

OVERSIGHT DESIGN

Approval gates: a specification for human decision authority in EMR automation

ABSTRACT

Clinical AI is governed almost everywhere by the assertion that a clinician reviews the output. This paper specifies what that review has to be. An approval gate is defined as a state transition the system cannot perform, carried out by a named person with authority, recorded as its own event, and defaulting to not proceeding. Criteria for placing gates, properties that make verification cheap, prohibited patterns, and instrumentation are set out, together with the cost the design imposes.

Type Original design specification **References** 27 **Reading time** 14 min

Last reviewed September 2026

1 The phrase that carries the weight

Nearly every deployment of clinical AI rests on one sentence: a clinician reviews the output before it reaches a patient. It is the answer given to regulators, to buyers, and to clinicians asking what happens when the model is wrong. It is also, as written, not a design. It names a person and a moment and specifies nothing about either.

The experimental literature is unkind to the unspecified version. In a randomized laboratory experiment, 120 final-year medical students prescribing with incorrect decision support made 86.6% more prescribing errors than with no decision support at all.¹ Twenty-seven radiologists lost 36.8 to 59.9 accuracy points on mammograms where the AI-supplied BI-RADS category was wrong, least among the most experienced readers.² Forty-four physicians, all of whom had received AI-literacy training, scored 14.0 percentage points lower on diagnostic reasoning when given deliberately flawed model output (95% CI −19.7 to −8.3).³ And in the one documentation pipeline measured end to end, review cut speech-recognition errors roughly twentyfold in volume while the clinically significant share of the survivors did not fall: 5.7%, then 8.9%, then 6.4%.⁴

A sibling paper appraises that evidence in detail. This one takes it as given and asks the engineering question: if the reviewing step is the control, what must the control be?

Two findings shape the answer. Over-reliance produces two distinct failures, omission, missing what the system did not flag, and commission, acting on what it wrongly flagged; a gate has to be built against both.⁵ And automation bias, across 40 studies drawn from nine databases, appeared in single tasks rather than multitasking ones, typically diagnosis, and where verification complexity was high.⁶ The reviewer's problem is not divided attention. It is that checking is expensive. The opposite failure is equally documented: acceptance of decision-support alerts fell about 30% per additional reminder in the same encounter, 87.9% of repeats of an already-overridden alert were overridden too, and override rates showed no association with general workload.⁷ Asking for confirmation more often does not buy more confirmation.

86.6%more prescribing errors when the suggestion was wrong¹**-14.0 pts**diagnostic reasoning, physicians trained in AI literacy³**30%**fall in alert acceptance per repeat in one encounter⁷

2 What an approval gate is

An approval gate is a state transition that (a) cannot be performed by the system, (b) is performed by a named person who holds the authority to perform it, (c) is recorded as its own event, and (d) has a default of not proceeding. All four clauses are load-bearing. A mechanism missing any one is something else, usually a notification.

The system cannot perform it. The transition has to be unavailable to the software under every condition, including the ones it is most confident about. A system capable of the transition that merely declines to has implemented a policy, and policies live in configuration, which drifts. In ChartVoyant the drafted note is not a note: it is a draft that waits, nothing signs itself, and the draft stays a draft until the confirmation step. What holds the gate shut is the state the entry is in, not a rule about who presses what.

A named person with the authority performs it. Attribution to a role, a queue, a workstation or a service account is not attribution. The signing event in ChartVoyant's audit trail carries the signing clinician and the timestamp, the minimum being an identity that resolves to one human being. Authority is a separate requirement, because that person must also be permitted to make this transition on this record. California's AB 3030 shows the legal form of the same idea: its disclosure duties for generative AI used in patient communications about clinical information do not apply where the communication is "read and reviewed by a human licensed or certified health care provider."⁸ The exemption turns on a licensed person, not on a review having nominally occurred.

It is recorded as its own event. A gate that leaves no trace distinct from the transition it authorizes cannot be examined afterwards, and a control nobody can examine is an assertion. The approval has to be a first-class record: who, when, against which version, with which outcome. ChartVoyant writes every action, including the AI's, to a log that can only be added to, each entry locked to the one before it, so an altered history fails verification. Drafting and signing are separate entries, and the drafting entry links to the transcript it came from.

