


The Problem-Oriented Medical Record, 60 Years Later

Can AI finally deliver what Larry Weed imagined?



Echo Health AI

Perspectives on connected care

October 2026 · Australian focus

© 2026 Visualisation Hub Pty Ltd trading as Echo Health AI. All rights reserved.

EXECUTIVE SUMMARY

The data may be structured. The clinical story often isn't.

In 1968 Lawrence Weed argued that the medical record should be organised around the problems clinicians are trying to understand and manage. Medicine adopted much of that idea: problem lists, SOAP notes and problem-oriented assessment and planning are now everywhere, and modern record systems include sophisticated problem-oriented tools.

Yet the fuller model has proved hard to sustain. A clinician can usually find a patient's diagnoses, medicines, results and notes. What the record rarely holds is the relationship between them: which result informed which problem, why a medicine was started or stopped, and what happened next.

Australia has confronted this for fifty years. Its College built a national problem-oriented record in the 1970s, and its BEACH program showed what problem-linked data make possible. When BEACH closed in 2016, its authors warned that electronic records still lacked the structure that links management to the problem managed.

The obstacle was not that the idea was wrong. The structure was expensive in clinician attention to create and maintain. This paper asks whether AI changes that cost: not by extracting facts, which is the easier part, but by proposing the **relationships** between facts and problems for a clinician to confirm as care happens.

That is a hypothesis, not a finding. We set out the evidence, the risks and what would have to be true.

AI may both depend on the problem-oriented record and help to maintain it.

IN THIS PAPER

THE IDEA

1. An idea about reasoning, not a form	03
2. What medicine adopted, and what stayed hard	04

THE AUSTRALIAN RECORD

3. Australia engaged early	05
4. Problem-oriented, problem-linked data	06
5. BEACH: what linkage made possible	07

THE GAP

6. Structured data, unstructured story	08
7. The record by data type, and the condition story	09
8. Three levels of the clinical record	10

COMPUTING THE PROBLEM

9. Reconstructing the story afterwards	11
10. Extraction is not relationship creation	12
11. The AI paradox	13

THE OPPORTUNITY

12. From BEACH to the routine record	14
13. What Echo One is building	15
14. Designing proposals that stay proposals	16
15. Unified clinical storytelling	17
16. What must be true, and what we do not yet know	18
17. Conclusion: the record that guides	19

SOURCES

References	20
------------	----

ABOUT THIS PAPER

A position paper on clinical-record design. It does not report outcomes of any Echo One deployment.

Echo One features are described only as shown in its published product preview and screenshots. Larry Cox is an illustrative patient.

03 / PERSPECTIVES ON CONNECTED CARE

1. An idea about reasoning, not a form

Larry Weed's contribution is often remembered as the SOAP note. His argument was larger: the record should be organised around the patient's problems, so that it could guide reasoning, support continuity and make care auditable.

Weed set out the problem-oriented medical record (POMR) in a two-part *New England Journal of Medicine* article in 1968, "Medical records that guide and teach", and in a book the following year. [1, 2, 3] The record had four parts: a defined database of what is known about the patient, a complete problem list, initial plans for each problem, and progress notes organised by problem, written in the structure that became known as SOAP. [4]



Figure 1. The problem-oriented record as Weed described it. The problem is the organising principle of the whole record, not a heading inside one note.

For Weed, the record was a clinical instrument. He called it the doctor's "scientific notebook", [4] and later wrote that a problem-oriented record makes it possible "to trace not only what was done, but why it was done." [5] The POMR, he told an interviewer in 2009, "cannot change the multiplicity of problems that physicians face... But the POMR enables a highly organized approach to that complexity." [6]

He also saw early that paper would not be enough. At the University of Vermont his team built PROMIS, the Problem-Oriented Medical Information System, described as the first clinical information system to use a touch-screen terminal. [4, 7] In a 1970s demonstration it preserved "the logic of medical action" and provided feedback loops for correction. [8] It also met resistance: in a contemporary study, physicians found it more time-consuming than paper and felt coerced rather than attracted into using it. [9]

A lesson from PROMIS. Structure that asks clinicians for more time than it returns is resisted, however sound the idea. Any new attempt has to be judged by the attention it costs as well as the structure it creates.

04 / PERSPECTIVES ON CONNECTED CARE

2. What medicine adopted, and what stayed hard

The POMR did not fail. It changed the medical record. Problem lists are a standard part of electronic records; SOAP structures clinical notes in many settings; problem-oriented assessment and planning is taught to most medical students. Modern record systems offer problem-oriented charting, in which assessments and plans are written under each problem rather than in one note. Where it has been introduced, problem lists are used more: in one intensive care study, new problems recorded per encounter rose from 3.6 to 6.2 after problem-based charting began. [10] One US health system built a problem-oriented documentation toolkit within a commercial record that is used in more than 92,000 visits a month. [11]

But the fuller model, a record that is accurate, current and richly linked to problems over time, has remained hard to keep. One editorial judgement is that the system was never fully implemented and that SOAP notes were what survived. [12] The measured evidence is more specific:

- **Completeness varies widely.** Across ten organisations on three continents, diabetes appeared on the problem list for between 60.2% and 99.4% of patients whose results indicated it. [13]
- **Lists accumulate.** At one US centre, for every 100 encounters problems were added 43.7 times and deleted 2.1 times; 10.5% of patients had coded duplicates. [14]
- **Ownership is unclear.** Clinicians disagree about who maintains the list, what belongs on it and when it should be updated. [15, 16]
- **The problem persists.** A 2026 scoping review of 103 studies found 57% reported problem lists that were incomplete, outdated, duplicated or bypassed. [17]

These are not failures of clinicians. Researchers describe them as sociotechnical: time pressure, workflow, usability, terminology, governance and shared care all contribute. [17, 18] A problem list is shared work with no single owner, and clinical relationships are many-to-many: one medicine can treat two problems, one result can inform three.

