# Enhancing University Research with IBM watsonx

## Contents





11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

# Executive Summary

Filipino universities are rich in scholarly content but often lack efficient tools to mine institutional research (journals, theses, dissertations) for insights. IBM's watsonx platform offers a solution by using AI-powered search and generative models on campus data. By ingesting and vectorizing a university's corpus into a knowledge base, faculty and students can pose natural-language queries and receive summarized, cited answers drawn from the actual documents. watsonx's Retrieval-Augmented Generation (RAG) framework connects foundation language models to institutional data, enabling up-to-date, trustworthy responses without retraining models. The expected outcomes include faster discovery of relevant literature, automated summarization of complex findings, and cross-referencing of related work across departments. For example, La Trobe University built a RAG pipeline with watsonx (using Watson Discovery and watsonx.ai) on 100+ research papers and "efficiently access university-approved knowledge, saving valuable time" in research tasks.

- **Consolidate Institutional Knowledge:** A watsonx-based system ingests all existing publications into a unified AI-ready repository, breaking silos across departments.
- **AI-powered Search and Summaries:** Users can ask research questions in natural language. The system retrieves relevant documents (via embeddings) and uses generative AI to extract and summarize answers (with citations).
- **Cross-Reference and Collaboration:** By linking concepts across thousands of pages, watsonx helps cross-reference methodologies and findings, fostering interdisciplinary insights.
- **Responsible AI Governance:** watsonx.governance tools ensure the AI's outputs are tracked, explainable, and compliant with university policies.

Modernizing research: AI-driven search can transform how students and faculty find and analyze knowledge from university libraries and archives.



11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

## The Research Challenge in Academia

Universities today generate vast amounts of scholarly work, from graduate theses to faculty publications. However, this wealth of information is often locked in disparate systems or document repositories, making it time-consuming to discover relevant research. Traditional keyword search misses context across domains, and sifting through PDFs for specific insights is laborious. Decision-makers must find ways to harness this "dark" knowledge base to improve research efficiency and collaboration.

Generative AI models like GPT are powerful at summarizing text, but their fixed training data can become outdated. **Retrieval-Augmented Generation (RAG)** solves this by connecting AI models to an external knowledge base. In a RAG setup, a language model receives information retrieved from an indexed document store, ensuring answers are grounded in the latest data. For universities, RAG means answers can be drawn directly from *institutional research documents*, not just the model's original training. This gives researchers access to a conversational assistant that cites actual papers and findings, reducing the risk of "hallucinations" (AI fabricating facts).



11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

## IBM watsonx Platform Overview

IBM watsonx is an integrated AI and data platform designed for enterprise use. Its three core components align well with university needs:

- **watsonx.ai:** An AI development studio with tools for building and deploying generative AI and machine learning models. It provides a catalog of foundation models (including open-source and IBM models) that can be used for summarization, classification, or question-answering on text. For example, faculty can fine-tune or prompt-tune a model for summarizing research articles, or use it to generate draft abstracts from data.
- **watsonx.data:** A lakehouse environment that unifies disparate data sources into a single, secure repository. Using watsonx.data integration, universities can build scalable ETL pipelines to ingest unstructured files (PDFs of theses, Word docs, etc.) alongside structured research data. The platform then transforms and indexes these documents, enabling efficient search across the combined dataset.
- **watsonx.governance:** A governance layer for ethical AI. It tracks model usage, monitors outputs for bias, and maintains audit trails. In an academic setting, this helps ensure AI-generated outputs are reliable and that any adaptations comply with data protection policies (e.g. student privacy).

Together, these components let IT teams set up a **secure, end-to-end RAG pipeline**: unifying data (via watsonx.data), building and querying models (via watsonx.ai), and overseeing quality and compliance (via watsonx.governance).

## Retrieval-Augmented Generation (RAG) for University Knowledge

Retrieval-Augmented Generation is a powerful approach for university research support. In this architecture, when a user poses a query (e.g. "What are the common findings on renewable energy policy in recent Philippine studies?"), the system first retrieves relevant document excerpts from the university's repository using semantic search. These snippets, grounded in actual research text, are then fed as context into a generative AI model, which formulates a concise, coherent answer. Crucially, the AI can cite which documents contributed to the answer, so scholars can verify and follow up.