The default is not proceeding. The transition does not occur on a timer, a queue drain, a session expiry, a retry or a deploy. If nobody approves, nothing happens, and the non-event is visible as a pending item. Default-proceed gates are the most common failure in this class, because they are indistinguishable from working gates on every day that nothing goes wrong.

3 Where gates belong, and where they do not

The central trade is not whether to have gates but how many. A gate on every action reproduces the alert-fatigue pattern, and the mechanism is specific: repetition, not busyness, destroys attention.⁷ Across 17 studies, drug-safety alert override rates ran from 49% to 96%, with the reviewers explicit that overriding is frequently justified and the error-producing conditions systemic rather than individual.⁹ Excessive warnings lead clinicians to ignore the vital ones.¹⁰ The foundational clinical-decision-support design paper put speed first and told implementers to ask for additional information only when they truly need it.¹¹ Too few gates produce the opposite failure: consequential actions leaving the practice with nobody answerable for them.

Four criteria decide the question. **Reversibility:** does undoing the action restore the prior state, or only append a correction? **Blast radius:** how many records, people and downstream systems does it touch? **Egress:** does it leave the practice, reaching a payer, a patient, another provider or a registry? **External accountability:** is someone answerable for it outside the practice, under a license, a signature, or a submission made in their name? A yes on either of the last two is close to dispositive; the first two decide the rest.

TABLE 1

The four criteria applied to surfaces that exist in ChartVoyant. The last column states what the criteria imply, not what ships; only the note, where the draft waits, and automatic payment posting describe published behavior.

Surface	Reversible	Blast radius	Leaves the practice	Answerable outside	What the criteria imply
Clinical note	No; corrected by addendum	One patient, every later reader	Yes, on release	The signing clinician, by license	Gate (the draft waits)
Diagnosis and procedure codes	Yes, until submitted	Claim and medical-necessity record	With the claim	The billing provider	Gate, or joined to the note gate
Claim to the clearinghouse	Only by void or corrected claim	Payer, patient balance	Yes	The billing provider	Gate
Referral and fax filing	Yes, but a misfile is a disclosure	Another patient's chart	No, unless misfiled	The practice	Gate on chart assignment
Scheduling an appointment	Yes	One slot	Only as a notification	No one	Record, do not gate
Patient message on clinical information	No, once sent	One patient, acting on it	Yes	The reviewing clinician	Gate
Posting an insurer payment	Yes, by reversal	The ledger	No	No one at posting	Record, do not gate (posted automatically)

For referral and fax intake the gate belongs on one decision inside the action: filing is reversible, but assigning a document to the wrong patient is a disclosure that refiling does not undo. Patient-facing messages carry two legal hooks, AB 3030's exemption for provider-reviewed communications⁸ and the Texas requirement of a clear, conspicuous, plain-language disclosure that AI is being used in the provision of health care services.¹² Both presume a record of which human being stood behind the message.

WHAT THE CRITERIA ARE NOT

These four criteria are engineering judgment, not a measured result. No trial has compared gate placements in an EMR, and the evidence base is thinnest where the criteria say gates matter most: of 519 health-LLM studies, 44.5% evaluated on medical examination questions, 5% on real patient care data, and 0.2% each on billing codes and on prescriptions.¹³ Criteria derived by reasoning should be labeled as such, and revised when somebody measures.

4 Making verification cheap

This section matters most, and it follows from the finding that automation bias tracks verification complexity rather than divided attention.⁶ If checking one assertion is expensive, the gate produces confirmation without verification, the failure it was added to prevent. Four properties make checking cheap.

The source is one movement away from the assertion. For any statement in a draft, the material it came from has to be reachable without leaving the review surface and without a search. In ambient documentation that material is the transcript, and the link has to be recorded rather than reconstructed: ChartVoyant's drafting event links to the transcript it drew on. One movement is the target, because the alternative is a reviewer who checks the first two claims and accepts the rest. In 97 outpatient encounters with 388 blinded paired specialist reviews, 31% of ambient notes contained hallucinations against 20% of physician-drafted notes, at near-parity overall quality ratings.¹⁴ That combination is exactly the condition under which unaided reading fails.

