A group of business professionals in a meeting room, looking at documents and screens. The scene is dimly lit with a purple hue, suggesting a modern, high-tech environment. Several people are seated around a long table, some holding papers and others looking at large digital displays in the background. The overall atmosphere is professional and focused.

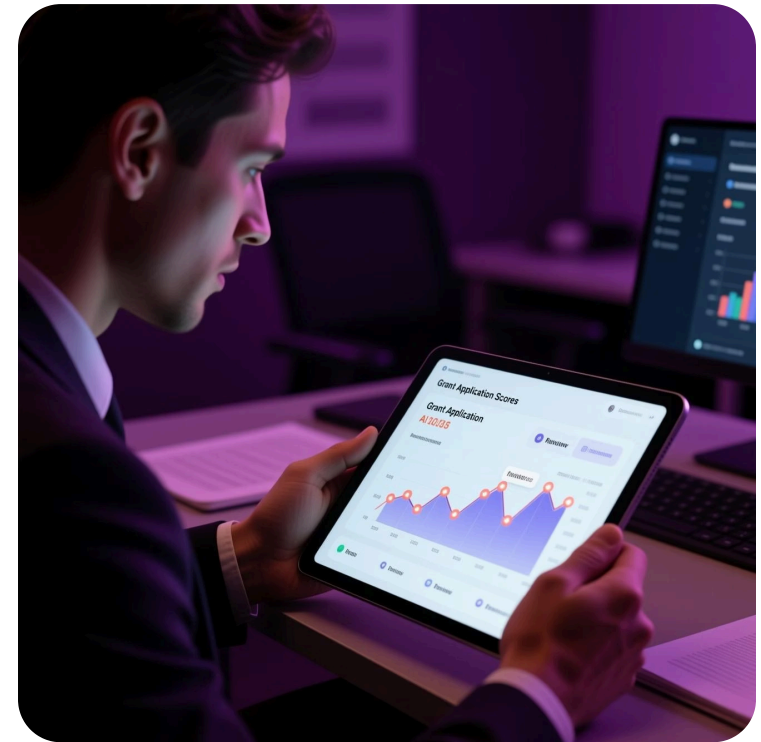
Keeping Humans in the Loop: The Future of AI-Augmented Peer Review for Competitive Grants

As government agencies embrace digital transformation, AI offers powerful capabilities to streamline grant application peer reviews. Yet the most effective solutions maintain human judgment at their core. This document explores how human-in-the-loop approaches to AI implementation ensure accountability, fairness, and optimal outcomes in the grant review process, with specific focus on Revelo Software's eReviewer system as a model of responsible AI augmentation in government decision-making.

The Challenge: Balancing AI Efficiency with Human Judgment

Government agencies face mounting pressure to process increasing volumes of grant applications with limited resources. AI promises to alleviate this burden through automation, but introduces new challenges. When AI guidance precedes human judgment, a phenomenon called **automation bias** can occur, where reviewers defer too heavily to algorithmic suggestions, potentially compromising decision quality.

Studies reveal that organizations across sectors have scaled back AI investments due to trust issues. Government contexts present even higher stakes, where AI systems must meet strict standards for transparency, auditability, and compliance with federal regulations. The January 2025 White House memorandum (M-25-22) emphasized the need for trustworthy, human-supervised AI in federal acquisitions.



The critical challenge is implementing AI in ways that enhance human decision-making rather than replacing it, particularly in high-stakes contexts like competitive grant reviews where nuance and ethical judgment remain essential.

Speed vs. Quality

Agencies need faster reviews without sacrificing thoroughness or accuracy.

