THE REFERENCE GUIDE

100 practical pearls

for healthcare AI decisions

Ten categories for better questions, useful team discussions and clearer judgment.

Harvey Castro MD · DR GPT™

Editorial synthesis · September 2026
Source map and review limitations included

Reference edition · September 2026

Start with a decision

This guide brings together 100 practical pearls from the DR GPT™ Podcast Intelligence collection. Use it to prepare a team discussion, examine an AI proposal or develop a teaching session. The pearls are editorial lessons and suggested applications, not a scientific ranking or a set of treatment recommendations.

The underlying review covered 41 local files, with substantive September briefs prioritized and selected archive sections screened. The source keys preserve where an idea originated. Many pearls are editorial adaptations, not statements made by the named guest. Original episode links have not all been independently verified. Check the source before quoting or presenting a specific factual claim.

How to use the collection

01 / Clinical AI and patient care

A clinical AI decision includes the people, information and handoffs around the tool. Use these pearls when planning a pilot or reviewing an existing workflow. They prompt a team to ask who checks the output, whether it fits the intended setting and what happens when the tool is unavailable. A claim about faster work still needs to be distinguished from evidence about patient care.

1. Evaluate the clinician and AI together; a strong model can still produce a weak workflow. [F1]
2. Define who reviews an AI recommendation and what happens when they disagree. [JEV]
3. Measure whether documentation tools return usable time to clinicians and patients. [SAT]
4. Fix broken intake forms and handoffs alongside introducing new AI. [TB]
5. A narrow clinical task still needs validation for its intended patients and setting. [F1]
6. Make uncertainty visible before a clinician acts on an output. [F1]
7. Give clinicians enough context to check an answer without reconstructing the entire case. [F1]
8. Separate administrative time savings from evidence of better clinical outcomes. [SAT]
9. Keep a workable fallback for interruptions in an AI service. [KEN]
10. Compare the risks of adoption with the consequences of leaving the current workflow unchanged. [DEB]

PUT IT TO WORK

Take one proposed pilot and name the reviewer, the recipient of the final output and the fallback owner.

Category 1 of 10 · Source keys refer to the source map

02 / Evaluation and measurement

Evaluation begins with a definition of useful work. A demonstration may leave out difficult cases, the time needed to correct an answer or the effort required from the person supervising it. These questions help a team describe what it wants to measure before a pilot starts. Use them to compare an AI proposal with the current process under conditions that make the comparison meaningful.

11. Choose the success measure before choosing the model. [JW]
12. Test on representative work from the setting where the tool will be used. [F1]
13. Include review time and correction costs when calculating time saved. [LIAM]
14. Measure completed tasks rather than impressive demonstrations. [MW]
15. Inspect consequential errors even when average performance looks strong. [JEV]
16. Check whether a benchmark rewards shortcuts that undermine the real objective. [WES]
17. Distinguish faster execution from better reasoning. [AWG]
18. Compare systems under similar budgets, tools, and access to human help. [AWG]
19. Seek independent evidence and local evaluation before relying on a vendor’s performance claims. [ALL]
20. Treat a system update as a reason to reassess performance, not an automatic improvement. [F1]

PUT IT TO WORK

Write down the current process, the proposed comparison and the outcome you will measure before testing.

Category 2 of 10 · Source keys refer to the source map

03 / Product design and adoption

Adoption requires attention to the task a person is trying to finish. Before adding a feature, examine where work begins, what information is needed and which interruptions make completion difficult. These pearls connect product decisions to the experience of the people using the tool. They are useful prompts for an interview, a workflow observation or a pilot review.

21. Map the existing workflow before proposing automation. [LIAM]
22. Start with a recurring task that has a clear owner and visible cost. [GREG]
23. Choose the least expensive model that reliably meets the task's requirements. [JW]
24. Use structured classification when the job needs a decision category rather than prose. [JEV]
25. Place assistance in the communication channels people already use. [GI]
26. Remove unnecessary app switching from the path to completing work. [GI]
27. Make defaults useful, with advanced controls available when needed. [ADB]
28. Measure repeat use after the pilot's initial excitement fades. [JW]
29. Give users a clear way to correct the system and report failures. [F1]
30. Ask whether a proposed feature fits what customers already trust you to do. [DARA]

PUT IT TO WORK

Observe one recurring task and mark every handoff or application change needed to finish it.