The diff is visible, and ordered by risk rather than size. A reviewer shown a finished document reads for plausibility. A reviewer shown what changed reads for correctness. The ordering has a measurable target: in 12,999 clinician-annotated sentences the overall hallucination rate was 1.47%, of which 44% were graded major, and the major ones clustered in the Plan section at 21%.¹⁵ A diff that gives a reworded history of present illness the same prominence as a changed plan has sorted attention by text volume instead of by consequence.

The gate presents what changed, not only the result. Confident, fluent, internally consistent output is the normal case, and models do not signal when they extend a false premise. In 300 physician-validated vignettes each seeded with one fabricated element, six models elaborated on the fabrication rather than rejecting it at rates from 50% to 82.7%, 66% overall; a mitigation prompt reduced that to 44%, and setting temperature to zero produced no significant improvement.¹⁶ A supported assertion and an unsupported one therefore have to look different on the screen, because they do not sound

different on the page. The same holds upstream, where transcription systems have produced fabricated phrases with no counterpart in the audio, disproportionately after long non-vocal pauses, though the measurements come from aphasia and control speech corpora rather than medical dictation.¹⁷

The reviewer can see when a value is unavailable. A blank is ambiguous: it may mean normal, not measured, not retrieved, or not applicable. Federal certification rules agree. The Decision Support Interventions criterion at 45 C.F.R. § 170.315(b)(11) requires, at paragraph (v), that identified users be able to access, record and modify source attributes and to see when a source attribute value is unavailable,¹⁸ a requirement adopted in HTI-1.¹⁹ It binds health IT developers rather than practices, and it is a transparency obligation, not a validation one: no federal body evaluates model performance under it. Its future is unsettled, since a December 2025 proposed rule would remove the artificial intelligence model-card requirements entirely, with no final rule located as of September 2026.²⁰ The design argument does not depend on the rule surviving.

5 What a gate must not do

No pre-checked consent. A checkbox that arrives checked has already made the decision and relabeled it as the reviewer's. It fails clause (d) while appearing to satisfy the definition, and it manufactures a record of an approval nobody gave.

No single confident recommendation presented without its basis. The statutory test for decision support that stays outside device regulation has four elements, and the fourth requires software “enabling such health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient.”²¹ The current guidance is FDA's final version of 29 January 2026, which superseded the September 2022 guidance that most secondary commentary still describes.²² Law-firm analyses of the revision report that it accepts a single recommendation where only one is clinically appropriate; the objection here is not to singularity but to a recommendation whose basis is absent. The element is about the basis being independently reviewable, not about a click: adding a button saying a clinician must confirm does not by itself satisfy it. The radiologists in the mammography study were shown a BI-RADS category, a conclusion with no visible reasoning attached.² A ranked set with its supporting and contradicting evidence exposed is reviewable. A bare label is not, however many times it is confirmed.

No gate whose only options are accept and accept-later. Rejection has to be reachable in one action, recorded as an outcome in its own right, and has to leave the draft available for correction rather than destroying it. A surface offering only approval and deferral is a queue, and queues get cleared.

No batching that lets one action approve many. One act of attention must not produce many approvals. The temptation comes from the same place as everything in section 3: repetition exhausts reviewers, so a hundred pending items invites a control that clears a hundred.⁷ A batch approval records

what did not happen, which is worse than an honest record of an automated action.

6 Measuring whether a gate is working

A gate is a control, and an unmeasured control is a belief. Four families follow, all derivable from an append-only event record without additional monitoring of the clinician. First, time-to-decision distributions, read at the fast end rather than the median: a mode at the minimum possible interaction time is the signature of clicking, and it appears long before it appears in outcomes. Second, edit rates, the proportion of drafts altered before approval, segmented by surface, clinician and draft length. Third, the location of those edits, the operational version of the finding that dangerous errors concentrate in particular sections of a note.¹⁵ Fourth, rejection rates, which should be neither zero nor noise. A rejection rate of zero means the gate is not a decision; a rate indistinguishable from random means the drafts are not worth reviewing.

