



# Energy Technology Cost Estimation Agent

James Tourigny, Tyler Simone  
Computer Science, University of New Hampshire, Durham, NH 03824

## Introduction

- Energy companies need to compare and forecast energy technology costs for budgeting and planning decisions. This agent lets them query a massive dataset using natural language instead of manually filtering through hundreds of thousands of rows.
- Developed an LLM-powered agent for estimating energy technology costs.
- Covers wind, solar, nuclear, geothermal, natural gas, and more.
- Returns key cost metrics from the NREL ATB 2024 dataset: LCOE (Levelized Cost of Energy) - CAPEX (Capital Expenditure) - Fixed & Variable O&M (Operations Costs) - Capacity Factor, etc.
- Answers natural language queries with cited, data-backed responses from the dataset.

## Methodology

- Tech and Tools used:
- API:** OpenRouter
  - Model:** stepfun/step-3.5-flash
  - Embeddings:** HuggingFace (all-MiniLM-L6-v2)
  - Vector Store:** FAISS (via LangChain)
- Tools Implemented:
- Retrieve:** Semantically search over the FAISS index – used when the query is vague or agent needs to explore available metrics.
  - Estimate:** Filters the DataFrame by technology, metric, scenario, and year, then returns summary statistics for a point-in-time lookup
  - Forecast:** Fits a linear trend across all historical metric-years for a given slice and projects to target year, reporting R squared and extrapolation distance.
- RAG retrieval method:
- Each CSV row embedded as a text chunk in FAISS
  - Approximate nearest-neighbor similarity search
  - Grounds LLM in real data to reduce hallucinations

## Data

- Source:** NREL Annual Technology baseline (ATB) 2024
  - Format:** CSV with ~570k+ rows
  - Each row represents a specific measurement or projection for an energy technology under certain conditions
- The dataset combines:
- Technology type** (solar, wind, nuclear, storage)
  - Scenario Assumptions** (Conservative, Moderate, Advanced)
  - Core metric being measured** (CAPEX, LCOE, Fixed O&M, Capacity Factor, etc.)
  - Metric year for historical and projected values**
  - Case filters** (Market vs R&D, tax credit scenarios)
  - Contextual details** (Resource quality, maturity level, scale)
  - All monetary values reported in 2022 USD**

## Forecast/Estimate

```
You: What will the CAPEX for offshore wind be in 2055?

[Tool call] {
  "tool": "forecast",
  "args": {
    "technology_contains": "offshore wind",
    "metric": "CAPEX",
    "target_year": 2055
  }
}

[Tool result] (forecast): {"target_year": 2055, "projected_value": 2984.3745...

Agent: Based on the forecast model for offshore wind CAPEX, the projected value in 2055 is **2984.37 $/kW (2022 USD)**.
```

**\*\*Important notes:\*\***

- This is an **\*\*extrapolation\*\*** beyond the available data range (2022-2050), extending 5 years past the last observed data point. Such projections are inherently less reliable.
- The model has a low R<sup>2</sup> value of **\*\*0.1756\*\***, indicating that the linear trend explains only about 18% of the variance in the data. The fit is weak, so the projection should be interpreted with caution.
- The fitted slope is **\*\* -120.63 \$/kW per year\*\***, reflecting a downward cost trend over the historical period.

If you need a different scenario or additional details, please let me know.

## Results

While working on this project, we learned combining an LLM with structured tool calls and vectorRAG can generate accurate predictions and estimations across the given dataset.

The example query on the left shows:

- Agent's ability to properly make a query from natural language
- Agent's ability to make a proper tool call and receive a result
- Agent's ability to use linear regression to predict capital cost of offshore wind in 2055 based on the dataset.

## Conclusion/Next Steps

We learned the fundamentals of AI agent development and were able to use internal components to develop a product capable of data analysis with natural language querying.