Category 3 of 10 · Source keys refer to the source map

04 / Agents, permissions, and operational control

An agent that can act needs clearly defined authority. Consider which records it may access, what it may change and which decisions still belong to a person. These pearls help frame an operational review before connecting an agent to important systems. Access, spending, records and recovery should be explicit enough for another team member to understand how the assignment is controlled.

31. Grant agents only the access needed for the assigned task. [WES]
32. Set spending limits before allowing an agent to purchase or consume paid resources. [GI]
33. Define which consequential actions require a person’s approval. [JEV]
34. Treat retrieved documents and stored memories as possible sources of misleading instructions. [WES]
35. Keep an audit trail of the agent’s access and actions without retaining unnecessary sensitive information. [TB]
36. Check what happens to stored data when an integration is disconnected. [GI]
37. Plan how to stop an agent and recover from a failed action. [WES]
38. Assess the coordination cost before adding more agents. [AWG]
39. Specify what successful completion means before giving an agent a long assignment. [AWG]
40. Separate controlled safety test failures from confirmed incidents in ordinary use. [ROV]

PUT IT TO WORK

Draw the boundary between what an agent may read, recommend and change.

05 / Business and institutional strategy

A useful institutional strategy connects a proposed investment with a problem the organization can describe. These pearls invite leaders to examine dependencies, incentives and the work required after a purchase. Use them when discussing a vendor, a new service or a change in staffing. They are questions for local analysis, not predictions that a particular business model will succeed.

41. Identify the main constraint before spending on a larger technical solution. [MOON]
42. Test demand for a specific service before expanding its feature list. [GREG]
43. Price a proposed service around a measurable customer benefit. [LIAM]
44. Include implementation, support, and oversight in the cost of delivering AI. [JW]
45. Ask who benefits financially when an AI tool succeeds. [MOON]
46. Document the cost and difficulty of switching vendors before becoming dependent on one. [KEN]
47. Make data portability part of the procurement discussion. [KEN]
48. Examine an evaluator’s incentives as well as their technical expertise. [TB]
49. Assess how a proposed safety rule affects competition and entry costs. [IT]
50. When automation changes jobs, plan for the transition rather than assuming new demand will resolve it. [SHIFT]

PUT IT TO WORK

Add implementation, support and exit costs to one vendor discussion.

Category 5 of 10 · Source keys refer to the source map

06 / Creativity, communication, and professional judgment

The library can help develop a talk, article or teaching session when the audience and purpose are clear. These pearls focus on selecting ideas, explaining them concretely and protecting the distinction between original thought and borrowed material. Use them during preparation and editing. A consistent voice should leave room for different formats and for the uncertainty that belongs to the subject.

51. Define the point you want to make before asking AI to draft it. [RUB]

52. Preparation can matter more than the speed of the final execution. [RUB]

53. Use AI to generate options; apply your own criteria when selecting the final work. [RUB]

54. Judge a piece by its value to the audience, not the volume produced. [JULIA]

55. Give creative work enough time to develop beyond immediate engagement metrics. [RUB]

56. Do not let every urgent request displace your most important work. [RUB]

57. Keep a recognizable visual and editorial style across a content series. [MW]

58. Use a concrete example to explain an idea before adding more terminology. [RUB]

59. Disclose synthetic presenters when viewers could mistake them for a real professional. [JULIA]

60. Check a memorable line against the original recording before quoting it publicly. [READ]

PUT IT TO WORK

Select one idea and explain it with a concrete example your intended audience will recognize.

Category 6 of 10 · Source keys refer to the source map

07 / Evidence and media judgment

A compelling conversation can generate a useful question without settling it. These pearls help readers examine where a claim came from, what kind of evidence supports it and what remains unresolved. They also support fair treatment of disagreement. Several guests repeating a story does not create several independent findings, and a speaker's expertise should be considered in relation to the claim.

61. Count original evidence sources; several summaries of one story are still one source. [READ]

62. Separate an expert's forecast from an observed result. [DEB]

63. Check whether a striking number describes a population, subgroup, or selected example. [HUB]

64. Distinguish statistical significance from a change people would notice or value. [SCULL]

65. Look for explanations of an association before treating it as cause and effect. [BACH]

66. Read beyond the headline before deciding what an article establishes. [AF]

67. Assess reliability and political perspective separately. [AF]

68. Keep expertise specific to the subject being discussed. [ROV]

69. Ask whether a proposed explanation makes a testable prediction. [SAB]

70. Keep uncertainty attached to a claim when turning it into slides, audio, or social posts. [READ]

PUT IT TO WORK

Trace one widely repeated claim back to its original evidence.