WHAT THESE MEASURES DO NOT ESTABLISH

A gate with a 100% approval rate is indistinguishable, from the outside, from no gate at all, and none of these measures shows that a reviewer understood what they approved. They measure the shape of attention, not its content, and the end-to-end documentation study is the analogue: review cut error volume roughly twentyfold without reducing the clinically significant share.⁴ Showing that a gate catches the errors that matter requires evaluation on real cases between offline validation and large comparative trials, the gap DECIDE-AI was built for with its 27 reporting items.²³ This section specifies what to collect; no measurements of ChartVoyant's own gate behavior have been published, and nothing here reports any.

7 The cost

Gates are friction, and the friction is not a rounding error. A systematic review and meta-analysis of human-LLM collaboration named the supervision burden a “vigilance tax”; across 10 peer-reviewed studies, composite diagnostic and management scores improved by 4.88 percentage points (95% CI 0.65 to 9.12) with a 95% prediction interval crossing null, while factual error rates remained around 26% to 36%.²⁴ In a randomized trial reported as a preprint, physicians with model access took 119.3 seconds longer per case (95% CI 17.4 to 221.2).²⁵ In a randomized trial of diagnostic reasoning, physicians using an LLM scored no better than physicians without one, an adjusted difference of 2 points (95% CI

−4 to 8; $P = .60$), while the model alone scored 16 points above the control arm.²⁶ A later trial found that where the model sat in the sequence of work changed the result.²⁷ Review time is real work, and buying it back by weakening the review is the failure this specification exists to prevent.

The honest statement is this: a system with proper approval gates is slower than the same system without them, and that is the intended trade rather than a defect to be optimized away. The tension with the design literature's first commandment, that speed matters most,¹¹ is why the placement criteria exist. The response to friction is fewer gates in better places, with cheaper verification at each. It is never a gate that proceeds on its own.

8 The properties, as a test

The specification reduces to ten questions, answerable about any system that claims a human is in the loop.

1. Can the system perform the transition itself under any condition, including a timeout, a retry or a deploy? If so, there is no gate.
2. Does the record name one person rather than a role, a queue or a service account?
3. Did that person hold the authority for this transition on this record, and is that checked rather than assumed?
4. Is the approval a separate event from the transition it authorizes, on a record that cannot be rewritten?
5. What happens if nobody acts? If the action proceeds, the default is wrong.
6. How many seconds and how many movements does verifying one assertion against its source take?
7. Does the surface show what changed since the last approved version, ordered by consequence, or only the current result?
8. Can the reviewer distinguish a missing value from a normal one?
9. Is the basis of a recommendation reviewable independently of the recommendation?
10. Is rejection a first-class outcome with its own record, and what is the observed rejection rate?

None of these asks whether the model is good. A better model does not remove the need for a gate: the failures above all occurred with output good enough to be believed. A gate is a statement about who is answerable when the output is wrong, plus a specification for keeping that person's attention affordable. The difference between a real gate and a ceremonial one is visible from outside, in the ten answers.

REFERENCES

Entries 1 to 7 and 13 to 17 are the empirical studies the specification rests on, with entries 24 to 27 the trials and meta-analysis behind section 7. Entries 9 to 11 are the older clinical-decision-support design literature the placement criteria answer to, and entries 8, 12 and 18 to 22 are statutes, regulations and agency guidance cited for what they require rather than as evidence of what works. Entry 25 is a preprint and entry 3 is a journal article whose volume, issue and pages could not be confirmed; both are flagged in place.

