AI Literacy in Healthcare

8-Question Assessment (With Answer Guidance)

Purpose:

Use this handout to help clinicians and leaders identify their current AI literacy level and pinpoint targeted next steps. Work through the questions in order. When you reach a question you can't confidently answer, note your level and review the guidance.

How to use:

For each question, read the prompt and the "What a solid answer includes" guidance. If you can address most bullets with examples relevant to your setting, move to the next question.

1. AI vs. Traditional Automation
2. Predictive AI vs. Generative AI
3. Hallucinations in AI Tools
4. Data Bias and Performance Transfer
5. Evaluating Vendor Claims
6. "No-Regrets" AI Investments
7. Diagnosing Why Pilots Don't Scale
8. Real ROI vs. Hype in Healthcare AI

Source: The Health Management Academy

1. AI vs. Traditional Automation

Prompt:

How would you explain the difference between "AI" and "traditional automation" to a colleague?

What a solid answer includes:

- **Traditional Automation:**
  Deterministic, rule-based workflows (if-then logic), does exactly what it's programmed to do, no learning.
- **AI:**
  Models learn patterns from data, handle ambiguity, make probabilistic predictions, improve with new data.

Examples:

- **Rules-based:**
  EHR alert triggers when potassium > 6.0.
- **AI:**
  Model predicts risk of readmission using multi-variable patterns learned from historical data.

---

2. Predictive AI vs. Generative AI

Prompt:

A colleague says, "We're using predictive AI to identify sepsis risk and generative AI to draft clinical notes." Why are these fundamentally different?

What a solid answer includes:

- **Predictive AI:**
  estimates likelihood of an outcome (classification/regression).
  - ***Inputs***: structured/unstructured features
  - ***Output***: probabilities/scores (e.g., sepsis risk within 12–24 hours).
- **Generative AI:**
  creates new content (text/audio/images) conditioned on prompts and context (e.g., clinical note drafts).

Different objectives, evaluation methods, and risks:

- **Predictive:**
  measured by AUROC, calibration, sensitivity/specificity.
- **Generative:**
  measured by accuracy to source, coherence, factuality, style compliance; prone to hallucinations.

3. Hallucinations in AI Tools

Prompt:

A documentation assistant claimed a symptom never mentioned during the visit. What likely went wrong, and why is this unique to AI tools?

What a solid answer includes:

- **Hallucination:**
  Model outputs plausible but false information, not grounded in source data.

How to prevent this:

- **Causes*:***
  Overgeneralization, weak grounding to transcripts/EHR, prompt or context errors, token truncation, model attempting to "fill gaps."
- **Mitigations*:***
  Retrieval-augmented generation (RAG) with strict citation to source, constrained templates, clinician-in-the-loop review, confidence flags, audit logs.

---

4. Data Bias and Performance Transfer

Prompt:

A sepsis predictor trained at AMCs performs poorly at a rural hospital. What data issues could explain this?

What a solid answer includes:

- **Dataset shift:**
  differences in patient demographics, comorbidities, prevalence, care pathways, lab availability, documentation practices.
- **Sampling Bias:**
  AMC data may overrepresent complex cases; rural data may differ in acuity/resource patterns.
- **Labeling Differences:**
  sepsis definitions/coding vary; gold standards may not match.
- **Feature Distribution Drift:**
  lab timing/frequency, vital sign completeness, EHR vendor differences.

How to prevent this:

- **Mitigations*:***
  local validation, recalibration, transfer learning, federated/representative training data, ongoing monitoring.

5. Evaluating Vendor Claims

Prompt:

Your organization is vetting a vendor's clinical decision support tool. What questions test credibility?

What a solid answer includes:

- **Training Data:**
  Size, sources, representativeness, timeframes, inclusion/exclusion criteria, PHI handling.
- **Validation*:***
  External, multi-site validation; prospective vs. retrospective; metrics (AUROC, PPV/NPV, calibration); subgroup performance (equity).
- **Real-World Performance*:***
  Workflow fit, measured clinical/operational outcomes, silent trial results, degradation over time, monitoring plan.
- **Explainability and Safety:**
  Model type, feature importance, error modes, human-in-the-loop, override, auditability.
- **Governance and Compliance*:***
  FDA status (if applicable), MDR/IVDR, HIPAA, SOC2/ISO 27001, bias testing, incident response.
- **Implementation*:***
  Integration (EHR APIs, FHIR), change management, training, total cost of ownership, ROI.