Trust & Transparency

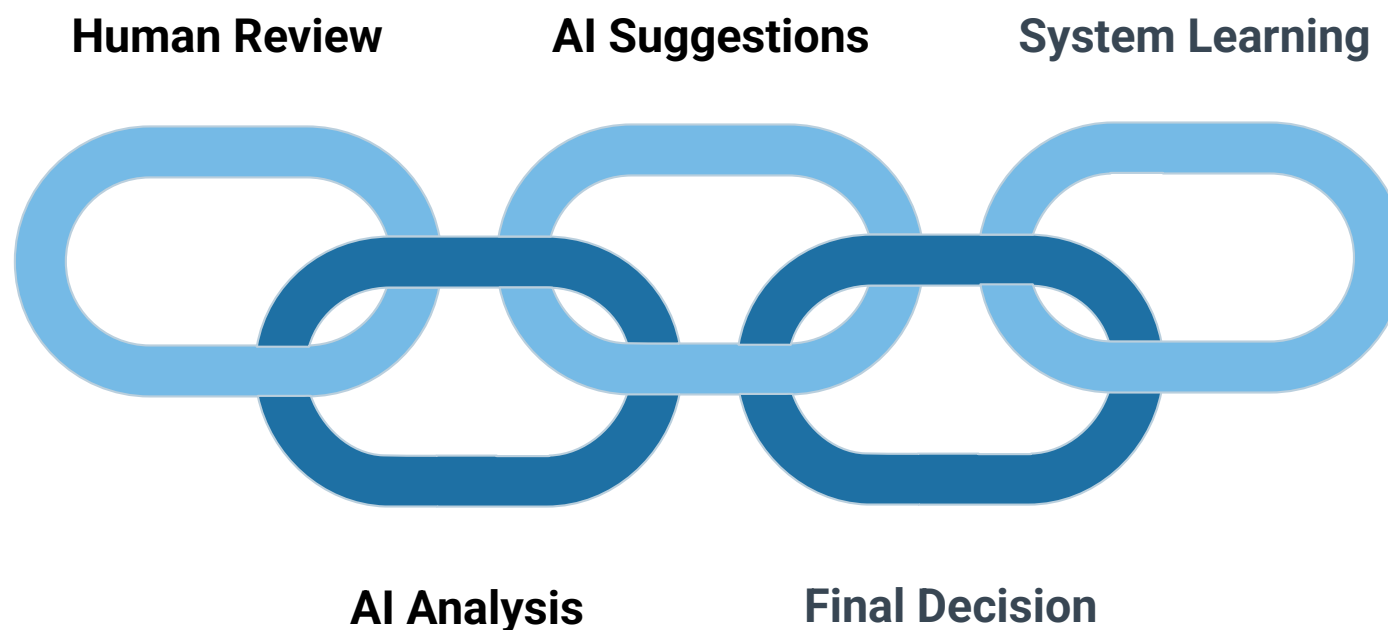
AI systems must be explainable and accountable, with clear audit trails.

Regulatory Compliance

Solutions must align with federal guidance on responsible AI implementation.

Understanding Human-in-the-Loop (HIL) AI

Human-in-the-Loop (HIL) AI represents a collaborative approach where human intelligence remains central to the decision-making process, with AI serving as a supporting tool rather than an autonomous decision-maker. This framework proves especially valuable in contexts requiring nuanced judgment, ethical considerations, and accountability—all hallmarks of the grant review process.



Unlike fully automated approaches, HIL systems recognize the complementary strengths of human and artificial intelligence. Humans excel at contextual understanding, ethical reasoning, and applying domain expertise to novel situations. AI excels at processing large volumes of data, identifying patterns, and maintaining consistency across many similar cases.

When properly implemented, HIL systems create a virtuous cycle of improvement. Human reviewers maintain decision authority while gaining insights from AI analysis. Meanwhile, the AI system learns from human feedback to refine its suggestions over time, becoming increasingly aligned with expert judgment without assuming control of the process.

"The most effective AI systems are those that respect the complexity of human decision-making. The key question is no longer 'Can AI do this?' but rather, 'How can AI help humans do this better?'"

This approach directly addresses automation bias by positioning AI as an advisor rather than an authority, preserving human agency while still leveraging computational advantages. For government grant programs specifically, HIL designs ensure compliance with regulatory requirements for transparency, explainability, and appropriate oversight of automated systems.

The Science Behind Human-AI Collaboration

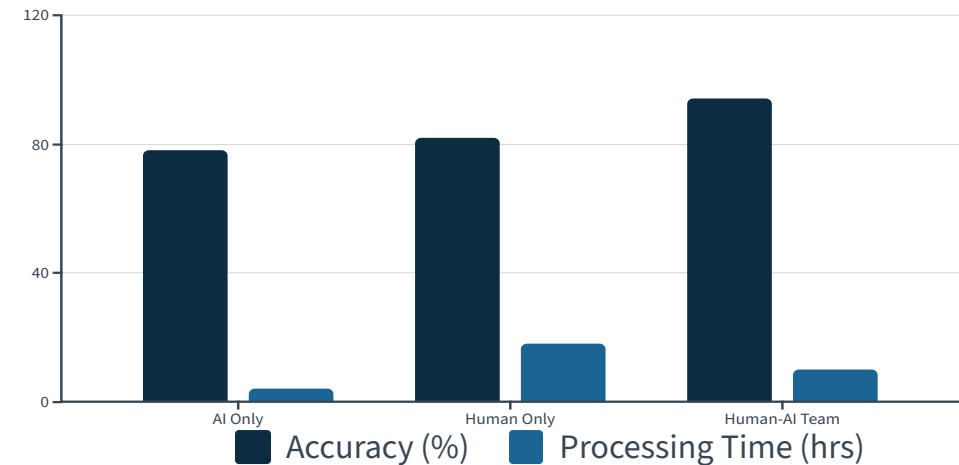
Recent research into human-AI collaboration reveals compelling evidence for the superiority of hybrid approaches over either standalone AI or unaided human judgment. Studies in Human-in-the-Loop Reinforcement Learning (HITL-RL) demonstrate that AI systems guided by human input—especially during ambiguous or low-confidence tasks—consistently outperform autonomous alternatives.

