

AI and the Clinician: Navigating Responsible Use in Rehabilitation

Aimee E. Perron, PT, DPT, NCS, CEEAA, FNAP



Learning Outcomes

- Distinguish the current and emerging roles of artificial intelligence in healthcare and rehabilitation practice, including common applications, opportunities, and limitations
- Deduce key factors that influence clinical decision-making when using artificial intelligence, including bias, data quality, transparency, and the potential for overreliance on generated outputs
- Integrate core principles that guide responsible and human-centered use of artificial intelligence, including fairness, accountability, transparency, safety, and respect for human autonomy
- Apply practical strategies to evaluate and responsibly integrate artificial intelligence into clinical practice, including AI tools that support ADL performance, functional mobility, and patient education, while maintaining clinical judgment, professional accountability, and patient-centered care

Chapter 1

AI in Clinical Practice: Why This Matters Now



What Is AI? A Quick Primer

- **Artificial Intelligence (AI):** Systems designed to perform tasks that typically require human intelligence
- **Machine Learning (ML):** AI that learns patterns from data without being explicitly programmed
- **Large Language Models (LLMs):** AI trained on massive text datasets to generate human-like text (e.g., ChatGPT)
- **Key point:** Most clinicians are already interacting with AI
 - EHR auto-suggestions, ambient documentation tools, predictive alerts

Machine Learning vs. Generative AI in Practice

Machine Learning

- Looks for patterns in numbers to predict, sort, or score something
- The output is a number or a label, never a full sentence
 - Fall risk scores
 - Return-to-work prediction
 - Swallow/aspiration risk scoring
 - Early deterioration alerts
- However it was trained, the question that matters clinically is the same: Has it been validated?

Generative AI

- Takes a prompt and writes new text, images, or audio (not a number)
- The output is a best guess at what sounds right, but not a checked fact
 - SOAP notes
 - HEP drafts
 - OT: ADL/goal notes
 - Shift notes, care plan drafts
- Bad training data means confident-sounding but wrong output
 - Always check before you use it

Why Generative AI Sounds Confident Even When It's Wrong

- **Hallucinations:** fluent, plausible-sounding output that is factually wrong or entirely fabricated
- LLMs predict the statistically most likely next words—they don't verify facts or “know” when they're wrong
- Sometimes called **the “stochastic parrot”** problem: mimicking the pattern of language without understanding its meaning¹
- Hallucination rates vary by tool and task
- The only reliable safeguard is expert human review of the actual output

1. Bender et al., 2021

A Rehab-Specific Simple Example: When the Output Sounds Right but Isn't

- Researchers gave ChatGPT-4 a textbook stroke case and asked it to generate a rehabilitation prescription and ICF codes¹
 - The model produced a complete, well-organized prescription and rationale in seconds
 - Broader and more detailed than the textbook's own standard answer
 - But it made a real error: It misclassified the ICF category for the case, even though the ICF codes themselves were correct

1. Zhang et al., 2023

A Note on Terminology

- **“AI Ethics,” “Responsible AI,”** and **“Trustworthy AI”** get used interchangeably, but they aren’t the same thing:
 - **AI Ethics:** The narrower question of “what is the right thing to do?”
 - Human judgment applied to AI
 - **Responsible AI:** The broader practice of building structures, roles, and accountability around developing and using AI well
 - AI ethics is one part of this larger picture
 - **Trustworthy AI:** A more contested term
 - Critics note it can imply a level of dependability AI hasn’t earned, shifting the burden onto the user to “trust” rather than onto developers and clinicians to earn that trust
 - **Responsible use** is the practices and judgment YOU bring to AI, not a claim that the technology itself is inherently trustworthy

Responsible AI: A Bigger Conversation

- “Responsible AI” isn’t healthcare-specific
 - Set of themes applied across every sector that uses AI
- Core themes
 - Fairness
 - Transparency
 - Privacy
 - Accountability
 - Safety
 - Human autonomy

Responsible AI: A Bigger Conversation (cont.)

- **Criminal justice:** predictive policing and facial recognition raise the same bias questions discussed in this course
- **Hiring:** resume-screening algorithms have faced their own well-documented bias problems
- **Finance:** credit-scoring models face the same fairness scrutiny as clinical risk tools
- **Education, government benefits, media platforms, and law** are having this same conversation right now

Where AI Touches the Episode of Care

Documentation: Ambient scribes, auto-notes, coding assist



Screening and Assessment: Movement analysis, wearables, differential screening



Treatment Planning: HEP generation, CDS tools, protocol matching



Monitoring: Telerehab, remote sensors, predictive alerts



Patient Education: Chatbots, communication aids, translated materials

A 2025 systematic mapping review identified 240 studies applying AI in rehabilitation across these touchpoints, most concentrated in neurological (57.9%) and orthopedic (22.7%) conditions¹