---

6. "No-Regrets" AI Investments

Prompt:

What foundational investments pay off even if specific AI tools become obsolete?

What a solid answer includes:

- **Data Quality and Interoperability:**
  Standardized data (FHIR, terminologies), metadata, lineage, de-duplication, robust ETL pipelines.
- **Governance:**
  Clear AI use policies, risk review boards, model lifecycle management, bias and safety monitoring, documentation standards.
- **Workforce AI Literacy:**
  Role-based training for clinicians, informatics, IT, leadership; safe-use practices; prompt hygiene; oversight.
- **Vendor Evaluation Processes:**
  Standardized due diligence checklist, pilot-to-scale criteria, post-deployment monitoring.

Source: The Health Management Academy

- **Infrastructure:**
  Secure data platforms, MLOps, access controls, audit logging, API integration capabilities.

7. Diagnosing Why Pilots Don't Scale

Prompt:

The CEO asks, "We keep piloting AI but nothing scales. What's our biggest bottleneck?" How do you diagnose it?

What a solid answer includes:

- **Data Infrastructure Maturity:**
  Reliability of feeds, interoperability, latency, environments for development/ testing/ production, monitoring.
- **Governance Structures:**
  Decision rights, risk approvals, clinical safety sign-off, model change control, accountability.
- **Change Management:**
  Workflow redesign, clinician adoption, training, incentives, measurement of impact, support channels.
- **Strategic Alignment and Business Case:**
  Clear problem statements, value metrics, funding, owner identified, scale plan, IT capacity.
- **Integration Readiness:**
  EHR/API integration, single sign-on, embedding in existing workflows, alert fatigue management.

8. Real ROI vs. Hype in Healthcare AI

Prompt:

Where is AI creating real ROI today, and where is it mostly hype? Give specific examples.

What a solid answer includes:

- **Proven ROI (examples):**
  - **Ambient Clinical Documentation:** reduced clinician burden, improved note quality, time savings.
  - **Operational Optimization:** patient flow, staffing forecasts, denials management, supply chain.
  - **Imaging Triage and Workflow Support:** prioritization for critical findings (site-dependent regulatory clearance).
  - **Contact Center / Administrative Automation:** prior auth assistance, scheduling support.
- **Emerging/hype (examples to treat cautiously):**
  - Fully autonomous diagnosis without clinician oversight

Source: The Health Management Academy

- Broad, unsupervised generative summaries inserted into the EHR without grounding or review.
- Complex inpatient decision support without rigorous multi-site prospective validation.

- **Evidence Checklist:**
  - Peer-reviewed studies, multi-site outcomes, regulatory status (where applicable), clinician satisfaction, sustained performance over time.

Fluency Levels and Next Steps:

☐ **AI Novice (Q1–2):**
Build foundational literacy; start with AI Fundamentals and Glossary.

☐ **AI Aware (Q3–4):**
Deepen understanding of AI-specific risks/benefits; review AI Governance and risk slides.

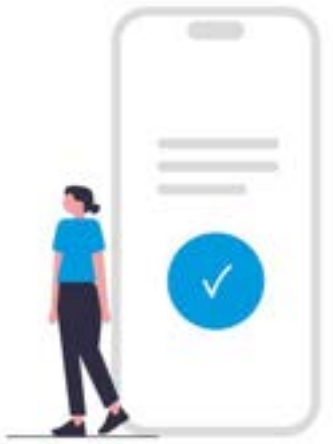

☐ **AI Informed (Q5–6):**
Focus on implementation and governance; use the FUMBLE framework and "no-regrets" investments.

☐ **Healthcare AI Competent (Q7–8):**
Lead strategy; apply the AI Maturity Roadmap and Use Case Library.

☐ **Healthcare AI Fluent (All 8):**
Maintain fluency; mentor peers; subscribe to AI Catalyst Pulse for updates.

Quick Tips for Participants:

✓ Anchor AI tools to clinical/operational goals with measurable outcomes.

✓ Require local validation and ongoing monitoring before scaling.

✓ Keep humans in the loop; design for safety, transparency, and accountability.

✓ Invest in durable foundations (data, governance, literacy) to ensure ROI as tools evolve.

Source: The Health Management Academy