11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

**Key benefits of RAG include:**

- *Access to current domain data:* Models stay up-to-date by drawing on the latest papers and reports, rather than relying on static pre-training.
- *Reduced hallucination:* By anchoring responses in real documents, RAG lowers the chance of AI making up facts].
- *Increased trust:* Automatic citation of sources (e.g. "As reported in University Journal 2023") lets users check validity.
- *Cost efficiency:* Universities avoid the expense of retraining large models; instead, they supply data at query-time.

For example, La Trobe University built a watsonx-based RAG solution over 100+ autism research papers. Researchers could quickly ask complex questions and get answers synthesized from those papers, dramatically speeding literature review. This resulted in an 8.7× cost saving compared to traditional development (largely by using watsonx.ai and watsonx Assistant in-house).

## Technical Workflow: From Documents to Insights

The end-to-end workflow with watsonx involves several steps:

1. **Document Ingestion:** Connect to institutional repositories (e.g. file servers, document databases, cloud storage). Using *watsonx.data integration* (IBM's unstructured data ETL), batch-upload all research documents – theses, journal PDFs, conference papers – into the lakehouse. The platform's connectors and SDK support common formats, with no-code options for librarians.
2. **Preprocessing and Transform:** The pipeline automatically extracts text from each document (including OCR for scanned pages) and applies cleaning steps (remove duplicates, redact sensitive info). watsonx integration can leverage open-source tools (like Tesseract or LangChain) under the hood to handle diverse file types.



11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

3. **Chunking and Embedding:** Long documents are split into semantically coherent passages or "chunks." Each chunk is fed into an embedding model (a type of neural network) to generate a numerical vector representing its meaning. Pre-built operators in watsonx's ETL take care of this at scale. The resulting vectors are stored in an enterprise vector database (for example, Milvus, which IBM's platform integrates with). This allows **semantic retrieval**: user queries are also converted to vectors and matched to similar document chunks.

4. **Query and RAG Processing:** When a user asks a question through a chat interface or search box, the system runs a two-step AI process. First, it finds the top-k relevant chunks from the vector index (fast nearest-neighbor search). Second, it feeds those chunks into a foundation language model via watsonx.ai. The model synthesizes an answer and can even highlight or cite the specific excerpts used.

5. **Summarization and Analysis:** Beyond Q&A, watsonx.ai tools can generate summaries of longer reports, extract key entities, or translate content (e.g. Tagalog-English) to make research more accessible. For instance, a professor could select several related papers and ask watsonx to "summarize key differences in methodology," receiving a concise comparative summary.

Throughout this pipeline, watsonx.governance tracks model usage. Every query can be logged, and results evaluated for bias or accuracy, ensuring the system remains reliable and auditable.

