

# Case Study: Enhancing Clinical Research Ethics Review with AI – The REAI Project

**Authors:** Rahul Mahurkar & Bibekananda Pahi

**Date of publication:** 25<sup>th</sup> April, 2025

## Key Stakeholders:

- **Client, A South African NGO** – Project Owner and Platform Provider
- **Zepto Digital Labs** – Technology Partner responsible for AI model development, integration, and deployment
- **Research Ethics Committees (RECs)** – End users and beneficiaries of the REAI solution
- **Clinical Researchers / Applicants** – Indirect stakeholders impacted by improved application turnaround
- **Data Protection Authorities** – Compliance stakeholders ensuring GDPR, POPIA, and GCP standards adherence

## 1. Executive Summary

The client, in partnership with technology experts (Zepto Digital Labs, India), has developed and deployed **REAI**—a pioneering Artificial Intelligence (AI) solution integrated into the **RHInnO Ethics** platform. REAI supports Research Ethics Committees (RECs) in reviewing clinical research applications by automating document analysis, validating applicant data, and generating reviewer-ready summaries and suggestions aiding the Research Ethics Committees (RECs) to arrive at a final approval decision. Designed to empower RECs in **low- and middle-income countries (LMICs)**, this innovation significantly enhances the quality,

speed, and consistency of ethical reviews, enabling better decision-making while reducing workload.

## 2. Background / Context

The **RHInnO Ethics** platform is a cloud-based solution that digitizes and streamlines the end-to-end ethics review process for clinical research. Its mission is to strengthen ethical review systems, particularly in **low- and middle-income countries (LMICs)**, by making them more efficient, transparent, and accessible.

Despite RHInnO's success, Research Ethics Committees continued to face:

- An increasing **volume and complexity** of research submissions.
- A **manual, time-consuming review process**, arranging between **2 to 6 months** for each new application processing, prone to variation.
- Difficulty in **validating consistency** between applicant-provided answers and supporting documents.
- High **cognitive load** for reviewers, especially when reviewing lengthy protocols and attachments.
- High **cost** of RECs.

Recognizing this, Client initiated the **REAI project** (Review Ethics Artificial Intelligence) to enhance RHInno Ethics with intelligent capabilities that support—not replace—human reviewers.

### 3. The Problem / Need

#### Key Issues Identified:

- Reviewers spent significant time navigating lengthy submissions, extracting relevant information manually.
- Submissions often included inconsistencies between applicant answers and documents, which were hard to detect.
- There was no easy way to provide a high-level overview or summary of each submission, making triage and prioritization difficult.
- Structured review forms required repetitive manual input, slowing down the process and introducing the risk of oversight.

The goal was to develop an AI solution that:

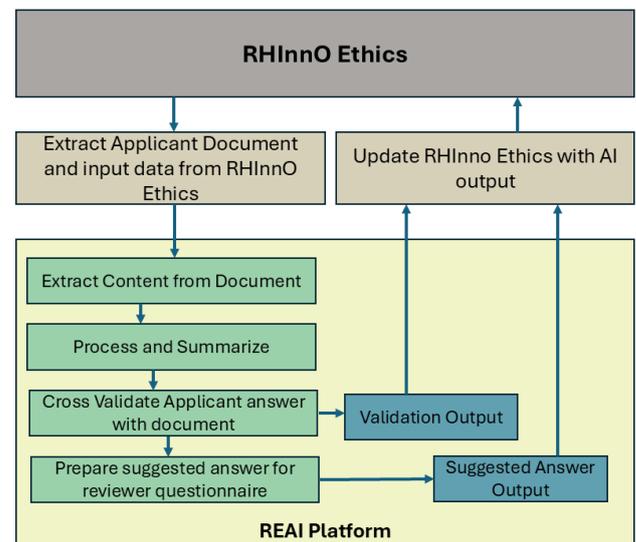
- **Improves review efficiency** and consistency.

- **Assists reviewers** with validated, structured information.
- **Supports standardization** across review committees.
- Upholds the **ethical and regulatory integrity** of the review process.