The research by Natarajan et al. (2025) found that human-AI teams achieved 23% better accuracy than either humans or AI working independently when evaluating complex applications with multiple evaluation criteria. This performance gain stems from complementary capabilities: humans provide contextual understanding and ethical judgment that AI lacks, while AI contributes consistency and can identify patterns across large datasets.

Agudo et al. (2024) further demonstrated that the *sequence* of human-AI interaction significantly impacts outcomes. When humans form independent judgments before seeing AI suggestions, the final decisions show greater accuracy and less automation bias than when AI guidance precedes human review.

Critically, Alon-Barkat and Busuioc (2023) found that government decision-makers display "selective adherence" to algorithmic advice—following it when it aligns with their expertise but critically evaluating it when it contradicts their judgment. This selective adherence represents an ideal middle ground between wholesale rejection of AI assistance and uncritical acceptance of all algorithmic outputs.

These findings provide a scientific foundation for designing peer review systems that maximize the benefits of both human expertise and AI capabilities while minimizing the risks of either working in isolation.



Comparative Performance in Complex Decision Tasks

Based on Natarajan et al. (2025)

eReviewer: A Model HIL Implementation for Grant Review

Revelo Software's eReviewer exemplifies effective Human-in-the-Loop design principles in practice. This Peer Review Management System (PRMS) was purpose-built for government grant competitions with HIL philosophy embedded in its core architecture.



Independent Review

Reviewers first conduct independent assessments and scoring without AI influence, preserving authentic human judgment.



AI Analysis

System analyzes review patterns to identify potential inconsistencies, outliers, or areas needing further examination.



Insights Presentation

AI insights are presented as suggestions, never directives, with clear explanations of the reasoning behind each observation.



Human Deliberation

Reviewers and panels discuss AI insights alongside their own assessments, maintaining full authority over final decisions.

Unlike systems that position AI as the primary decision-maker with humans merely providing oversight, eReviewer maintains reviewer autonomy throughout the process. All AI insights are presented *after* initial scoring to avoid anchoring bias, and reviewers have complete discretion to accept or reject these insights based on their expertise.

The system incorporates comprehensive logging and transparency features, allowing program officers to track the rationale behind every decision—whether driven by human judgment, AI suggestion, or a combination of both. This audit trail ensures compliance with federal requirements for explainable AI while providing valuable data for continuous improvement.

Key Features of HIL-Based Peer Review Systems



Sequential Processing

Human reviewers form independent judgments before AI analysis is presented, preventing automation bias and preserving authentic expert assessment.



Reviewer Authority

All AI suggestions can be overridden, with intuitive interfaces that make human judgment the default authority in the review process.



Pattern Recognition

AI analyzes across applications to identify inconsistencies, trends, and outliers that might escape human notice when reviewing applications individually.

Transparent Reasoning

Every AI insight is accompanied by clear explanations of the underlying logic and data sources, building reviewer trust through transparency.

Comprehensive Logging

The system maintains detailed records of all human decisions, AI suggestions, and reviewer responses, creating an audit trail for compliance and continuous improvement.

Learning Loop

Reviewer feedback on AI suggestions feeds back into the system, improving future rounds of review while preserving human oversight of the learning process.

Effective HIL peer review systems emphasize collaboration rather than replacement. They provide tools that enhance reviewer capabilities—such as simplified comparison views across multiple applications or automatic checking for scoring consistency—while ensuring humans retain control over substantive judgments about application quality and merit.

These systems also acknowledge the social dimension of peer review by supporting panel deliberation with useful visualizations and insights, rather than attempting to bypass the valuable consensus-building that occurs through reviewer discussion. The goal is augmenting the collective intelligence of review panels, not supplanting it with algorithmic alternatives.

