

chatNSR: An AI-Enhanced Nuclear Science References Knowledge Base

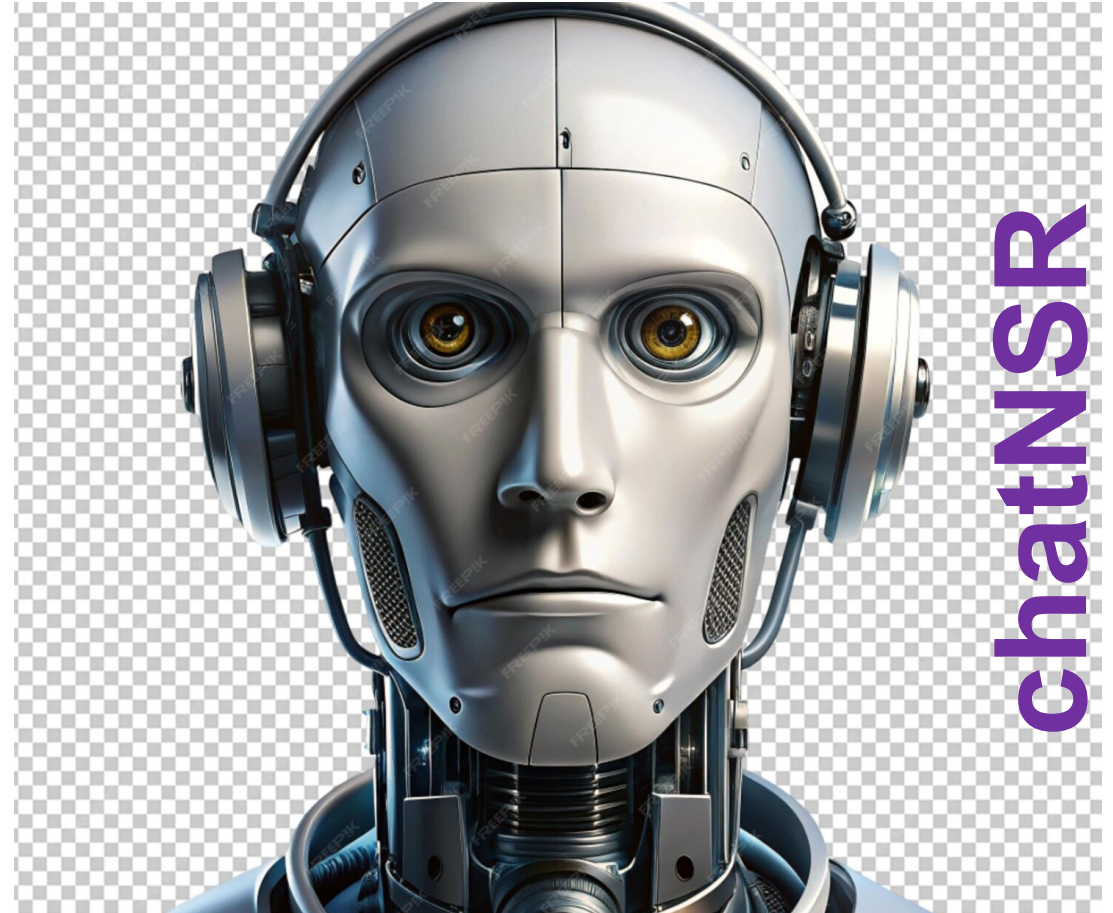
R. Arcilla and D. Brown

National Nuclear Data Center, Brookhaven National Laboratory, Upton, NY

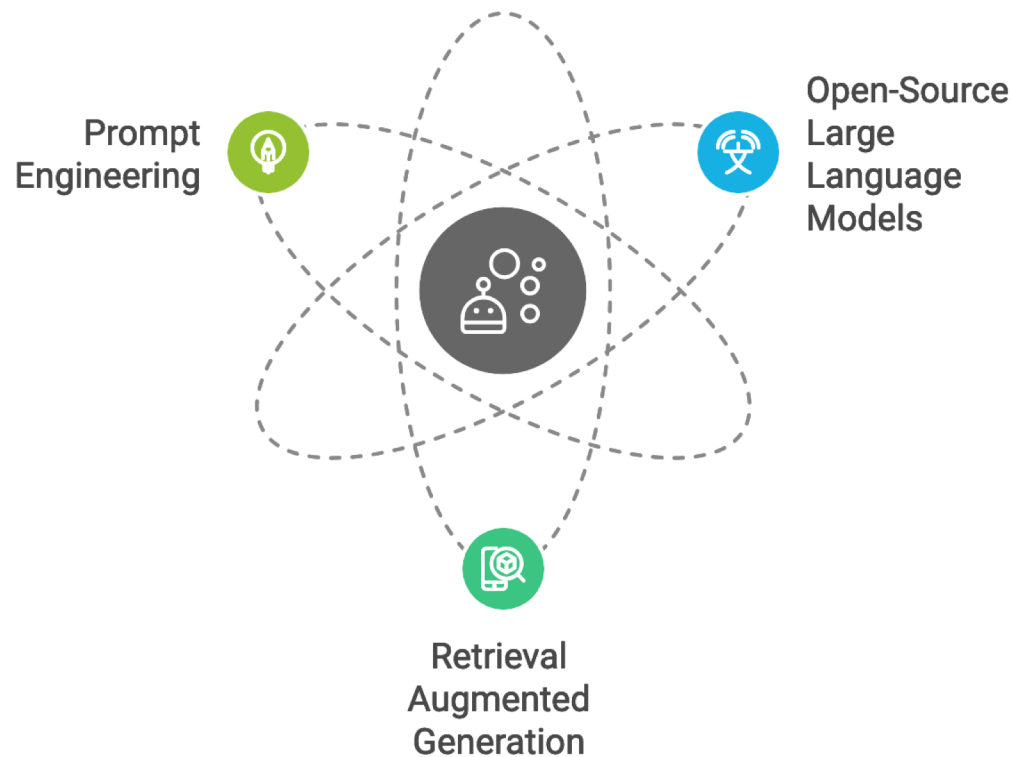
IAEA Meeting, November 11, 2024

Motivation

- Automate summarization and keywording of articles
- Enhance timeliness and scope of NSR's database contents
- Simplify retrieval interface using natural (conversational) language (currently English only)