Category 7 of 10 · Source keys refer to the source map

08 / Health, longevity, and research interpretation

Health discussions often move quickly from a mechanism or measurement to an implied benefit. This category provides questions for reading research and preparing educational content. Ask what was measured, who was studied and how long follow up lasted. These are interpretation prompts, not treatment instructions. Specific decisions about care require the relevant evidence and the patient's clinical context.

71. Ask whether a biomarker change has been shown to predict a meaningful health benefit. [F2]

72. Keep a plausible biological mechanism separate from demonstrated treatment benefit. [F2]

73. Evaluate benefits alongside adverse effects and treatment discontinuation. [HUB]

74. Check how closely a study's participants resemble the people you want to help. [HUB]

75. Do not infer lifelong benefit from a short follow up period. [F2]

76. For weight related interventions, ask what changed in function and body composition as well as weight. [HUB]

77. Use several relevant measures when assessing fitness; one score cannot answer every question. [GAL]

78. Distinguish a coaching target from a validated clinical threshold. [GAL]

79. A product's availability does not establish its effectiveness. [KHAN]

80. Revise a health hypothesis when stronger evidence contradicts it. [KHAN]

PUT IT TO WORK

Read one study and distinguish its measured result from the benefit implied by the headline.

Category 8 of 10 · Source keys refer to the source map

09 / Infrastructure, robotics, and space

Technical capability is one part of deploying a system. These pearls broaden the discussion to power, maintenance, supply chains and the conditions under which a demonstration took place. Some topics extend beyond clinical AI; their value here is to improve questions about readiness and feasibility. Use the space and robotics items for strategic discussion rather than as claims about healthcare effectiveness.

81. Assess power, connectivity, and maintenance alongside model capability. [MOON]

82. Include infrastructure lead times when estimating a deployment schedule. [MOON]

83. Compare local and cloud deployment for the actual workload and operating environment. [JW]

84. Require a full privacy and security assessment regardless of where a model runs. [JW]

85. Ask how much human intervention occurred during an autonomous robot demonstration. [MW]

86. Distinguish a prototype, a pilot, and a system operating reliably at scale. [MW]

87. Consider hazardous tasks when identifying useful applications for robotics. [ALL]

88. Evaluate orbital computing as a complete infrastructure proposal, including launch, heat rejection and repair. [AWG]

89. Examine labor and material supply chains when discussing the costs of AI. [KARA]

90. Look for useful adjacent applications of a technology without assuming they share identical economics. [ALL]

PUT IT TO WORK

Ask which operating conditions and human interventions were omitted from one demonstration.

Category 9 of 10 · Source keys refer to the source map

10 / Turning information into better decisions

Useful information should change how a question is examined or a decision is prepared. These final ten pearls describe editorial and learning practices for doing that: preserving the source, checking understanding and learning from a limited test. They are suggested working habits, not clinical findings or verified claims of improved learning. Use them to make the rest of the collection easier to inspect and apply.

91. Store the original source, date and exact passage beside each insight so another reader can check it. [READ]

92. Label what was said, what was paraphrased and what you concluded as separate parts of the record. [READ]

93. Mark the current reviewed version and preserve the correction history so old claims do not quietly return. [READ]

94. When speakers disagree, identify the assumption or evidence behind the disagreement before choosing a view. [DEB]

95. Record a prediction’s deadline and success measure before checking whether it came true. [ZOIN]

96. Recall the main idea without reopening the brief, then check what you remembered against the source. [ZOIN]

97. Before using an idea in a talk, ask what evidence would weaken it and where it might not apply. [ZOIN]

98. Record what remains unknown and what evidence would be needed to answer it. [READ]

99. Give each useful insight a specific possible use, such as a teaching example or a question for a pilot review. [READ]

100. Test one idea in a low risk setting, record the result and decide what to change before trying it again. [ZOIN]

PUT IT TO WORK

Choose one insight to test and record the result beside its source.

