

Diagnostic Inference in Vat Photopolymerisation: Discriminating Between Competing Defect Mechanisms Under Uncertainty

Juan Segurola

3Dresyns by Resyner Technologies S.L.

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

ER-184 | *Engineering Review* | Version 0.4 | 11 September 2026

Statement of Significance

Visible defects in vat photopolymerisation are rarely unique signatures of a single mechanism. Curl, separation, weak interfaces, dimensional drift, incomplete layers and surface defects can arise from different combinations of optical, chemical, rheological, interfacial, geometric and post-processing states. Treating an observation as its own root cause creates false certainty and can make a corrective action appear successful without establishing why the failure occurred. ER-184 defines a bounded evidence architecture for diagnostic inference. It separates observation, classification, localisation, mechanism attribution, root-cause attribution and corrective-action verification, and it makes an unresolved result an acceptable scientific outcome.

Keywords: vat photopolymerisation; diagnostic inference; root-cause analysis; failure analysis; uncertainty; discriminating measurement; VPP qualification

Abstract

ER-184 proposes a diagnostic logic for vat photopolymerisation and has not been validated as a diagnostic system. The framework starts from an observable signature and the state of the specimen and process, constructs a set of competing mechanism hypotheses, derives predictions that differ among those hypotheses, selects measurements capable of discriminating among them, and updates the candidate set only to the extent justified by the evidence. Controlled Engineering Reviews already show why this separation is necessary: cure shrinkage is not equivalent to warpage [1]; bottom-up separation depends on interface, resin, exposure and geometry state [2]; filled-resin behaviour couples rheology, recoating, dispersion and optical state [3]; native interlayer failure requires verified failure location and cannot be inferred from optical overlap alone [5]; and photocuring labels operate on distinct mechanistic axes that should not be collapsed [6,7]. The resulting architecture is not a fault-code table and does not assign diagnostic probabilities, sensitivity, specificity or universal thresholds. Its intended output is a traceable evidence state: one mechanism retained, several mechanisms retained, one or more mechanisms eliminated by an adequate test, or insufficient evidence to assign a unique cause. Corrective action is treated as an intervention that must change the predicted discriminator and the failure outcome before it strengthens a causal attribution.

1. Purpose and scope

The purpose of ER-184 is to define how evidence should be organised when more than one physical mechanism can explain an observed VPP defect. This six-stage evidence-to-decision architecture is consistent with the controlled corpus framework [9]. It is a diagnostic-inference paper, not a catalogue of defects and not a replacement for the specialist Engineering Reviews that own the underlying mechanisms.

ER-031 retains cure shrinkage, residual stress and warpage [1]. ER-034 retains bottom-up separation and peel mechanics [2]. ER-036 retains rheology, recoating and dispersion of highly filled suspensions [3]. ER-037 retains debinding defects in photopolymer-derived ceramic and metal green bodies [4]. ER-073 retains native same-material interlayer cohesion and anisotropy [5]. ER-084 retains the mechanistic atlas of photocuring pathways [7]. ER-184 uses these papers only to define competing branches, predictions and discriminating evidence.

The scope begins after an observable deviation has been detected. It ends when the available evidence justifies retaining or eliminating candidate mechanisms, or when it is insufficient to assign a unique cause. Product release, regulatory disposition and service qualification remain separate decisions.

2. Why symptom lookup is not diagnosis

A symptom is an observation. A mechanism is a physical or chemical process capable of producing an observation. A root cause is the specific mechanism, state and causal history that produced the observed event in the investigated system. These are different claims and require different evidence.

ER-031 provides a direct example. Polymerisation shrinkage is a local material response, whereas residual stress and warpage are structural outcomes that depend on nonuniform conversion, evolving stiffness and constraint [1]. Observing curl therefore does not identify a bulk shrinkage value as the root cause. ER-034 provides another example: a separation-force trace depends on interface state, resin state, exposure history, temperature, projected area and drainage path [2]. A visible separation cannot identify one of those variables without additional evidence.