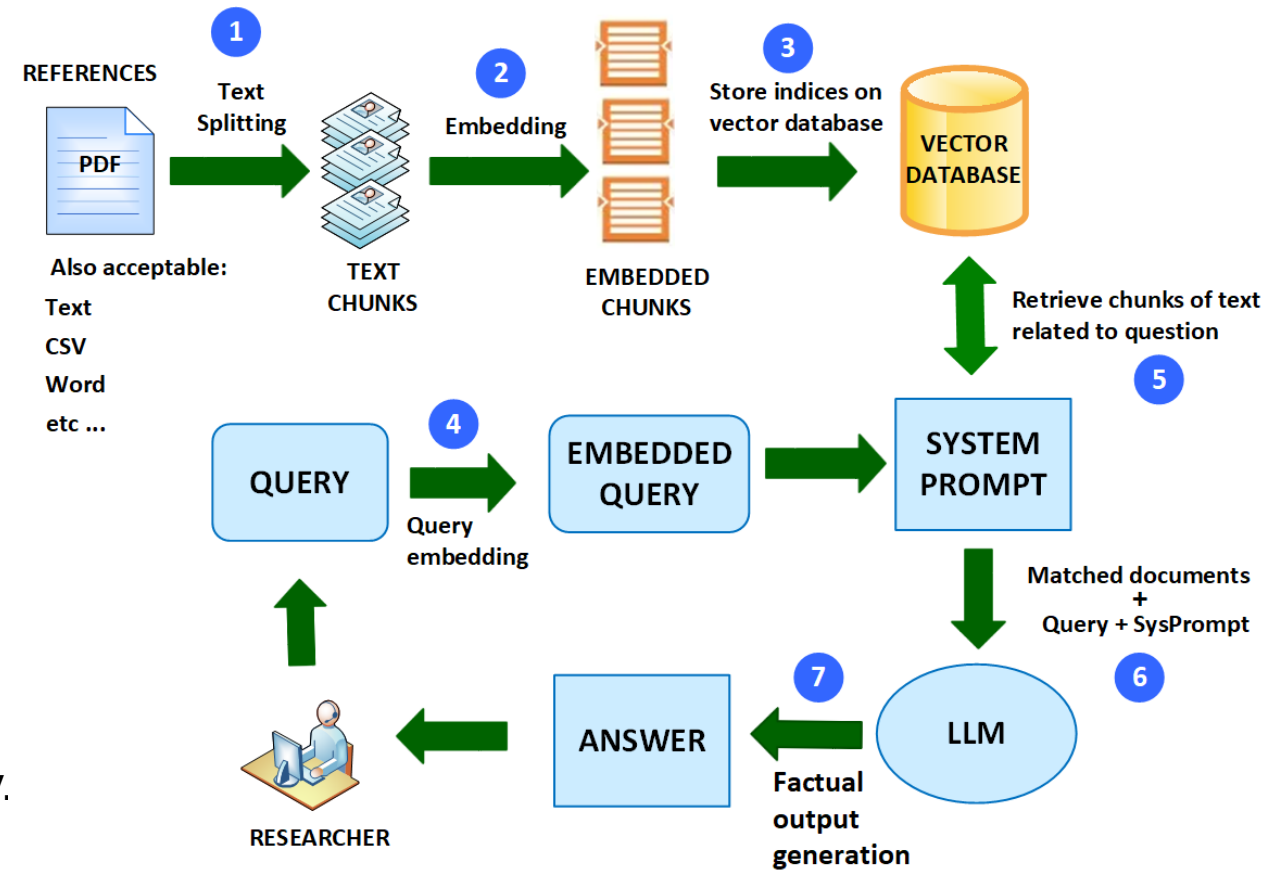
AI Components



- Large Language Models (LLM): Enables understanding user's query and generation of human-like response.
- Retrieval-Augmented Generation (RAG): retrieves relevant documents from vector database to ensure provided information is contextually appropriate and accurate.
- System prompt engineering (SYSPROMPT): crafting prompts to guide LLMs in generating quality and relevant information.

Workflow

1. Chunking: Document is split into smaller segments for easier processing later.
2. Embedding: Text chunks are converted into vectors for fast search.
3. Vectorized texts and their indices are stored on a vector database.
4. Researcher's query is converted to vectors and appended to the system prompt for use in search.
5. System searches vector database for matching vector chunks.
6. Matched vector chunks, user query and system prompt are merged for LLM's use.
7. LLM generates factual response to user's query.



