

Ontologies for Photopolymer Materials, Processes and Qualification Evidence: Terms, Relations and Interoperability

Turning material, process and evidence descriptors into a shared semantic layer without confusing vocabulary control with scientific truth

Juan Segurola

3Dresyns by Resyner Technologies S.L.

© 2026 Juan Segurola. All Rights Reserved.

Correspondence: info@3Dresyns.com · www.3Dresyns.com

ER-391 · Engineering Review · Version 0.1 · 5 October 2026

Abstract

Photopolymer engineering data are difficult to combine when the same word denotes different states, when identical quantities are reported under incompatible conditions, or when relationships are hidden in prose. This review defines a practical ontology architecture for photopolymer materials, vat-photopolymerisation processes and qualification evidence. It separates identifiers, classes, relations, quantities and provenance; distinguishes controlled vocabulary from formal semantics; and shows how ontology terms can connect formulation state, printer state, specimen state, test method and engineering claim. The objective is not to create a universal chemistry ontology, but to specify the minimum semantic commitments needed for interoperable engineering evidence.

Keywords: ontology; photopolymer; additive manufacturing; interoperability; provenance; semantic data; qualification

1. Ontology begins where labels stop being sufficient

A controlled term list can prevent spelling variants, but it does not state whether a resin batch is an instance of a formulation, whether a post-cure is an activity, or whether a tensile result was generated from a particular specimen state. Gruber defined ontology as an explicit specification of a conceptualisation, which is useful here because the engineering problem is relational rather than lexical [1]. For photopolymer qualification, the semantic layer must therefore encode what an entity is, how it changed and which evidence object supports which claim.

2. Three semantic layers should remain separate

A robust implementation separates domain classes, measurement semantics and provenance. Domain classes describe material, machine, process step, geometry, specimen and environment. Measurement semantics bind a value to a quantity kind, unit and method; unit ontologies such as QUDT illustrate why a number without a quantity type is not interoperable [2]. Provenance records which activity generated or transformed an entity, using relations equivalent to wasGeneratedBy, used and wasDerivedFrom in PROV-O [3]. Mixing these layers produces brittle schemas because every new test method forces changes to chemistry or process classes.

Table 1. Minimal semantic blocks for photopolymer evidence.

Block	Core entities	Required relation
material	formulation, batch, constituent	batch instantiates formulation; constituent belongs to formulation
process	exposure, wash, post-cure, ageing	activity transforms one material state into another
specimen	geometry, orientation, region	specimen derived from state and build context
measurement	quantity, value, unit, method	measurement generated from specimen under method
evidence	source, claim, version	source supports, qualifies or contradicts claim

3. State is a first-class entity

Photopolymer evidence is strongly state dependent. Liquid formulation, exposed green part, washed part, thermally post-cured part and aged specimen are not interchangeable representations of one object. Each state transition should be represented as an activity with inputs, parameters, time and outputs. The resulting graph makes it possible to ask whether two modulus values refer to specimens with comparable cure history rather than merely comparable material names.

4. Relations carry more engineering meaning than class depth

Deep taxonomies are tempting but often fragile. A shallower class system with precise relations is more useful: hasConstituent, processedBy, printedOn, hasExposureCondition, derivedFromBatch, testedBy, hasOrientation, supportsClaim and contradictsClaim. Ontology-design practice in other scientific domains emphasises coordinated identifiers and reusable relations rather than isolated local hierarchies [4]. The same principle reduces duplication between chemistry, manufacturing and evidence vocabularies.

Table 2. Interoperability failures that an ontology should expose.

Failure	Example	Semantic remedy
homonym	cure used for exposure and post-cure	separate process classes and preferred labels
state collapse	green and post-cured modulus merged	state-transition entities
unit ambiguity	dose reported without area basis	quantity kind plus unit ontology
identity ambiguity	commercial name treated as fixed composition	versioned material identifier
provenance loss	claim detached from source version	PROV-style derivation links

5. Interoperability requires explicit boundary rules

A photopolymer ontology should not silently infer chemical identity from trade names, infer cure conversion from exposure time, or equate nominal printer settings with delivered optical dose. Such mappings require evidence objects. FAIR principles motivate machine-actionable identifiers and metadata [5], but FAIRness does not guarantee semantic correctness. Interoperability is achieved when systems can exchange meaning while preserving uncertainty and state boundaries.

6. Qualification use case

For qualification, the ontology is successful when a query can retrieve all evidence supporting a service claim and expose the state mismatch that would invalidate transfer. The semantic graph should answer: which specimen states support the claim, which measurements were made under which methods, what transformations occurred before measurement, and which source or version supplied each assertion. This is sufficient for evidence navigation without requiring the ontology to decide whether the claim is true.

Table 3. Qualification queries that should be machine answerable.