Category 10 of 10 · Source keys refer to the source map

How to use the final ten

Pearls 91 to 100 help you apply any idea from the first 90. Use the steps relevant to your decision. A prediction needs a deadline; an editorial observation may not.

Hypothetical example: Your team wants to explore whether an AI documentation tool saves time. Pearl 13 prompts you to include review and correction time in the comparison. No results are assumed below.

91 to 93 · Keep a checkable record

Save the source link, date and relevant passage. Separate the speaker's claim from your question: Does this tool reduce total documentation time in our workflow? Date your notes and retain corrections.

94 to 95 · Define the comparison

If accounts disagree, check whether they count review time or use different tasks. If testing a prediction, state the expected result, measurement and deadline before the test.

96 to 98 · Check your understanding

Explain the idea in your own words, then compare it with the source. Identify what could weaken it, such as longer correction time. Record missing evidence, including whether the demonstration resembles your setting.

99 to 100 · Try a limited application

Use fictional cases in an approved test environment. Compare the current process with the proposed workflow, including review and correction. Record errors as well as total time. Decide whether to revise the test, stop or pursue further evaluation.

Keep one brief record: source, question, uncertainties, test, result and next decision. A small simulation can inform further evaluation. It does not establish patient benefit or readiness for clinical use.

Worked example · Applying pearls 91 to 100

Sources and review limits

Source notes preserve the September 22 review. FDA transparency guidance was reopened September 23. Other source links remain at the review depth stated below. No quotation is reproduced here.

F1. FDA, transparency principles for machine learning enabled medical devices

Primary source checked during this review. Supports human and AI team evaluation, intended use, limitations, and communication. Some operational pearls are my applications.

Open original source

F2. FDA, surrogate endpoint resources

Primary source checked during this review. Supports the distinction between biomarkers, surrogate endpoints, and clinical benefit. Study design questions are editorial applications.

Open original source

JEV. Greg Isenberg with Ryan Vogel, Jev discussion

Reported in the September creator brief. Pearls are applications, not validation of the product.

Open original source

SAT. All In, Satya Nadella on Microsoft and AI

Reported in the September long form brief.

Open original source

TB. TBPN, High Interest Rate Phenomenon and Who Votes on AI Safety, September 15

Third party transcript link recorded in the September brief. Operational controls are my applications.

Open original source

KEN. Ken Rutkowski, The Tech Industrial Complex

Healthcare procurement lessons are adaptations of the dependency argument.

Open original source

Source map · Inspiration is distinct from verified evidence

DEB. Diary of a CEO, The Great AI Debate, September 17

Reported in the September brief. Numerical risk predictions excluded.

Open original source

JW. Jordan Wilson, Everyday AI, episodes 864 and 865

The corpus supplies a show website and feed rather than a verified episode permalink.

Open original source

LIAM. Liam Ottley, accounting automation discussion

September creator brief plus April workflow material. Build speed claims excluded.

Open original source

MW. Matt Wolfe, robotics discussion

September creator brief, with archive and retrospective synthesis where indicated by file references.

Open original source

WES. Wes Roth, AI security discussion

September creator brief and April evaluation themes. Specific incident allegations not asserted here.

Open original source

AWG. Alex Wissner Gross, The Innermost Loop, September 22

Brief reports several dated posts. Quantitative performance, revenue, and orbital claims excluded.

Open original source

ALL. All In, Musk and Shotwell on AI risks, peer review, and Starship

September long form brief; other All In interviews within the same brief inform infrastructure applications.

Open original source

GREG. Greg Isenberg, creator and founder material

April briefs and September creator roundup. No verified episode permalink supplied for some archive ideas.

Open original source

GI. Greg Isenberg, Instinct discussion

Reported in September creator brief. Data deletion item is a diligence question, not an allegation.

Open original source

ADB. The AI Daily Brief, personal agents

Reported in September creator brief.

Open original source

DARA. Moonshots, episode 244 with Dara Khosrowshahi

April 6 brief supplies show index, not an episode permalink. Platform fit is an adaptation; numerical claims excluded.

Open original source

ROV. Joe Rogan Experience, episode 2554 with Carlo Rovelli

Third party transcript link reported in September long form brief.

Open original source

MOON. Moonshots, April infrastructure episodes and deep dive

Episodes 241 and 246 and the April deep dive. Forecasts and market size figures excluded; pearls are planning applications.

