

Governing Surface Treatment Knowledge for Dimensional Assessment in Precision Machining

Mohd Soufhwee Abd Rahman^{1*}, Muhamad Akbar Osman Bin Md Hashim², Hambali bin Boejang³,
Amir Hamzah bin Abdul Rasib⁴

^{1,3,4}Senior lecturer at Faculty of Industrial and Manufacturing Technology and Engineering, Universiti Teknikal Malaysia Melaka, Durian Tunggal, Melaka, Malaysia

²Master student in Faculty of Industrial and Manufacturing Technology and Engineering, Universiti Teknikal Malaysia Melaka, Durian Tunggal, Melaka, Malaysia

***Corresponding Author**

DOI: <https://doi.org/10.47772/IJRISS.2026.100900034>

Received: 12 September 2026; Accepted: 17 September 2026; Published: 28 September 2026

ABSTRACT

A treatment name found on a drawing is not, by itself, enough to justify using a dimensional parameter. This paper examines how surface-treatment terminology can remain searchable while analytical use is limited to records whose source and applicability are known. The industrial review covered 4,612 active production drawings; 185 contained treatment-related information, giving 186 recorded treatment occurrences. The terminology register retained 35 categories, whereas 11 records were taken forward for detailed analysis. The governance structure keeps source wording, controlled terminology and analytical records separate, so a parameter is considered only after identity, revision, meaning, applicability and lifecycle status have been resolved. The 11 selected records accounted for 118 of 186 occurrences (63.44%); an 11-record frequency-only selection could reach 120 occurrences (64.52%), a difference of two occurrences or 1.08 percentage points. The remaining 24 categories stay in a controlled expansion queue rather than being treated as analytically equivalent to the selected set. The paper also specifies prospective validation for independent review, inter-reviewer agreement and completeness, staged category qualification, external-site replication and longitudinal production outcomes; these are not reported as completed experiments. The contribution is a source-preserving method for admitting surface-treatment knowledge to dimensional assessment, together with an inspectable register comparison and explicit boundaries on curation reliability, generalisability and production-effect claims.

Keywords: knowledge-based engineering; surface treatment; knowledge governance; engineering provenance; applicability; precision machining

INTRODUCTION

In practice, surface treatment often appears on a machining drawing as a short note. The note may give a process name, a specification, or a class, yet none of those items alone tells an engineer what dimensional shift should be used. The usable meaning also depends on the material, treatment condition and feature being assessed. A repository can therefore retrieve the right terminology and still provide an incomplete basis for calculation. The engineering question is not simply whether the treatment can be recognised, but whether the retrieved information is sufficiently qualified for use.

This concern is familiar in knowledge-based engineering (KBE). Verhagen et al. (2012) identified transparency and knowledge traceability among the continuing challenges in KBE development, while Kügler et al. (2023) later highlighted weak links between knowledge sources and automation as a barrier to reuse. The issue becomes especially important when a numerical constant is embedded in a compensation routine. The value may remain in the system long after the specification, revision or conditions that originally justified it have become unclear.

Existing manufacturing research shows several ways of formalising engineering knowledge. Salaets et al. (2022), for example, express design rules through surface properties and bounds, whereas Zhang and Lobov (2024) use semantic representations to support engineering reasoning. The concern addressed here sits one step earlier. Before a treatment parameter is allowed into a feature-level calculation, the repository must establish what the record represents, where it came from and under which conditions it is valid.

Surface-treatment literature provides a practical reason for making that distinction. Electropolishing studies examine dimensional response under defined processing conditions (Beamud-González et al., 2023). Anodic-layer behaviour varies with substrate and process conditions (Razzouk et al., 2024), and work on electroless Ni-P coatings distinguishes phosphorus content and treatment history when characterising the coating (Genova et al., 2023). Such findings support keeping process and condition differences visible. They do not, however, justify copying a published value directly into a customer-specific machining instruction.