THE POSITION OF THIS PAPER

The obstacle to the full problem-oriented record may have been less that the idea was wrong and more that its structure was expensive, in clinician attention, to create and maintain.

Fairness. Commercial records invest heavily here. Epic, for example, offers problem-oriented charting and AI features that summarise a patient's history and "the story behind a patient's medications" (vendor description). [19] The question this paper asks is not which interface is better. It is who, or what, maintains the clinical relationships.

05 / PERSPECTIVES ON CONNECTED CARE

3. Australia engaged early

Australian general practice did not just read about Weed. In 1972 he gave a keynote address in Melbourne at the Fifth World Conference on General Practice, the meeting at which WONCA was inaugurated, and drew overflow audiences. [20, 21] The College's Practice Management Committee then resolved to develop a problem-oriented record for Australian general practice, building on local experiments that had already begun. A prototype was trialled in Melbourne practices, debated at a national workshop in Canberra in 1974, redesigned, and launched as the RACGP Health Record in 1975. The College adopted standards for computerised primary-care records in 1980, revised in 1988. [20]

Hospital medicine was asking the same question. In 1976 R. P. Lambert and Ian Mackay reported a problem-orientated outpatient record that had run for two and a half years and 5,000 attendances. Their patients averaged 3.1 concurrently treated problems, and 18 doctors who used it rated it favourably. [22] Their conclusion speaks directly to this paper:

Such systems “may help to resolve the present paradox of the clinical chart becoming decreasingly useful with increasing complexity of illness.”

Lambert & Mackay, *Medical Journal of Australia*, 1976. [22]




Figure 2. Selected Australian milestones. Each is verified in the sources cited; it is not a complete history.

The College's own history estimates that about 40% of practices used its record by 1981 and about 60% of primary-care doctors by 1985. [20] These are a participant's estimates, not measured figures, but the direction is clear: problem orientation was mainstream in Australian general practice before most practices owned a computer.

06 / PERSPECTIVES ON CONNECTED CARE

4. Problem-oriented, problem-linked data

A problem list says what is wrong. Linkage says what was done about each thing that is wrong. Australian primary-care researchers understood the difference early, and for decades collected data built on it.

From 1969, a national survey asked general practitioners to record encounters on structured forms, later classifying the problems managed with the international classification WONCA was developing. [23] In 2001, researchers at the University of Sydney described a form for collecting what they called, in the title of their paper, “problem oriented, problem linked data”. Used across 10,937 encounters and completed for 87% of eligible patients, it linked each service to its problem “so that it is possible to address questions about patterns of service provision”, including clinical costing and compliance with guidelines. [27]

THE PHRASE THAT MATTERS

Problem-oriented, problem-linked data. Echo One is not proposing a new idea. Australian researchers were explicitly pursuing problem-linked information decades ago. What is new is the possibility that software could maintain those links during routine care without asking the clinician to create every one by hand.

Computerisation did not, by itself, deliver the links. When general practice moved onto clinical software, records gained structured fields for diagnoses, prescriptions and requests, but the relationship between them remained optional. In a 2005 national survey, only 65% of GPs who used clinical software regularly recorded the reason for prescribing a medicine electronically. [26] Without the reason, a prescription is an event. With it, the prescription becomes part of a problem’s history.

Coding helped. ICPC-2 PLUS, developed from the free-text language Australian GPs actually use, gave clinical systems a terminology closer to practice than classification alone. [25] Codes, though, answer “what is this?” They do not answer “what is this *for*?” That second question is the one the problem-oriented record was designed to answer, and it is the subject of the next section.

Evidence boundary. The 2001 instrument was a research collection, not a clinical record. The 65% figure is self-reported and predates current software. Both establish what researchers valued, not what current records achieve.

07 / PERSPECTIVES ON CONNECTED CARE

5. BEACH: what linkage made possible

Bettering the Evaluation and Care of Health (BEACH) ran at the University of Sydney from April 1998 to March 2016. Each year about 1,000 randomly sampled general practitioners each recorded 100 consecutive encounters on a structured paper form, and by its close BEACH held almost 1.8 million encounter records from 10,789 individual GPs. [28] Each encounter recorded up to four problems managed; problems were coded with ICPC-2 PLUS and medicines with the ATC classification.

BEACH did not merely count medicines, referrals and tests. It linked each management action to the problem it managed. Its final report is explicit about how: “The structure itself forces linkage of actions to the problem being managed.” [28]




Figure 3. Problem-linked management in BEACH. Every action was recorded against the problem it addressed, so data could show management per problem, not just activity per visit.