1. Morone et al., 2025

Opportunities

- Personalized treatment recommendations
- Remote monitoring and telerehabilitation
- Improved access to care in underserved areas
- Reduced documentation burden
- Efficiency in literature search and evidence synthesis
- Support for digital and health literacy³
- Reduced administrative burden⁴
- AI-assisted coding and claims review⁴
- Pairing a clinician with AI (rather than either alone) has shown improved accuracy, fewer errors, and greater clinician trust in early healthcare research⁵

1. Hao et al., 2025

2. Adnan et al., 2025

3. “Generative AI Rewriting,”
2024

4. “AI in Health Care,” 2025

5. Olawade et al., 2026

Limitations

- Clinical effects remain inconsistent across RTC
- Most studies use small datasets from single sites
- Most published studies don't report who's in the training data
 - 84% of clinical AI/ML studies worldwide didn't disclose the racial composition of their data, and this evidence gap is consistent with what we see in rehabilitation specifically²
- External validation applied in only ~6% of rehabilitation AI studies¹
- Only ~10% of rehabilitation AI studies report any explainability method, and over half include no comparator group at all¹
- Limited explainability, as many systems are “black boxes”
- Implementation barriers: cost, technology literacy, reliability, user fatigue

1. Morone et al., 2025

2. Otokiti et al., 2025

Example of Why External Validation Matters

- A 30-day mortality-risk model for older surgical patients was derived and internally validated on ~50,000 patients—strong performance (AUC 0.85)¹
- Tested on a different country's health system, that same model's accuracy dropped to only adequate (AUC 0.70)—a real, meaningful decline¹
- The cause: the local patient population had a different underlying data distribution than the national training set
 - “Out-of-distribution” problem
- A large, diverse training dataset does not guarantee a tool will work at your specific site with your specific patients
- **Practical takeaway:** Ask whether a tool has been validated on a population like yours, not just validated at all

1. Feng et al., 2022

2. Morone et al., 2025

Relational Risks

- Evidence so far has focused on whether AI's outputs are accurate
- Overreliance on automation can quietly erode a clinician's own hands-on assessment and reasoning skills over time
 - **Deskilling:** experienced clinicians lose skills they once had through disuse
 - **Never-skilling:** students and new clinicians never build the skill because AI does it from the start
 - **Mis-skilling:** clinicians absorb errors or biases from AI output as correct practice
- Patients report feeling the difference when a clinician is dividing attention between them and a screen (shorter, less personal-feeling encounters)
- Qualitative research links heavy AI use to a felt loss of empathy and connection in the visit, even when clinical accuracy holds steady¹
- Left unaddressed, this can erode both the patient-therapist relationship and the PT's own sense of professional identity

1. Horvath & Lorincz, 2026

Why This Matters Now: The Practice Gap

- AI tools are proliferating faster than guidance for clinicians
 - AI publications in rehabilitation medicine grew from 22 in 2015 to 1,449 in 2025¹
- Only 8% of WHO European Region member states have a health-specific AI strategy²
- 63% of physical therapists report no experience with AI applications at work³
- Clinicians need frameworks to evaluate these tools, not just use them

1. Rabinowitz et al., 2026

2. WHO/Europe, 2026

3. *Int J Environ Res Public Health*, 2022

Chapter Summary

- **AI is not one thing:** Machine learning finds patterns in data; generative AI predicts plausible language, and that difference sets what you can trust an output for
- **Fluency is not accuracy:** Generative models are built to sound right, so errors arrive in confident, well-formed prose
- **AI already touches the full episode of care:** documentation, screening and assessment, treatment planning, monitoring, and patient education.
- **Opportunity and limitation travel together:** real time savings alongside bias, thin external validation, and relational cost at the point of contact
- **Tools are outpacing governance and evidence:** Clinical judgment and professional accountability stay with the clinician, not the model

Chapter 2

What Shapes AI-Supported Clinical Decisions



Bias in AI: Where Does It Come From?