1. Lyell D, Magrabi F, Raban MZ, et al. *Automation bias in electronic prescribing*. BMC Medical Informatics and Decision Making. 2017;17:28. [doi:10.1186/s12911-017-0425-5](https://doi.org/10.1186/s12911-017-0425-5) RCT
2. Dratsch T, Chen X, Rezazade Mehrizi M, et al. *Automation Bias in Mammography: The Impact of Artificial Intelligence BI-RADS Suggestions on Reader Performance*. Radiology. 2023;307(4). [doi:10.1148/radiol.222176](https://doi.org/10.1148/radiol.222176) READER STUDY
3. Qazi IA, Ali A, Khawaja AU, et al. *Automation Bias in Large Language Model–Assisted Diagnostic Reasoning among Physicians Trained in AI Literacy — A Randomized Clinical Trial*. NEJM AI. Volume, issue and pages not confirmed. [doi:10.1056/Aloa2501001](https://doi.org/10.1056/Aloa2501001). Preprint version: medRxiv, posted 26 Aug. 2025, [doi:10.1101/2025.08.23.25334280](https://doi.org/10.1101/2025.08.23.25334280) RCT
4. Zhou L, Blackley SV, Kowalski L, et al. *Analysis of Errors in Dictated Clinical Documents Assisted by Speech Recognition Software and Professional Transcriptionists*. JAMA Network Open. 2018;1(3):e180530. [doi:10.1001/jamanetworkopen.2018.0530](https://doi.org/10.1001/jamanetworkopen.2018.0530) CROSS-SECTIONAL
5. Goddard K, Roudsari A, Wyatt JC. *Automation bias: a systematic review of frequency, effect mediators, and mitigators*. Journal of the American Medical Informatics Association. 2012;19(1):121–127. [doi:10.1136/amiajnl-2011-000089](https://doi.org/10.1136/amiajnl-2011-000089) SYSTEMATIC REVIEW
6. Lyell D, Coiera E. *Automation bias and verification complexity: a systematic review*. Journal of the American Medical Informatics Association. 2017;24(2):423–431. [doi:10.1093/jamia/ocw105](https://doi.org/10.1093/jamia/ocw105) SYSTEMATIC REVIEW
7. Ancker JS, Edwards A, Nosal S, et al. *Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system*. BMC Medical Informatics and Decision Making. 2017;17:36. [doi:10.1186/s12911-017-0430-8](https://doi.org/10.1186/s12911-017-0430-8) COHORT
8. California Legislature. *Assembly Bill No. 3030, Health care services: artificial intelligence*. Stats. 2024, ch. 848; filed 28 Sept. 2024; operative 1 Jan. 2025; codified at Cal. Health & Safety Code §§ 1339.75–1339.78. [leginfo.legislature.ca.gov](https://leginfo.ca.gov) STATUTE
9. van der Sijs H, Aarts J, Vulto A, Berg M. *Overriding of Drug Safety Alerts in Computerized Physician Order Entry*. Journal of the American Medical Informatics Association. 2006;13(2):138–147. [doi:10.1197/jamia.M1809](https://doi.org/10.1197/jamia.M1809) SYSTEMATIC REVIEW

10. Kesselheim AS, Cresswell K, Phansalkar S, Bates DW, Sheikh A. *Clinical Decision Support Systems Could Be Modified To Reduce 'Alert Fatigue' While Still Minimizing The Risk Of Litigation*. Health Affairs. 2011;30(12). Page range not confirmed. [doi:10.1377/hlthaff.2010.1111](https://doi.org/10.1377/hlthaff.2010.1111) POSITION PAPER
11. Bates DW, Kuperman GJ, Wang S, et al. *Ten Commandments for Effective Clinical Decision Support: Making the Practice of Evidence-based Medicine a Reality*. Journal of the American Medical Informatics Association. 2003;10(6):523–530. [doi:10.1197/jamia.M1370](https://doi.org/10.1197/jamia.M1370) POSITION PAPER
12. Texas Legislature. *Texas Responsible Artificial Intelligence Governance Act*. H.B. 149, 89th Leg., R.S. (2025); Tex. Bus. & Com. Code tit. 11, subtit. D, chs. 551–554, health-care disclosure at § 552.051(f); effective 1 Jan. 2026. capitol.texas.gov STATUTE
13. Bedi S, Liu Y, Orr-Ewing L, et al. *Testing and Evaluation of Health Care Applications of Large Language Models: A Systematic Review*. JAMA. 2025;333(4):319–328. [doi:10.1001/jama.2024.21700](https://doi.org/10.1001/jama.2024.21700) SYSTEMATIC REVIEW
14. Palm E, Manikantan A, Mahal H, et al. *Assessing the quality of AI-generated clinical notes: validated evaluation of a large language model ambient scribe*. Frontiers in Artificial Intelligence. 2025;8:1691499. [doi:10.3389/frai.2025.1691499](https://doi.org/10.3389/frai.2025.1691499) COMPARATIVE STUDY
15. Asgari E, Montaña-Brown N, Dubois M, et al. *A framework to assess clinical safety and hallucination rates of LLMs for medical text summarisation*. npj Digital Medicine. 2025;8:274. [doi:10.1038/s41746-025-01670-7](https://doi.org/10.1038/s41746-025-01670-7) BENCHMARK
16. Omar M, Sorin V, Collins JD, et al. *Multi-model assurance analysis showing large language models are highly vulnerable to adversarial hallucination attacks during clinical decision support*. Communications Medicine. 2025;5:330. [doi:10.1038/s43856-025-01021-3](https://doi.org/10.1038/s43856-025-01021-3) BENCHMARK
17. Koenecke A, Choi ASG, Mei KX, Schellmann H, Sloane M. *Careless Whisper: Speech-to-Text Hallucination Harms*. Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency; arXiv:2402.08021. arxiv.org/abs/2402.08021 BENCHMARK
18. Office of the National Coordinator for Health Information Technology. *Decision support interventions*. 45 C.F.R. § 170.315(b)(11) (2026). ecfr.gov REGULATION
19. Office of the National Coordinator for Health Information Technology. *Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing*. Final rule. 89 Fed. Reg. 1192–1438 (Jan. 9, 2024); Doc. No. 2023-28857; RIN 0955-AA03; effective Feb. 8, 2024. federalregister.gov REGULATION
20. Assistant Secretary for Technology Policy / Office of the National Coordinator for Health Information Technology. *Health Data, Technology, and Interoperability: ASTP/ONC Deregulatory Actions To Unleash Prosperity*. Proposed rule. 90 Fed. Reg. 60970–61034 (Dec. 29, 2025); Doc. No. 2025-23896; RIN 0955-AA09; comments closed Feb. 27, 2026. federalregister.gov REGULATION

