



Use of AI for Interpreting Technical Specifications for Power Upgrades in Nuclear Power Plants

June 2025

Changing the World's Energy Future

John Norman Mapes



DISCLAIMER

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

Use of AI for Interpreting Technical Specifications for Power Uprates in Nuclear Power Plants

John Norman Mapes

June 2025

**Idaho National Laboratory
Idaho Falls, Idaho 83415**

<http://www.inl.gov>

**Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517**



**Presented by: Vivek
Agarwal, Ph.D.**

**Technical Lead: John
Mapes Ph.D.**

06/18/2025

Use of AI for Interpreting Technical Specifications for Power Uprates in Nuclear Power Plants

Presented for NPIC & HMIT 2025

LRS #:INL/CON-25-84815



Organization of Presentation

- Background information provided on power plant uprates.
- Discussion of the current and proposed approaches to power plant uprates.
- Explanation of what data is used to draft a LAR.
- Methods such as retrieval augmented generation (RAG) and fine-tuning are discussed.
- Use case analysis is performed.
- Different failure types are examined.
- Conclusions are drawn from the analysis.
- Future work is proposed.

Background on Power Plant Uprates

- Power uprates have been approved since the 1970s to bring approximately 8 GW of additional power to the existing fleet.
- As of January 2022, the NRC has approved 171 uprates.
- An uprate provides additional revenue from increased power generation.

From 2000

NRC Review of License Amendment Request

The NRC requires commercial nuclear power plants to operate at or below the approved maximum power level. The power level is used to establish the safety of the plant and is available in the license and Technical Specifications (TSs) for the plant. The plant may submit a license amendment request (LAR) and revised TSs to the NRC for approval when seeking a power uprate (U.S. NRC 2024).

U.S. NRC. “Power Uprates” U.S. Nuclear Regulatory Commission (Updated 2024)
<https://www.nrc.gov/reactors/operating/licensing/power-uprates/submit-requests.html>



Current Approach to Power Plant Uprates

- Not surprisingly, power uprate submissions are time-consuming and costly to prepare.
- They require extensive interpretation labor to comprehend the TSs and associated documents since it is unrealistic for any one person alone to understand the information.

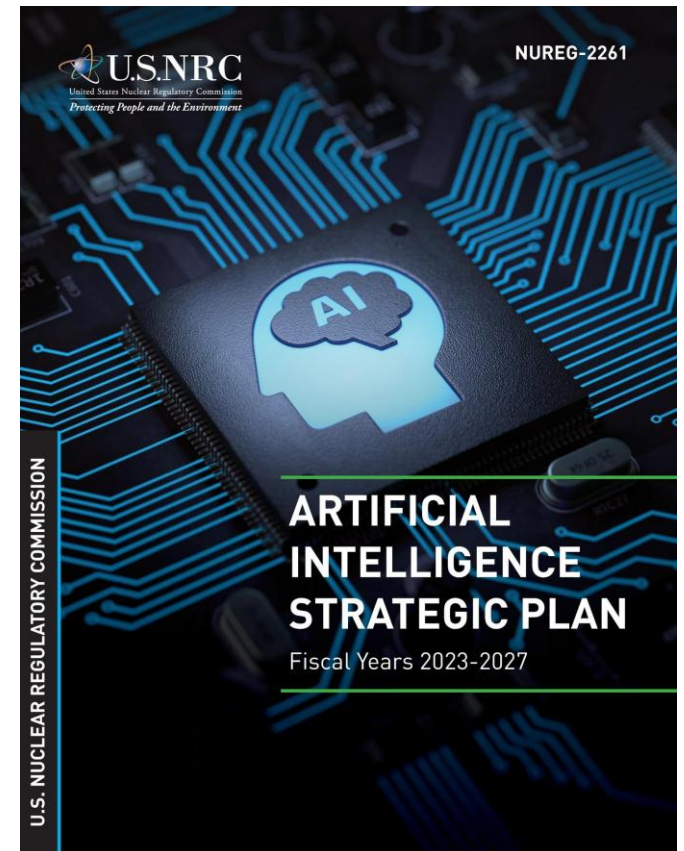


Proposed Approach to Power Plant Upgrades

- Thus, the goal of this research is to reduce the time needed to draft power upgrade submissions and their cost.
- Specifically, this work investigates using artificial intelligence (AI) models that can assimilate knowledge from TSs to enable the thorough evaluation of the impacts of power upgrades.
- Ultimately, AI will support the safe and effective implementation of power upgrades that meet all regulatory requirements.