- Training data bias: AI trained on nonrepresentative datasets produces biased recommendations
 - Example: A system trained on younger adults may generate inappropriate recommendations for older adults recovering from stroke
- **Algorithmic bias:** Model design choices embed human values at every stage¹
- **Selection bias:** Which patients and outcomes are included in the data?
- Landmark case: A widely used commercial health algorithm used cost as a proxy for illness
 - Systematically underestimated the health needs of Black patients relative to White patients with similar chronic illness burden²
- Bias-related harm takes two forms: distributive harm (one population gets a less accurate or less useful output) and representational harm (a tool's language or outputs stereotype or disrespect a group, even when access to the tool itself is equal)

1. Yu et al., 2024

2. Obermeyer et al., 2019

Data Quality and Representativeness

- “Garbage in, garbage out”
- Poor-quality or biased data → inaccurate or inequitable recommendations
- Rehabilitation-specific concern: Most AI studies use raw or tabular data from single sites
- Multimodal inputs used in only ~11% of rehabilitation AI studies
- AI systems can also be surprisingly fragile: small amounts of noise or corrupted input data can shift a model’s output meaningfully¹
- Key question to ask: “What data was this tool trained on?”
 - Related question, often skipped: Did the people in that data consent to it being used to train an AI? “Publicly available” (a photo, a note, a recording) is not the same as “free to use for training”

1. Feng et al., 2025

Addressing Bias: What Actually Helps

- **Bias detection:** audit training data for over- or underrepresented groups; test whether small input changes shift outputs disproportionately
- **Fairness metrics:** quantify whether a tool's accuracy or error rate is comparable across patient groups (e.g., age, race, sex)
- **Explainable AI (XAI):** techniques like SHAP and LIME can reveal which patient features are driving a given recommendation
- Bias mitigation acts
 - Before training (rebalancing the dataset)
 - During training (adjusting the model)
 - After training (correcting outputs post hoc)
- Beyond bias-specific fixes
 - Professional recommendations call for validating and certifying tools before deployment
 - National-level safety monitoring and adverse-event reporting

1. Jain et al., 2024

2. Labkoff et al., 2024

System Design Shapes Outcomes

- A tool doesn't have to be inaccurate to fail
- Separate logins, interrupted documentation, duplicate data entry¹
- Supports adoption: embedding directly in the EHR, minimal manual input, real-time support, customizable alerts, and clear visual outputs¹
- Timing matters: a recommendation that pops up immediately captures attention and increase overreliance³
- Organizational readiness: data management, change management, IT support
- Designing the tool together with the people who will use it^{1,2}

1. Nesa et al., 2025

2. Sahni & Carrus, 2023

3. Brunyé et al., 2025

Automation Bias and Overreliance

- Automation bias: The tendency to over-rely on automated outputs without critical evaluation
- Errors of commission: Following incorrect AI advice
- Errors of omission: Failing to act when AI does not prompt
- Clinicians were 10× more likely to be correct with correct AI, but accuracy decreased reciprocally with incorrect AI¹
- Those who stand to gain the most from AI are also the most susceptible to automation bias
- Long-term reliance can erode clinical skills over time, not just in the moment³
- Expert clinicians tend to benefit less from AI decision support than novices
- Clinician's baseline attitude toward AI predicts how much they'll rely on it²

1. Kücking et al., 2026

2. Jackson et al., 2026

3. Brunyé et al., 2025

Transparency: The “Black Box” Problem

- Many AI systems lack interpretability (cannot see how a recommendation was generated)
- Explainable AI (XAI) aims to make AI reasoning visible (but can introduce new biases)
- The goal isn't maximizing trust in AI, it's calibrating “appropriate trust,” so clinicians rely on it only when that reliance is actually warranted¹
- Explainability cuts both ways: seeing a rationale increases trust
 - Including trust in a wrong recommendation, which can make a bad call harder to catch^{1,2}
- Some tools try to ground responses in specific, current sources (retrieval-augmented generation) rather than relying on the model's static training data alone
 - Can reduce errors, but does not eliminate them
- **Key question:** “Can I explain to my patient how AI contributed to their care plan?”
 - If the answer is no, proceed with caution

1. Jackson et al., 2026

2. Brunyé et al., 2025

Chapter Summary

- **Bias enters early.** It comes from the data, the labels, and the population a model was trained on, not from the algorithm alone.
- **Representativeness is a clinical concern.** If your patients were not in the training data, the model's confidence does not transfer to them.
- **Mitigation is real but partial.** Better sampling, auditing, and reporting help, but fairness criteria conflict, so improving one can worsen another.
- **Design choices carry ethical weight.** Defaults, alert thresholds, and how output is displayed shape what clinicians actually do with it.
- **Automation bias is the human half.** Overreliance on a confident, opaque output is where transparency stops being an abstraction.

Chapter 3

What Does “Responsible AI” Actually Mean?



Latent Values

- Value-sensitive design (VSD): people often can't name their own values until an experience violates one
 - Example: You may not consciously think “I value knowing when I’m talking to a machine” until an undisclosed chatbot makes you feel misled, and the value surfaces as a small breach of trust
- This is exactly why ethical frameworks matter before something feels wrong
- Values worth naming ahead of time as you adopt AI tools: privacy, trust, autonomy, freedom from bias, accountability

Four Principles of Biomedical Ethics

Autonomy

Respect the patient's right to make informed decisions about their own care, including whether AI is used in it.