Against this background, the paper asks how an industrial treatment vocabulary can remain broad enough for retrieval while analytical use is restricted to records with a defined source and applicability, and what occurrence coverage is retained under that restriction. The case uses the repository documented by Hashim (2026) together with a retrospective fixed-capacity comparison. The framework combines preservation of original source wording, a governed admission decision before parameter use, and a direct comparison between the selected scope and a frequency-only alternative. It also describes how specification changes and category expansion can be managed without assuming that similar names or related specifications are interchangeable. The claims remain limited to knowledge governance and register coverage; curation reliability, source completeness, the remaining 24 categories, cross-site generalisability and production effect are prospective validation questions rather than results inferred from this dataset.

Case Study

Evidence units and industrial scope


The source dataset comprised 4,612 active production drawings from a Malaysian precision-machining operation. A drawing was treated as surface-treatment related when its governing information contained an explicit treatment call-out, a specification, or an associated note with possible dimensional relevance. On that basis, 185 drawings were identified. The register contains 186 treatment occurrences because one drawing may carry more than one treatment requirement.

A separate group of 26 drawings contained FN1, FN2 or FN3 functional-fit references. These drawings were not added to the treatment-related count because the documented register did not establish their exact intersection with the treatment set. The fit information was retained for engineering interpretation rather than treated as another coating or removal process. Keeping the three counting units separate avoids confusing drawing count, treatment occurrence count and functional-fit reference count.

All counts used in this paper come from the final thesis register and its documented control structure, particularly Section 3.4 and Appendix A of Hashim (2026). The calculations reported here regroup those existing counts; they are not the result of a second review of the original 4,612 drawings or a new annotation exercise. The original wording, record links and review decisions remain part of the research provenance established in the thesis.

Three layers of representation

The repository was organised in three layers: original source information, a terminology register and governed analytical records. The first layer preserves what the drawing or source actually says. The second gives stable identities and aliases so the terminology can be retrieved consistently. The third contains only those records for which authority, dimensional meaning and limitations have been established for the intended analytical use. Figure 1 summarises this separation.



Recognising a treatment name does not by itself authorise a parameter.

Figure 1: Separation of original source wording, terminology coverage and conditionally usable analytical records.

All 35 categories remained in the terminology register, but only 11 were selected for detailed DFM development. Inclusion in the initial 11-record analytical set depended on documented evidence readiness: a traceable source path, controlled identity, interpretable parameter basis and bounded applicability sufficient for deterministic research use; frequency alone was not an admission criterion. The historical selection was not performed as an independently scored dual-review experiment, so these factors describe the documented basis of selection rather than a quantified curation-accuracy result. The other categories were not deleted; they remained searchable in the register. Their presence in the vocabulary did not allow them to inherit the parameter or behaviour model of a neighbouring selected record.

METHODOLOGY

Identity preservation during normalisation

During normalisation, the source wording was kept alongside a controlled identity and any approved aliases. An alias serves retrieval only. It does not erase a class distinction or make two specifications equivalent. The record therefore keeps enough information to determine whether a retrieved category actually corresponds to the requirement invoked on the drawing, rather than merely sharing a similar name.

The resulting 35-category list should be read as an operational industrial register, not as a universal or mutually exclusive taxonomy of treatment physics. Some entries are branded processes, some refer to specification families, and others are broader treatment descriptions. Generic chemical-film or chromate-conversion wording, for instance, may be related to a named conversion-coating requirement, but the recorded occurrences should not be merged unless source-level evidence supports that decision. The reported counts therefore apply to the register at the level of normalisation used in the original study.

Source hierarchy and unresolved authority

For authority, the exact customer drawing and its applicable revision take priority, followed by the invoked specification, approved process evidence and controlled internal engineering evidence. That order is interpreted within the relevant engineering context. A newer generic document does not automatically rewrite an older requirement that is still invoked by the drawing. Where source values differ in meaning, unit, class or material condition, the difference has to be reconciled rather than averaged away.

When an in-scope authority question is unresolved, the record is held for Engineering Review. A condition that is explicitly incompatible or prohibited is marked Blocked, while a treatment or condition that lies outside the governed knowledge domain is marked Unsupported. These outcomes are useful because they preserve the reason automatic use was refused instead of reducing every failure to the same generic status. Table 1 lists the record fields used to maintain these distinctions during reuse.