The NRC is committed to continue to keep pace with technological innovations to ensure the safe and secure use of AI in NRC-regulated activities.

—NUREG-2261



Dataset

- What data contains information for drafting a LAR?
 - Technical Specifications (TS)
 - Technical Specification Bases (TS bases)
 - Standard Technical Specifications (STS)
 - Standard Technical Specification Bases (STS bases)
 - Updated Final Safety Analysis Report (UFSAR)



NUREG-1431
Volume 1
Revision 5.0

Standard Technical Specifications

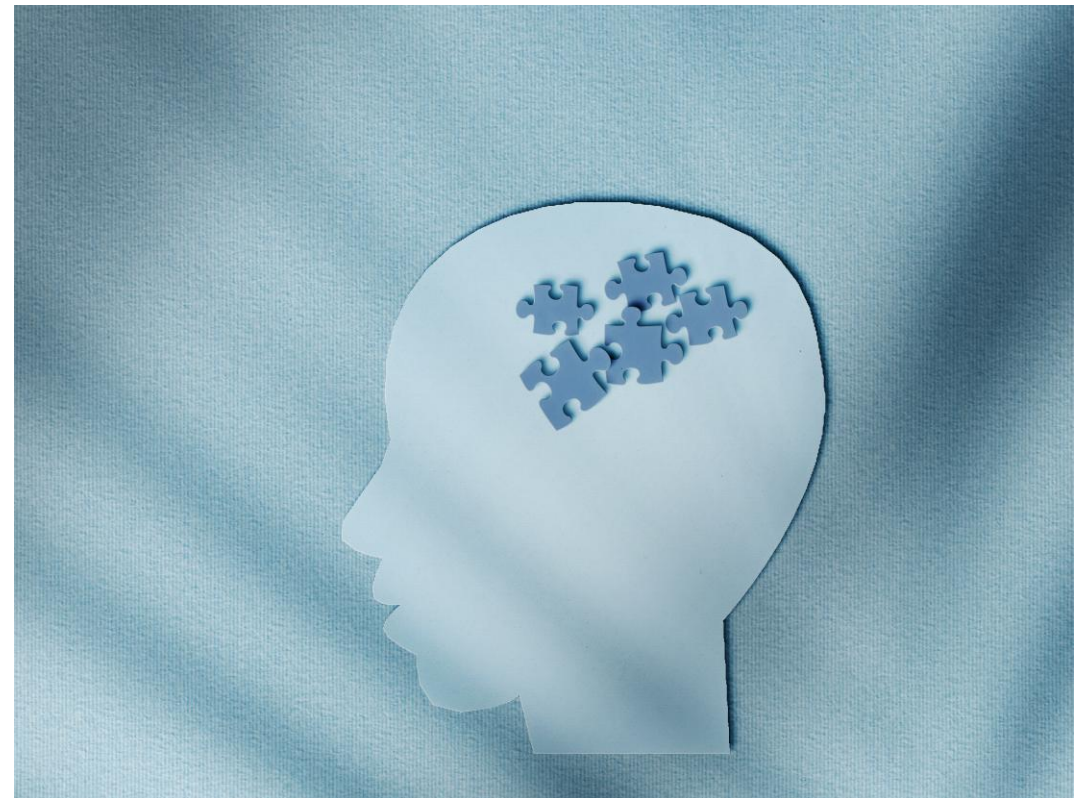
Westinghouse Plants

Revision 5.0

Volume 1, Specifications

Hypothesis

- The objective of this broader effort is to evaluate the hypothesis that generative AI has can answer questions to assist a LAR drafter by using RAG and fine-tuning LLMs to evaluate the TSs and associated documents.
- This is a stepping stone to automatically drafting the LAR.



Difficulty of Problem

- Hallucinations remain an unsolved problem of LLMs.
- Incorrect references remain an unsolved problem of RAG.
- The LLM declines to answer when it does not know how to provide the correct information.

