

Human judgment and automated action

When may AI assist, recommend or act?

Human oversight becomes meaningful when someone has the competence, time and authority to intervene.

Recommended starting position

Separate assistance, consequential recommendations and actions. Assign a responsible person, define what must be reviewed, and preserve a meaningful route to explanation, correction and human reconsideration. Authorize connected actions separately, within limited permissions.

A nominal human approval can become a rubber stamp if the reviewer cannot inspect the basis or correct the result. Faster decisions are not necessarily better decisions. Stakes rise when a system can change records, send messages or affect access to an opportunity.

What the evidence supports

Purdue limits sole reliance on AI for specified consequential decisions. North Shore includes human oversight and routes for questioning outcomes. NIST supplies a broader voluntary framework for managing AI risk.

Establish these foundations

- Name the accountable decision owner.
- Give reviewers the information and authority needed to challenge a result.
- Make correction, suspension and recovery possible.

The choices to make locally

Assist

Consider this when: AI drafts, explains or organizes material for a person.

Tradeoff: Still requires verification appropriate to the task and protection of the information involved.

Recommend

Consider this when: AI influences a decision affecting a person or institution.

Tradeoff: Define the evidence a reviewer needs, when review occurs and how a person can challenge an outcome.

Act

Consider this when: AI can send, change, spend, approve or otherwise affect the world.

Tradeoff: Use explicit permissions, confirmation for consequential actions, action records, limits and a way to stop and recover.

Illustrative campus case

An advising assistant can change a schedule.

The same product can explain course requirements, recommend an individual schedule and change an enrollment record.

Those functions require different approval and accountability arrangements.

1. For explanations, establish current approved information and access to an advisor.
2. For recommendations, define review of academic and financial consequences and a correction route.
3. For record changes, explicitly approve the action, permissions, confirmation, audit record and recovery process.

What the team produces: An authorization that distinguishes advice from action and identifies the responsible owner.

Who is responsible

Board and trustees: Ask who can explain, correct and stop consequential AI-influenced decisions.

Institutional leaders: Resource human review and ensure authority is matched to responsibility.

AI task force: Specify the boundaries between assistance, recommendations and actions for each use.

Illustrative language to adapt

Consequential AI-supported decisions must have a designated human decision owner and a defined review and reconsideration route. Reviewers must be able to challenge the output and obtain appropriate evidence. AI systems may take actions only within documented permissions, with controls for confirmation, monitoring, suspension and recovery appropriate to the consequences.

Read alongside model policy sections 6, 9, 11. Reconcile local authority and requirements before adoption.

Evidence of implementation

- A named reviewer with relevant expertise and authority.
- Records showing substantive review of representative difficult cases.
- A usable human reconsideration route and tested stop mechanism.

Reconsider when: Reassess when a system moves from advice to action, affects additional rights or services, or review becomes ineffective.

Supporting sources

[Purdue University / Artificial Intelligence Use](#)

[North Shore Community College / Artificial Intelligence Principles and Policies](#)

[National Institute of Standards and Technology / AI Risk Management Framework](#)