Table 1: Knowledge fields and the conditions they preserve

Field group	Retained information	Consequence for reuse
Identity and aliases	Controlled treatment identity, original wording, class and approved aliases	Retrieval does not erase the invoked requirement.
Source and revision	Applicable drawing/specification, issue and source linkage	Earlier decisions remain associated with their original basis.
Parameter meaning	Quantity, source unit, behaviour model and contribution basis	A value is interpretable only with its physical and documentary context.
Applicability	Material, process, geometry, exposure and limitations	In-scope uncertainty is resolved before automatic use.
Lifecycle	Draft, In Review, Effective, Suspended or Superseded	Only an applicable Effective record may participate in authorised automatic use.
Decision provenance	Knowledge snapshot and limitations associated with the engineering decision	Later source changes do not silently alter the historical basis.

Applicability and lifecycle

Lifecycle status answers a different question from treatment mechanism. A record may be technically plausible and still remain Draft or In Review because its authority has not been completed. Even an Effective record is usable only inside its stated applicability. If a record is later Suspended or Superseded, that change affects future use without erasing the basis of an earlier released decision. The control approach therefore creates successor revisions instead of silently replacing the historical record.

Equation (1) condenses these checks into an admission predicate. I represents resolved identity, S the applicable source authority, V the retained source or knowledge revision, A the material and process domain, E an Effective lifecycle state, and P an adequate parameter basis for the intended use. The expression is a compact description of the governance conditions, not a claim that a new software interface was implemented for this paper. Passing the predicate only makes the knowledge eligible for consideration; characteristic type, masking, tolerance and geometry still have to be assessed separately.

In practical review, six questions have to be answered: what source wording and class the record represents; which document authorises it; which revision is being retained; what material and process domain it covers; whether the lifecycle state permits use; and what quantity, unit and evidence support the parameter. A familiar treatment name cannot compensate for a blank parameter or unresolved class. Model D also requires a positive, bounded rationale for a no-significant-shift decision. These checks define the admission contract; they are not evidence that every field in every record was independently audited for completeness.

$$\text{Knowledge admission} = I \wedge S \wedge V \wedge A \wedge E \wedge P \quad (1)$$

The governed no-significant-shift category (Model D; see Table 3) deserves separate attention because a no-significant-shift decision can easily be misunderstood. It represents a positive engineering disposition for a bounded application, not a default value for missing information. Likewise, a recorded contribution factor of zero does not prove that no physical layer exists. The repository keeps a deliberate distinction between an authorised decision that the shift is not significant for the stated model and a situation in which the treatment response is simply unknown.

Research parameters and specification identity

VTR 01 is the frozen internal research baseline used for deterministic development and reference calculations. Its treatment magnitudes are research values; they are not universal constants inferred from the names of external standards. Specification identity and numerical parameter provenance are therefore stored separately. Any later production value still has to be justified by the applicable drawing, material and treatment evidence.

Electroless nickel makes this separation particularly clear. ASTM B733 describes coating classifications and properties associated with phosphorus content and subsequent heat treatment (ASTM International, 2022), while SAE AMS2404 and the AMSC26074 document family are separately identifiable specifications (SAE International, 2013, 2024). The thesis therefore keeps legacy MIL-C-26074 Class 3 and Class 4 as distinct identities and does not interpret the class numbers as thickness values. A current publication or later specification edition also does not retrospectively replace the authority that was invoked by an earlier drawing.

Descriptive analysis of the documented register

The register analysis was deliberately limited to arithmetic and structural checks. For each of the 35 rows, the recorded occurrence count and the selected or outside-scope status were retained, then the category total, occurrence total, selected-record count and status-based sums were recalculated. No extraction-accuracy score, inter-reviewer coefficient or quantitative completeness score was estimated because the thesis does not contain the independent labels needed to support those measures.

$$\text{Selected occurrence share} = 118 / 186 \times 100 = 63.44\% \quad (2)$$

Equation (2) is only a register-coverage measure. It should not be read as the proportion of production drawings or characteristics automatically authorised for compensation. A selected category can still be inapplicable to a particular feature. Equally, a category outside the initial analytical set is not necessarily rare, unimportant or technically unsuitable; it simply falls outside the bounded development scope used here. Appendix A reproduces the full count extract so that the aggregation can be checked directly.

