

Living Systematic Reviews for Rapidly Changing Additive-Manufacturing Evidence: Search, Updating and Version Governance

Defining when an engineering evidence synthesis should remain live, what triggers an update and how a changed conclusion is versioned without continuous editorial churn

Juan Segurola

3Dresyns by Resyner Technologies S.L.

© 2026 Juan Segurola. All Rights Reserved.

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

ER-395 · Engineering Review · Version 0.1 · 5 October 2026

Abstract

Fast-moving additive-manufacturing fields can outdate a static review before its engineering conclusions are acted upon. Living systematic reviews address this problem by maintaining an active surveillance and update process, but continuous searching does not imply continuous publication. This review adapts living-review principles to photopolymer engineering. It defines eligibility for living status, search cadence, update triggers, version governance, statistical cautions and stopping rules. The proposed model separates evidence surveillance from formal review release so that a corpus can remain current without generating unstable micro-versions.

Keywords: living systematic review; additive manufacturing; evidence surveillance; update trigger; version governance; systematic review

1. Living status is a resource commitment

Living systematic review methodology is intended for questions where evidence is important, uncertain and likely to change [1]. An engineering corpus should therefore not make every review live by default. The decision should be based on expected evidence velocity, decision consequence and whether new evidence could plausibly alter the engineering conclusion.

2. Search cadence and release cadence are different clocks

Automated or scheduled searches can run frequently, while formal releases occur only when a trigger is met. This distinction avoids micro-versioning. A search cycle creates a surveillance record; a release cycle creates a new review version. PRISMA 2020 remains relevant because each released version should transparently report the search and selection process [2].

Table 1. Eligibility for living-review status.

Criterion	Living favoured when	Static review favoured when
evidence velocity	new relevant studies appear frequently	evidence base changes slowly
decision impact	new evidence can change qualification	conclusion is descriptive only
uncertainty	material unresolved uncertainty remains	boundaries are stable
maintenance capacity	search and screening can be sustained	no durable maintenance route

3. Update triggers should be claim based

A new paper is not automatically an update trigger. Triggers include evidence that changes a boundary condition, introduces a previously missing material state, reverses a key conclusion, changes a standard used for qualification, or materially reduces uncertainty. Low-impact additions can remain queued until a scheduled consolidation.

4. Statistical updating needs repeated-look discipline

Repeated meta-analysis can inflate error rates or produce unstable inference if each update is treated as a fresh independent analysis. Living-review methodology has developed sequential approaches and cautions for repeated updating [3]. Many photopolymer questions are too heterogeneous for quantitative pooling; in those cases, the equivalent discipline is to preserve state stratification and avoid turning increased source count into false precision.

Table 2. Update triggers.

Trigger	Action
new state or failure mode	full differential review
standard or regulatory change	scope-specific update
quantitative shift within same conclusion	queue for consolidated release
duplicate confirmatory evidence	record in surveillance log
source correction or retraction	immediate dependency check

5. Version governance protects the engineering record

Every release should have a stable identifier, frozen search date, source manifest, change log and explicit statement of whether conclusions changed. Older releases remain citable. A new version should not silently overwrite the evidence base that supported a

historical decision. Living guideline methods likewise emphasise the relationship between updated evidence and updated recommendations [4].

6. Stopping rules are part of the protocol

A review can leave living status when evidence velocity falls, the engineering question stabilises, the technology is superseded or the decision no longer warrants continuous surveillance. A stopping decision should be documented just as an update decision is. This prevents indefinite maintenance commitments.

Table 3. Version-governance record.

Field	Required content
search cut-off	date and databases/sources
change set	added, removed, corrected evidence
claim delta	unchanged, narrowed, broadened, reversed
review authority	who approved the new version
supersession	which version is current and what remains citable

7. Engineering implementation

For photopolymer AM, the most useful living candidates are questions tied to rapidly changing printer architectures, new standards, regulatory evidence, or data-driven modelling. Fundamental mechanisms with stable evidence may be better served by scheduled review rather than continuous surveillance. The governing principle is decision relevance, not novelty.

8. Scope boundary and neighbouring reviews

ER-395 owns search surveillance, update triggers and version governance for living reviews. ER-394 owns provenance relationships between sources, claims and versions; ER-396 owns adjudication when updated evidence conflicts; and ER-400 consumes the resulting review states in corpus synthesis. ER-395 does not resolve scientific contradiction by itself.