The diagnostic error is not merely semantic. If a corrective action changes two or more state variables at once, success can be incorrectly attributed to the variable the operator expected to matter. ER-184 therefore requires the observation and each causal claim to be recorded at its own evidential level.

3. Six evidential stages

ER-184 separates six stages that are often collapsed in troubleshooting practice. Detection establishes that a deviation exists. Classification assigns an observable type. Localisation identifies where and when the deviation occurred. Mechanism attribution evaluates which physical or chemical processes remain compatible with the evidence. Root-cause attribution links a retained mechanism to the actual state and history that produced the event. Corrective-action verification establishes whether an intervention changed the predicted discriminator and the failure outcome.

Passing one stage does not automatically pass the next. A classifier may correctly label a delamination while providing no evidence about whether the initiating mechanism was insufficient native cohesion, excessive peel loading, support instability or another state. A localised crack may identify the failure plane while leaving multiple mechanisms compatible with that location. A corrective intervention may suppress the failure without proving that the inferred cause was unique.

Table 1. Evidential stages and permissible conclusions

Stage	Question answered	Permitted output	Shortcut prohibited
Detection	Is a deviation observable under the stated method?	Detected / not detected within method capability	Treating absence of observation as proof of absence
Classification	What observable class describes it?	Defined symptom class with measured descriptors	Using the class label as a cause
Localisation	Where and when did it arise?	Process stage, spatial location, layer/region, timing	Assuming the visible final location is the initiation site
Mechanism attribution	Which physical or chemical mechanisms remain compatible?	Retained / eliminated / unresolved candidate branches	One-to-one symptom-to-mechanism mapping
Root-cause attribution	Which mechanism plus state/history produced this event?	Causal attribution only to evidence-supported specificity	Selecting the familiar or frequent cause without discriminating evidence
Corrective-action verification	Did the intervention change the predicted discriminator and failure outcome?	Strengthened, weakened or unchanged causal support	Treating one successful fix as proof of a unique root cause

4. Stage 1: define the observable and the state

The diagnostic record begins with the specimen and process state, not with a cause label. The observable should be described using measurable features where possible: location, orientation, scale, timing, repetition, spatial distribution, layer index, surface or bulk character, dimensional vector, force-time signature, image morphology or another traceable descriptor.

The same visible feature can carry different meaning at different states. ER-073 distinguishes delamination that occurs during printing from final conditioned interlayer cohesion [5]. ER-031 requires post-cure drift to be separated from print-stage distortion by measurement at defined time points [1]. A diagnosis that omits the state boundary risks assigning a later observation to an earlier mechanism.

Minimum state variables depend on the branch being tested, but material identity, printer and optical configuration, exposure history, geometry, temperature, resin condition, relevant post-processing and measurement timing should be available whenever they can change the candidate set.

5. Stage 2: construct the candidate-mechanism set

Candidate mechanisms should be generated from controlled mechanism owners, not from familiarity or frequency. The set should be broad enough to include credible alternatives but bounded enough that each branch makes testable predictions.

A candidate set is conditional on the observed state. A bottom-up print-stage separation can legitimately open peel/interface, exposure/green-state, support and native-cohesion branches [2,5]. A filled-resin layer defect can open recoating/refill, local concentration or agglomeration, optical attenuation and exposure-state branches [3]. A defect that first appears during debinding must open the thermal/decomposition branch owned by ER-037 rather than being automatically back-attributed to print-stage cure [4].

A mechanism should not enter the candidate set merely because it exists elsewhere in the literature. It must be physically compatible with the material, architecture, process stage and observed state under investigation.

6. Stage 3: derive mechanism-specific predictions

A useful hypothesis predicts something that competing hypotheses do not all predict. The prediction can concern location, direction, response to a controlled perturbation, temporal evolution, force signature, concentration profile, conversion state, fracture path or another observable.