For the frequency-only comparator in Equation (3), the same capacity of 11 records was retained. The 35 occurrence counts were ranked from highest to lowest and the first 11 were summed. If every row is assumed to have the same one-record capacity cost and no source or applicability constraint is imposed, this produces the largest occurrence total obtainable from 11 rows of the register. Ties may yield different category combinations with the same total. The comparator is therefore a descriptive benchmark only; it is neither an approved alternative repository nor a test of physical treatment performance.

$$\text{Maximum frequency-only share} = 120 / 186 \times 100 = 64.52\% \quad (3)$$

Independent curation validation boundary

A limitation of the original drawing review is that it was not designed as a blinded dual-annotator study. The register therefore does not contain two independent pre-adjudication label sets from which Cohen's kappa, or another inter-reviewer reliability statistic, could validly be calculated. It also does not contain an independently measured category-level completeness score for the admission fields. For those reasons, neither a retrospective agreement value nor a retrospective completeness metric is reported as if it had been independently observed.

A follow-on curation study should address both reliability and completeness. Two reviewers should independently classify a stratified sample before adjudication, scoring treatment-related status, controlled identity and class, source or specification linkage, and analytical-scope eligibility separately. The six admission fields used here - identity and class, source authority, retained revision, applicability domain, lifecycle state and parameter basis - should also be audited for presence, traceability and internal consistency. Raw agreement and Cohen's kappa should be calculated before disagreements are resolved against governing evidence, and completeness should be reported at record and field level. These measures test curation consistency and documentation completeness, not physical parameter validity.

The current dataset likewise cannot support a causal estimate of downstream manufacturing benefit because no longitudinal before-and-after production trial was part of the original study. A valid production-effect study should freeze the knowledge version, define matched part and treatment strata together with baseline and intervention periods, and prespecify endpoints such as absolute post-treatment dimensional prediction error, first-pass acceptance, treatment-related nonconformance or scrap, and Engineering Review workload. No such effect size is estimated retrospectively in this paper.

RESULTS AND DISCUSSION

Coverage of the selected analytical set

Table 2: Recalculated coverage of the documented industrial register

Measure	Count	Interpretation
Production drawings reviewed	4,612	Underlying industrial review population
Treatment-related drawings	185	4.01% of the reviewed drawing population
Treatment occurrences	186	Counting unit for the category register
Selected analytical categories	11 of 35	31.43% of the category register
Occurrences in selected categories	118 of 186	63.44% of recorded occurrences
Categories outside initial analytical scope	24 of 35	68.57% of categories retained for traceability
Occurrences outside initial analytical scope	68 of 186	36.56% of recorded occurrences
Frequency-only maximum at 11 records	120 of 186	64.52%; ignores evidence readiness and applicability
Difference from the recorded selection	2 occurrences	1.08 percentage points of the 186-occurrence register

The main coverage result is straightforward. The 11 selected categories account for 118 of the 186 recorded occurrences, leaving 68 occurrences across the other 24 categories. Choosing 11 categories by frequency alone would raise the total only to 120 occurrences, or 64.52%, compared with 63.44% for the recorded selection. The gap is two occurrences, equivalent to 1.08 percentage points on the common 186-occurrence denominator shown in Figure 2. This comparison concerns descriptive coverage of the register and should not be interpreted as the share of production characteristics authorised for compensation.

Equal capacity: 11 records from the 35-category register




Figure 2: Occurrence coverage at an equal eleven-record capacity. Both bars use the 186-occurrence denominator; the frequency-only comparator ignores knowledge authority.