### 4. Solution Approach

The client and Zepto Digital Labs (its technology partners) designed REAI, a multi-component AI platform tightly integrated into the RHInno Ethics workflow. Here's how it works:

#### Workflow between RHInno Ethics and an AI-



#### powered REAI Platform

#### Core Capabilities:

- **Document Ingestion & OCR:** Converts scanned PDF documents into searchable text using Optical Character Recognition (OCR) techniques (e.g., Tesseract).
- **Content Extraction & Summarization:** Uses advanced NLP models (e.g., BART, T5) to generate concise summaries of long clinical documents such as protocols, consent forms, and investigator brochures.

- **Answer Validation:** Applies transformer-based QA models (e.g., BERT, DistilBERT) to verify whether the applicant’s questionnaire responses align with uploaded documents.
- **Reviewer Questionnaire Assistance:** Generates draft answers for reviewer forms using AI (e.g., Flan-T5), which can then be validated or edited by human reviewers.
- **Discrepancy Alerts:** Flags potential mismatches or missing data, drawing reviewer attention to issues.

#### Technical Stack:

- **Hugging Face Transformers** for NLP tasks (summarization, QA, classification).
- **LangChain** for managing AI task workflows and chaining multiple steps.
- Custom-built data pipelines for document ingestion, parsing, and storage.

### 5. AI Model Selection & Fine-Tuning Process

A key technical achievement in the REAI project was the deliberate and strategic selection, customization, and deployment of AI models that met the dual demands of performance and data privacy.

#### Model Selection Criteria

**To ensure practical utility without compromising security, models were selected based on the following criteria:**

- Performance on NLP tasks relevant to ethics review (summarization, question answering, classification)

- Resource efficiency, so models could run on internal, secure servers without the need for high-end GPUs
- Interpretability, so outputs could be reviewed and audited by human ethics committee members
- Fine-tuning readiness, enabling adaptation to the domain-specific language of clinical research protocols.

#### Selected Models

Function	Model(s) Used	Reason for Selection
Document Summarization	BART, T5	Strong abstractive summarization with fluent output
Answer Validation (Q&A)	BERT, DistilBERT	Efficient QA models trained for answer consistency tasks
Reviewer Form Answer Generation	LLama	Good balance of language generation and factual grounding
Workflow Management	LangChain	Allows chaining of multiple AI steps securely and modularly

#### Training & Fine-Tuning Strategy

- Historical application data (securely stored) was tokenized and used to fine-tune base models for high relevance and lower hallucination.
- Evaluation was done using custom QA benchmarks and domain-specific validation sets.
- Hyperparameters were optimized for fast inference and low memory usage, making deployment possible even on mid-tier server infrastructure.

- All training and testing were conducted within the client’s private environment, in compliance with GDPR, POPIA, and GCP.

## 6. AI Capabilities Breakdown

To clarify REAI's role in different stages of the review workflow, the table below outlines AI-supported capabilities and their implementation maturity:

Capability	Model Power	Human Involvement	Value Addition
OCR & Text Extraction	Tesseract	Minimal	Converts image-based PDFs into searchable text
Protocol Summarization	T5/BART	Reviewer oversight	Speeds up understanding of long documents
Answer-Document Validation	BERT QA	Reviewer confirmation	Flags inconsistencies quickly
Drafting Reviewer Responses	LLama	Must be reviewed	Reduces reviewer workload
Discrepancy Alerts	Rule engine + QA model	Reviewed manually	Improves data integrity
Reviewer Workflow Orchestration	LangChain	Backend only	Ensures AI modules work in sequence

## 7. Implementation

The implementation was structured in well-defined phases:

### Phase 1: Research & Design

- Conducted user research with REC members and administrators.
- Mapped existing review workflows.
- Identified pain points and AI integration opportunities.

### Phase 2: Prototype Development

- Created proof-of-concept modules for OCR, summarization, and validation.
- Tested AI outputs against historical RHInnO data to ensure relevance.

### Phase 3: Integration & Pilot

- Integrated AI modules into RHInnO Ethics.
- Conducted internal pilot with simulated data and reviewer feedback.
- Added UI components for displaying AI-generated summaries and suggested answers.