That linkage made BEACH data unusually useful. It allowed researchers to report antibiotic prescribing for each respiratory condition against guideline rates, [29] and governments and industry commissioned analyses for medication reviews and Pharmaceutical Benefits Scheme listing applications. BEACH produced about 178 refereed articles and more than 1,000 bespoke reports. [28, 30] As its final report notes, prescribing data alone hold no record of the problem being managed, and morbidity cannot be reliably inferred from medication. [28]

Its authors also saw the gap in routine records. In 2013 they wrote that “most electronic systems lack the structured problem orientation that provides reliable linkage of management actions to a patient problem”, citing Weed. [31] In 2015 they put it more briefly: “most EHRs lack a problem-oriented structure that enforces this linkage.” [30]

BEACH closed for lack of direct government support and falling industry support. [28] What was lost was not only a dataset but a method: the BEACH team warned that its “mandatory recorded relationships between problem management and the management actions taken” were being ignored as research turned to extracting data from routine records. [28]

What if the structure that made BEACH valuable could exist natively in routine care?

08 / PERSPECTIVES ON CONNECTED CARE

6. Structured data, unstructured story

Australian primary care now holds large longitudinal datasets. Patron, at the University of Melbourne, holds de-identified records for more than 3.5 million patients from 129 practices. [32] MedicineInsight, now held by the Australian Commission on Safety and Quality in Health Care, draws on hundreds of practices. [33] Researchers using them have measured, carefully, how much of the clinical picture reaches structured fields.

- **Diagnoses are often uncoded.** In 84 Victorian practices, 26.4% of diagnoses were recorded only as free text. Relying on codes alone would undercount asthma by 36.7% and chronic kidney disease by 22.6%. [34]
- **Conditions go unrecorded.** Of patients with pathology consistent with chronic kidney disease, 23.2% had it recorded in one national dataset, [35] and 28.1% in another. [36]
- **Reasons are missing or narrative.** Analyses of prescribing note that the reason for a prescription is often absent or held in free text and progress notes. [37, 38]
- **Structure is recovered afterwards.** To code free-text diagnoses across practices, the POLAR program ran language processing over some 57,000 records with four rounds of expert correction. [39]

FREE TEXT IS NOT THE PROBLEM

The clinician may know it. The note may contain it. The structured record may not. Narrative carries clinical meaning that codes cannot; the opportunity is to keep that richness while also creating structure.

There is a further distinction. A dataset may contain a patient's diabetes diagnosis, every HbA1c, each weight, metformin and empagliflozin, and years of notes, and still not say *what happened to the diabetes*: what was done, why, and what followed. Longitudinal data is not necessarily a longitudinal clinical story. POLAR asks whether computational methods can structure narrative afterwards. The question for routine care is different: can the structure be created as the information is created, with the clinician who knows the story?

Fairness. These datasets were not designed to hold clinical relationships, and their custodians document their limits openly. The figures describe what routine records contain, not the quality of the research built on them.

09 / PERSPECTIVES ON CONNECTED CARE

7. The record by data type, and the condition story

Most electronic records organise information by kind: problems in one place, medicines in another, results, observations, orders and notes elsewhere. That arrangement is reliable and familiar. It also leaves one task to the clinician, every time: the semantic join.




Figure 4. Two views of the same information. Illustrative, based on the demonstration record of Larry Cox, an illustrative patient. Neither view replaces the other; the right-hand view depends on relationships someone has to record.

The questions a clinician brings to a chronic condition are relational. Which result relates to which condition? Why was this medicine started, and why stopped? Which observation changed management, and what happened after treatment changed? A record organised by data type can answer each question, but only after a person reads across it and reconstructs the answer, often from free text.

Weed's problem-oriented record put those relationships at the centre. The difficulty has always been who records them, and at what cost. Each link is a small act of clinical judgement: this weight belongs with the diabetes; this dose change was because of a hypoglycaemic episode. Multiplied across every visit, those small acts are the maintenance burden that problem lists have struggled under.

10 / PERSPECTIVES ON CONNECTED CARE

8. Three levels of the clinical record

It helps to separate three questions a record can answer. They are complementary. A problem-oriented longitudinal record does not replace the encounter note or the patient summary; it adds a view neither was designed to give.




Figure 5. Three levels of the clinical record. Most records are strong at levels 1 and 2. Level 3 depends on relationships between items, not only on the items themselves.

Modern records do levels 1 and 2 well, and vendors are investing in level 3. Epic's patient summary column keeps problems, medicines and allergies visible across screens, [40] and its AI features summarise a patient's history, describe the story behind a patient's medicines, collect cancer staging from notes and turn recommendations in results into follow-ups (vendor descriptions). [19] Epic reports Insights summaries at more than 16 million uses a month. [41]

Level 3 views help clinicians. In a simulated retrieval study, a problem-oriented view that showed the results and medicines relevant to a chosen problem cut task time from 205 to 173 seconds and errors from 7.7% to 3.4%. [42] In a Dutch randomised crossover study, prescribers given a correctly structured problem list answered 56.3% of clinical questions correctly, against 33.5% with an uncurated list. [43]

A summary generated when a chart is opened and a relationship recorded during care are different things. The first reconstructs the story for a reader, on demand. The second makes the relationship part of the record, with provenance, so that the next clinician, the next query and the next decision-support rule can rely on it. Both have value. This paper is about the second.