Beneficence

Act in the patient's best interest, does this AI tool actually improve their care?

Nonmaleficence

"Do no harm." Could an unvalidated or biased AI output cause harm if followed uncritically?

Justice

Fair distribution of benefits and burdens. Does the tool work equally well across all patient populations?

1. Giacobello, 2025

2. Chiang et al., 2021

Ethical Frameworks for AI in Healthcare

Framework	Key Emphasis
WHO (2021)	Protect autonomy, promote well-being and safety, ensure transparency, foster accountability, ensure equity, promote sustainability
AMIA (2022)	Beneficence, nonmaleficence, autonomy, justice, explainability, interpretability, fairness, dependability, transparency, accountability
SHIFT Framework (2022)	Sustainability, human-centeredness, inclusiveness, fairness, transparency
OECD AI Principles (2019/2024)	Inclusive growth and well-being, human rights and fairness, transparency and explainability, robustness and safety, accountability, referenced by the EU, US, and UN
NIST AI Risk Management Framework (2023)	Trustworthy AI is valid/reliable, safe, secure/resilient, accountable/transparent, explainable, privacy-enhanced, and fair; managed through a Map → Measure → Manage → Govern cycle
EU Ethics Guidelines for Trustworthy AI (HLEG)	Human agency and oversight, technical robustness and safety, privacy and data governance, transparency, diversity and fairness, societal/environmental well-being, accountability
FAVES Framework	Stands for Fair, Appropriate, Valid, Effective, Safe—a foundation for equitable, patient-centered deployment with independent oversight

Not All AI Involvement Should Look the Same

- A 2026 nursing ethics framework proposes tiering how much AI should be involved based on how relationally dependent the care task is¹
- Low relational-dependency tasks (e.g., inventory tracking, equipment scheduling) can reasonably use more autonomous AI
- High relational-dependency tasks (e.g., breaking difficult news, building rapport with a new patient) call for minimal AI involvement, with the clinician clearly in the lead

1. Xia et al., 2026

Fairness

- AI should not perpetuate or amplify health disparities
- Fairness requires representative training data and equitable outcomes across populations
- Ask: “Does this tool perform equally well for all of my patients?”
- Connects to APTA Core Value: Social Responsibility



Accountability

- Clinicians remain responsible for ALL decisions informed by AI
- Responsibility does not transfer to the AI developer
- Professional liability remains with the treating clinician
- AI supports decision-making; it does not make decisions
- Connects to APTA Core Values: Accountability, Duty

Transparency and Explicability

- Transparency: Assumptions, limitations, and potential conflicts must be declared
 - Model card: the developer's summary of a tool's intended use, training data, performance across patient groups, and known limitations
 - Ask the vendor for one; a missing model card is a red flag
- **Explicability = Intelligibility + Accountability**
 - **Intelligibility:** Can the clinician understand how the AI works?
 - **Accountability:** Who is responsible for the outcomes?
- Patients deserve to know when and how AI is used in their care
- Connects to APTA Core Value: Integrity

Safety and Human Autonomy

- Safety includes “safe failure”; what happens when AI is wrong?
- **Human-in-the-loop:** Clinicians must remain actively involved in interpreting and validating AI outputs
- Human-centered AI supports clinician judgment and patient values, but it does not replace human interaction
- Patients want to know that a human clinician is making the final decisions
- Connects to APTA Core Values: Compassion and Caring, Altruism

The RIPS Model of Ethical Decision-Making

Realm

- **Individual:** This patient, this clinician
- **Organizational/Institutional:** Clinic policy, workflow
- **Societal:** Regulation, payer policy, public trust

Individual Process

- **Moral Sensitivity:** Do I notice the issue?
- **Moral Judgment:** What should be done?
- **Moral Motivation:** Do I prioritize it?
- **Moral Courage:** Will I act on it?

Situation

- **Issue:** Clear right answer, needs attention
- **Dilemma:** Competing “rights”
- **Distress:** Know the right action, barriers prevent it
- **Temptation:** Right vs. self-interest
- **Silence:** A violation is not named

APTA Code of Ethics and Core Values

APTA Core Value	How It Shows Up With AI
Accountability	Owning outcomes of AI-informed decisions, not deferring responsibility to the tool
Altruism	Prioritizing patient benefit over convenience or efficiency gains from AI
Collaboration	Involving patients and the care team in decisions about AI use
Compassion and Caring	Keeping the human relationship central even when a tool assists
Duty	Staying current on the tools you use and their limitations
Excellence	Holding AI-assisted care to the same standard as any other care
Integrity	Being transparent with patients about when and how AI is used
Social Responsibility	Considering whether a tool serves all patients equitably