### Phase 4: Reviewer Onboarding & Feedback Loop

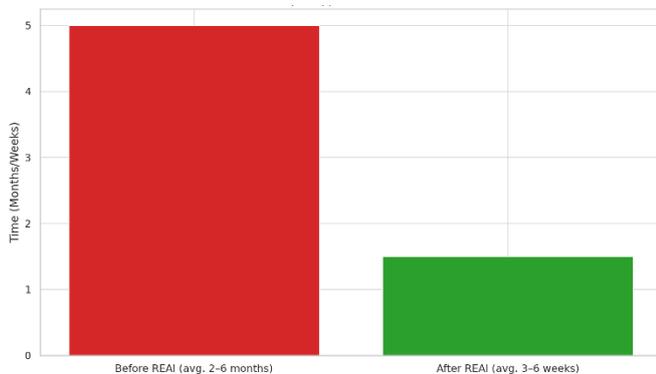
- Trained reviewers on using AI-assisted workflows.
- Set up structured feedback collection to fine-tune model behaviour.
- Iteratively refined outputs and added interpretability features to build trust.
- Challenges Addressed:
  - Ensured AI suggestions are clearly marked and editable by reviewers.
  - Addressed concerns around data privacy and compliance with ethical standards: A key challenge during implementation was the handling of sensitive applicant and research data. Due to the confidential nature of clinical research protocols and personal information involved, data security was non-negotiable. The systems and operational processes adhered strictly to regulatory standards including:
    - Good Clinical Practice (GCP) guidelines
    - EU General Data Protection Regulation (GDPR)
    - South Africa’s Protection of Personal Information Act (POPIA)

To maintain compliance, **no data was allowed to leave the client’s secured server infrastructure** for training or inference. This ruled out the use of large, cloud-based AI models or external APIs. Instead, smaller yet powerful models (e.g., BERT, T5, Flan-T5) were carefully selected for their balance of performance and resource efficiency. These models were then **fine-tuned on historical data entirely within secure internal servers**, ensuring full data control while still achieving high levels of accuracy and utility.

- Handled poor-quality scans using enhanced OCR models.

## 8. Results & Impact

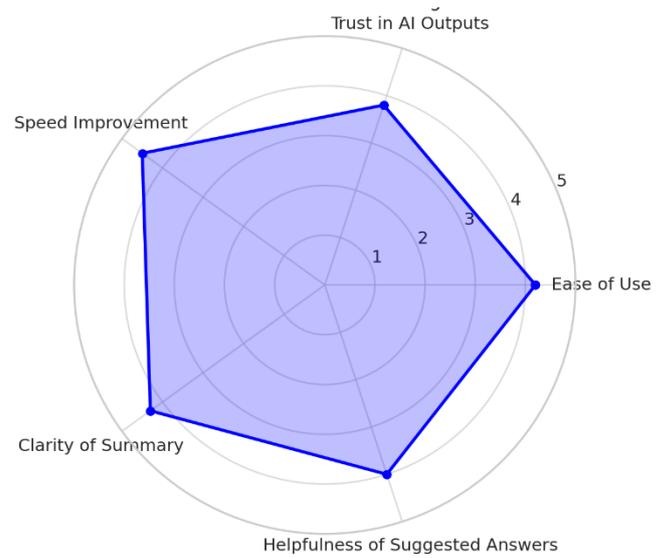
### Quantitative Impact (Initial Pilot):



### Reviewer Time Per Application Before Vs After REAI

- 60-70% reduction** in reviewer time per application.
- Consistent summaries** produced for all document types, including protocols and consent forms.
- Validation accuracy** of applicant answers above **85%**.

### Qualitative Feedback:

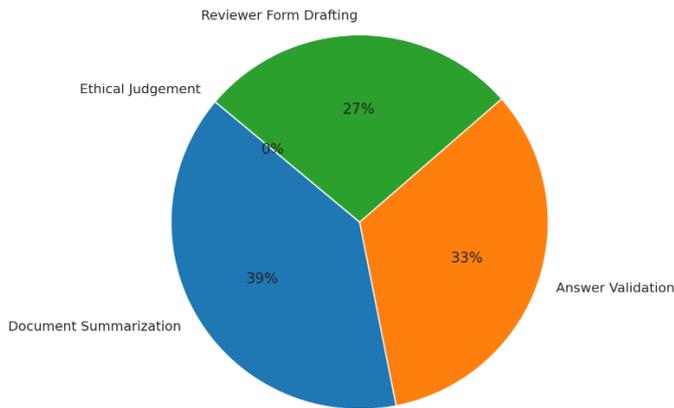


### Reviewer Feedback on AI Integration

- Reviewers appreciated the **structured summaries**, which helped them focus on core ethical issues faster.
- Suggested answers** to reviewer questionnaires saved time and acted as a quality check.
- Increased confidence in cross-validated applicant inputs.