9. Living mode parameters must be declared before surveillance starts

A review is not 'living' merely because authors intend to update it. The living mode requires operational parameters: search sources, search frequency, screening cadence, trigger criteria, analysis cadence, publication cadence and retirement conditions. PRISMA-LSR formalises this by adding reporting requirements for living-mode parameters, changes to methods, changes to results and author roles across versions [5]. For additive-manufacturing evidence, these parameters should also identify which technologies, standards and machine generations are monitored because a rapid hardware or materials change can alter applicability even when the underlying scientific question is stable.

10. Search triggers and publication triggers are different decisions

A monthly search does not imply a monthly publication. New records can be screened continuously while a public update is released only when a predefined trigger is met. Useful triggers include a new study that crosses a qualification boundary, a new standard revision, evidence that reverses a prior conclusion, a new material or machine class that expands scope, or accumulation of enough concordant evidence to change confidence. Separating surveillance from release prevents needless version churn while preserving currency. It also allows the update latency to be measured: the time from source availability to screening, from screening to synthesis and from synthesis to public version.

11. Version governance needs a structured delta record

Each living version should declare what changed and what did not. The delta record should separate search additions, inclusion/exclusion decisions, extracted data changes, risk-of-bias changes, synthesis changes, conclusion changes and editorial-only corrections. PRISMA-LSR specifically emphasises transparent reporting of changes between versions [5]. For engineering use, the record should additionally identify which qualification claims were re-opened and which remained valid by dependency analysis. This makes a living review compatible with controlled engineering decision records rather than turning it into an endlessly mutable webpage.

12. Retirement is a scientific state, not project abandonment

Living status consumes resources and should stop under explicit conditions. A methodological survey of living systematic reviews found that update practice and retirement reporting were inconsistent, including reviews that lagged far beyond planned update intervals [6]. Retirement can be justified when the question is no longer decision-critical, the evidence has stabilised sufficiently for the intended use, the technology has been superseded, or surveillance yield becomes too low to justify the resource. The final living version should record the last search date, retirement rationale and conditions that would justify reactivation.

13. Additive-manufacturing surveillance architecture

A practical AM implementation can use three queues. The first monitors peer-reviewed literature and preprints for new process/property evidence. The second monitors standards and regulatory documents whose revision changes method validity. The third monitors machine, material and software revisions that can create a new state even without new academic literature. Items are deduplicated but not collapsed: a machine firmware change and a new journal paper may trigger different reviews. The living system therefore maintains a source ledger, a claim-impact ledger and a publication ledger. This architecture keeps rapid surveillance separate from authorisation to change the public engineering conclusion.

14. Surveillance performance should be measured

A living process can fail by missing evidence as well as by updating too slowly. Search performance should therefore be monitored through sentinel records, known-key-study checks, database coverage and periodic review of terminology drift. Screening disagreement

and false exclusions should be sampled. When machine-learning prioritisation is used, the model version and human override process should be recorded. These controls matter because a living review creates an expectation of currency; an opaque surveillance pipeline can give false confidence that absence of an update means absence of new evidence.

15. Resource burden and update latency are engineering variables

Living review design is a resource-allocation problem. More frequent searching reduces detection latency but increases screening overhead; frequent public releases increase editorial burden and version noise. The protocol should therefore budget expected search yield, screening capacity, expert review time and publication delay. A trigger that cannot be acted on within the promised interval is not an effective trigger. Tracking median and worst-case latency from source publication to reviewed incorporation provides a direct operational measure of whether the living system is meeting its purpose.

16. Change classification protects downstream users

Not all updates deserve the same response from users. ER-395 proposes classifying releases as surveillance-only, evidence addition without conclusion change, boundary change, conclusion change or methodological change. The release record should state the class and the affected claims. Downstream qualification systems can then decide whether a new version requires immediate re-evaluation or simple archival replacement. This is especially important for an engineering corpus in which many later documents may depend on one living review.

17. Surveillance diagnostics and process degradation

A living review should monitor its own ability to detect relevant evidence. Useful indicators include database coverage, search failure rate, duplicate burden, proportion of records requiring full-text review, time from source publication to screening and time from eligibility to public update. Sudden changes in any indicator can reveal a broken search string, inaccessible database or staffing bottleneck. The purpose is not to optimise a bibliometric dashboard but to detect degradation in the review process before it becomes a hidden gap in evidence surveillance. [1,5,6]

18. Trigger calibration and false alarms

Update triggers can be too sensitive as well as too insensitive. A trigger based on any new publication may create excessive version churn, while a trigger based only on statistically significant reversal can miss important boundary changes. Calibration should therefore use historical examples: replay prior search updates and ask which trigger would have fired, how quickly and whether the resulting update would have changed an engineering decision. The selected trigger should state the tolerated false-alarm burden and the consequences of delayed detection. This makes update governance auditable rather than intuitive.