1. APTA, 2026

Chapter Summary

- **Technology is never value-neutral.** Design choices embed latent values long before a tool reaches the clinic. Surfacing them is the first ethical act.
- **Biomedical ethics still applies.** Autonomy, beneficence, nonmaleficence, and justice all carry over, but each takes a new form when a model mediates the decision.
- **Four commitments recur across frameworks:** fairness, accountability, transparency and explicability, and safety with preserved human autonomy
- **Match oversight to the stakes.** Not all AI involvement should look the same. A documentation assist and a triage recommendation warrant different human control.
- **The profession already has tools for this.** The RIPS Model and the APTA Code of Ethics and Core Values give PT a structured way to reason through these cases.

Chapter 4

Applying AI Responsibly in Clinical Practice



Evaluating AI Tools: A Practical Checklist

Before adopting any AI tool, ask

- Can the outputs be interpreted and explained in a clinical context?
- What are its known limitations?
- How does it fit within existing clinical workflows?
- Who is accountable if the output is wrong?
- What is the tool's purpose and intended use?
- What data was it trained on? Is it representative of my patients?
- Has it been validated externally? By whom?
- Is it FDA-cleared for this specific indication?

Regulatory Landscape

- FDA classifies AI as Software as a Medical Device (SaMD)
- Three clearance pathways: 510(k), De Novo, Premarket Approval (PMA)
- 1,450+ AI/ML devices authorized by FDA through 2025—most via 510(k)¹
- FDA clearance does **NOT** necessarily mean improved clinical outcomes have been demonstrated
- Using AI tools off-label introduces clinical and legal risks
- **Moving target:** HHS released an AI strategy in December 2025² and is actively soliciting comments on accelerating AI adoption in clinical care

1. U.S. Food and Drug Administration, 2026

2. U.S. Department of Health and Human Services, 2025

The Human-in-the-Loop Principle

AI Output: Create a draft or suggestion



Clinician Review: Verify against patient presentation



Clinical Judgment Applied: Modify, confirm, or override



Patient-Centered Decision: Choose with the patient, not for them



Documented Rationale: What was accepted, changed, or rejected

Keep this flow in mind through every case scenario that follows. The AI output is a starting point, never the final word.

When AI Errors and Human Shortcuts Compound

- A well-documented pattern outside healthcare: several people have been wrongfully arrested after a facial-recognition match, but investigators had also skipped their own required verification steps before acting on that match^{1,2}
- The lesson isn't "facial recognition is uniquely unreliable"
 - AI error and a skipped human safeguard compound each other
- Healthcare parallel: AI documentation or decision-support error becomes a patient safety event less often because the AI was wrong, and more often because a clinician skipped a verification step

1. American Civil Liberties Union, 2021

2. American Civil Liberties Union, 2023

AI-Generated Documentation

Scenario: A PT uses an AI ambient listening tool that auto-generates a SOAP note after a patient visit


- The clinician **MUST** verify accuracy and completeness before finalizing
- AI may miss nuances: patient affect, nonverbal cues, prior visit context
- Transcription errors and speaker attribution mistakes are common
- Professional accountability remains with the clinician
- HIPAA considerations
 - Where is the data going?
 - Who has access?



AI-Generated Home Exercise Program

Scenario: A PT uses AI to generate an HEP for a patient with MS. The patient reports significant fatigue and pain during the visit. The AI-generated program does not reflect these real-time symptoms.

- AI outputs may not reflect real-time patient status
- The clinician must modify the program based on patient presentation and clinical judgment
- This is where human-centered AI principles become actionable
- AI supports, but does not replace, individualized clinical reasoning



AI for Clinical Decision Support in Stroke Rehab

Scenario: An AI system trained on younger adult data recommends treatment progression for a 78-year-old patient recovering from a stroke

- Bias concern: Nonrepresentative training data
- The clinician must critically evaluate whether the recommendation is appropriate for this specific patient
- Ask: “What population was this tool trained on?”
- Clinical judgment must override AI when the recommendation does not fit the patient

Fall Risk AI vs. Clinical Judgment

Scenario: A home health PT's AI-based fall risk algorithm flags a patient as low risk. PT's exam, gait deviations, home hazards, and medication changes suggest otherwise. The agency's workflow nudges clinicians to "follow the score" to keep documentation efficient.