11 / PERSPECTIVES ON CONNECTED CARE

9. Reconstructing the story afterwards

Researchers have tried for twenty years to compute problem-oriented views from records that were not written that way. Their work is the foundation for anything new, and it shows both what is possible and where it falls short.

- **Language processing that proposes problems.** At Utah, a system read clinical notes and proposed problems for physicians to add. In a randomised trial in intensive care, problem-list sensitivity rose from about 9% to 41%. [44]
- **Rules that infer problems.** At Brigham and Women's Hospital, rules over medicines, results and billing data inferred 17 conditions; in a randomised trial, clinicians accepted 41.1% of the suggested problems. [45, 46] A later multi-site trial added problems at 4.6 times the control rate but did not change quality measures. [47]
- **Linking medicines to problems.** A knowledge base built from clinicians' own links connected about a third of previously unlinked medicines to problems. [48]
- **Problem-oriented summaries.** IBM's ProSPER generated problem lists and disease-specific views; problem-list generation reached an F1 of 0.74 on intrinsic evaluation, with a clinical evaluation still to come. [49, 50]
- **Problem-centred timelines.** Columbia's HARVEST combined a timeline with problems extracted from notes and was deployed for clinicians at NewYork-Presbyterian from 2013. Clinicians performed as well as with the standard system, and its authors noted that negated terms such as "no history of asthma" were still indexed. [51]

Retrospective reconstruction

Read the record after care, infer problems and links, present a view. Works on history that exists; inherits its gaps; inferences may be wrong and are rarely confirmed.

≠

Structuring during care

Propose the relationship as information is created, while the clinician who knows the reason is present to confirm, change or reject it.

Figure 6. Two ways to obtain a problem-oriented record. They are complementary: history arrives unlinked and must be reconstructed; new care can be linked as it happens.

Two lessons recur. Machine proposals find real gaps: in one pilot, a generated problem list identified at least one important problem physicians had missed in 89% of assessments. [50] And clinician review is a genuine filter, not a formality: well under half of suggestions are accepted in most trials. [46, 47] Completeness is also not the same as better care. More problems on a list did not, by itself, move quality measures. [47]

12 / PERSPECTIVES ON CONNECTED CARE

10. Extraction is not relationship creation

Much clinical AI is extraction: turning words into structured facts. Extraction matters, but it is the easier half. The harder half, and the half the problem-oriented record depends on, is relationship creation: stating which problem a fact informs, or which problem an action addresses.



Figure 7. The distinction this paper turns on. A weight is easy to extract. Knowing that this weight belongs with this patient's diabetes, and recording that link with its provenance, is what builds the condition's story.

The evidence suggests why the distinction matters. Language models summarise well: in one study, physicians rated adapted models' summaries, including problem lists from progress notes, as equivalent or better than experts' summaries in most cases. [52] They code exactly far less well: given code descriptions, GPT-4 returned the exact ICD-10-CM code 33.9% of the time. [53] In an international challenge to link clinical text to SNOMED CT, the organisers judged that no submission could yet be considered usable for clinical decision support. [54]

A relationship proposal asks a narrower question than open coding: given this patient's recorded problems, and this item's own words, which problem does it belong with, if any? It is still fallible. Relationships are many-to-many, context matters, and some links are clinical judgements that only a clinician can make. That is why a proposed relationship should remain a proposal, shown with the words it rests on, until a clinician accepts it.

The value is not that AI can extract a weight. It is that AI may help create and maintain the relationships between clinical facts and the problems they inform.

13 / PERSPECTIVES ON CONNECTED CARE

11. The AI paradox

The 2026 scoping review of electronic problem lists in primary care states the tension precisely: “Increasing use of PL for CDS, interoperability, and AI-enabled tools amplifies the clinical risks associated with poor data quality.” [17] As problem lists feed more decision support, data exchange and AI, their errors travel further.

AI systems need context. Language models’ clinical performance is sensitive to the quantity and order of the information they are given, [55] and a curated problem list changed human prescribers’ decisions in a controlled trial. [43] Problem-oriented structure may therefore matter more in the AI era, not less.




Figure 8. A possible feedback loop. Better maintained structure gives AI better context; AI-assisted maintenance gives clinicians better structure. The loop only works if every proposal is reviewable and wrong links are corrected, not propagated.

The same review is clear about the conditions. Automation and AI “can assist with PL curation and reconciliation but are most effective when embedded within clinician-led workflows and supported by clear governance”, and formal indicators such as marking a list “reviewed” did not reliably correspond to meaningful updates. [17]

The loop can also run the wrong way. A model given a stale problem list may propose links to the wrong problem; a clinician under time pressure may accept a plausible proposal; automation bias is well documented in clinical AI. [17] The design question is therefore not whether AI should touch the problem-oriented record, but under what controls.

Evidence boundary. No peer-reviewed study we found has yet evaluated AI-maintained problem relationships in routine longitudinal primary care against a clinician standard. The loop is a hypothesis to test.

14 / PERSPECTIVES ON CONNECTED CARE

12. From BEACH to the routine record

BEACH obtained problem-linked data by asking GPs to complete a structured form, outside their own records, for 100 encounters a year. That was its strength and its limit: the structure was excellent and the collection was deliberate, sampled and separate from care.

Modern AI suggests a different route. The consultation is already conducted in clinical language. If software can propose structured facts from that language, and propose the problem each fact or action belongs with, the clinician's task changes from creating every link to reviewing proposed ones.