For warpage, a structural-constraint hypothesis predicts dependence on geometry, support release or constraint history rather than on one bulk shrinkage value alone [1]. For bottom-up separation, an interface-state hypothesis predicts a response to controlled film or surface condition, while a geometry-driven hydrodynamic contribution predicts dependence on projected area, orientation or drainage path under a fixed interface state [2]. For a filled suspension, a local concentration or agglomeration hypothesis predicts spatial or time-dependent differences that a single bulk viscosity measurement may not detect [3].

The purpose of prediction is discrimination, not narrative completeness. A mechanism that produces no distinguishable prediction under the available measurements cannot be uniquely attributed from those measurements.

7. Stage 4: select discriminating measurements

A discriminating measurement is one whose result can change the status of at least one candidate mechanism under the stated conditions. Measurements should therefore be selected after the candidate set and predictions are defined, not simply because an instrument is available.

ER-034 identifies calibrated force-time traces, interface condition, resin temperature and viscosity, green-part state, projected area, orientation and drainage path as relevant evidence for separation mechanics [2]. ER-036 requires time-resolved concentration or height information for sedimentation claims and warns that surface and vat-bottom samples are not automatically equivalent [3]. ER-073 requires failure-location evidence because a lower mechanical result can arise from grip, surface, void, geometry or interface effects [5].

A negative result eliminates a mechanism only when the method had adequate capability to detect the predicted state. Absence of a visible feature, absence of a signal below an unspecified detection capability, or a coupon that passes at one location is not general evidence that the mechanism is absent.

Table 2. Example candidate branches and discriminating evidence

Observable / state	Competing branches that may remain open	Discriminating evidence examples	Mechanism owner
Curl or warpage	Conversion gradient; evolving stiffness; structural constraint; post-cure drift	Time-resolved geometry; support-release response; spatial cure/state map; boundary-condition record	ER-031
Bottom-up print separation	Release-interface state; refill/hydrodynamics; green-state exposure; geometry/drainage; supports; native cohesion	Force-time trace; film state; resin temperature/viscosity; exposure history; projected area; failure location	ER-034 / ER-073
Filled-resin layer defect	Flow resistance; insufficient recovery; local concentration/agglomeration; trapped gas; optical attenuation mismatch	Rheology under matched history; time/position sampling; layer observation; optical/cure response	ER-036
Final interlayer failure	Native boundary weakness; bulk defect; surface/void origin; earlier print-stage peel damage	Interface-targeted mechanics; fracture path; image-based failure-location classification; matched bulk control	ER-073
Defect first appearing during debinding	Thermal/decomposition branch versus inherited green-state defect	Green-state record; debinding programme and atmosphere; defect timing; stage-resolved inspection	ER-037
Photocuring pathway attribution	Trigger, active species, network-forming reaction, spatial/temporal control and structure-evolution states can be different axes	Axis-specific chemical/optical/state measurements; specialist pathway evidence	ER-084

8. Stage 5: update hypotheses without manufacturing certainty

Evidence updates the candidate set. A result may eliminate one branch, retain one branch, retain several branches or fail to discriminate. ER-184 deliberately permits the result 'insufficient evidence to assign a unique cause'. That output is scientifically stronger than forcing a unique label when two or more mechanisms remain compatible with the observations.

No prior probability is assigned in this review. Terms such as 'most likely', diagnostic probability or confidence percentage require a declared empirical or statistical basis. Frequency in an operator's experience can guide which test is efficient to perform first, but it is not evidence that the mechanism produced the investigated event.

When several branches remain, the next action is another discriminating measurement or a controlled intervention, not a stronger adjective.

9. Worked architecture: curl, warpage and dimensional drift

Curl or warpage is a useful example because the observable can be generated by multiple states. ER-031 distinguishes chemical shrinkage from residual stress and geometry-dependent warpage and states that spatial dose gradients can create conversion gradients whose dimensional consequence depends on geometry and constraint [1].

A diagnostic record should first localise when the distortion appears: during printing, immediately after support removal, after washing, after post-cure or after conditioning. Distortion appearing specifically after support removal supports a released-constraint branch but does not identify the earlier source of the stored stress. A spatial correlation with dose or thickness can support a conversion-gradient branch, while a geometry-specific pattern can support a constraint or stiffness-distribution branch. The discriminator is the measured relation between state change and dimensional response, not the word 'warpage'.