The 11 selected records were not simply the 11 most frequent entries. Eight categories with at least seven occurrences contribute 102 occurrences. Four more categories each occur six times: Alodine Class 3, ASTM B733 nickel, High Phosphorus NiP and Black Oxide/Blackening. Any three of those four would bring a purely frequency-based 11-record set to 120 occurrences. The documented selection contains two of the six-occurrence categories and Passivation, which occurs four times. Replacing Passivation with either excluded six-occurrence category would add two occurrences, but frequency by itself gives no basis for authorising that replacement. The result is consistent with evidence readiness having influenced selection, although it does not independently prove that the curation decision was correct.

Table 3: Selected treatment records, dimensional models and occurrence counts

Controlled identity	Model	Count
Clear Anodize Type II	A Conversion/build-up	25
Black Anodize Type II	A Conversion/build-up	18
Hard Anodize Type III	A Conversion/build-up	14
Electropolishing	B Removal	12
Electroless Nickel MIL-C-26074 Class 3	C Deposition	10
Electroless Nickel MIL-C-26074 Class 4	C Deposition	8
Alodine Class 1A	D Governed no-significant-shift	8
Electroless Nickel SAE AMS2404	C Deposition	7
Electroless Nickel ASTM B733	C Deposition	6
Electroless Nickel High Phosphorus NiP	C Deposition	6
Passivation	D Governed no-significant-shift	4

Grouping the selected records by mechanism gives 57 occurrences across the three Model A identities, 12 for the single Model B identity, 37 across the five Model C identities and 12 across the two Model D identities. These totals describe only the selected analytical records. The outside-scope categories were not retrospectively assigned to a mechanism for this calculation, as doing so would extend the documented scope. Nor does membership in the same model make the underlying specifications interchangeable.

Cases that explain the governance requirement

Alodine provides a simple example of why the distinction matters. Class 1A is one of the selected analytical records, whereas Class 3 remains a separate terminology entry outside that detailed scope. Treating the common process name as sufficient evidence of equivalence would remove the class boundary retained in the source. The repository can recognise both terms without granting them the same analytical authority.

Electroless nickel gives a second example. The selected MIL-C-26074 classes, SAE AMS2404, ASTM B733 and High Phosphorus NiP records all use a deposition model, yet their documentary basis and material or process conditions are not identical. Published work on Ni-P coatings reports differences associated with chemistry and processing (Genova et al., 2023), which is consistent with keeping those identities separate rather than averaging apparently related thickness information.

Lifecycle response to a specification update

Consider a released drawing and an existing decision that invoke MIL-C-26074 Class 3. If an engineering change later proposes AMS2404 for a future requirement, the repository does not treat the two specification identities as synonyms and does not copy the earlier parameter automatically. The released MIL-C-26074 record and its decision snapshot remain traceable. AMS2404 is handled as a separate candidate record in Draft and is reviewed against the exact revision, material and process domain, and parameter basis. It can become Effective for future use only when the drawing or an approved engineering change explicitly invokes AMS2404 and the full admission conditions have been met.

If that proposed change is not yet approved, or if its applicability is still uncertain, the candidate remains In Review or Engineering Review while the earlier released basis is left intact. An explicitly incompatible requirement can be marked Blocked. In this way, document currency and applicability are handled separately, and a newer publication cannot silently replace an older drawing-invoked authority. The example describes the lifecycle logic only; it does not claim that MIL-C-26074 and AMS2404 are technically equivalent.

Scaling path beyond the initial eleven records

The remaining 24 categories are still present in the terminology register, so they form a controlled expansion queue rather than discarded knowledge. They represent 68.57% of the categories and 36.56% of the recorded treatment occurrences and remain outside the initial analytical scope; the paper does not relabel them as modelled or approved. Among the higher-occurrence outside-scope entries are Alodine Class 3 and Black Oxide/Blackening, with six occurrences each; Chrome Based Coating, DLC Coating and Everslik Coating, with five each; and Chem Film per MIL DTL 5541 and Medium Phosphorus NiP, with four each. Occurrence count can help decide where review effort is spent, but it has to be considered together with engineering consequence and evidence readiness. It cannot authorise a record on its own.