Query	Minimum graph path
Which data support the claim?	claim -> evidence -> measurement -> specimen
Are specimens comparable?	specimen -> state -> process history
Which values share a quantity definition?	measurement -> quantity kind -> unit
What changed after a source update?	claim -> source version -> derivation
Where is uncertainty explicit?	measurement or claim -> uncertainty descriptor

7. Scope boundary and neighbouring reviews

ER-391 owns semantic identity, relations and interoperability. ER-392 owns the minimum information that must be reported for an additive-manufacturing study, while ER-394 owns source-to-claim provenance and version dependency. ER-391 therefore defines meaning and relationship semantics; it does not decide reporting completeness or whether a claim is sufficiently supported.

8. Identity, namespaces and version governance

Interoperability fails when identifiers are treated as decoration rather than engineering controls. A formulation family, a commercial grade, a batch, a specimen and a test result require different identifiers because they have different lifecycles. The ontology should therefore distinguish persistent concept identifiers from version identifiers and physical-instance identifiers. A concept such as a resin grade can persist while a formulation revision changes; a batch is an instance of one revision; a printed specimen is derived from that batch through a process history; and a test result is generated from the specimen by a measurement activity. PROV-O supplies a general provenance vocabulary for these relations, while QUDT supplies quantity and unit semantics [2,3]. Materials-specific ontologies such as the Materials Design Ontology show how modular ontology design can support integration across heterogeneous materials databases rather than forcing all concepts into one monolithic class tree [6].

9. Competency questions are the practical validation test

An ontology is not validated by the number of classes it contains. It is validated by whether the intended engineering questions can be answered without manual reinterpretation. For photopolymer qualification, representative competency questions include: Which modulus results were generated from specimens printed from the same formulation revision but on different machines? Which exposure conditions produced a given green-state geometry? Which values were measured after a defined post-cure and conditioning history? Which claim depends on a source that has been superseded? Which results use the same quantity kind but incompatible methods? If these questions require free-text reading, the semantic commitments are incomplete. Each competency question should therefore be converted into expected graph patterns and tested against example records before the ontology is released.

10. Constraint validation and semantic error classes

Formal semantics do not automatically guarantee valid engineering data. A graph can be syntactically legal yet physically incoherent. Constraint checks are therefore needed above the ontology layer. Examples include a tensile-strength value with a temperature unit, a post-cure activity that precedes printing in the provenance chain, a specimen linked simultaneously to two incompatible material revisions, or a claim that cites a result without the specimen state from which it was derived. These errors are better treated as validation failures than as ontology redesign problems. A practical system separates four layers: class and relation semantics, quantity/unit validation, provenance-order validation and domain rules. This separation allows each layer to evolve without destabilising the others.

11. Crosswalks, deprecation and semantic change

Long-lived engineering corpora need controlled change. Terms will be split, merged or deprecated as evidence and practice evolve. Every such change should retain the old identifier, state the replacement relation and declare whether the mapping is exact, broader, narrower or context-dependent. Silent renaming is particularly dangerous because historical records may appear current while their meaning has

shifted. The OBO Foundry experience demonstrates the value of shared principles for coordinated ontology evolution [4], while FAIR principles make persistence and machine accessibility part of reuse [5]. For ER-series use, a semantic release therefore requires a versioned term set, change log, validation tests and crosswalks sufficient to re-evaluate older evidence without rewriting the source record.

12. Worked photopolymer evidence graph

Consider a tensile result from a UV-post-cured VPP specimen. The graph should not store only material name, value and unit. It should connect the physical specimen to its build, the build to printer and parameter set, the parameter set to exposure wavelength and layer thickness, the specimen to washing and post-cure activities, the result to the test method and environment, and the engineering claim to the result through an explicit support relation. If another result differs, the graph allows the conflict to be stratified by process or specimen state before a scientific contradiction is asserted. This worked structure is the minimum level at which ontology becomes an engineering tool rather than a vocabulary exercise.

13. Mapping external vocabularies without semantic collapse

Interoperability often requires mapping a local photopolymer vocabulary to external schemas. The mapping should preserve uncertainty about equivalence. An exactMatch relation is appropriate only when the two terms have the same extension and intended use. A broader or narrower mapping is safer when one vocabulary collapses states that another distinguishes. For example, a generic 'curing' concept may cover thermal, photo and dual-cure operations, while an engineering graph may need exposure wavelength, atmosphere and post-cure temperature as distinct activities. The crosswalk should therefore be versioned and tested by competency questions rather than accepted because labels look similar. This is especially important when importing datasets whose original semantics cannot be changed.

14. Machine-actionable constraints and SHACL-like validation logic

A production ontology benefits from explicit constraints that can be evaluated automatically. Typical rules include: every measured value must have a quantity kind and unit; every specimen must derive from one declared material revision; every measurement must identify the specimen state and method; every process activity must have a temporal position consistent with the lineage; and every qualification claim must resolve to at least one evidence object. The exact constraint language can vary, but the validation report should return an actionable error class rather than a generic schema failure. Constraint violations are then reviewed as data-quality findings, not silently coerced into a valid graph.