A digital compensation that reduces dimensional error does not by itself validate a causal mechanism because compensation can mask the outcome without removing the underlying material/process state [1].

10. Worked architecture: separation and delamination

The word delamination can describe different events. During bottom-up VPP, separation from the vat interface and refill mechanics can damage a green structure [2]. In the final conditioned part, native same-material interlayer cohesion is a different property [5]. ER-184 therefore requires the failure stage and location to be established before a mechanism branch is selected.

For a print-stage event, candidate branches can include release-interface state, viscous refill or hydrodynamic loading, exposure-dependent green state, projected area and drainage, support instability and native-cohesion limitations [2,5]. Force-time traces, film condition, resin temperature and viscosity, exposure history, geometry, support state and image-based failure location can discriminate among subsets of these branches.

Cure depth exceeding layer thickness is not proof of adequate interlayer strength. ER-073 explicitly treats optical overlap as distinct from interface strength and requires interface-sensitive mechanics and failure-location evidence [5]. A failed layer can therefore remain diagnostically unresolved even when a working-curve measurement appears acceptable.

11. Worked architecture: filled-resin layer defects

Highly filled photopolymer suspensions add a diagnostic problem because bulk measurements can miss local or time-dependent state. ER-036 states that surface and vat-bottom samples are not assumed equivalent and that stable bulk viscosity does not exclude local concentration or agglomerate changes [3].

For an incomplete or nonuniform layer, candidate branches can include excessive flow resistance, insufficient recovery time, recoating geometry, trapped gas, local concentration change, agglomeration, altered optical attenuation and exposure-state mismatch [3]. A useful sequence pairs rheology with position- and time-resolved sampling, layer formation observation and optical/cure measurement under the same material history. A working curve measured immediately after mixing cannot automatically close a branch for a resin that has rested, been repeatedly sheared or changed temperature [3].

This branch illustrates a general rule: a diagnostic measurement should reproduce the state about which the causal claim is made.

12. Process-stage routing: debinding and secondary transformation

When a defect first appears after printing during debinding of a photopolymer-derived ceramic or metal green body, the candidate set must include the thermal/decomposition mechanisms owned by ER-037 [4]. ER-184 does not restate those mechanisms because the detailed source text is outside the present sentence-level evidence pack.

The diagnostic rule is nevertheless important: a later process-stage defect must not be assigned to print-stage cure merely because print parameters are easy to change. The investigation should preserve the green-state record, debinding programme, atmosphere and defect timing so that print-stage and debinding-stage hypotheses can be discriminated. If those records are absent, an unresolved cross-stage attribution is the correct result.

13. Mechanism classification is itself an evidential problem

Photocuring mechanism labels can create a second diagnostic error: several labels may describe different scientific levels rather than competing mechanisms. ER-083 separates trigger, active species, network-forming transformation, temporal or spatial control, structure-evolution overlay and qualification endpoint [6]. ER-084 owns the detailed mechanistic atlas [7].

A formulation described as dual cure, cationic, phase-separating or reactivatable is therefore not diagnosed by one label alone. The evidence must identify which level is being attributed and which measurement addresses that level. Conversion does not by itself identify active-species supply, morphology, dark reaction or reactivation [6].

ER-184 uses this layered classification only to prevent category errors in diagnosis. Detailed pathway evidence remains with ER-084 and the specialist chemistry reviews.

14. Corrective action as a controlled causal test

A corrective action can provide causal evidence only if it is framed as an intervention with a prediction. Before the change, the record should state which candidate mechanism the intervention is expected to affect, which discriminator should change and which competing branches should remain approximately controlled.

If a modified parameter removes the visible defect but also changes exposure, temperature, peel history and support geometry, the intervention is not specific enough to establish which mechanism was causal. Conversely, if a targeted intervention changes the predicted discriminator in the expected direction and suppresses the failure while relevant competing conditions remain controlled, the mechanism attribution is strengthened.