Admission of additional categories should retain the original source, identity, revision, applicability, parameter-basis and lifecycle checks, while adding independent curation review and a category-level completeness audit as strengthened controls for future expansion. Each candidate should therefore pass source and identity adjudication, revision confirmation, applicability definition, parameter-basis qualification, independent review and completeness checking, and lifecycle approval before analytical activation. A lower-frequency category may reasonably be addressed before a higher-frequency one if its engineering consequence is greater or its evidence package is more mature. Until qualification is complete, the remaining categories stay retrievable but analytically inactive. This provides a controlled route beyond the initial 11 records without weakening the distinction between recognition and authorised analytical use.

Table 4: Documented distinctions and the inference they prevent

Case	Retained distinction	Inference that is not authorised
Alodine classes	Class 1A and Class 3 remain separate	A similar name grants the same analytical authority.
Electroless nickel records	Specification, class and phosphorus condition remain identifiable	All deposition records share one interchangeable parameter.
Model D	Positive bounded no-significant-shift disposition	Unknown treatment or missing data means zero effect.
Source revision	Invoked authority and research freeze have separate identities	A current publication silently replaces an earlier requirement.
Outside-scope terminology	Recognised categories remain retrievable	Recognition alone permits automatic calculation.

Table 4 is intended as a documentary check on the method rather than a new physical experiment or an independently scored validation set. Its purpose is to show where an ordinary name-based lookup could lose information that matters to a later calculation. By keeping those distinctions visible, the repository exposes the point at which knowledge is not yet safe to use instead of allowing the ambiguity to pass unnoticed into the arithmetic.

Contribution and transferability

The practical contribution lies in separating a broad industrial vocabulary from the smaller set of records that are actually eligible for analytical use. Provenance, applicability and lifecycle control are not new ideas by themselves. Here they are brought together as an admission contract for treatment-sensitive dimensional assessment and examined alongside a fixed-capacity coverage comparison. In the documented register, the evidence-readiness selection retains 63.44% of occurrences, compared with a 64.52% frequency-only ceiling at the same capacity of 11 records. That comparison describes scope; it does not measure curation quality or prove the physical accuracy of any parameter.

This control sits upstream of the reasoning method used later. A rule engine may apply its logic consistently and still produce a poor engineering decision if the input value came from the wrong source or outside its valid domain. Keeping the provenance and eligibility of that value visible is therefore useful whether the subsequent implementation uses procedural code, constraints or a semantic representation. No comparison between those implementation technologies is attempted here, and no claim of superiority is made.

The frequency comparator should also be interpreted cautiously. It holds record capacity constant and changes only the selection rule. Under the category definitions used in this register, the recorded selection is two occurrences below the frequency-only ceiling. That difference could change if the categories were normalised differently, if some records required substantially more curation effort, or if low-frequency requirements carried higher engineering risk. Frequency can guide planning, but it is not a sufficient basis for release authority.

Transferability and numerical generalisability are therefore different issues. Another manufacturing site could adopt the sequence of source preservation, terminology normalisation, applicability review, lifecycle control and coverage accounting, but it would still need to build its own register from the drawings and specifications actually invoked there. A proper replication should define the receiving site's drawing population, apply the same screening and admission rules, use independent review for category mapping, report field-level completeness and inter-reviewer agreement, and compare occurrence coverage and governance outcomes with this case. Only after such replication could cross-site robustness be assessed. The present 35-category distribution, the 11 selected identities and the research values should not be assumed to carry over unchanged.

Limitations

The evidence comes from one industrial population. Its category frequencies, the selected 11-record set and the numerical research values are therefore specific to that setting and cannot be taken as the distribution of surface treatments across precision machining generally. What is potentially transferable is the governance method, not the frequency profile or numerical values observed at this site. External replication using the same screening, independent review and admission controls is required before cross-site robustness can be claimed.

The original review also did not include a separately reported blinded dual-annotation exercise, an inter-reviewer agreement statistic or a category-level audit of source completeness. Without independent pre-adjudication labels, a retrospective Cohen's kappa would not be valid, and without a pre-defined completeness audit the paper cannot manufacture a quantitative completeness score after the fact. The prospective validation protocol defined in the Methodology specifies how agreement and record- and field-level completeness should be measured in a follow-on study, but no numerical value is claimed for data that were not collected.