- **Realm:** Is this an individual clinical call, or is organizational workflow pressure shaping the decision?
- **Individual Process:** Does the PT notice the conflict (sensitivity)? Do they judge it worth acting on (judgment)? Do they have the motivation and courage to document a higher risk level than the algorithm?
- **Situation:** Is this a dilemma (two legitimate pulls—efficiency vs. safety), or is it moral distress (the PT knows the right call but feels institutional pressure not to make it)?
- **Responsible action:** Document your own clinical risk determination and rationale; the algorithm informs, it does not decide



The Undisclosed AI Chatbot

Scenario: A clinic rolls out an AI chatbot that handles patient intake questions and sends home program reminders. Patients are not told they are interacting with AI rather than clinic staff.

- **Integrity:** Is withholding that AI is involved a failure of honesty with patients?
- **Accountability:** If the chatbot gives a patient incorrect information, who owns that error?
- **Autonomy (bioethics) and Collaboration (APTA Core Value):** Patients can't make an informed choice about engaging with the tool if they don't know it exists
- Hidden risk: Conversational AI often uses first-person language that feels person-like
- Ask your vendor: Are patient conversations used to further train the model, or shared with any third party?
- Responsible action: Disclose AI use clearly, include opt-in option where appropriate, and give patients a clear path to reach a human

AI-Driven Utilization Review

Scenario: An insurer's AI utilization-management tool flags a SNF patient for early discharge based on a cost-optimization model, ahead of what the treating PT believes is medically appropriate

- **Beneficence vs. Justice:** Is the tool optimizing for population-level cost fairness at the expense of this patient's best interest?
- **Nonmaleficence:** What is the harm risk of premature discharge for this specific patient?
- **Autonomy:** Has the patient been included in this conversation, or is the decision happening around them?
- **Responsible action:** Document the clinical rationale for continued care and use the payer's appeal process

Talking to Patients About AI

- Patients are increasingly aware of AI in healthcare
- Best practice: Explain the role of the tool while emphasizing that clinical judgment drives decision-making
- Transparency builds trust; vagueness erodes it
- Suggested language: “We used a technology tool to help organize some information, but all treatment decisions are made by me based on your specific needs and goals.”
- Informed consent for ambient documentation: Use an opt-in approach¹

1. Lawrence et al., 2025

Strategies for Responsible Integration

- Maintain a “human-in-the-loop” approach at all times
- Always review and edit AI-generated outputs before acting on them
- Evaluate tools before adoption using a structured checklist
- Use RIPS or the Four Principles when a case feels ethically ambiguous
- Stay informed about evolving regulations and institutional policies
- Document when and how AI tools contributed to clinical decisions

Getting Better Output: Prompt Engineering Basics

- Prompt engineering: refining how you ask a generative AI tool for something to get a more useful and accurate response
- Be specific: State the format, criteria, and context you want, rather than a vague, open-ended ask
- Break complex requests into steps: Smaller, well-defined prompts produce more reliable output than one large request
- Ask the tool to state its confidence or flag missing information, rather than letting it fill gaps with a best guess
- This does not replace clinical review; a better-engineered prompt still produces a draft that needs the same verification as any other AI output
- If you use these tools for patient education materials or documentation support, the same review standard from Case 1 and Case 2 still applies

Chapter Summary

- **Evaluate before you adopt.** Validation population, transparency, oversight, and data handling are what separate a usable tool from a liability.
- **Regulation lags practice.** Most tools reach the clinic without FDA review, so appraisal falls to the clinician and the organization.
- **Keep a human in the loop.** AI output is a draft and requires review; judgment, decision-making, and a documented rationale remain the clinician's job.
- **Errors compound quietly.** A flawed output plus a human shortcut (automation bias, time pressure) is more dangerous than either alone.
- **Integration is a skill.** Disclosure, workflow strategy, and better prompting all change the quality of what you get back.

Course Summary

- AI is already in rehabilitation practice
 - The question is not whether to engage, but how to engage responsibly
- Bias, overreliance, and lack of transparency are real risks requiring active mitigation
- General bioethics (Four Principles) and healthcare AI frameworks (WHO, AMIA, SHIFT) converge on the same core ideas
- Professional specific tools (i.e. RIPS Model and the APTA Code of Ethics) give you a way to apply these principles case by case
- Clinicians remain responsible for all patient care decisions, regardless of AI involvement