One successful correction is not proof of uniqueness. Repeatability, reversal or an appropriate comparator may be needed when the claim matters to release or qualification.

15. Minimum diagnostic evidence record

A diagnostic record should be sufficient for another qualified investigator to reconstruct the evidence state. It should identify the observable, process stage and specimen state; the candidate mechanisms and their controlling sources; predictions; measurements and method capability; results; eliminated and retained branches; unresolved alternatives; intervention, if any; and the evidence linking the intervention to both the discriminator and the failure outcome.

The record should also preserve negative evidence with its capability limits. 'No defect seen' is not equivalent to a quantified inspection result. 'No signal detected' is incomplete without the relevant method capability. 'Coupon passed' is not evidence for an untested location, geometry or material state.

Table 3. Minimum diagnostic evidence record

Record block	Minimum content	Reason
Case identity	Material/version/lot, printer/optics, geometry, build file, date and operator or automated cell identity	Defines the investigated system
Observable	Measured signature, images/traces, location, orientation, scale and timing	Separates observation from interpretation
State history	Exposure, temperature, resin state, supports, washing, post-cure, conditioning and relevant transitions	Preserves causal chronology
Candidate set	Each mechanism, owner/source and physical compatibility rationale	Makes alternatives explicit
Predictions	Expected distinguishing result for each branch	Turns narratives into testable hypotheses
Measurements	Method, calibration/capability, sampling position/time and acceptance logic	Defines what a positive or negative result can mean
Evidence update	Retained, eliminated and unresolved branches with rationale	Prevents forced certainty
Intervention	Targeted change, predicted discriminator response and controlled co-variables	Makes corrective action a causal test
Verification	Replicates, discriminator response, failure outcome and residual alternatives	Separates improvement from root-cause proof

16. What validation of the diagnostic system would require

The architecture in ER-184 is proposed and not validated. Validation would require a set of cases with independently established mechanism states or sufficiently strong reference evidence, a prespecified diagnostic procedure, blinded or otherwise bias-controlled application of that procedure, and reporting of both correct and unresolved outcomes.

Only such a programme could support claims such as diagnostic sensitivity, specificity, error rate, predictive value or calibrated probability. A collection of successful troubleshooting anecdotes cannot supply those performance measures because case selection, ground truth and failed diagnoses are uncontrolled.

Until that work exists, the framework should be judged by traceability and evidential discipline rather than by an asserted diagnostic accuracy.

17. Limitations and evidence gaps

ER-184 is built from a controlled mechanism corpus rather than from a prospective validation dataset. It therefore defines an inference architecture but cannot quantify how often the architecture reaches the correct mechanism or how frequently it remains unresolved.

The mechanism coverage is also intentionally incomplete. It is bounded to branches supported by the current controlled corpus. New failure modes or mechanisms should enter only after source verification and ownership assignment. ER-037 and ER-084 are used here mainly as routing boundaries because their exact sentence-level manuscript evidence was not required for the present examples.

The most important research need is a deliberately constructed VPP diagnostic benchmark containing paired process histories, state-resolved measurements, known perturbations, failure-location evidence and independent mechanism adjudication. Such a dataset would allow the proposed architecture to be tested rather than merely argued.

18. Conclusions

Diagnostic inference in vat photopolymerisation should not begin with a fault-code table. It should begin with a state-resolved observation and a set of physically credible competing hypotheses. The investigation then asks what each hypothesis predicts, which measurements can discriminate among those predictions, what evidence can legitimately eliminate a branch and when uncertainty must remain explicit.

The controlled Engineering Review corpus demonstrates why this is necessary. Warpage is not equivalent to shrinkage [1]; separation is coupled to interface, exposure, resin and geometry state [2]; filled-resin defects can evade bulk measurements [3]; interlayer failure requires failure-location evidence [5]; and photocuring classifications occupy different mechanistic levels [6,7]. These mechanism papers remain the owners of their science. ER-184 contributes the cross-mechanism evidence architecture.

