FluBroad-Voice: An AI Presentation Agent for Broadly Neutralizing Antibody Research

Technical White Paper – Version 2.0
March 2026
A product of ailink.ca — in partnership with McGill University

---

1. Introduction

The exponential growth of virology literature and public databases (PubMed, GISAID, MAAD) has created an information paradox: while more data is available than ever, synthesizing it into actionable scientific insight remains a major bottleneck. Researchers often spend days preparing literature reviews, slide decks, and presentations—time that could be devoted to experimental design and discovery.

FluBroad-Voice is an open-source AI agent developed by ailink in close collaboration with McGill University. It automates the end-to-end process of literature retrieval, knowledge synthesis, and presentation generation for research on broadly neutralizing antibodies (bnAbs) against influenza. By combining Retrieval-Augmented Generation (RAG), domain-specific ontologies, and modular output pipelines, FluBroad-Voice produces a fully narrated, citation-backed 10-minute presentation in under five minutes.

This white paper describes the system's architecture, core modules, performance metrics, and deployment options, with the goal of fostering adoption and further co-development with the virology community.

---

2. System Overview & Core Capabilities

FluBroad-Voice is designed around the "three connections and four transformations" principle:

- Three connections:

    - Internal experimental data (optional)
    - Public databases (PubMed, GISAID, MAAD, PDB)
    - Scientific literature
- Four transformations:
    - Structured literature review
    - Academic slide deck (PPTX / LaTeX Beamer)
    - Spoken narration with synchronized visuals
    - Interactive Q&A (planned)

The system is built to be modular, locally deployable, and privacy-preserving—all components can run on a laboratory's own infrastructure without sending data to external cloud services.

---

3. Technical Architecture

The architecture follows a layered, agent-based design:

Layer	Components
Data Sources	PubMed (via Entrez), GISAID (authorized API), MAAD, PDB, optional internal data
Knowledge Layer	Vector store (Chroma/FAISS) + Neo4j knowledge graph (antibody–epitope–sequence–literature)
Processing Layer	Literature fetcher, RAG pipeline, narrative generator (LLM), information extractor (Pydantic schema)
Output Layer	PPT generator (python-pptx / LaTeX), speech synthesizer (Edge-TTS), video composer (FFmpeg)

Interaction Layer	Gradio interface / REST API; future conversational module

All modules are containerized (Docker) and can be orchestrated with a single `docker-compose` command.

---

4. Core Modules & Implementation

4.1 Literature Retriever

- PubMed Fetcher: Biopython Entrez wrapper with rate limiting (≤3 requests/sec) and API key support.
- Query: Configurable keywords, date filters, and result limits.
- Extensible: Pre-integrated with bioRxiv (planned) and Europe PMC.

4.2 RAG Knowledge Base

- Embeddings: OpenAI `text-embedding-3-small` (alternative local models: MedCPT, BioBERT)
- Vector Store: Chroma with persistent storage; chunk size = 1000, overlap = 200.
- Metadata: PMID, title, year, journal stored for traceability.

4.3 Narrative Generator (LLM)

- Model: GPT-4 (temperature = 0.1) with optional local deployment via LLaMA-3.
- PMRC Framework: Prompts enforce a Problem–Motivation–Result–Conclusion structure tailored to bnAb research.
- Citation Forcing: Every claim must be accompanied by a PMID.

4.4 Information Extractor

- Schema: Pydantic models for antibody name, target protein (HA/NA/M2e), epitope region, neutralization spectrum, IGHV gene, clinical phase, and PMID.
- Method: LLM JSON mode + post-validation; accuracy verified against expert annotations.

4.5 Presentation & Speech Synthesis

- PPT Generator: Uses `python-pptx` with lab-customizable templates; automatically inserts text, tables, and basic charts.
- Speech Synthesizer: Edge-TTS (free) or Azure TTS with realistic prosody; generates a timed audio track.
- Video Assembly: FFmpeg combines slide images and audio; optional mouse-cursor highlighting for key figures.

---

5. Performance Evaluation

We evaluated FluBroad-Voice on a test set of 200 recent papers (2023–2025) related to influenza broadly neutralizing antibodies. Three virology experts provided reference annotations.

Metric	Result	Benchmark
Literature Retrieval Recall	94.2%	vs. manual systematic review
Literature Retrieval Precision	91.5%	vs. expert screening
Review Logical Structure (1–5)	4.7	Expert rating
Review Content Accuracy (1–5)	4.5	Expert rating
Citation Traceability	100%	All claims linked to PMID

Time Saved (vs. manual)	97%	26 hours → 50 minutes
Cohen's κ (GPT-4 vs. expert)	0.89	Near-human consistency in relevance classification

Table 1: Key performance indicators from the validation study.

The system consistently produces outputs that domain experts consider suitable for use as a draft for grant proposals, lab meetings, or teaching materials.

---

6. Deployment & Privacy

FluBroad-Voice is designed with a local-first deployment model:

- Docker containers encapsulate all services (vector DB, API, frontend).
- Offline mode: Once models are downloaded, the system can run without internet access.
- Data sovereignty: No proprietary lab data leaves the local environment.
- Access control: LDAP/SSO integration is available for multi-user labs.

Minimal hardware requirements:

- 8 CPU cores, 16 GB RAM, 100 GB storage
- Recommended: 32 GB RAM + NVIDIA T4 GPU (for faster embedding and optional local LLM)

---

7. Open-Source Core & Non-Profit Mission

ailink is a registered non-profit organization. The core framework of FluBroad-Voice is released under the MIT License and available at:
🔗 https://github.com/yiweixidu/ailink

Domain-specific components (e.g., fine-tuned prompts, GISAID adapters, antibody extraction schemas) are made freely available to all academic collaborators and contribute to the sustainability of the project. Revenue from commercial deployments is reinvested into community development, documentation, and support for early-career researchers.

---

8. Future Directions

Planned enhancements include:

- Multi-virus support: Expand the ontology to HIV, SARS-CoV-2, and RSV.
- Interactive Q&A: Allow researchers to query the knowledge base conversationally, with full citation traces.
- Deep omics integration: Automatically ingest BCR-seq data and incorporate evolutionary dynamics.
- Collaborative annotation: Enable teams to refine and correct generated outputs, feeding into continuous improvement of the models.

---

9. Conclusion

FluBroad-Voice demonstrates that a carefully designed AI agent can meaningfully reduce the administrative burden on virology researchers while maintaining scientific accuracy and traceability. By open-sourcing the core framework and maintaining a non-profit model, ailink aims to build a sustainable ecosystem where labs worldwide can adapt and extend the tool for their own needs.

We invite the virology community to explore the code, contribute improvements, and partner with us to shape the next generation of AI-assisted scientific communication.

Contact
ailink.ca
Email: info@ailink.ca
GitHub: github.com/yiweixidu/ailink