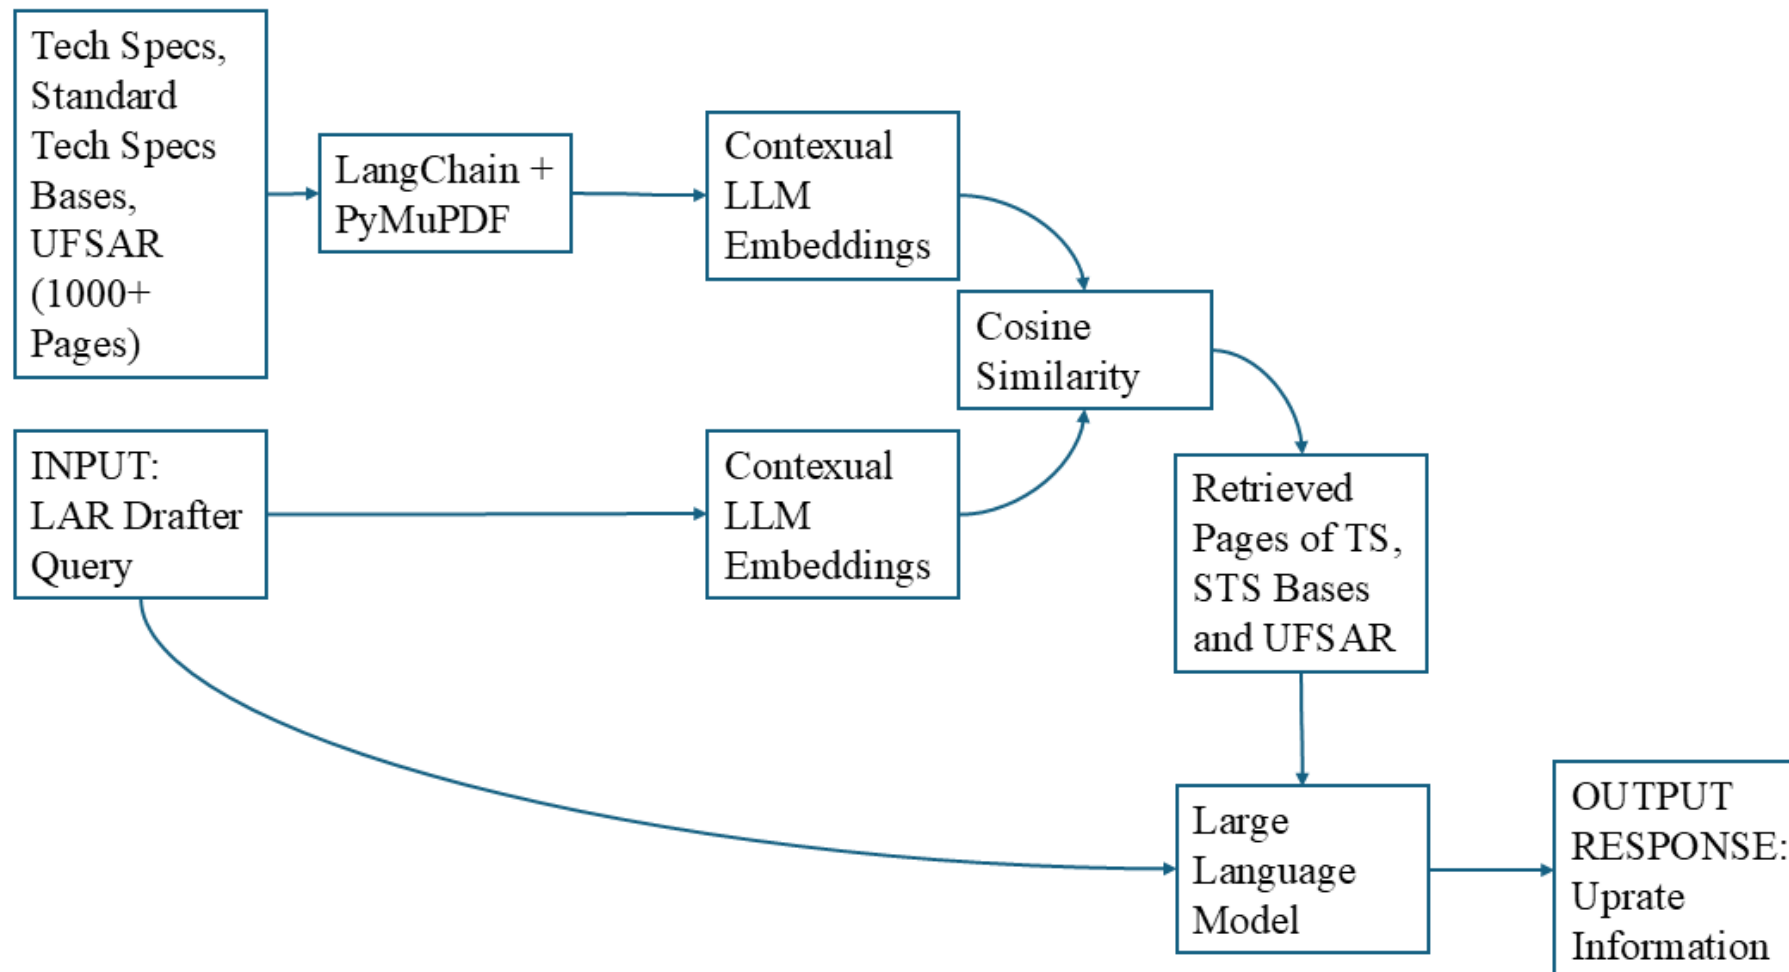


Methods—Fine-Tuning

- Fine-tuning means hyperfocusing an existing LLM on specific data to accomplish a desired task.
- In this work we fine-tuned a Microsoft Phi-3.5 model to attempt to understand approximately 2 million NRC documents stored in the Agencywide Documents Access and Management System (ADAMS) database.
- Then we did additional fine-tuning on TSs, STS bases, a UFSAR, and related LARs.
- The goal is to take an LLM with intern capabilities and convert it into an expert LLM that has years of nuclear experience.
- This is what is needed to draft LARs.



Methods—RAG Flowchart



How LAR Use Cases Are Evaluated

QUESTION

- A LAR drafter asks the question about information known to exist in LARs that is not common knowledge.

AI-RETRIEVED CONTEXT

- Information is retrieved based on the question.

PROMPT

- The retrieved information is combined with the question.

AI RESPONSE

- The AI produces an answer to the LAR drafter's question.

SME REVIEW

- **A subject-matter expert (SME) evaluates the quality of the LLM response.**

RAG Baseline Case

I found a document with this information:...The QPTR limit ensures that the gross radial power distribution remains consistent with the design values used in the safety analyses....

Yes, the QPTR (Quadrant Power Tilt Ratio) limit ensures that the gross radial power distribution remains consistent with the design values used in the safety analyses.