Bibliography

Medbridge

AI and the Clinician: Navigating Responsible Use in Rehabilitation

Aimee E. Perron, PT, DPT, NCS, CEEAA, FNAP

- Adegunle, F., Chhatwal, K., Arab, S., Alabduljabar, M. S., Raslan, M. A., Sayed, O., & Goldsweig, A. M. (2026). Bias and oversight in clinical AI: A review of decision support tools and equity frameworks. *Journal of General Internal Medicine*, 41(7), 1957–1968. <https://doi.org/10.1007/s11606-026-10229-5>
- Adnan, T., Abdelkader, A., Liu, Z., Hossain, E., Park, S., Islam, M. S., & Hoque, E. (2025). AI-powered speech-based screening tool for Parkinson's disease. *npj Parkinson's Disease*.
- AI in health care: Closing the revenue cycle gap. (2025). *The American Journal of Managed Care*, 31(4), 161–162.
- Alsobhi, M., Sachdev, H. S., Chevidikunann, M. F., Basuodan, R., Dhanesh Kumar, K. U., & Khan, F. (2022). Facilitators and barriers of artificial intelligence applications in rehabilitation: A mixed-method approach. *International Journal of Environmental Research and Public Health*, 19(23), Article 15919. <https://doi.org/10.3390/ijerph192315919>
- American Civil Liberties Union. (2021). *Williams v. City of Detroit*.
- American Civil Liberties Union. (2023, February). *Porcha Woodruff wrongful arrest, Detroit*.
- American Physical Therapy Association. (2026). *Code of ethics for the physical therapy profession*.
- APTA House of Delegates. (2024a). *Digital health technologies, digital therapeutics, and AI in physical therapist practice* (P07-24-12-11).
- APTA House of Delegates. (2024b). *Ethical and effective integration of artificial intelligence across PT practice, education and research* (P07-24-10-09).
- APTA House of Delegates. (2024c). *Inappropriate use of artificial intelligence by payers* (P07-24-11-10).
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In

Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623). <https://doi.org/10.1145/3442188.3445922>

- Branco, J. P., Graça, D. T., Costa, E., Branco, C. A., Sampaio, F., Barroso, J., Pinheiro, J. P., Cantista, P., Nunes, R., & Lains, J. (2026). AI-based rehabilitation: The way forward in addressing unmet needs in musculoskeletal disease. *Frontiers in Public Health*, 14, Article 1773733. <https://doi.org/10.3389/fpubh.2026.1773733>
- Brunyé, T. T., Mitroff, S. R., & Elmore, J. G. (2025). Artificial intelligence and computer-aided diagnosis in diagnostic decisions: 5 questions for medical informatics and human-computer interface research. *Journal of the American Medical Informatics Association*.
- Budhu, J. A., Anderson, N., Branson, C., et al. (2025). Health equity considerations in the age of artificial intelligence. *Neurology*.
- Chiang, S., Picard, R. W., Chiong, W., et al. (2021). Guidelines for conducting ethical artificial intelligence research in neurology. *Neurology*.
- Feng, Q. J., Harte, M., Carey, B., et al. (2025). The risks of artificial intelligence: A narrative review and ethical reflection from an oral medicine group. *Oral Diseases*.
- Feng, S., van Walraven, C., Lalu, M. M., Moloo, H., Musselman, R., & McIsaac, D. I. (2022). Derivation and external validation of a 30-day mortality risk prediction model for older patients having emergency general surgery. *British Journal of Anaesthesia*, 129(1), 33–40.
- Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
- García-Magariño, I., Plaza, I., & Lacuesta, R. (2023). A governance and assessment model for ethical artificial intelligence in healthcare. *ERCIM News*, 134, 18–19. <https://ercim-news.ercim.eu/en134/special/a-governance-and-assessment-model-for-ethical-artificial-intelligence-in-healthcare>
- Giacobello, M. L. (2025). Informed consent and bioethical advances in clinical settings. *Frontiers in Psychology*.
- Hao, J., Yao, Z., Tang, Y., Remis, A., Wu, K., & Siu, K. C. (2025). Artificial intelligence in physical therapy education: Evaluating clinical reasoning performance in musculoskeletal care using ChatGPT. *Musculoskeletal Care*.
- Horvath, D., & Lorincz, N. S. (2026). Artificial intelligence in patient-centered care and macro-, meso-, and micro-level determinants of rehumanization and