Figure 9. The bridge. A design opportunity, not a demonstrated result.

If this works, routine records could hold some of the problem-linked structure that BEACH collected by design, for every patient rather than a sample, and the clinician who knows the reason would be the one confirming it.

Evidence boundary. It has not been shown that AI-assisted linking can produce BEACH-quality data in routine care. BEACH's value came from standard definitions, trained coders and random sampling, none of which a consultation tool provides by itself. Whether proposed links are accurate, whether review is meaningful under time pressure, and whether linked data are comparable across practices are open questions.

15 / PERSPECTIVES ON CONNECTED CARE

13. What Echo One is building

Echo One is being built around problem-linked records. Its condition history brings together, on one time axis, what has been recorded against a problem: each medicine's doses and changes with the recorded reason, the measurements linked to the problem, referrals and the follow-up still open. Gaps with no readings are shown as gaps.





Figure 10. Condition history in Echo One. Larry Cox's type 2 diabetes: medicines for this problem with dose changes and recorded reasons (hypoglycaemia; recurrent genital mycotic infection), and HbA1c and eGFR linked to it. Larry Cox is an illustrative patient; demonstration data.



A change opened: the recorded reason beside the nearest readings either side, which the product labels “for comparison only; nothing here says the change caused a difference.”

PRODUCT PREVIEW

In a published product preview, a diabetes review is transcribed and Echo One proposes what to record, each proposal shown with the exact words it came from. A weight the GP said belonged with the diabetes is prefilled to link to that problem; the clinician can change or clear the link, and nothing is linked until it is recorded. [56]

Status. The condition history shows relationships clinicians have recorded. Consultation proposals are shown as a product preview; the consultation pipeline is not switched on for use. No outcome or accuracy claim is made.

16 / PERSPECTIVES ON CONNECTED CARE

14. Designing proposals that stay proposals

Imported history arrives with no links between items and problems. To help fill the history, Echo One's design includes a linking agent that proposes where an existing medicine, test, consult, document or referral belongs among a patient's recorded problems. Its design rules are the paper's argument in practice, and they apply to proposals from any source. [57]

A proposal is not a link

Proposals are stored separately. A link is made only when a clinician accepts, one proposal at a time, through the same path and checks as a link made by hand. There is no accept-all.

Every proposal shows its evidence

A model's proposal must quote the item's own text exactly. If any reference or quote cannot be verified, the whole run is refused.

It files; it does not judge

The proposal says where an item belongs, with a general reason ("metformin is used to treat type 2 diabetes"), never an assessment of the patient.

Deterministic first

Where a person stated the problem, such as a referral's reason, it is matched to the problem list exactly or not at all. A bounded model call is used only for what remains.

Started by a person

A clinician starts it for one patient. There is no background sweep across the practice.

Confidential stays out

Confidential conditions are not candidates and restricted notes are not read. A rejected link is kept, so it is not proposed again.

These rules follow the pattern Echo One applies to all model output: the model proposes, deterministic checks verify, the clinician decides, and provenance stays with the record. [58] They respond directly to the evidence in sections 9 to 11: proposals find real gaps, clinician review is a real filter, exact coding is unreliable, and wrong links must not silently propagate.

Status and boundary. The linking agent is part of Echo One's current design and development; it is not described here as in routine use. Its accuracy, the rate at which clinicians accept or reject proposals, and the effect on clinician time have not been evaluated, and we make no claim about them.

17 / PERSPECTIVES ON CONNECTED CARE

15. Unified clinical storytelling

We use the phrase **unified clinical storytelling** for a conceptual destination, not a product feature. We define it as a longitudinal representation of a patient's care in which each problem carries its own history, assembled from relationships that clinicians have confirmed, so that what happened, why and what followed can be read without reconstructing it from separate lists.

Unified, because the story draws on every source in the record and across the care team: the consultation, results, letters, referrals and medicines. Clinical, because its relationships are clinical judgements, made or confirmed by clinicians. Storytelling, because medicine is understood over time, and a record that holds the sequence and its reasons supports the reasoning Weed wanted the record to teach.

01 Relationships are first-class

An item's link to a problem is recorded, with who made it and why, not inferred each time it is read.

02 The clinician holds authority

Software may propose; a clinician confirms, changes or rejects every link.

03 Provenance travels

Every proposed fact or link keeps the words or source it came from.

04 Comparison, not causation

Readings before and after a change are shown side by side, never presented as an effect.

05 Gaps stay visible

Missing readings and unlinked items are shown as missing, not smoothed over.

06 Many-to-many is normal

One medicine may serve two problems; one result may inform several.

07 Complementary views

The condition story sits beside the encounter note and the patient summary, not instead of them.

08 Narrative is kept

Structure is added to clinical language, never in place of it.

Each principle answers a known failure: lists without owners, inferences without evidence, summaries that imply cause, and structure that costs more attention than it returns.

18 / PERSPECTIVES ON CONNECTED CARE

16. What must be true, and what we do not yet know

The opportunity carries real risks. For AI-assisted relationship creation to improve the record rather than degrade it, several things must be true.