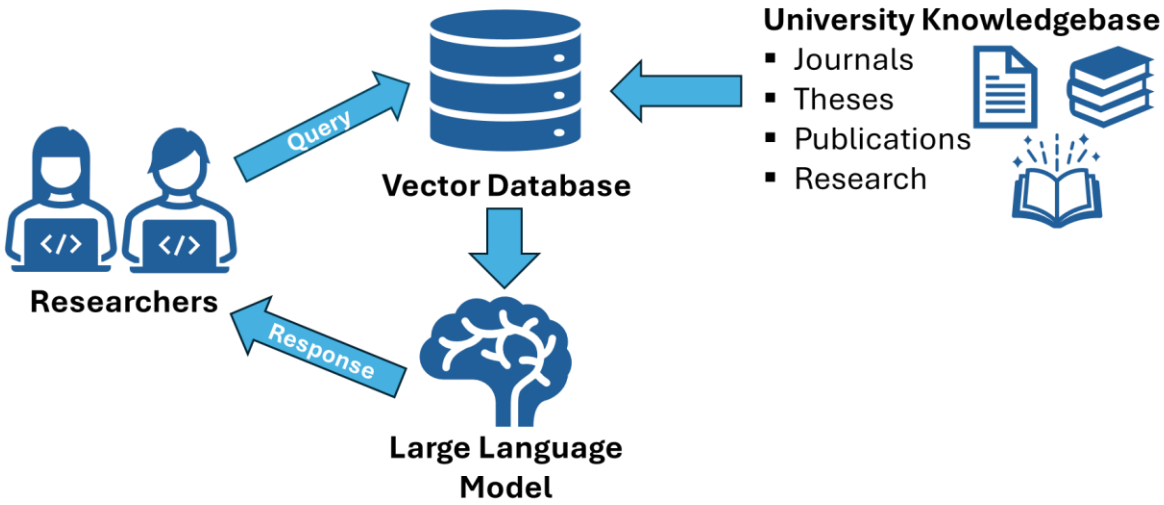




11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

## Expected Outcomes for Universities

Implementing watsonx and RAG is a strategic investment with measurable returns:

- **Faster Research Discovery:** Students and faculty spend less time searching manually. Preliminary studies (like at La Trobe) show researchers can complete literature surveys in weeks instead of months, freeing time for analysis.
- **Higher Research Quality:** AI-generated summaries and cross-references help ensure no relevant study is overlooked. Interdisciplinary connections become visible (e.g. linking engineering papers with social science surveys on the same topic).
- **Democratized Knowledge:** Non-experts (students, policy-makers) can query complex technical documents in plain language, broadening research impact. For example, administrative staff might query institutional data policies or grant summaries.
- **Cost Savings:** By using pre-built IBM AI tools, the university avoids hiring large AI teams. The La Trobe example achieved an **8.7× cost reduction** by leveraging watsonx Assistant and in-house expertise.
- **Responsible AI Oversight:** watsonx.governance provides transparency. If needed, the university can review why the AI gave a certain answer, or retrain models on local feedback. Compliance officers can ensure student data and unpublished drafts remain private.

In the Philippine context, these outcomes translate to a competitive edge: more efficient research production, improved rankings, and innovation ready for commercialization (supporting CHED's push for research outputs and industry partnerships).

## Implementation Considerations

- **Infrastructure:** watsonx is typically deployed on IBM Cloud Pak for Data (can run on-premises or cloud). Philippine institutions should assess whether to host on their own servers or collaborate with a cloud provider. Consider starting with a pilot on cloud to reduce upfront costs.
- **Data Preparation:** Ensuring high-quality ingestion is key. Libraries may need to digitize old theses or standardize metadata. watsonx's unstructured data tools (IBM's new ETL for documents) are designed to scale and handle Philippine languages if needed.



11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com

- **User Training:** While non-experts can use the query interface easily, training workshops for faculty and librarians will maximize adoption. Show how to phrase queries and validate AI summaries. Emphasize that the AI assistant complements, not replaces, expert judgment.
- **Governance and Policy:** Establish guidelines for use (e.g. verifying AI-suggested citations). Use watsonx.governance to monitor outputs for accuracy. Institute controls so embargoed or sensitive documents are only accessible to authorized users.
- **Cost and Licensing:** IBM offers academic pricing on watsonx; partnerships (like IBM Skills Academy) may provide training grants. Total cost depends on scale (number of documents, users, and compute usage). However, the model approach means costs grow slower than retraining massive AI models.

## References


**What is RAG (Retrieval Augmented Generation)? | IBM**

https://www.ibm.com/think/topics/retrieval-augmented-generation

**La Trobe University | IBM**

https://www.ibm.com/case-studies/la-trobe-university

**Overview of IBM watsonx.ai and IBM watsonx.governance software**

https://www.ibm.com/docs/en/watsonx/w-and-w/2.0.0?topic=overview-watsonx

**Academic Software partners with IBM watsonx: AI for Education**

https://academicsoftware.com/en-gb/news/ibm-watsonx-academic-software-partners

**Unstructured Data Processing | IBM**

https://www.ibm.com/products/watsonx-data-integration/unstructured-data-processing




11MN CyberOne Bldg. Eastwood Cyberpark, Quezon City Philippines 1110
Tel. +632 8912 6745 or +632 7210372 Email. info@eacomm.com