While still in progressive deployment, early impact indicators are promising:

Metric	Before REAI	After REAI	Improvement
<b>Reviewer Time per Application</b>	2–6 months	3–6 weeks (approx.)	60–70% faster
<b>Answer Validation Accuracy</b>	Manual cross-checks (~unknown)	>85% automated validation accuracy	Significant improvement
<b>Application Consistency Check</b>	Manual, inconsistent	Automated, flagged discrepancies	Improved consistency
<b>Reviewer Satisfaction Score</b>	(Baseline: estimated low due to workload)	High (based on pilot feedback surveys)	Notably improved
<b>Cost Reduction for RECs</b>	High administrative cost	Reduced effort = lower operational cost	Economical gain



AI Contribution in Reviewer Tasks

### Strategic Impact:

- Supports **standardization** across RECs by reducing variability in assessments.
- Contributes to **capacity building**, enabling newer or smaller RECs to handle complex reviews more efficiently.
- Lays the groundwork for **data-driven improvements** in ethical oversight.

## 9. Lessons Learned

### Key Insights:

- Secure AI deployment is achievable but requires trade-offs: Training AI models within a privacy-first infrastructure meant we had to be strategic in choosing models that were lightweight enough to run on internal servers, yet robust enough to deliver meaningful insights. This required a balance of model performance, processing overhead, and compliance with global data protection laws.
- AI is most effective when it **augments human expertise**, not replaces it.
- Human-AI collaboration requires **transparent model outputs** and **editability**.

- Domain-specific knowledge is essential in **fine-tuning NLP models** for healthcare ethics.
- Ethical review processes require **deep trust**—AI suggestions must be explainable and optional.

### Challenges Overcome:

- Ensuring AI outputs are **contextually relevant** in diverse types of research protocols.
- Handling **multilingual documents** or unstructured data.
- Designing a **user experience** that integrates AI without overwhelming reviewers.

## 10. Reviewers Quote

A few real reviewers feedback on the REAI Platform-

“ The AI-generated summaries helped me cut review time in half and focus on ethical issues ”  
– Senior Reviewer, West Africa REC

“ The REAI has eliminated the subjectivity in the answers and reduced the overall review time drastically. ” – Senior Reviewer, South Africa REC

## 11. Future Plans

The client plans to build on the REAI success with the following expansions:

- **Rollout to more countries** and RECs across Africa and Asia.
- Addition of **multilingual support** for French, Portuguese, and local languages.

- Integration of **active learning loops** to improve AI accuracy through reviewer corrections.
- Enhanced visual dashboards for **insights, trends, and reviewer analytics**.
- Exploration of use in **non-clinical research reviews**, such as social or behavioral studies.

The vision is for REAI to become a **trusted co-pilot** for ethical reviewers worldwide.

## 12. Ethical AI Commitment

- REAI was built to enhance—not replace—human judgement.
- All outputs are transparent, explainable, and 100% editable by reviewers.
- Our models are trained and deployed in compliance with GDPR, POPIA, and GCP standards.

## 13. Conclusion

The REAI project represents a landmark innovation in the field of clinical research ethics review. By thoughtfully integrating Artificial Intelligence into the RHInnO Ethics platform, Client and Zepto Digital Labs have demonstrated that technological advancement and ethical rigor can go hand-in-hand.

Through strategic model selection, strict adherence to international data privacy laws, and a strong focus on human-AI collaboration, REAI has significantly improved the quality, consistency, and speed of ethics reviews.

As REAI expands across geographies and evolves through active learning, it holds the promise of becoming a trusted partner to Research Ethics Committees worldwide—enabling faster, fairer, and more informed ethical decision-making.