- dehumanization: Qualitative interview study. *Journal of Medical Internet Research*.
- Jackson, N. J., Brown, K. E., Miller, R., et al. (2026). Factors influencing the effectiveness of artificial intelligence-assisted decision-making in medicine: A scoping review. *Journal of the American Medical Informatics Association*.
- Jain, S. S., Elias, P., Poterucha, T., et al. (2024). Artificial intelligence in cardiovascular care, Part 2: Applications. *Journal of the American College of Cardiology*.
- Ke, Y., Jin, L., Ong, J. C. L., et al. (2026). AI-induced never-skilling in medical education. *Nature Medicine*, 32(6), 1997–2006.
<https://doi.org/10.1038/s41591-026-04438-y>
- Kücking, F., et al. (2026). Impact of AI recommendation correctness on diagnostic accuracy in clinical decision-making. *International Journal of Medical Informatics*, 207, Article 106223.
- Labkoff, S., Oladimeji, B., Kannry, J., et al. (2024). Toward a responsible future: Recommendations for AI-enabled clinical decision support. *Journal of the American Medical Informatics Association*.
- Lawrence, K., et al. (2025). Informed consent for ambient documentation using generative AI in ambulatory care. *JAMA Network Open*.
- Modi, T. B. (2023). Artificial intelligence ethics and fairness: A study to address bias and fairness issues in AI systems, and the ethical implications of AI applications. *Revista Review Index Journal of Multidisciplinary*, 3(2), 24–35.
<https://doi.org/10.31305/rrijm2023.v03.n02.004>
- Morone, G., Carbonetti, R., Martino Cinnera, A., Malfitano, C., Bisirri, A., Negrini, F., & REHALISE Group. (2025). Artificial intelligence in rehabilitation: A living systematic mapping review—First release. *European Journal of Physical and Rehabilitation Medicine*, 61(5), 694–708.
<https://doi.org/10.23736/S1973-9087.25.09331-1>
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1).
<https://doi.org/10.6028/NIST.AI.100-1>
- Nesa, L., Rony, M. K. K., Chowdhury, S., et al. (2025). Artificial intelligence in healthcare: A scoping review of medical professionals' acceptance and institutional challenges in implementation. *Journal of Evaluation in Clinical Practice*.
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>

- Olawade, D. B., Plabon, S. B., Ojo, A., et al. (2026). Human in the loop artificial intelligence in healthcare: Applications, outcomes, and implementation challenges. *International Journal of Medical Informatics*.
- Organisation for Economic Co-operation and Development. (2024). *Recommendation of the Council on artificial intelligence* (OECD/LEGAL/0449). (Original work published 2019)
- Osnat, B. (2025). Patient perspectives on artificial intelligence in healthcare: A global scoping review of benefits, ethical concerns, and implementation strategies. *International Journal of Medical Informatics*, 203, Article 106007. <https://doi.org/10.1016/j.ijmedinf.2025.106007>
- Otokiti, A. U., Shih, H. J., & Williams, K. S. (2025). Gender and racial bias unveiled: Clinical AI/ML algorithms are fanning the flames of inequity. *Oxford Open Digital Health*, 3, Article oqaf027.
- Palmieri, S., Robertson, C. T., & Cohen, I. G. (2026). New guidance on responsible use of AI. *JAMA*, 335(3), 207–208. <https://doi.org/10.1001/jama.2025.23059>
- Rabinowitz, A., Trauger, M., & Williams, M. (2026). Reining in unbridled AI enthusiasm: Protecting the integrity of rehabilitation science and clinical care. *Archives of Physical Medicine and Rehabilitation*.
- Sahni, N. R., & Carrus, B. (2023). Artificial intelligence in U.S. health care delivery. *The New England Journal of Medicine*.
- Shah, N. H., Halamka, J. D., Saria, S., et al. (2024). A nationwide network of health AI assurance laboratories. *JAMA*.
- Siala, H., & Wang, Y. (2022). SHIFTing artificial intelligence to be responsible in healthcare: A systematic review. *Social Science & Medicine*, 296, Article 114782. <https://doi.org/10.1016/j.socscimed.2022.114782>
- Solomonides, A. E., Koski, E., Atabaki, S. M., Weinberg, S., McGreevey, J. D., Kannry, J. L., Petersen, C., & Lehmann, C. U. (2022). Defining AMIA's artificial intelligence principles. *Journal of the American Medical Informatics Association*, 29(4), 585–591. <https://doi.org/10.1093/jamia/ocac006>
- Swisher, L. L., Arslanian, L. E., & Davis, C. M. (2005). The Realm-Individual Process-Situation (RIPS) model of ethical decision-making. *HPA Resource*, 5(3), 1, 3–8.
- U.S. Department of Health and Human Services. (2025, December 4). *HHS artificial intelligence strategy*.
- U.S. Food and Drug Administration. (2026). *Artificial intelligence-enabled medical devices*.

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*.

World Health Organization Regional Office for Europe. (2026). *Statement: Govern AI in health before the gaps become irreversible; WHO European Region AI readiness assessment*.

Xia, L., Wang, Y., Zhang, Y., et al. (2026). The Relational Care-AI Alignment Framework: An ethical model for AI involvement in person-centred fundamental care. *Journal of Advanced Nursing*.

Yu, K. H., Healey, E., Leong, T. Y., Kohane, I. S., & Manrai, A. K. (2024). Medical artificial intelligence and human values. *The New England Journal of Medicine*.

Zhang, L., Tashiro, S., Mukaino, M., & Yamada, S. (2023). Use of AI large language models as a clinical tool in rehabilitation medicine: A comparative test case. *Journal of Rehabilitation Medicine*, 55, Article jrm13373.