A further limitation is analytical scope. Twenty-four of the 35 categories, representing 68 of 186 occurrences, remain outside the initial analytical set. They are retrievable but are not authorised for analytical use until the same source, applicability, parameter and lifecycle requirements are satisfied. The staged expansion protocol

provides a controlled path for qualifying them, but the present analytical scope remains 11 records until that work is completed.

Finally, the study did not include a longitudinal production trial comparing compensation error, first-pass acceptance, treatment-related nonconformance or scrap before and after use of the admission method. No downstream reduction in dimensional error, nonconformance or scrap is therefore claimed. The research values have also not been physically qualified across all suppliers, materials or treatment lots, so a record may be adequate for deterministic software verification while still lacking the evidence needed for production authority.

CONCLUSION

Recognition of a treatment name and authority to use its parameter are different decisions. In the documented register, 35 categories account for 186 occurrences; the 11 selected analytical records account for 118, while an 11-record frequency-only selection could reach 120. The method preserves source identity, revision, parameter meaning, applicability and lifecycle as admission conditions. The specification-change example preserves the earlier decision basis, and the staged expansion protocol provides a controlled route for qualifying the other 24 categories. The present evidence supports the governance structure and descriptive register comparison only; it does not establish inter-reviewer reliability, quantified completeness, cross-site generalisability, qualification of the remaining categories, or a causal reduction in dimensional error or scrap. Those outcomes require the prospective validation defined above.

Future Research

Further work should first establish curation reliability and completeness through independent dual review. Treatment identity, class, source linkage and analytical eligibility should be scored before adjudication so raw agreement and Cohen's kappa can be reported, while the six admission fields are audited at record and field level for completeness.

The remaining 24 categories should then be qualified through the staged admission protocol. Replication at another site should repeat screening, mapping, independent review and completeness auditing on that site's own drawings and specifications, and report its population, coverage, agreement, completeness and proportion of candidates reaching Effective status. Numerical treatment values should not be transferred by default.

A separate longitudinal study should test production effect using a frozen knowledge version, matched part and treatment strata, defined baseline and intervention periods, and prespecified measures such as post-treatment dimensional prediction error, first-pass acceptance, treatment-related nonconformance or scrap, and Engineering Review workload. Physical trials remain necessary before any research parameter is promoted to production authority.

ACKNOWLEDGEMENT

The authors acknowledge Saifudin Hafiz Bin Yahaya for supervising the underlying master's research, Universiti Teknikal Malaysia Melaka for academic support, and the participating precision-machining organisation for access to the industrial information used to construct the treatment register.

Data and Research Provenance

The underlying source is Hashim (2026), particularly Section 3.4 and Appendix A. Appendix A of this paper reproduces the recorded category counts and development status. The values 118/186 and 11/35, the 11-record frequency-only maximum of 120/186, and the resulting two-occurrence difference were calculated retrospectively from that register. The frequency comparator was added for this article; it does not represent a new drawing survey or a separate experiment from the original thesis. Confidential customer drawings are not reproduced.

APPENDIX A

Table 5 reproduces the 35 categories recorded in the industrial register. "Selected" means that the category was included in the 11-record research analytical scope. "Outside scope" means that the category remained in the terminology register without that analytical designation. Neither status should be interpreted as unrestricted production approval. Source: Hashim (2026), Appendix A, Table A.2.

Table 5: Complete count extract used for the descriptive analysis

Category	Count	Research status
Clear Anodize Type II	25	Selected
Black Anodize Type II	18	Selected
Hard Anodize Type III	14	Selected
Electropolishing	12	Selected
Electroless Nickel Plating per MIL C 26074 Class 3	10	Selected
Electroless Nickel Plating per MIL C 26074 Class 4	8	Selected
Alodine Class 1A	8	Selected
Electroless Nickel Plating per SAE AMS2404	7	Selected
Alodine Class 3	6	Outside scope
Electroless Nickel Plating per ASTM B733	6	Selected
Electroless Nickel Plating (High Phosphorus NiP)	6	Selected
Black Oxide / Blackening	6	Outside scope
Chrome Based Coating	5	Outside scope
DLC Coating	5	Outside scope
Chem Film per MIL DTL 5541	4	Outside scope
Electroless Nickel Plating (Medium Phosphorus NiP)	4	Outside scope
Everslik Coating	5	Outside scope
Passivation	4	Selected
Titanium Nitride (TiN)	3	Outside scope
Cadmium Plating	3	Outside scope
ASM White Coating	3	Outside scope
Dry Film Lubricant	3	Outside scope
Annealing	3	Outside scope