Real-World Impact: Case Studies in HIL Implementation

Practical implementations of Human-in-the-Loop AI in grant peer review have demonstrated significant improvements in efficiency, consistency, and reviewer satisfaction across various government agencies. The following case studies illustrate these benefits with measurable outcomes.

Federal Science Agency

Implemented eReviewer for a competitive research grants program with 1,200 annual applications and a panel of 80 expert reviewers.

- Reduced total review cycle time by 37%
- Improved inter-reviewer scoring consistency by 28%
- 89% of reviewers reported greater confidence in final decisions
- Identified 15% more potential conflicts of interest than manual processes

State Economic Development Office

Adopted HIL approach for small business innovation grants program processing 800 applications annually.

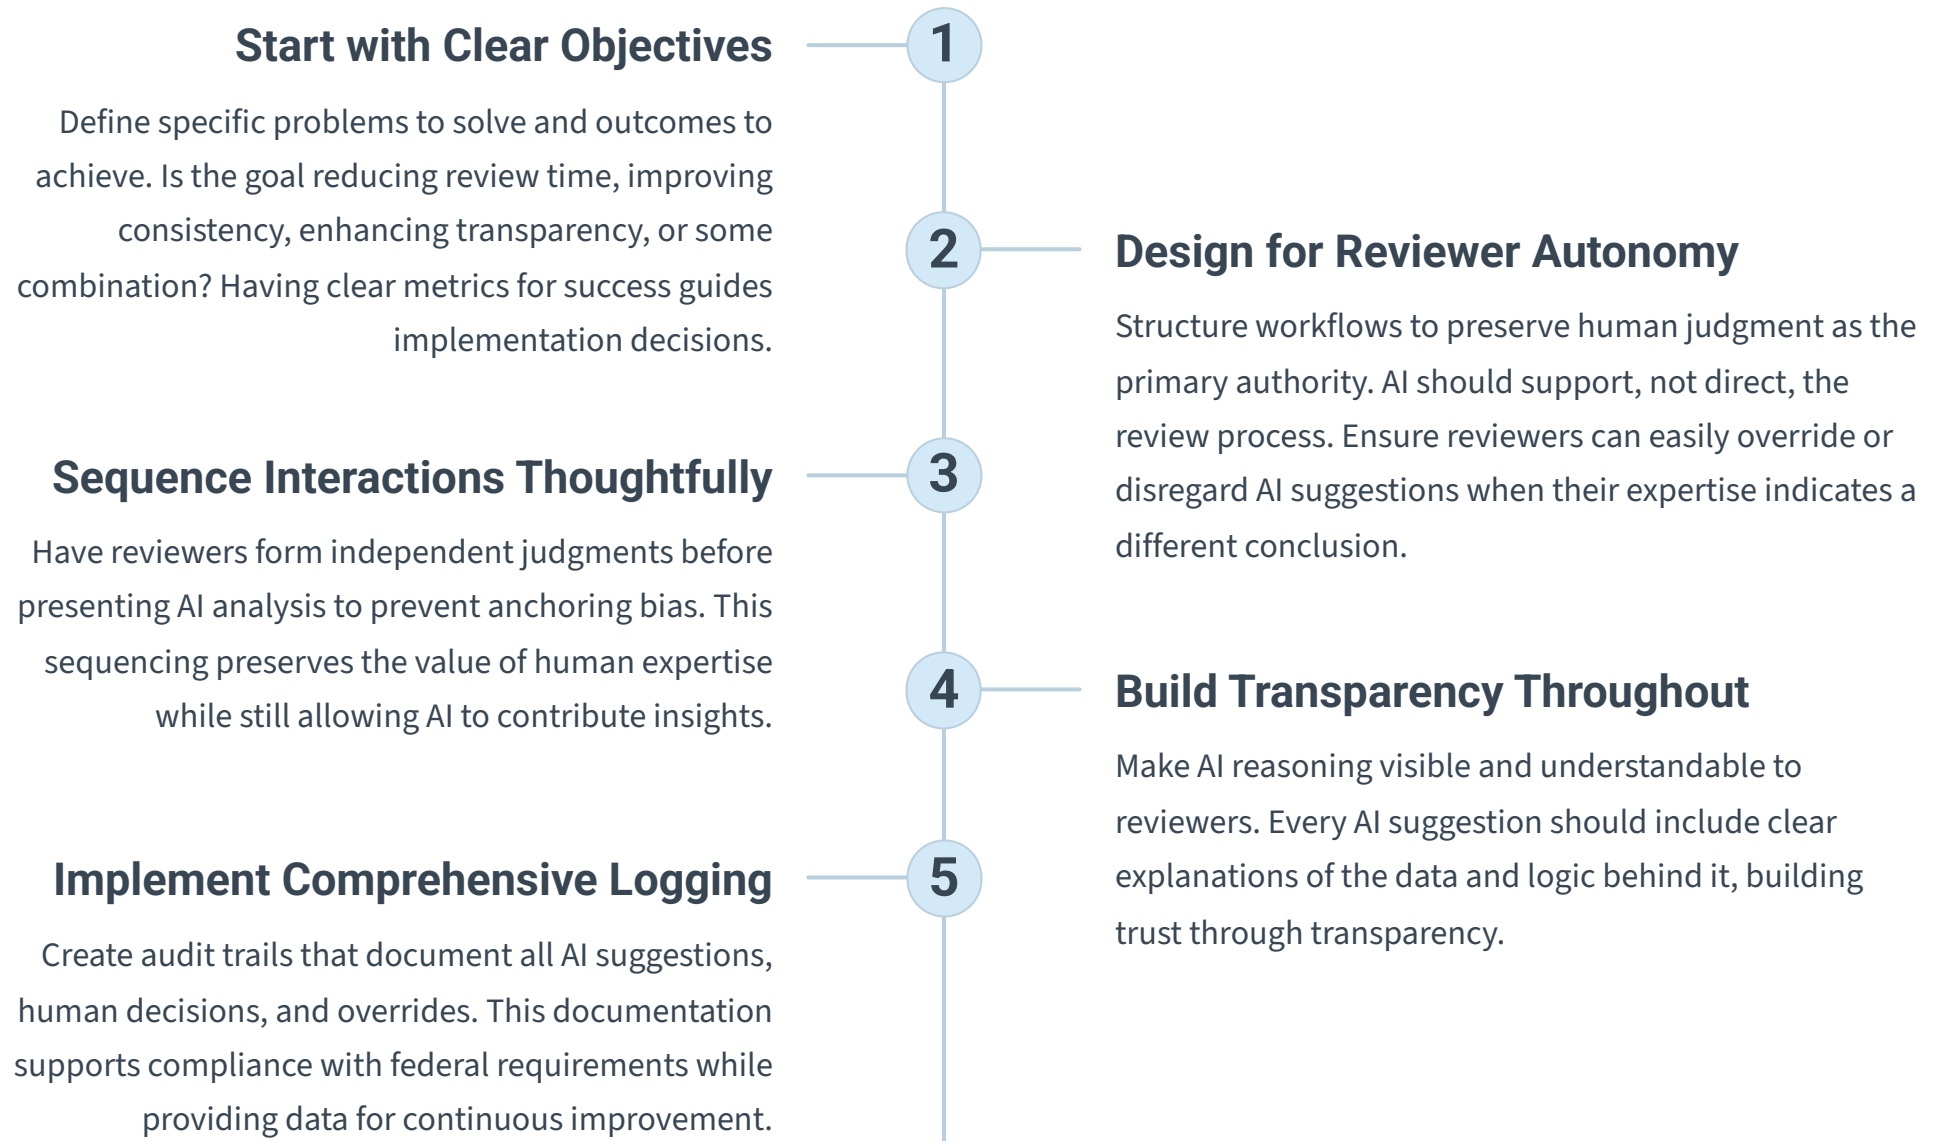
- Reduced administrative burden by 42 hours per review cycle
- Improved applicant feedback quality through AI-assisted response generation
- Identified scoring drift patterns that led to revised evaluation criteria
- Maintained full compliance with state transparency requirements

A key finding across implementations was that reviewer autonomy and proper sequencing of human judgment before AI analysis were critical success factors. In one case study, a pilot that presented AI-generated scores before human review resulted in significant clustering around AI suggestions—demonstrating automation bias in action. When the process was reversed to prioritize independent human judgment first, the final decisions showed greater diversity and ultimately higher quality outcomes.

Equally important was the system's ability to learn from reviewer feedback. Over multiple funding cycles, the AI components became increasingly aligned with expert consensus, reducing the frequency of overrides while maintaining human control of the process. This continuous improvement loop created a virtuous cycle where both human reviewers and AI capabilities strengthened over time.