Open original source

IT. Impact Theory, The Real Reason AI Companies Are Desperate to Be Regulated

September brief had show notes only. Competition item is a question to examine, not an assertion about motives.

Open original source

SHIFT. The Shift, Salim Ismail discussion

September science brief. Labor transition item is an application, not a guaranteed economic outcome.

Open original source

RUB. Diary of a CEO with Rick Rubin, September 21

September brief used an automatic transcript. No verbatim quotations reproduced.

Open original source

JULIA. Julia McCoy, avatar and content production discussion

September science brief. Disclosure and audience value items are editorial recommendations.

Open original source

READ. DR GPT Podcast Intelligence editorial and collection practices

My synthesis of the folder instructions, extraction workflow, revisions, and logs. This is not an external evidentiary citation.

HUB. Huberman Lab with Chris Thompson, gut health and weight loss

September brief. Evaluation questions retained; doses, thresholds, treatment advice, and unverified effect sizes excluded.

Open original source

SCULL. Lex Fridman, episode 502 with Andrew Scull

September long form brief. General evidence interpretation retained; clinical treatment claims not repeated.

Open original source

BACH. Diary of a CEO, David Bach replay on home ownership

September brief. Association and confounding lesson only; no financial recommendation.

Open original source

AF. Ad Fontes Media, podcast chart

September science brief. Ratings are assessments rather than proof of individual claims.

Open original source

SAB. Sabine Hossenfelder, science coverage

September science brief synthesizes multiple items without full new transcripts. Testability item is my general evidence principle, not a quotation.

GAL. Huberman Lab Essentials with Andy Galpin, assessing fitness

September brief identifies reused earlier material. No training prescription or numerical cutoff included.

Open original source

KHAN. Adeel Khan reference guide, evidence and access essays

Reference guide includes conflicting levels of certainty. Only research interpretation lessons retained. No personal protocol or treatment recommendation reproduced.

KARA. Joe Rogan Experience, episode 2556 with Siddharth Kara

September long form brief. Labor and material costs are an analytical lens; unverified statistics excluded.

Open original source

ZOIN. Mihailo Zoin, NotebookLM workflow reference

Reference guide used for proposed learning practices. Performance percentages and feature availability claims excluded.

Trace an item to the archive

These file numbers refer to the companion Podcast Intelligence Source Review inventory. They preserve the origin of the synthesis and do not add independent corroboration. The complete source inventory is supplied separately.

PEARL : FILES	PEARL : FILES	PEARL : FILES	PEARL : FILES
1: 19, 31	26: 19	51: 25	76: 27
2: 19	27: 19	52: 25	77: 27
3: 31, 41	28: 10, 19	53: 25	78: 27
4: 41	29: 19, 31	54: 21	79: 37
5: 25, 19	30: 33	55: 25	80: 37
6: 19, 27	31: 17, 19	56: 25	81: 34, 35, 23
7: 19, 31	32: 19	57: 13	82: 34, 15
8: 31, 19	33: 19	58: 25, 13	83: 19, 35
9: 29	34: 19	59: 21	84: 19, 1
10: 25	35: 41	60: 7, 13	85: 19, 35
11: 10, 19	36: 19	61: 1, 2, 7	86: 19, 34
12: 19, 31	37: 19	62: 25, 35	87: 31, 19
13: 11, 19	38: 23	63: 27	88: 23, 31
14: 13, 19	39: 23, 19	64: 31	89: 31
15: 19	40: 31, 23	65: 25	90: 31, 33
16: 16, 17, 21	41: 34, 35	66: 21	91: 7, 14
17: 23	42: 8, 9	67: 21	92: 7, 13
18: 23	43: 11, 19	68: 31	93: 7, 20, 21, 30, 31
19: 31, 41	44: 10, 19	69: 21	94: 25, 31, 21
20: 1, 35	45: 35, 1	70: 7, 6, 21	95: 38, 35
21: 11, 19, 41	46: 29	71: 37, 27	96: 38, 2
22: 8, 9, 19	47: 29	72: 37, 24	97: 38, 6
23: 19	48: 41, 31	73: 27	98: 3, 6, 7
24: 19	49: 31	74: 27, 37	99: 1, 14
25: 19	50: 21, 33, 35	75: 37, 27	100: 38, 2

Provenance · Archive file numbers