Red Anodize Type II	2	Outside scope
Chromate Conversion Coating	2	Outside scope
Nitriding	2	Outside scope
Zinc Plating + Chromate	2	Outside scope
Aluminium Titanium Nitride (AlTiN)	2	Outside scope
Chemical Etching	2	Outside scope
Electroless Nickel Plating + PTFE (Nickelflon)	1	Outside scope
Nickel Plating	1	Outside scope
Hybrid Coating (Ni PTFE)	1	Outside scope
Vinyl Wrap	1	Outside scope
Precious Metal Plating (Gold)	1	Outside scope
Teflon Coating	1	Outside scope

REFERENCES

1. ASTM International. (2022). ASTM B733-22: Standard specification for autocatalytic (electroless) nickel-phosphorus coatings on metal. ASTM International. <https://store.astm.org/b0733-22.html>
2. Beamud-González, E. M., Núñez-López, P. J., & García-Plaza, E. (2023). Electropolishing stainless steel optimization using surface quality, dimensional accuracy, and electrical consumption criteria. *Materials*, 16(5), 1770. <https://doi.org/10.3390/ma16051770>
3. Genova, V., Paglia, L., Pulci, G., Pedrizzetti, G., Pranzetti, A., Romanelli, M., & Marra, F. (2023). Medium and high phosphorous Ni-P coatings obtained via an electroless approach: Optimization of solution formulation and characterization of coatings. *Coatings*, 13(9), 1490. <https://doi.org/10.3390/coatings13091490>
4. Hashim, M. A. O. B. M. (2026). Development and validation of a knowledge based design for manufacturability (DFM) compensation engine for precision machining applications. Master of Manufacturing Engineering (Quality System Engineering) thesis, Universiti Teknikal Malaysia Melaka. Final submission dated 13 September 2026.
5. Kügler, P., Dworschak, F., Schleich, B., & Wartzack, S. (2023). The evolution of knowledge-based engineering from a design research perspective: Literature review 2012–2021. *Advanced Engineering Informatics*, 55, 101892. <https://doi.org/10.1016/j.aei.2023.101892>
6. Razzouk, E., Koncz-Horváth, D., & Török, T. I. (2024). Critical challenges in the anodizing process of aluminium–silicon cast alloys—A review. *Crystals*, 14(7), 617. <https://doi.org/10.3390/cryst14070617>
7. SAE International. (2013). AMSC26074D: Electroless nickel coatings. Stabilized June 2013. SAE International. <https://doi.org/10.4271/AMSC26074D>
8. SAE International. (2024). AMS2404K: Plating, electroless nickel-phosphorous. SAE International. <https://doi.org/10.4271/AMS2404K>
9. Salaets, R., Decraemer, B., Eyckens, P., Boudewyns, W., Van Houdt, W., & Beyers, K. (2022). A flexible approach for design rule formalization and evaluation. *Procedia CIRP*, 109, 556–561. <https://doi.org/10.1016/j.procir.2022.05.294>
10. Verhagen, W. J. C., Bermell-Garcia, P., van Dijk, R. E. C., & Curran, R. (2012). A critical review of Knowledge-Based Engineering: An identification of research challenges. *Advanced Engineering Informatics*, 26(1), 5–15. <https://doi.org/10.1016/j.aei.2011.06.004>

-
11. Zhang, L., & Lobov, A. (2024). Semantic Web Rule Language-based approach for implementing Knowledge-Based Engineering systems. *Advanced Engineering Informatics*, 62, 102587. <https://doi.org/10.1016/j.aei.2024.102587>