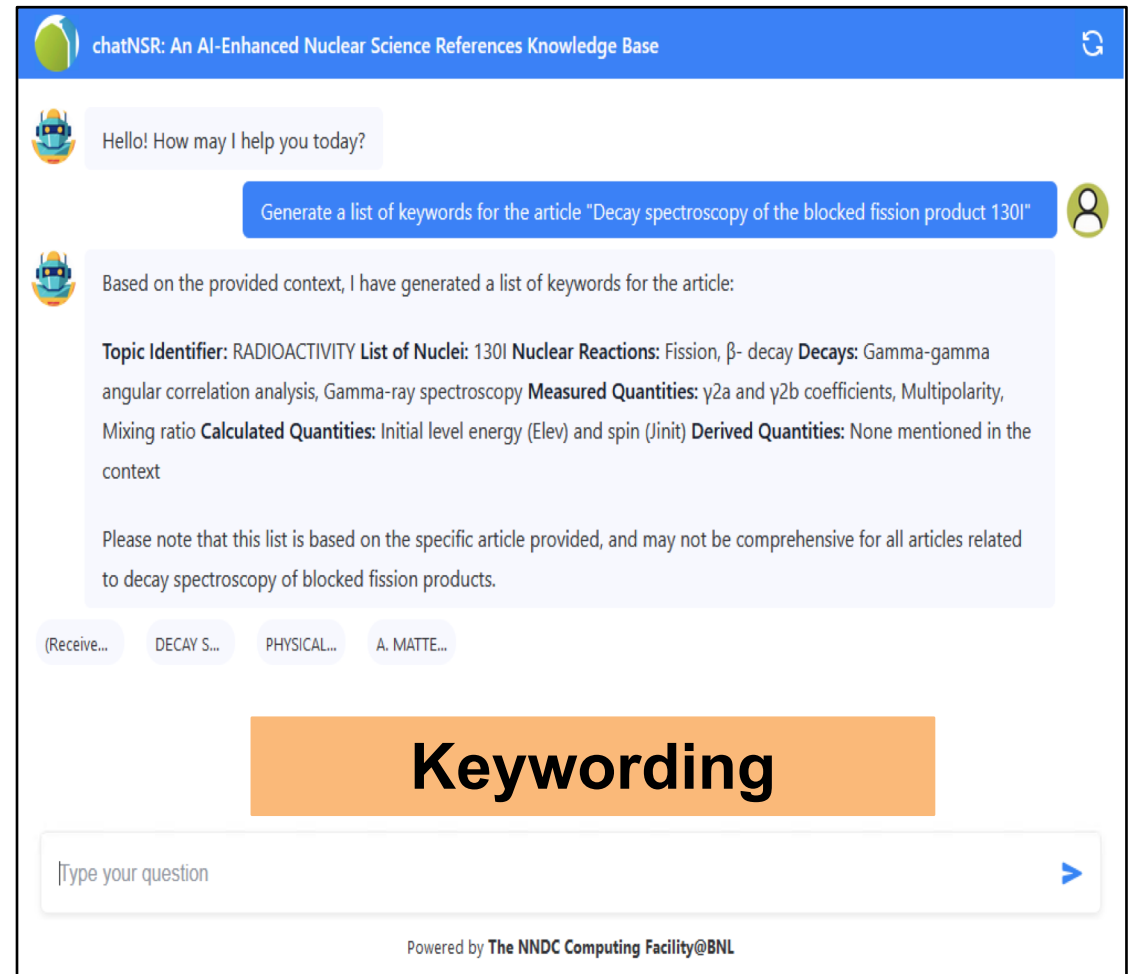
Current Status

- Prototype being tested using Llama-3.2 model (instruction-tuned) with 3 billion parameters
- ~400 Physical Review C and Elsevier PDF articles ingested, vectorized and stored on a local vector database.
- Metadata (JSON format) of top 4 most relevant sources are shown per response.
- B. Pritychenko's (NSR Manager) Feedback: Summarization quality is already acceptable.

The screenshot displays the chatNSR interface, titled "chatNSR: An AI-Enhanced Nuclear Science References Knowledge Base". It shows a user's query: "Summarize the article 'Calculations of Maxwellian-averaged cross sections and astrophysical reaction rates using the ENDF/B-VII.0, JEFF-3.1, JENDL-3.3, and ENDF/B-VI.8 evaluated nuclear reaction data libraries'". The AI response provides a detailed summary of the article, including the methods used (four evaluated nuclear reaction data libraries: ENDF/B-VII.0, JEFF-3.1, JENDL-3.3, and ENDF/B-VI.8), the purpose of the calculations (to compare with existing data and provide a comprehensive evaluation for astrophysical applications), and the authors' findings (comparison with other sources like JENDL-3.3 and experimental data, highlighting strengths and limitations). Below the summary, there are four small text boxes containing the start of the article's text: "Calculati...", "cross sec...", "cross sec...", and "grals, Ma...". A large orange button labeled "Summarization" is positioned below the summary. At the bottom, there is a text input field with the placeholder "Type your question" and a blue arrow button. The footer of the interface reads "Powered by The NNDC Computing Facility@BNL".

Current Status (continued)

- B. Pritychenko's (NSR Manager) Feedback: Keywording still needs a lot of work.
- More advanced system prompts being developed to improve keywording capability.