19. Retirement, archival state and handover

A living review cannot be assumed to remain live indefinitely. Retirement criteria should be defined in advance, for example when the topic stabilises, surveillance resources cease, the review is superseded by a broader corpus or no stakeholder remains responsible for maintenance. Retirement does not delete the work. The final release should state the last search date, unresolved screening backlog, successor record if any and the status of downstream dependencies. A dormant review without such a handover is dangerous because users may interpret a recent-looking document as actively surveilled evidence when no monitoring is occurring. [4-6]

20. Screening quality control

Living surveillance can fail even when the search runs on schedule. Duplicate resolution, title/abstract screening and full-text eligibility decisions should therefore be sampled for consistency. A second reviewer or adjudication sample can estimate disagreement and identify ambiguous eligibility rules. If disagreements cluster around a topic boundary, the protocol should clarify the rule rather than simply retrain the screener. Automated classifiers may reduce workload, but their threshold, training data and human override process should be versioned because a model update can change which records reach review. [1,2,5]

21. Evidence accumulation and repeated looks

Repeated updating creates a risk of overreacting to early or noisy evidence. Where quantitative synthesis is used, the review should state how repeated analyses are interpreted and whether thresholds for changing a conclusion differ from thresholds used in a one-off review. In sparse engineering literatures, formal sequential methods may be impractical, but the principle remains: a single new study should not automatically reverse a mature conclusion if it is small, indirect or methodologically discordant. Conversely, a well-designed study that crosses a previously identified boundary can justify an immediate scope change even without a large meta-analysis. [3,4]

22. Downstream dependency notification

A living review becomes infrastructure when other documents depend on it. Each release should therefore expose a machine-readable change summary identifying claims added, removed, narrowed or strengthened. Downstream reviews can subscribe to those claim-level changes rather than re-reading the whole document after every surveillance cycle. A conclusion-changing update should trigger explicit impact assessment; a surveillance-only release may require no downstream action. This keeps a living evidence system responsive without creating unnecessary re-review of stable dependent documents.

23. Limitations and research priorities

Living-review methods were developed largely in evidence domains with mature indexing and review infrastructure; additive-manufacturing evidence is more fragmented across journals, standards, conference literature and manufacturer documentation. The important research question is therefore operational: which surveillance architecture detects decision-relevant change with acceptable latency and workload. Priorities include retrospective replay studies of search triggers, measurement of missed eligible studies, classifier audits where machine screening is used, and dependency-aware notification trials for downstream engineering documents. Resource requirements and retirement behaviour should be reported alongside scientific outcomes. Until these practices are tested in AM corpora,

ER-395 supports transparent living-review governance but does not imply that high-frequency updating is always superior to a well-controlled periodic review. [1-6]

References

1. Elliott JH, Synnot A, Turner T, et al. Living systematic review: 1. Introduction-the why, what, when, and how. *Journal of Clinical Epidemiology*. 2017;91:23-30. doi:10.1016/j.jclinepi.2017.08.010.
2. Page MJ, McKenzie JE, Bossuyt PM, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*. 2021;372:n71. doi:10.1136/bmj.n71.
3. Simmonds M, Salanti G, McKenzie J, Elliott J. Living systematic reviews: 3. Statistical methods for updating meta-analyses. *Journal of Clinical Epidemiology*. 2017;91:38-46. doi:10.1016/j.jclinepi.2017.08.008.
4. Akl EA, Meerpohl JJ, Elliott J, Kahale LA, Schunemann HJ. Living systematic reviews: 4. Living guideline recommendations. *Journal of Clinical Epidemiology*. 2017;91:47-53. doi:10.1016/j.jclinepi.2017.08.009.
5. Akl EA, Khabsa J, Iannizzi C, et al. Extension of the PRISMA 2020 statement for living systematic reviews (PRISMA-LSR): checklist and explanation. *BMJ*. 2024;387:e079183. doi:10.1136/bmj-2024-079183.
6. Akl EA, El Khoury R, Khamis AM, et al. The life and death of living systematic reviews: a methodological survey. *Journal of Clinical Epidemiology*. 2023;156:11-21. doi:10.1016/j.jclinepi.2023.02.005.

Competing interests

The author is Founder, CEO and CTO of 3Dresyns by Resyner Technologies S.L.; ER-395 is a methodological engineering review and does not provide a comparative commercial recommendation.

Data availability

ER-395 is a literature-based engineering review; no new experimental dataset was generated.