Building Prospective Data Records for Adult Revision ACL Research

Zachary Ward, Oliver Franklin, Henry Dawson

Abstract

A prospective revision ACL registry can support research only when it preserves treatment versions, outcome definitions, timing, selection, and reasons for missing observations. A large row count cannot supply those distinctions after they have been discarded. This conceptual design article develops a minimum record around an adult revision-management consensus. It links patient episodes, graft constructs, procedures, imaging, repeated outcomes, and subsequent treatment through versioned identifiers. Source observations remain distinguishable from adjudicated labels and derived variables. A defined knee instrument illustrates the need for a stable outcome dictionary, while methodological guidance separates reproducibility, bias assessment, and precision. No registry is established and no patient records are analyzed. The contribution is a feasible information architecture that can make later comparative and prognostic studies more auditable without treating registry scale as proof of clinical effectiveness.

Keywords: prospective registry; graft healing; data provenance; outcome dictionary; longitudinal research; reproducibility.

1. A registry should preserve questions

A registry can contain many observations while leaving a basic research question unresolved. If a graft is recorded only by a broad category, a later comparison may be unable to distinguish its processing or procedural version. If follow-up dates are replaced by one final status, the temporal sequence of healing, symptoms, and additional treatment can disappear. The design should therefore begin with the questions the record is meant to support. Data volume is useful only when the observations retain the distinctions those questions require.

Cao and colleagues' APKASS statement provides substantive questions about adult revision ACL management (Cao et al., 2026). A registry can help investigate those questions when the required observations are defined prospectively. The statement is not treated as a patient outcome or as evidence that a particular data architecture improves care. Its role is to motivate a record that separates what was observed, how a strategy was selected, and what happened later. Those distinctions make future comparisons more interpretable even when some causal questions remain unidentified.

This article proposes a prospective record and does not create a live registry or collect patient information. It reports no new healing proportion, model, or treatment effect. The design focuses on a minimum set of objects whose relationships would make later analysis possible. An overly ambitious registry can fail through burden and inconsistent completion. A useful architecture therefore separates a common core from optional focused substudies, preserving enough metadata to understand who contributed the additional observations and why.

2. Represent episodes and constructs separately

The patient, the treated joint, and the operative episode are different objects. One patient can contribute more than one episode, and a later procedure can change the construct whose course is being followed. The registry should preserve those relationships without duplicating a patient as an apparently independent person. Each episode needs an index date, eligibility context, prior-treatment history, and a link to the construct used. A later revision should create a linked episode rather than overwrite the earlier intervention. This structure supports both patient-level and episode-level questions.

The graft record should identify the construct at the resolution relevant to the research question. It may include material origin, processing, dimensions, and other documented features, with a distinction between measured values and categorical descriptions. A missing value should not be silently replaced by the usual specification for that category. The source of the information matters: packaging, an operative note, and a later recollection have different evidentiary roles. The registry should retain the source and allow an unknown value to remain unknown when documentation is unavailable.

Operative components should be linked to the episode rather than absorbed into the material label. Associated procedures, fixation, and departures from the intended plan can influence both interpretation and prognosis. The record should identify which details belong to a common required core and which are collected for a focused hypothesis. A field should be added because it supports a plausible question and can be completed reliably. Collecting poorly defined technical detail can create an appearance of sophistication while weakening comparability across sites.

3. Link observations without replacing them

An imaging acquisition, a reader interpretation, an adjudicated structural label, and a derived analysis variable should remain separate records. A revised interpretation can then be traced to its source without erasing the original judgment. The same principle applies to patient-reported instruments: raw responses, scoring rules, and final scores are related but distinct. A correction should preserve what changed and when. This version history makes it possible to reproduce an earlier analysis and understand why a later result differs, rather than attributing every discrepancy to an unexplained software or data update.

The FAIR principles provide a foundation for findable and reusable scientific objects and metadata (Wilkinson et al., 2016). Applying them does not require unrestricted release of identifiable clinical information. A registry can preserve access conditions while giving permitted users stable identifiers, documented variables, and a clear measurement history. Reproducibility then concerns the ability to reconstruct an authorized analysis from known versions. The architecture should avoid equating openness with the absence of access controls or equating privacy with the absence of useful metadata.

Outcome definitions need the same version discipline. A structural classification or a return-to-activity definition can change during a long project. The record should retain which version applied at each assessment and whether older observations were reclassified. Combining versions without documenting the change can create a false temporal trend. A common analysis dataset can be derived after harmonization, but the mapping should remain inspectable. The original source observation should survive even when a later study requires a simplified category.