15. Interoperability test suite

A semantic release should ship with test records representing ordinary, boundary and deliberately invalid cases. Ordinary cases verify routine queries. Boundary cases test multimaterial parts, reprocessed specimens, repeated post-cure, superseded sources and measurements with mixed unit systems. Invalid cases verify that the validator catches impossible state order, missing quantity semantics and ambiguous revision identity. The test suite should be rerun after every ontology or crosswalk change. This makes semantic interoperability reproducible in the same way that unit tests make software behaviour reproducible. The ontology becomes infrastructure only when a change can be shown not to break previously valid engineering queries.

16. Governance roles and semantic change control

An engineering ontology needs ownership rules comparable to those of controlled technical documentation. A term proposal should identify the problem it solves, examples that require it and whether an existing concept can be reused. A semantic change is then classified as additive, clarifying, deprecating or breaking. Additive changes can often preserve existing queries; a changed definition or relation cardinality may invalidate previous inferences and therefore requires a new ontology version plus migration notes. The release record should name the affected classes, properties, constraints and competency questions. This avoids a common failure mode in which the URI remains stable while its intended meaning drifts. In qualification evidence, semantic drift is equivalent to changing a test definition without updating the report because the same graph can then yield a different engineering interpretation after an ontology update. [1,3,4,6]

17. Worked interoperability failure modes

Several failure modes are especially relevant to photopolymer evidence. Unit-compatible values can still be semantically incompatible when one describes liquid resin and another the cured solid. Two 'post-cure' activities can differ in wavelength, irradiance, temperature and duration. A specimen labelled 'aged' may represent thermal, humid, radiative or chemical exposure. A generic material identifier may hide a formulation revision. A qualification claim may cite a source but omit the exact version or page supporting it. These are not merely metadata defects; they can cause invalid joins, duplicate entities or false comparisons. A robust ontology should make the incompatibility representable and allow validators to stop an automated merge rather than silently normalise the records.

18. Release gate for an ontology used in qualification

Before an ontology version is accepted for production use, ER-391 recommends a release gate with four independent checks: lexical review of terms and definitions; logical consistency and constraint validation; execution of the competency-question suite; and migration of a representative historical dataset. The migration test is critical because it exposes changes that look harmless on an empty schema but break real records. Release should be blocked if valid prior records become uninterpretable without an explicit migration path, or if invalid examples cease to be detected. The output is a versioned semantic contract that downstream ERs can cite with confidence rather than a vocabulary that changes invisibly over time.

19. Limitations and research priorities

Ontology quality is difficult to infer from class count or logical consistency alone. The main evidence gap is whether the semantic model improves real engineering retrieval, comparison and change control across independent datasets. Priorities are therefore shared competency-question suites, inter-ontology crosswalk challenges, migration tests on historical photopolymer records and user studies that compare graph-assisted review with conventional document search. Another priority is version-aware treatment of confidential composition and proprietary process information, because industrial interoperability often fails where disclosure boundaries are ignored. ER-391 therefore treats ontology release as an empirical interoperability claim: the schema is only successful if independent systems can

exchange records without losing the material state, process lineage, measurement meaning or qualification boundary encoded by the producer. [1-6]

References

1. Gruber TR. A translation approach to portable ontology specifications. *Knowledge Acquisition*. 1993;5(2):199-220. doi:10.1006/knac.1993.1008.
2. QUDT.org. Quantities, Units, Dimensions and Types (QUDT) schema and vocabularies. Living standard. <https://www.qudt.org/catalog/qudt-catalog.html>. Accessed 4 October 2026.
3. Lebo T, Sahoo S, McGuinness D, et al. PROV-O: The PROV Ontology. W3C Recommendation; 30 April 2013. <https://www.w3.org/TR/prov-o/>.
4. Smith B, Ashburner M, Rosse C, et al. The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology*. 2007;25:1251-1255. doi:10.1038/nbt1346.
5. Wilkinson MD, Dumontier M, Aalbersberg IJJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*. 2016;3:160018. doi:10.1038/sdata.2016.18.
6. Aameri B, Poveda-Villalon M, Sanfilippo EM, Terkaj W, Lambrix P, Armiento R, et al. The Materials Design Ontology. *Semantic Web*. 2024;15(2):481-515. doi:10.3233/SW-233340.

Competing interests

The author is Founder, CEO and CTO of 3Dresyns by Resyner Technologies S.L.; ER-391 is a methodological engineering review and does not provide a comparative commercial recommendation.

Data availability

ER-391 is a literature-based engineering review; no new experimental dataset was generated.