- **Proposals must be accurate enough to be worth reviewing.** Too many wrong proposals waste attention; too few find nothing.
- **Review must be meaningful.** Automation bias and time pressure can turn confirmation into a click. Designs should make rejection as easy as acceptance and avoid bulk acceptance.
- **The problem list itself must be maintained.** Links to a stale, duplicated or incorrect problem spread the error.
- **Ownership must be clear in shared care,** across general practice, specialists and hospitals.
- **Terminology must be shared** if linked data are to be comparable across practices and usable in exchange, as BEACH's standard coding was.
- **Privacy must be respected,** with confidential material excluded from automated processing.
- **Regulation must be considered.** Filing an item under a problem is different from assessing a patient; designs should keep that boundary explicit.

What we do not yet know

We found no published evaluation of AI-proposed problem relationships created during routine primary-care consultations. We do not know how accurate such proposals are across conditions and settings, how often clinicians accept wrong ones, whether the work saves or costs time, whether linked data would be comparable to BEACH's, or whether better-linked records improve care. The trials of problem-list suggestions are a caution: more complete lists did not by themselves improve quality measures. [47]

These are answerable questions. They need prospective evaluation in real practices, with clinician review recorded, errors analysed and patient outcomes measured, and Australian general practice, with its history of problem-linked research, is well placed to answer them.

19 / PERSPECTIVES ON CONNECTED CARE

17. Conclusion: the record that guides

Sixty years ago Larry Weed asked for medical records that guide and teach. Medicine took his problem list and his SOAP note and made them universal. The rest, a record whose every fact and action is connected to the problem it serves, proved costly to sustain, not because clinicians did not value it but because the connections cost attention to make.

Australia understood this early. Its College built a national problem-oriented record within three years of hearing Weed speak. Its clinicians described the chart becoming less useful as illness becomes more complex. Its researchers collected problem-linked data for half a century, and BEACH showed what that structure made possible before it closed.

AI does not settle the question. It may change the economics. Extraction is becoming cheap; relationship creation is the harder and more valuable task. If software can propose the relationships, show the evidence for each, and leave every decision with the clinician, the problem-oriented record may become something routine care can maintain rather than something research has to collect.

QUESTIONS FOR THE FIELD

- How should a proposed clinical relationship be represented, reviewed and audited?
- What standard of accuracy should a relationship proposal meet before it is shown?
- How can problem-linked data from routine care be made comparable across practices?
- What would a BEACH-like national dataset look like if its linkage came from routine records?

The data may be structured. The clinical story often isn't. The next generation of records should be judged by whether they can tell it.