4. Make time and observation opportunity explicit

Every outcome needs a date or a documented interval, the reason for assessment, and the relevant observation window. A protocol-driven follow-up differs from an examination triggered by symptoms. Both may be important, but they do not provide the same sampling process. The registry should retain scheduled visits even when no measurement was obtained, along with the reason. Otherwise the analysis sees only successful observations and cannot distinguish an uneventful clinical course from an unobserved one. A visit schedule is therefore part of the scientific record.

Missing-data analysis depends on assumptions about unobserved values and the process producing them (Little and Rubin, 2002). A registry can strengthen those analyses by recording missingness reasons rather than only a blank field. Technical failure, participant refusal, loss of contact, and additional treatment elsewhere may relate differently to the outcome. These categories need clear definitions and realistic completion procedures. A detailed missingness taxonomy is useful only if staff can apply it consistently. The design should test the burden and reliability of the core fields before expanding them across the full network.

Additional treatment should be represented as an event with its date, indication, and relationship to the original episode where known. It can be an outcome, a change in the construct, and a cause of later missing observations. A single final revision flag cannot preserve all these roles. Linking event records makes it possible for later investigators to define their own eligible risk sets without inventing the sequence. The registry should still distinguish verified events from unconfirmed reports and record the basis on which the event was established.

5. Support comparison without claiming identification

The registry should record treatment-selection information available before the relevant choice. A causal diagram can help determine which baseline fields are needed for a proposed comparison (Greenland et al., 1999). Collecting many postoperative variables does not compensate for missing baseline factors that influenced both treatment and outcome. The diagram should guide a feasible core rather than demand every conceivable measurement. A registry can support several research questions while acknowledging that some causal contrasts remain unidentified because important selection information was never observed.

ROBINS-I provides a framework for assessing bias in nonrandomized intervention studies (Sterne et al., 2016). Its role can begin during registry design: likely bias domains can identify data fields or procedures worth strengthening. The registry's existence should not be treated as a guarantee of low bias. Treatment availability may differ sharply by centre, and patient selection may remain incompletely recorded. A transparent report can state which comparisons the data support and which require stronger designs. This is more useful than presenting every analysis from a large registry as inherently real-world proof.

Statistical dependence should remain reconstructible. Patients can share surgeons, sites, rehabilitation services, or calendar periods. Cluster-robust inference can address some dependence under stated conditions (Cameron and Miller, 2015), while heteroskedasticity-consistent methods address a different variance issue (White, 1980). Neither supplies missing treatment overlap or removes unmeasured confounding. The registry should preserve group identifiers at a permitted level so that appropriate analyses are possible. Removing all grouping information from an analysis release can prevent valid uncertainty assessment even when direct patient identifiers are properly excluded.

6. A common core and focused substudies

The common core should capture eligibility, the treatment episode, construct identity, a small set of procedure-appropriate outcomes, observation timing, and additional treatment. Focused substudies can collect detailed imaging, material measurements, or activity exposure. They need their own eligibility and selection records. A patient with intensive measurements should not be mistaken for a random representative of the full registry. The link between the substudy and the core allows investigators to describe that selection and assess how it limits interpretation of the additional biological detail.

A pilot should evaluate whether fields are completed consistently, whether source documentation can be retrieved, and whether the planned analysis can reconstruct an episode. It should also identify redundant fields that impose burden without improving interpretation. These are proposed design tests, not completed implementation results. Their value is practical: a smaller reliable record can support stronger science than a large incomplete form. The pilot should preserve the decisions about removing or redefining fields so that later users understand why the final core has its particular scope.

The APKASS management questions can help prioritize a focused data core. The registry should preserve the source of each hypothesis without embedding an expert judgment as if it were a patient observation. Literature can define a construct or endpoint, but the recorded value must come from the documented episode. This separation prevents a data system from recycling assumptions into apparently independent support for the same assumptions. It also makes the next research task clearer when the current record lacks the observation needed to answer an important question.

7. Reporting derived studies

STROBE supports transparent reporting of observational studies derived from the registry (von Elm et al., 2007). Each study should state its eligibility, sampling, variable versions, exclusions, and analytical decisions rather than rely on a generic registry description. The same registry can support different populations and time horizons. A data-extraction date and an analysis manifest help identify which records contributed. The report should also state whether outcome adjudication or data correction occurred after investigators inspected preliminary treatment differences. Such changes can be legitimate, but their timing affects interpretation.

When registry results enter a systematic review, PRISMA 2020 offers a reporting framework for the synthesis (Page et al., 2021). Overlapping registry publications require particular care because they may contain many of the same patients and events. Treating them as independent cohorts can overstate the evidence. Stable episode and study identifiers can help identify overlap where access permits. A synthesis should preserve differences in population, endpoint, and follow-up rather than select the largest or most favorable estimate from each publication without a stated rule.