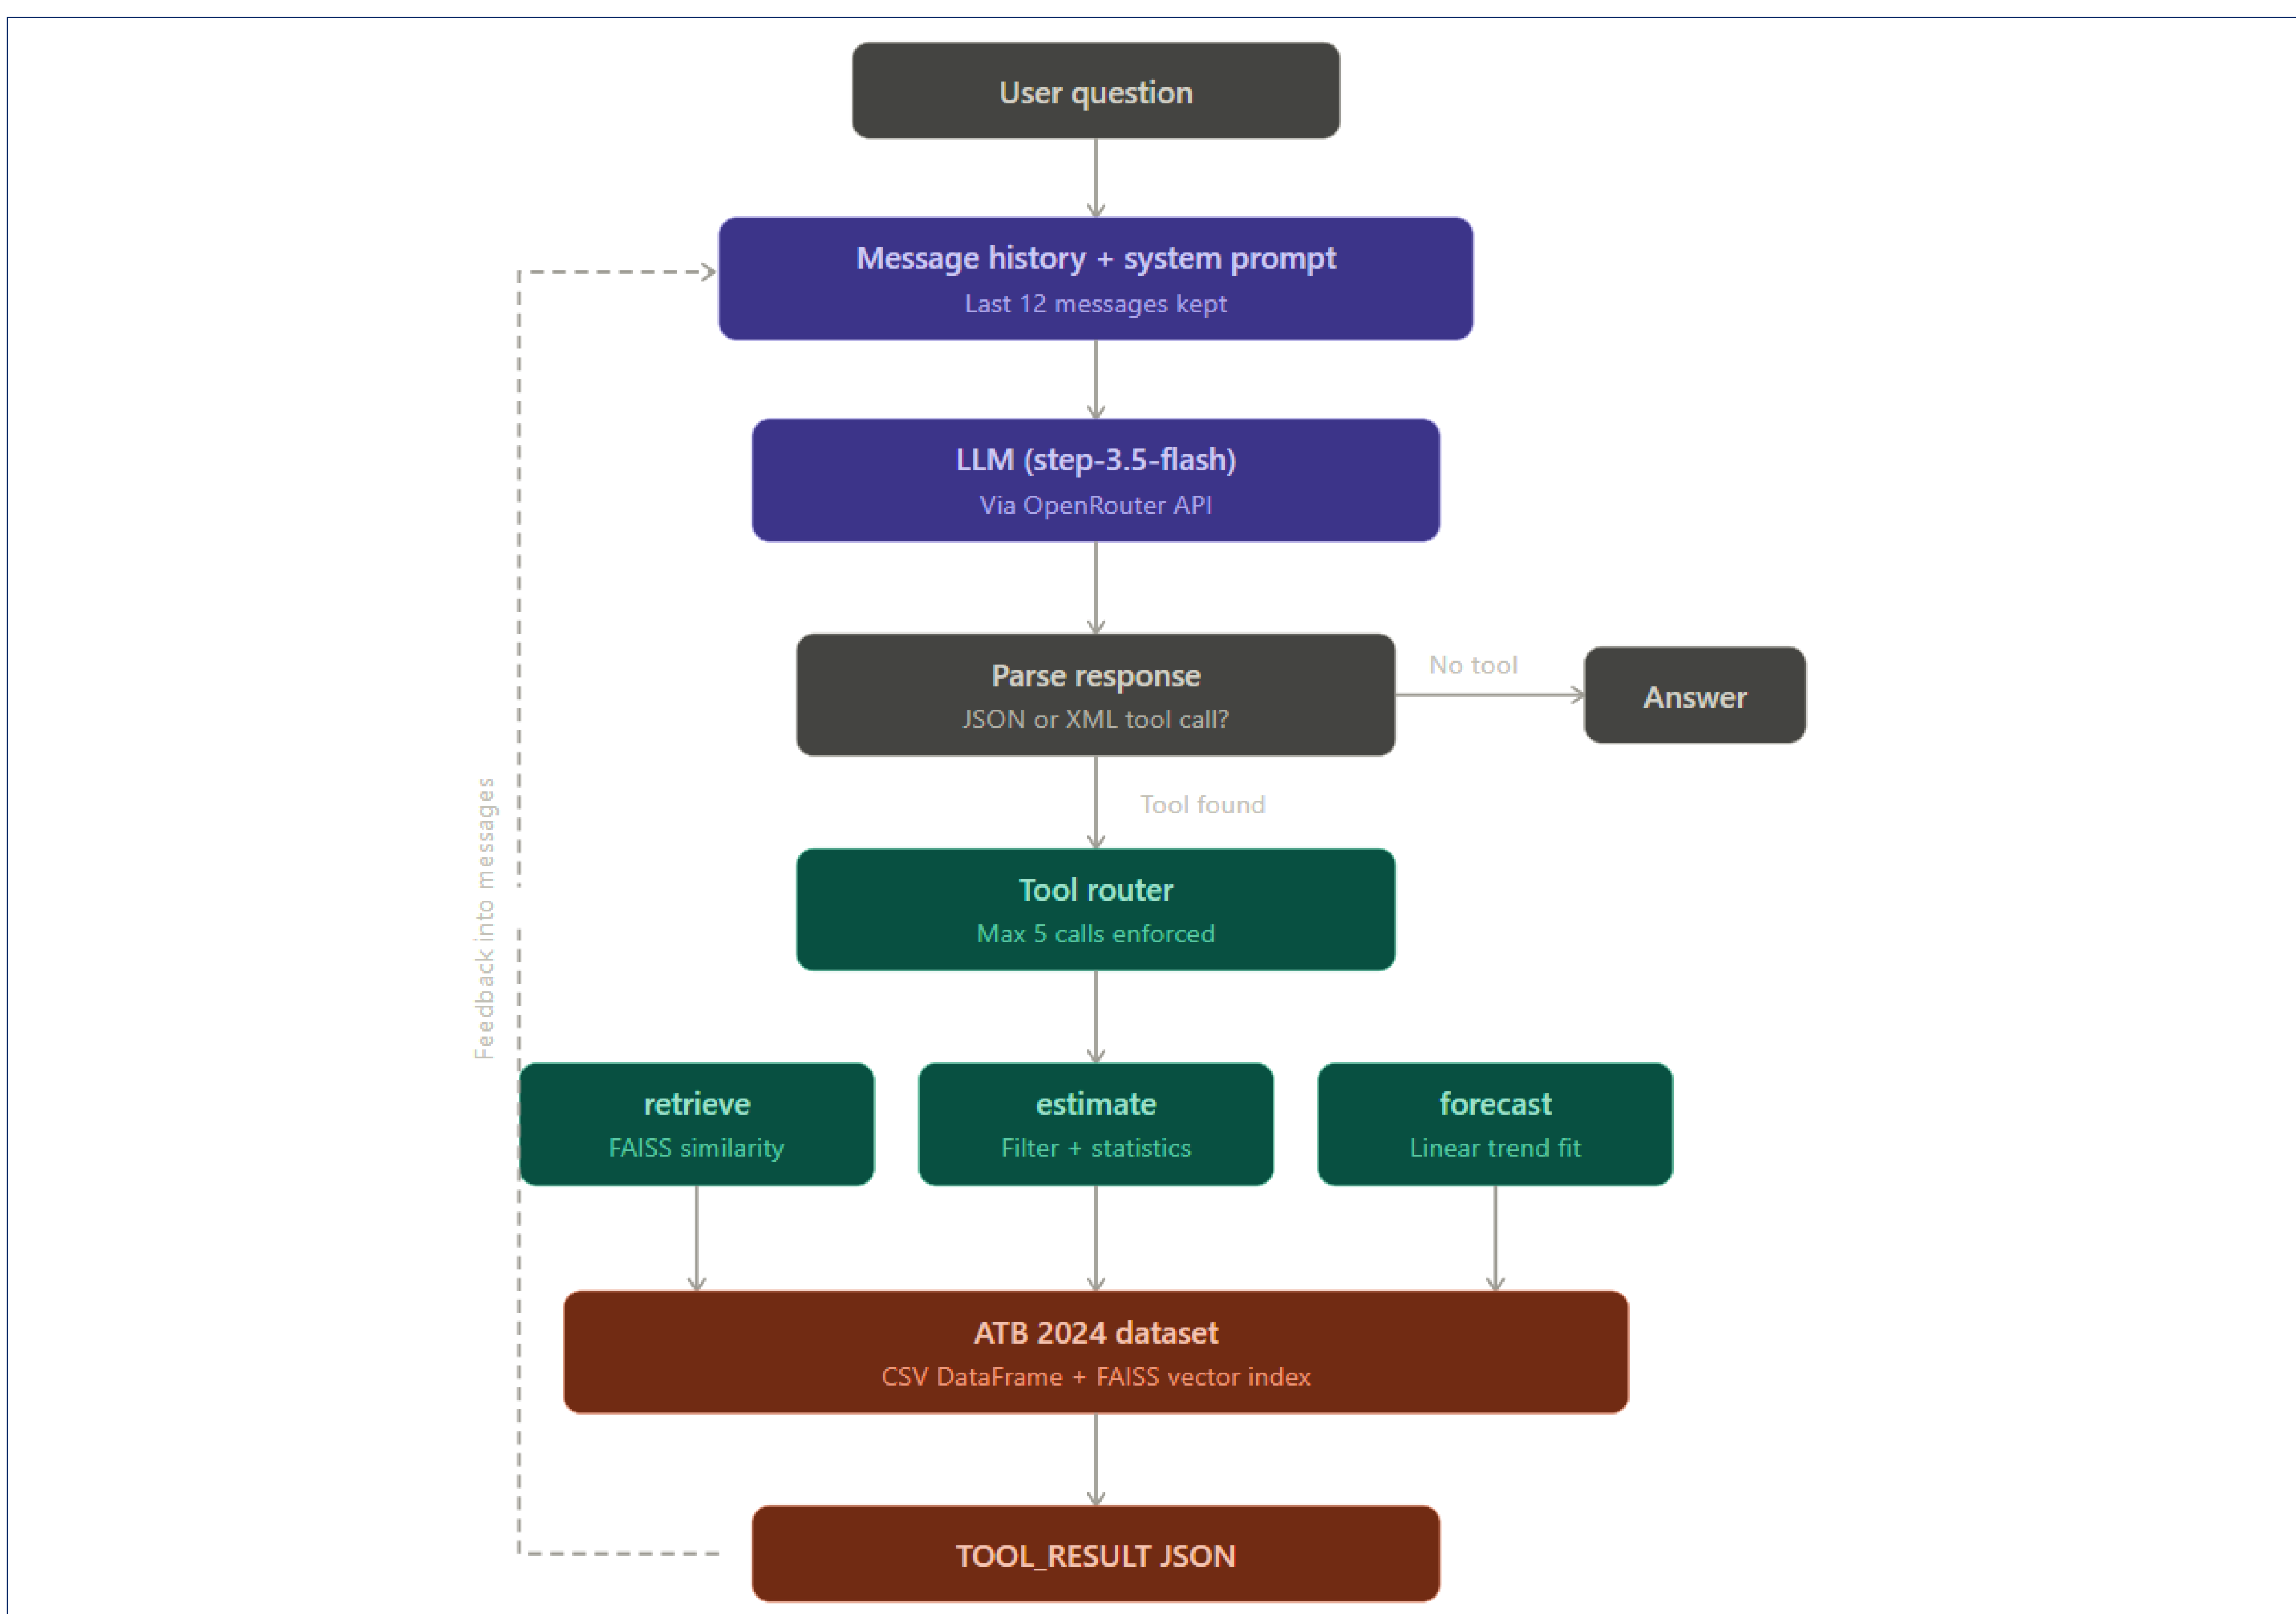
**Some details to note:**

- Free tier AI model would frequently rate limit, making it near unusable during high usage times
- Linear trend fitting for forecasting often show weaker R squared values.

**Future work ideas:**

- Package the agent as software with a user interface
- Add polynomial or learning-curve models to improve forecast accuracy
- Support multi-technology comparison queries (e.g. "compare solar vs wind CAPEX through 2040")

## Flow Chart Diagram for Agent Learning Loop



## Acknowledgements

Special thanks to Professor Matthew Magnusson for his aid in understanding and developing our agent.

## References

- NREL Annual Technology Baseline (ATB): <https://atb.nrel.gov/>  
 ATB 2024 Dataset (Open Energy Data): <https://data.openei.org/submissions/6006>  
 ATB 2024 Metric Definitions & Units: <https://atb.nrel.gov/electricity/2024/definitions>