

How to Build Your Own Literature-Review Agentic AI

Moji Ghadimi | Head of AI and Quantum Algorithms, QCIF

Members & Partners



Agentic AI in Research (What & Why)



Definition & Concept

- Agentic AI = LLM-based systems that **plan, reason, and act** across multi-step tasks — not just respond.
- In research → automates steps like searching, reading, synthesising, hypothesising, coding.
- Still, **human-in-the-loop is needed** — AI acts as an *assistant*, not an autonomous researcher.

Benefits

- Tames information overload (> 100k papers/week).
- Speeds literature review → hypothesis → analysis cycle.
- Enables reproducible, traceable, interactive pipelines.

Challenges / Cautions

- Hallucination & coverage gaps → need retrieval grounding.
- Data/privacy, IP, and transparency issues.
- Overselling autonomy risks loss of trust; human oversight remains essential.

Where Agentic AI Is Used in Research



Category	Purpose	Major Examples (Commercial)	Free / Academic Alternatives
Literature Review & Discovery	Search → summarise → map topics → draft reviews	<i>Web of Science Research Assistant</i> (Clarivate); <i>Nature Research Assistant</i> (Springer Nature)	<i>scienceOS</i> ; <i>OpenAlex agents</i> ; open repos like <i>agentic-scholar</i>
Hypothesis & Experiment Planning	Generate research questions, design experiments	<i>Owkin K Navigator</i> (biomedical AI)	<i>AIssistant</i> (open-source research collab tool)
Data Analysis & Manuscript Prep	Code generation, visualisation, writing	<i>Denario</i> , <i>Elicit Pro</i>	<i>Agentic-Research-Assistant</i> (GitHub)
Research Intelligence / Trend Mapping	Field mapping, citation networks & front detection	<i>Clarivate Agentic AI for Research & Learning</i>	<i>Semantic Scholar</i> / <i>OpenAlex</i> bibliometric agents
National R&D Initiatives	End-to-end scientific workflow automation	CSIRO Sciansa : multi-agent AI for scientists (idea → experiment → analysis)	—

Takeaway: Literature review is the most mature and widely adopted agentic AI use case today but not the only one.

A Personal AI Assistant



- **Small-scale, modular research assistant**, not a fully autonomous agent.
- Built during NCI training for Australian researchers (qcif-training.github.io/llm_workshop).
- Focus: reproducibility, grounded retrieval, transparent synthesis.
- Even for Agentic AI, human control is recommended → aligns with responsible research practice.

	title	summary	published	link	pdf	authors	summary_llm	structured
0	Local-projective-measurement-enhanced quantum ...	Quantum batteries have significant potential a...	2024-05-06		http://arxiv.org/pdf/2405.03093v1	Tinggui Zhang, Hong Yang, Shao-Ming Fei	Researchers investigated methods to enhance th...	{'methodology': 'Local projective measurements...
1	Entanglement and energy transportation in the ...	Quantum battery exploits the principle of quan...	2025-02-11		http://arxiv.org/pdf/2502.07513v1	Fan Liu, Hui-Yu Yang, Shuai-Li Wang, Jun-Zhong...	This study explores the concept of a quantum b...	{'methodology': 'Invariant subspace method', '...
2	High-performance Kerr quantum battery	We propose and investigate the performance of ...	2023-05-04		http://arxiv.org/pdf/2305.03202v2	Muhammad Shoufie Ukhtary, Ahmad R. T. Nugraha,...	Researchers propose a "Kerr quantum battery" c...	{'methodology': 'Investigation of a hybrid qua...

Problem & Motivation



Why automate literature review?

- Exponential growth of research output — manual review $\approx$ impossible
- Researchers need **custom, private, reproducible** pipelines (not generic ChatGPT)
- AI Assistant can:
 - Search and filter across APIs, databases, and PDFs
 - Extract **structured insights**
 - Write summaries **with traceable references**
- The contents of the course also cover practical tips for crafting **structured prompts** that yield reliable, **well-organised responses** while **minimising hallucinations**.

Structured Multi-Criteria Scoring

Another approach is to define **specific evaluation criteria** and let the LLM (or a human reviewer) score each abstract numerically.

For example, score each paper from **1–5** on: - *Method novelty*

- *Data similarity to your question*

- *Domain alignment*

- *Practical applicability*

Core Architecture



How it works (without the code):

- **Frontend:** Python notebooks (OpenAI API / local model) + simple chat interface
- **Memory/context management:** stores history for recall
- **Retrieval pipeline:**
 - Query → Semantic search → Top-k documents
 - LLM synthesises & cites evidence

Think of it as a *modular research assistant*, not a monolithic chatbot.