- 21.** United States Congress. *21st Century Cures Act, § 3060(a), Clarifying Medical Software Regulation*. Pub. L. No. 114-255, 130 Stat. 1033, 1130 (Dec. 13, 2016), codified at 21 U.S.C. § 360j(o)(1)(E). [uscode.house.gov](https://www.uscode.house.gov) **STATUTE**
- 22.** U.S. Food and Drug Administration. *Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff*. Final guidance issued 29 Jan. 2026, superseding the version issued 6 Jan. 2026 and the final guidance of 28 Sept. 2022; Docket No. FDA-2017-D-6569. [fda.gov](https://www.fda.gov) **GUIDANCE**
- 23.** Vasey B, Nagendran M, Campbell B, et al. *Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI*. *Nature Medicine*. 2022;28(5):924–933. [doi:10.1038/s41591-022-01772-9](https://doi.org/10.1038/s41591-022-01772-9) **REPORTING GUIDELINE**
- 24.** Wang G, Zhang K, Jiang J, et al. *Human–large language model collaboration in clinical medicine: a systematic review and meta-analysis*. *npj Digital Medicine*. 2026;9:195. [doi:10.1038/s41746-026-02382-2](https://doi.org/10.1038/s41746-026-02382-2) **META-ANALYSIS**
- 25.** Goh E, Gallo R, Strong E, et al. *Large Language Model Influence on Management Reasoning: A Randomized Controlled Trial*. medRxiv preprint, 2024. Not peer-reviewed at the time of writing. [doi:10.1101/2024.08.05.24311485](https://doi.org/10.1101/2024.08.05.24311485) **PREPRINT**
- 26.** Goh E, Gallo R, Hom J, et al. *Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial*. *JAMA Network Open*. 2024;7(10):e2440969. [doi:10.1001/jamanetworkopen.2024.40969](https://doi.org/10.1001/jamanetworkopen.2024.40969) **RCT**
- 27.** Everett SS, Bunning BJ, Jain P, et al. *From tool to teammate in a randomized controlled trial of clinician-AI collaborative workflows for diagnosis*. *npj Digital Medicine*. 2026;9:409. [doi:10.1038/s41746-026-02545-1](https://doi.org/10.1038/s41746-026-02545-1) **RCT**

RELATED PAPERS

Human oversight of clinical AI

The experimental evidence on how reviewers fail, in detail.

A write-path architecture for an AI-native EMR

The states a gate moves an entry between.

Regulating AI in certified EHR technology

Where oversight is a regulatory requirement rather than a design choice.