The screenshot displays the chatNSR interface, titled "chatNSR: An AI-Enhanced Nuclear Science References Knowledge Base". The chat history shows a user prompt: "Generate a list of keywords for the article 'Decay spectroscopy of the blocked fission product 130I'". The AI response provides a list of keywords categorized by type: Topic Identifier, List of Nuclei, Nuclear Reactions, Decays, Measured Quantities, Calculated Quantities, and Derived Quantities. A large orange box with the text "Keywording" is overlaid on the bottom half of the chat window. The interface also includes a search bar at the bottom with the placeholder text "Type your question" and a blue arrow button.

chatNSR: An AI-Enhanced Nuclear Science References Knowledge Base

Hello! How may I help you today?

Generate a list of keywords for the article "Decay spectroscopy of the blocked fission product 130I"

Based on the provided context, I have generated a list of keywords for the article:

Topic Identifier: RADIOACTIVITY **List of Nuclei:** 130I **Nuclear Reactions:** Fission, β - decay **Decays:** Gamma-gamma angular correlation analysis, Gamma-ray spectroscopy **Measured Quantities:** γ_{2a} and γ_{2b} coefficients, Multipolarity, Mixing ratio **Calculated Quantities:** Initial level energy (Elev) and spin (Jinit) **Derived Quantities:** None mentioned in the context

Please note that this list is based on the specific article provided, and may not be comprehensive for all articles related to decay spectroscopy of blocked fission products.

(Receive... DECA... PHYSICAL... A. MATTE...

Keywording

Type your question

Powered by The NNDC Computing Facility@BNL

Conclusion and Outlook

- LLMs, RAG and SYSPROMPT were successfully used to build a chatNSR prototype
- Summarization capability is already acceptable but keywording still needs a lot of work.
- More advanced LLMs (70 billion parameters) and more (100k's) high-quality references need to be ingested to improve chatNSR's performance (contextual accuracy).
- Future Work: Add advanced vision capability to interpret embedded charts and plots in articles.