- **Future Agentic capabilities:**
 - Plans → Executes sub-tasks (search, summarise, critique) → Returns final synthesis → Finds gaps → Runs Experiments 🧪

Example structured abstract summary:

```
{'methodology': 'Local projective measurements on bipartite quantum systems', 'data': 'Two-qubit Bell-diagonal and X-type states', 'finding': 'Local measurements can improve quantum battery capacity', 'limitations': 'Theoretical analysis, experimental development needed'}
```

Key Functions to Highlight



- **chat(), clear_chat()** → Easily send queries and manage chat history.
- **search_arxiv(query)** → retrieves relevant papers using API (Can be replaced with any other database).
- **extract_structured_info(abstract)** → produces structured output: {'methodology': '...', 'data': '...', 'finding': '...', 'limitations': '...'}
- **select_and_fetch_pdfs(abstracts, question)** → Selects the most relevant papers and fetches full text. This will be used to write a review.

Emphasise: “LLM ≠ knowledge base — retrieval provides grounding

Sample short review output



title	summary	pdf
0 Local-projective-measurement-enhanced quantum ...	Quantum batteries have significant potential a...	http://arxiv.org/pdf/2405.03093v1
1 Optimizing quantum battery performance by redu...	Quantum batteries have emerged as promising de...	http://arxiv.org/pdf/2505.08029v1
2 Efficient charging of multiple open quantum ba...	We explore a protocol that efficiently charges...	http://arxiv.org/pdf/2410.19303v1

The influence of local projective measurements on the charging capacity of quantum batteries has been a subject of interest in recent studies. Research has shown that local projective measurements can enhance the capacity of quantum batteries, particularly when considering bipartite quantum systems (<http://arxiv.org/pdf/2405.03093v1>). By applying local projective measurements on a subsystem of the quantum state, the battery capacity with respect to the whole system or a subsystem can be improved. This has been demonstrated using two-qubit Bell-diagonal states and X-type states as examples, highlighting the potential for local projective measurements to optimize quantum battery performance (<http://arxiv.org/pdf/2405.03093v1>).

The optimization of quantum battery performance is a complex task that involves understanding the interplay between the battery and charger Hamiltonians. To minimize the battery's influence during charging dynamics, a control parameter can be introduced to suppress the battery's contribution (<http://arxiv.org/pdf/2505.08029v1>). Furthermore, efficient charging of multiple open quantum batteries can be achieved through dissipation and pumping, where collective coupling of the charger and the battery to the same thermal reservoir enables super-extensive charging (<http://arxiv.org/pdf/2410.19303v1>). The incorporation of both dissipation and incoherent collective pumping on the charger can also facilitate efficient parallel charging of many quantum batteries, even when each battery is coupled to a different thermal reservoir (<http://arxiv.org/pdf/2410.19303v1>). These findings highlight the importance of considering local projective measurements and charging dynamics in the optimization of quantum battery capacity.

Insights on Using LLMs Properly



Insights shared in the course:

- **Prompt engineering ≠ magic** — clarity and structure matter more than clever phrasing
- **Use structured outputs** (JSON, tables, bullet formats) to make responses consistent and machine-readable
- **LLMs hallucinate when unguided** → retrieval grounding is essential
- **Reproducibility:** log prompts, model versions, and data sources
- **Ethics & IP:** always check license terms, privacy, and data sensitivity before ingestion

Future Directions / Takeaways



- **Democratising research assistance** → lightweight, local, and domain-tuned agents that anyone can build safely
- **Planning stages of research** → linking literature analysis with project management, dataset discovery, and method selection
- **Identifying research gaps** → using summaries to flag under-explored questions, missing datasets, or weak evidence bases
- **Multi-agent collaboration** → specialised agents for search, critique, synthesis, and debate working together
- **Human-in-the-loop** → researcher defines goals, verifies reasoning, and curates final outputs
- **Integration with lab and data systems** → agents connecting to Electronic Lab Notebooks, simulation pipelines, and FAIR data repositories
- **Responsible agent design** → auditability, reproducibility, privacy, and ethical guardrails baked into agent frameworks

Resources: Large Language Models for Research Workshop

https://qcif-training.github.io/llm_workshop