20 / PERSPECTIVES ON CONNECTED CARE

References

- [1] Weed LL. Medical records that guide and teach. *N Engl J Med* 1968;278(11):593–600. doi:10.1056/NEJM196803142781105
- [2] Weed LL. Medical records that guide and teach (concluded). *N Engl J Med* 1968;278(12):652–657. doi:10.1056/NEJM196803212781204
- [3] Weed LL. *Medical Records, Medical Education, and Patient Care: The Problem-Oriented Record as a Basic Tool*. Cleveland: Press of Case Western Reserve University; 1969.
- [4] Wright A, Sittig DF, McGowan J, Ash JS, Weed LL. Bringing science to medicine: an interview with Larry Weed, inventor of the problem-oriented medical record. *J Am Med Inform Assoc* 2014;21(6):964–968. doi:10.1136/amiajnl-2014-002776
- [5] Weed LL. New premises and new tools for medical care and medical education. *Methods Inf Med* 1989;28(4):207–214. PMID 2622370.
- [6] Jacobs L. Interview with Lawrence Weed, MD: the father of the problem-oriented medical record looks ahead. *Perm J* 2009;13(3):84–89. doi:10.7812/TPP-09-068
- [7] Schultz JR. A history of the PROMIS technology: an effective human interface. In: *Proceedings of the ACM Conference on the History of Personal Workstations*; 1986:159–182. doi:10.1145/12178.12188
- [8] Gilroy G, Ellinoy BJ, Nelson GE, Cantrill SV. Integration of pharmacy into the computerized problem-oriented medical information system (PROMIS): a demonstration project. *Am J Hosp Pharm* 1977;34(2):155–162. PMID 842545.
- [9] Fischer PJ, Stratmann WC, Lundsgaarde HP, Steele DJ. User reaction to PROMIS: issues related to acceptability of medical innovations. In: Anderson JG, Jay SJ, eds. *Use and Impact of Computers in Clinical Medicine*. Computers and Medicine series. Springer; 1980:284–301. doi:10.1007/978-1-4613-8674-2_18
- [10] Li RC, Garg T, Cun T, et al. Impact of problem-based charting on the utilization and accuracy of the electronic problem list. *J Am Med Inform Assoc* 2018;25(5):548–554. doi:10.1093/jamia/ocx154
- [11] Altman RL, Lin CT, Earnest M. Problem-oriented documentation: design and widespread adoption of a novel toolkit in a commercial EHR. *JAMIA Open* 2023;6(1):o0ad005. doi:10.1093/jamiaopen/o0ad005
- [12] Dunea G. Problems with medical records. *Hektoen International*, Summer 2015. Editorial essay. hekint.org
- [13] Wright A, McCoy AB, Hickman TT, et al. Problem list completeness in electronic health records: a multi-site study and assessment of success factors. *Int J Med Inform* 2015;84(10):784–790. doi:10.1016/j.ijmedinf.2015.06.011
- [14] Brown AR, McCoy AB, Wright A, Nelson SD. Decluttering the problem list in electronic health records. *Am J Health Syst Pharm* 2022;79(Suppl 1):S8–S12. doi:10.1093/ajhp/zxab381
- [15] Wright A, Maloney FL, Feblowitz JC. Clinician attitudes toward and use of electronic problem lists: a thematic analysis. *BMC Med Inform Decis Mak* 2011;11:36. doi:10.1186/1472-6947-11-36
- [16] Holmes C, Brown M, St Hilaire D, Wright A. Healthcare provider attitudes towards the problem list in an electronic health record: a mixed-methods qualitative study. *BMC Med Inform Decis Mak* 2012;12:127. doi:10.1186/1472-6947-12-127
- [17] Nair R, Bekker I, Singer A, Lau F. Current state of electronic problems lists in primary care: a rapid scoping review. *Fam Pract* 2026;43(4). doi:10.1093/fampra/cmago36
- [18] Simons SMJ, Cillessen FHJM, Hazelzet JA. Determinants of a successful problem list to support the implementation of the problem-oriented medical record according to recent literature. *BMC Med Inform Decis Mak* 2016;16:102. doi:10.1186/s12911-016-0341-0
- [19] Epic Systems. Art for clinicians. Vendor description, accessed 1 October 2026. epic.com/software/art
- [20] Fry F. *The RACGP medical record system: a short account of the development of the record*. Royal Australian College of General Practitioners, history series. racgp.org.au
- [21] WONCA. A brief history: the first twenty years. World Organization of Family Doctors. globalfamilydoctor.com
- [22] Lambert RP, Mackay IR. A problem-orientated outpatient record. *Med J Aust* 1976;2(10):368–372. doi:10.5694/j.1326-5377.1976.tb130257.x
- [23] Britt H, Miller GC. Measuring general practice activity in Australia: a brief history. *Aust Fam Physician* 2017;46(5). racgp.org.au/afp
- [24] RACGP. Dr John North OAM: biography. RACGP history. racgp.org.au
- [25] Britt H. A new coding tool for computerised clinical systems in primary care: ICPC plus. *Aust Fam Physician* 1997;26 Suppl 2:S79–S82. PMID 9254947.
- [26] McInnes DK, Saltman DC, Kidd MR. General practitioners' use of computers for prescribing and electronic health records: results from a national survey. *Med J Aust* 2006;185(2):88–91. doi:10.5694/j.1326-5377.2006.tb00479.x
- [27] Bolton P, Prior G, Mira M, Karr M, Usher H. An instrument for collecting problem oriented, problem linked data in primary medical care. *Aust Fam Physician* 2001;30(12):1190–1194. PMID 11838403.
- [28] Britt H, Miller GC, Henderson J, et al. *General practice activity in Australia 2015–16*. General practice series no. 40. Sydney: Sydney University Press; 2016. hdl.handle.net/2123/15514
- [29] McCullough AR, Pollack AJ, Plejdrup Hansen M, Glasziou PP, et al. Antibiotics for acute respiratory infections in general practice: comparison of prescribing rates with guideline recommendations. *Med J Aust* 2017;207(2):65–69. doi:10.5694/mja16.01042
- [30] Britt H, Miller G. BEACH program update. *Aust Fam Physician* 2015;44(6):411–414. PMID 26209994.
- [31] Britt HC, Miller GC. The Bettering the Evaluation and Care of Health (BEACH) program: where to from here? *Med J Aust* 2013;198(3):125–126. doi:10.5694/mja13.10012
- [32] Manski-Nankervis JA, Canaway R, Chidgey C, et al. Data resource profile: Primary Care Audit, Teaching and Research Open Network (Patron). *Int J Epidemiol* 2024;53(1):dyae002. doi:10.1093/ije/dyae002