Implementation Best Practices

Successfully integrating Human-in-the-Loop AI into grant peer review requires careful planning and execution. These best practices, derived from both research and practical experience, can guide government agencies toward effective implementation.



Change management is equally important for successful implementation. Reviewers may initially resist AI assistance due to skepticism or concerns about their role. Effective training that emphasizes how AI augments rather than replaces human judgment can address these concerns. Demonstrating early wins—such as reduced administrative burden or helpful pattern recognition—builds acceptance among review panels.

Pro Tip: Consider a phased implementation approach, starting with AI features that have the clearest value proposition and lowest risk, such as checking for scoring consistency or identifying potential conflicts of interest. This builds reviewer trust before introducing more sophisticated capabilities.

Finally, ensure ongoing evaluation of the system's impact. Regularly assess both quantitative metrics (review times, consistency scores) and qualitative feedback from reviewers and program officers to guide refinements to the human-AI collaboration model.

Regulatory Compliance and Ethical Considerations

Government agencies implementing AI-augmented peer review must navigate a complex regulatory landscape. The January 2025 White House memorandum (M-25-22) establishes specific requirements for federal AI acquisitions, emphasizing cross-functional oversight, privacy protection, and explainability in mission-critical systems like grant evaluation.

Human-in-the-Loop approaches directly address these requirements by maintaining clear human accountability for decisions while providing comprehensive documentation of how AI insights influence the process. This transparency is essential for both compliance and building public trust in government decision-making.

Beyond regulatory requirements, ethical considerations should guide implementation. Grant decisions often impact historically underserved communities, making fairness and bias mitigation paramount. HIL systems allow human reviewers to identify and correct potential biases in AI suggestions, providing an essential safeguard for equitable decision-making.

Data governance represents another critical compliance area. Grant applications often contain sensitive information about research plans, financial details, and intellectual property. Systems must implement appropriate access controls, encryption, and data minimization practices to protect this information while still enabling effective review.

Revelo Software's eReviewer addresses these requirements through several mechanisms:

- Role-based access controls that limit data exposure based on reviewer assignments
- Comprehensive audit logs that document all system activities for compliance verification
- Explainable AI features that make algorithmic reasoning transparent to users
- Bias detection tools that flag potential disparities in scoring patterns for human review
- FedRAMP compliance for cloud deployments, ensuring alignment with federal security standards

By maintaining human oversight throughout the process while implementing these technical safeguards, agencies can leverage AI benefits while meeting their regulatory obligations and upholding the public trust.



Privacy Protection

Systems must safeguard applicant and reviewer personal information with appropriate technical and procedural controls.



Explainability

AI reasoning must be transparent and understandable to both reviewers and applicants if requested.



Fairness

Agencies must evaluate systems for potential bias and implement appropriate mitigation measures.

Conclusion: The Future of Human-AI Partnership in Grant Review

As government agencies continue their digital transformation journey, the integration of AI into peer review processes represents both an opportunity and a responsibility. The evidence clearly demonstrates that Human-in-the-Loop approaches offer the optimal balance between efficiency and judgment, between innovation and accountability.

37%

Reduction in Review Time

HIL implementation streamlines administrative tasks while preserving quality deliberation

28%

Improved Consistency

AI-assisted review helps identify and address scoring variations across reviewers

89%

Reviewer Confidence

Reviewers report greater confidence in final decisions with AI support

The future of grant peer review lies not in replacing human judgment with artificial intelligence, but in creating thoughtful partnerships that leverage the strengths of both. By keeping humans firmly in the loop—with AI serving as a tool for insight rather than an autonomous decision-maker—agencies can achieve faster, more consistent reviews while maintaining the critical elements of human wisdom, ethical judgment, and contextual understanding.

Systems like eReviewer demonstrate how this partnership can work in practice: preserving reviewer autonomy while providing valuable AI-powered insights, maintaining complete transparency in how algorithms contribute to the process, and creating virtuous cycles of improvement as humans and AI learn from each other over time.

For government program officers and grant managers evaluating AI-assisted peer review options, the key consideration should not be how much human judgment can be automated away, but rather how effectively AI can be integrated to support and enhance human decision-making. The most successful implementations will be those that view AI as a collaborator rather than a replacement, recognizing that the combined strengths of human and artificial intelligence far exceed what either could achieve alone.

By embracing Human-in-the-Loop approaches, government agencies can confidently move forward with AI-augmented peer review systems that deliver on the promise of digital transformation while upholding the highest standards of fairness, accountability, and sound judgment in the allocation of public resources.