The strongest permitted conclusion is therefore not always a named root cause. It may be that one mechanism remains supported, that several remain compatible, that a mechanism has been eliminated by an adequate test, or that the available evidence is insufficient to assign a unique cause. That distinction is the basis for a traceable diagnostic record and for later validation of any automated or human diagnostic system.

Related Engineering Review boundaries

ER-031 owns cure shrinkage, residual stress and warpage.

ER-034 owns bottom-up separation and peel mechanics.

ER-036 owns rheology, recoating and dispersion in highly filled suspensions.

ER-037 owns debinding defect formation.

ER-073 owns native same-material interlayer cohesion and mechanical anisotropy.

ER-084 owns detailed photocuring pathway mechanisms.

ER-184 owns cross-mechanism diagnostic inference and evidential discrimination only. Its corpus-level limitations are interpreted within the broader Engineering Reviews evidence architecture [8].

References

- [1] Segurolo, J. Cure Shrinkage, Residual Stress and Warpage in VPP: A Qualification and Minimum-Reporting Note. ER-031. Zenodo (2026). Concept DOI: [10.5281/zenodo.22048036](https://doi.org/10.5281/zenodo.22048036).
- [2] Segurolo, J. Separation and Peel Mechanics in Bottom-Up VPP: A Qualification and Minimum-Reporting Note. ER-034. Zenodo (2026). Concept DOI: [10.5281/zenodo.22048042](https://doi.org/10.5281/zenodo.22048042).
- [3] Segurolo, J. Rheology, Recoating and Dispersion Control of Highly Filled VPP Suspensions. ER-036. Zenodo (2026). Concept DOI: [10.5281/zenodo.22048044](https://doi.org/10.5281/zenodo.22048044).
- [4] Segurolo, J. Debinding Defect Formation in Photopolymer-Derived Ceramic and Metal Green Bodies. ER-037. Zenodo (2026). Concept DOI: [10.5281/zenodo.22048046](https://doi.org/10.5281/zenodo.22048046).
- [5] Segurolo, J. Interlayer Cohesion and Mechanical Anisotropy in Vat-Photopolymerised Thermosets: Cure Overlap, Network Continuity and Qualification. ER-073. Zenodo (2026). Concept DOI: [10.5281/zenodo.22123351](https://doi.org/10.5281/zenodo.22123351).
- [6] Segurolo, J. Photocuring Pathway Taxonomy for Vat Photopolymerisation: Active Species, Network-Formation Modes, Spatial Gating and Qualification Boundaries. ER-083. Zenodo (2026). Concept DOI: [10.5281/zenodo.22123372](https://doi.org/10.5281/zenodo.22123372).
- [7] Segurolo, J. Mechanistic Atlas of Photocuring Pathways in Vat Photopolymerisation: From Optical Trigger to Network Formation and Structural Evolution. ER-084. Zenodo (2026). Concept DOI: [10.5281/zenodo.22123376](https://doi.org/10.5281/zenodo.22123376).
- [8] Segurolo, J. From Individual Reviews to a Qualification Corpus: Evidence Architecture Across Engineering Reviews ER-01 to ER-99. ER-100. Zenodo (2026). Concept DOI: [10.5281/zenodo.22161173](https://doi.org/10.5281/zenodo.22161173).
- [9] Segurolo, J. Evidence-to-Decision Qualification Framework for Vat Photopolymerisation. RP-02. Zenodo (2026). Concept DOI: [10.5281/zenodo.22048032](https://doi.org/10.5281/zenodo.22048032).

Competing interests

The author is Founder, CEO and CTO of 3Dresyns by Resyner Technologies S.L. The paper is an internal methodological synthesis of a controlled Engineering Review corpus and is not independent peer review of that corpus. No comparative product recommendation is made.

Data availability and AI-assistance statement

No new experimental dataset was generated. The proposed diagnostic architecture is derived from the cited controlled corpus and has not been prospectively validated. Generative AI assisted drafting and corpus organisation but is not an author. Juan Segurolo retains responsibility for the final article.