REFERENCES, CONTINUED

- [33] Daniels B, Havard A, Myton R, Lee C, Chidwick K. Evaluating the accuracy of data extracted from electronic health records into MedicinesInsight. *Int J Popul Data Sci* 2022;7(1):1713. doi:10.23889/ijpds.v7i1.1713
- [34] Canaway R, Chidgey C, Hallinan CM, Capurro D, Boyle DI. Undercounting diagnoses in Australian general practice: a data quality study with implications for population health reporting. *BMC Med Inform Decis Mak* 2024;24:155. doi:10.1186/s12911-024-02560-w
- [35] Kitsos A, Peterson GM, Jose MD, Khanam MA, et al. Variation in documenting diagnosable chronic kidney disease in general medical practice: implications for quality improvement and research. *J Prim Care Community Health* 2019;10. doi:10.1177/2150132719833298
- [36] Petzke D, Hallinan CM, Trevena J, Manski-Nankervis JA. Chronic kidney disease in Australian general practice using electronic medical record data. *BMC Nephrol* 2025;26:405. doi:10.1186/s12882-025-04345-3
- [37] Drug Utilisation Sub-Committee. Ocular lubricants: MedicinesInsight analysis. Public release document, Pharmaceutical Benefits Scheme, February 2021. pbs.gov.au
- [38] Gonzalez-Chica D, Begum M, Bernardo C, Hoon E, et al. Trends and patterns of benzodiazepines and Z-drugs prescriptions in Australian general practice: a national study (2011–2018). *Drug Alcohol Rev* 2023;42:427–438. doi:10.1111/dar.13561
- [39] Pearce C, McLeod A, Patrick J, et al. Coding and classifying GP data: the POLAR project. *BMJ Health Care Inform* 2019;26(1):e100009. doi:10.1136/bmjhci-2019-100009
- [40] Alberta Health Services, Connect Care. Storyboard (customer documentation of Epic). Accessed 1 October 2026. ehealth.connect-care.ca
- [41] Epic Systems. Epic AI Charting rolls out alongside an expanding set of built-in AI capabilities. 4 February 2026. Vendor statement. epic.com
- [42] Semanik MG, Kleinschmidt PC, Wright A, et al. Impact of a problem-oriented view on clinical data retrieval. *J Am Med Inform Assoc* 2021;28(5):899–906. doi:10.1093/jamia/ocaa332
- [43] Klappe ES, Heijmans J, Groen K, et al. Correctly structured problem lists lead to better and faster clinical decision-making in electronic health records compared to non-curated problem lists: a single-blinded crossover randomized controlled trial. *Int J Med Inform* 2023;180:105264. doi:10.1016/j.ijmedinf.2023.105264
- [44] Meystre SM, Haug PJ. Randomized controlled trial of an automated problem list with improved sensitivity. *Int J Med Inform* 2008;77(9):602–612. doi:10.1016/j.ijmedinf.2007.12.001
- [45] Wright A, Pang J, Feblowitz JC, et al. A method and knowledge base for automated inference of patient problems from structured data in an electronic medical record. *J Am Med Inform Assoc* 2011;18(6):859–867. doi:10.1136/amiajnl-2011-000121
- [46] Wright A, Pang J, Feblowitz JC, et al. Improving completeness of electronic problem lists through clinical decision support: a randomized, controlled trial. *J Am Med Inform Assoc* 2012;19(4):555–561. doi:10.1136/amiajnl-2011-000521
- [47] Wright A, Schreiber R, Bates DW, et al. A multi-site randomized trial of a clinical decision support intervention to improve problem list completeness. *J Am Med Inform Assoc* 2023;30(5):899–906. doi:10.1093/jamia/ocad020
- [48] McCoy AB, Wright A, Laxmisan A, et al. Development and evaluation of a crowdsourcing methodology for knowledge base construction: identifying relationships between clinical problems and medications. *J Am Med Inform Assoc* 2012;19(5):713–718. doi:10.1136/amiajnl-2012-000852
- [49] Liang JJ, Tsou CH, Dandala B, et al. Reducing physicians' cognitive load during chart review: a problem-oriented summary of the patient electronic record. *AMIA Annu Symp Proc* 2021;2021:763–772. PMID 35308927.
- [50] Devarakonda MV, Mehta N, Tsou CH, et al. Automated problem list generation and physicians perspective from a pilot study. *Int J Med Inform* 2017;105:121–129. doi:10.1016/j.ijmedinf.2017.05.015
- [51] Hirsch JS, Tanenbaum JS, Lipsky Gorman S, et al. HARVEST, a longitudinal patient record summarizer. *J Am Med Inform Assoc* 2015;22(2):263–274. doi:10.1136/amiajnl-2014-002945
- [52] Van Veen D, Van Uden C, Blankemeier L, et al. Adapted large language models can outperform medical experts in clinical text summarization. *Nat Med* 2024;30(4):1134–1142. doi:10.1038/s41591-024-02855-5
- [53] Soroush A, Glicksberg BS, Zimlichman E, et al. Large language models are poor medical coders: benchmarking of medical code querying. *NEJM AI* 2024;1(5). doi:10.1056/Aldbp2300040
- [54] Davidson R, Hardman W, Amit G, et al. SNOMED CT entity linking challenge. *J Am Med Inform Assoc* 2025;32(9):1397–1406. doi:10.1093/jamia/ocaf104
- [55] Hager P, Jungmann F, Holland R, et al. Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nat Med* 2024;30(9):2613–2622. doi:10.1038/s41591-024-03097-1
- [56] Echo Health AI. *From the consult to the record*. Product preview film, 88 seconds; narration approved 1 October 2026. Larry Cox is an illustrative patient. echo-one.health
- [57] Echo Health AI. Condition link proposals: design (Echo One product design, October 2026). Describes a design in development, not a deployed or evaluated service.
- [58] Echo Health AI. *After the Scribe: where the durable value in clinical AI lives*. 2026. echo-one.health/resources

PUBLICATION DETAILS

Published by Visualisation Hub Pty Ltd trading as Echo Health AI, October 2026. This is a position paper on clinical-record design, not a systematic review, clinical guidance or an evaluation of any product. Sources support the statements to which they are attached; the three-level model, the distinction between extraction and relationship creation, and unified clinical storytelling are Echo Health AI's argument. Vendor features are described as vendors describe them. Quotations are brief and attributed. Larry Cox is an illustrative patient and the screens show demonstration data; no patient information is used.

Companion papers at echo-one.health/resources: *The Cost of Standing Still*, *The Digital Front Door—and the Forgotten Back Door*, *After the Scribe and Informed Consent and Shared Benefit*. © 2026 Visualisation Hub Pty Ltd trading as Echo Health AI. All rights reserved.