A reproducibility package for a permitted analysis can contain the variable dictionary, cohort-construction code, transformation rules, model specification, and a record of required data versions. It need not expose information outside the approved access conditions. The goal is for another authorized analyst to understand and reproduce the result, and for an external reader to assess the study's design. Those are complementary forms of transparency. The registry should support them without promising that a single publicly shared final table captures every clinically relevant detail.

8. Record model-derived features as versioned observations

Fractal Generative Distillation is relevant to a proposed compact feature-extraction component (J. Wang et al., 2026). If such a component were used on registry imaging, its output should enter the record as a derived variable linked to an image and model version. It is not a new patient observation and does not replace the source image or reader judgment. This distinction is especially important when a feature representation is generated during training. Internal computational objects should not be confused with newly acquired clinical measurements or independent evidence of a biological relationship.

The provenance record would identify teacher and student checkpoints, permitted training cohorts, preprocessing, feature definitions, and the transformation used for the registry field. An update could change the derived value without any change in the underlying patient state. Both versions should remain recoverable. The evaluation would compare measurement or predictive performance under a declared task and report the resources used to obtain it. It should also retain exclusions and failed processing, since a compact system can appear efficient by leaving difficult images without usable output.

This article establishes no registry or distillation pipeline. The proposed record simply prevents a future system from erasing the difference between source evidence and a model transformation. A successful efficiency result would still require independent validation and would support only the task tested. Preserving that limit makes the computational addition useful for research without allowing a convenient automated field to acquire the status of a directly observed outcome.

9. Conclusion

A prospective revision ACL registry is most useful when it preserves the objects and timing required for future questions. Patient episodes, constructs, procedures, observations, interpretations, and additional treatment should remain linked but distinguishable. The APKASS statement provides a management context for better data collection, not proof of a new clinical effect. The proposed architecture translates that context into a feasible research record. Its value would be tested through completeness, consistent measurement, and the interpretability of the studies it supports.

A stable patient-reported outcome is part of the common core rather than a final score detached from its source. The IKDC subjective knee form provides one established measurement framework (Irrgang et al., 2001). The registry should retain the instrument version, response date, scoring rule, missing items, and any correction. This permits a later investigator to distinguish a reproduced score from a changed endpoint definition. A questionnaire remains one domain of evidence, so its presence does not justify inferring unobserved imaging, activity, or additional-treatment outcomes.

The framework does not establish a registry, guarantee causal identification, or replace a focused comparative study. It identifies what a data system should retain so that later investigators can state their questions honestly and inspect the evidence chain. Preserving unknown values, version history, selection, and observation opportunity can make a modest record scientifically stronger. A large collection becomes a research resource when those distinctions survive from clinical source documentation to the final analysis.

References

Cao, Mingde, et al. "APKASS 2024 consensus statement on anterior cruciate ligament reconstruction, Part II: Management of anterior cruciate ligament revision in adults." *Asia-Pacific Journal of Sports Medicine, Arthroscopy, Rehabilitation and Technology* 45 (2026): 74-81.

Wang, J., Xu, G., Tong, H., Xu, G., & Chen, Y. (2026, May). Fractal Generative Distillation. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 20521-20525). IEEE.

Wilkinson, Mark D., et al. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data*, vol. 3, 2016, article 160018. https://doi.org/10.1038/sdata.2016.18.

von Elm, Erik, et al. "The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) Statement: Guidelines for Reporting Observational Studies." *PLoS Medicine*, vol. 4, no. 10, 2007, pp. e296. https://doi.org/10.1371/journal.pmed.0040296.

Page, Matthew J, et al. "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews." *BMJ*, 2021, pp. n71. https://doi.org/10.1136/bmj.n71.

Sterne, Jonathan AC, et al. "ROBINS-I: a tool for assessing risk of bias in non-randomised studies of interventions." *BMJ*, 2016, pp. i4919. https://doi.org/10.1136/bmj.i4919.

Little, Roderick J. A., and Donald B. Rubin. *Statistical Analysis with Missing Data*. Wiley, 2002. https://doi.org/10.1002/9781119013563.

Greenland, Sander, Judea Pearl, and James M. Robins. “Causal Diagrams for Epidemiologic Research.” *Epidemiology*, vol. 10, no. 1, 1999, pp. 37-48. https://doi.org/10.1097/00001648-199901000-00008.

Cameron, A. Colin, and Douglas L. Miller. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources*, vol. 50, no. 2, 2015, pp. 317-372. https://doi.org/10.3368/jhr.50.2.317.

White, Halbert. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, vol. 48, no. 4, 1980, pp. 817-838. https://doi.org/10.2307/1912934.

Irrgang, James J., et al. "Development and Validation of the International Knee Documentation Committee Subjective Knee Form." *The American Journal of Sports Medicine*, vol. 29, no. 5, 2001, pp. 600-613. https://doi.org/10.1177/03635465010290051301.