Does the QPTR limit ensure that the gross radial power distribution remains consistent with the design values used in the safety analyses?

You are an AI assistant specifically designed to support a nuclear power plant... Use the retrieved information to answer the question.

SME review: The answer is correct; the AI performed as expected.

Overview of the Next Five Slides

| Failure Type | Details | Use Cases (Tasks) |
|--------------|---|--|
| FT 1 | The retriever retrieves the incorrect document. | Provide the design pressure of the containment. |
| FT 2 | The LLM incorrectly interprets the retrieved document. | Provide the design pressure of the containment. |
| FT 3 | LLM Uses Internal Knowledge Instead of Retrieved Knowledge. | Explain CSS function—i.e., limits (decreases) the peak containment pressure. |
| FT 4 | The LLM hallucinates and gives incorrect facts. | Explain CSS function—i.e., limits (decreases) the peak containment pressure. |
| FT 5 | The LLM declines to answer or responds, “I don’t know.” | Define and describe the acronym MSLB. |

Failure Type 1—Retriever Retrieves Incorrect Document.

I found a document with this information:... The containment design pressure, 55 psig...

The peak design pressure of the containment is 55 psig.

What is the peak design pressure of the containment?

You are an AI assistant specifically designed to support a nuclear power plant... Use the retrieved information to answer the question.

SME review: In this case an incorrect answer was given (55 psig instead of 54) because the RAG engine failed to identify the correct document for answering the question.

Failure Type 2—LLM Incorrectly Interprets the Retrieved Document

I found a document with this information:... The design pressure of 54 psig, as shown in table 6.2-6...

The peak design pressure for the containment of the nuclear power plant is 54 psig, as indicated in table 6.2-6.

Peak design pressure containment.

You are an AI assistant specifically designed to support a nuclear power plant... Use the retrieved information to answer the question.

SME review: The psig value is correct, but the mistake is that the LLM reports that Table 6.2-6 contains the peak design containment pressure, which it does not.

Failure Type 3—LLM Uses Internal Knowledge Instead of Retrieved Knowledge

*I found a document with this information:
[IRRELEVANT DOCUMENT].*

*The CSS (Containment Spray System) is
designed to limit the peak containment pressure.*

*Does the CSS operate to limit peak containment
pressure to less than the design pressure?*

*You are an AI assistant specifically designed to support
a nuclear power plant... Use the retrieved information to
answer the question.*

***SME review: Here the answer is correct; however, the retrieved
context is irrelevant to the question per an SME.***

Failure Type 4—LLM Hallucinates

*I found a document with this information:...
Containment Spray System was analyzed to
determine the resulting reduction in containment
pressure....*

*The Containment Spray System (CSS) is
designed to decrease peak containment
pressure during a Loss of Coolant Accident
(LOCA) or a Minor Stationary Leak (MSLB).*

*Does the CSS increase or decrease peak containment
pressure during a LOCA or MSLB?*

*You are an AI assistant specifically designed to support
a nuclear power plant... Use the retrieved information to
answer the question.*

***SME review: Here the answer is correct; however, the LLM
hallucinated that MSLB means “minor stationary leak.”***

Failure Type 5—LLM Declines to Answer

*I found a document with this information:
[IRRELEVANT DOCUMENT].*

Please define and describe the acronym MSLB.

*I don't know. The provided context does not
contain information about the acronym MSLB.*

*You are an AI assistant specifically designed to support
a nuclear power plant... Use the retrieved information to
answer the question.*

***SME review: Sometimes the LLM correctly reads the context
retrieved and concludes that it cannot answer the question posed.***

Conclusions

1. The hypothesis was partly confirmed because in certain use cases the LLM does correctly evaluate the TSs and associated documents.
2. Since the documents and page numbers are provided when using RAG, this provides a degree of explainability.
3. However, as shown in five uses cases, there are pitfalls associated with the system, such as incorrect document retrieval, LLM misinterpretations, reliance on internal knowledge, hallucinations, and refusal to answer. Knowing these limitations provides realistic expectations for how AI can and cannot assist in the power uprate drafting process.
4. Overall, the use cases indicate that RAG and fine-tuned LLMs may present a new opportunity for increased productivity during power uprate drafting.

Future Work

1. Experimentation with more recent models (for example, OpenAI Frontier models).
2. Investigate OpenAI RAG embeddings for locating reference material to address each LAR section requirement.
3. Use AI to draft sections of a LAR and combine them to form the full LAR.
4. Quantitative evaluations such as cosine similarity.



Sustaining National Nuclear Assets

lwrs.inl.gov

Norman.mapes@inl.gov

BACKUP SLIDES

John Mapes, Ph.D., Idaho National Laboratory (INL)

- Research and Development Scientist
- Expertise is in generative AI and natural language processing (NLP). NLP is the analysis of text data using machine learning techniques such as classification, regression, and clustering.
- Over 9 publications with approximately 90 cites total per Google Scholar; first author on many papers.
- Gen AI and NLP expert who has worked on projects for outage optimization, MIRACLE, and an INL document search tool.
- Experience with text analysis of nuclear power plant condition reports, operator logs, work requests, and work instructions.
- Themes of research include text classification, text search, generative AI, topic models, text summarization, fine-tuning of large-language models, and document embeddings.
- Ph.D. Louisiana Tech University (2018) in engineering with a concentration in cyber.



E-mail = norman.mapes@inl.gov
Phone = +1 225 246 3319

Elsevier Pure