Guide to Conducting a Literature Review in Robotics

A Step-by-Step Student Handbook

1. Introduction and Purpose

This guide applies to literature reviews in academic robotics research and is specifically designed for the purpose of the seminar. You are not required to follow these guidelines in detail. They are meant as supporting suggestions. In parts they will go beyond what will be required in the seminar.

2. Planning Your Literature Review

2.1 Define Your Research Question

Before searching for papers, formulate a clear, focused research question. Broad topics such as "robot navigation" will yield thousands of irrelevant results. Instead, specify dimensions such as:

- Environment: indoor, outdoor, underwater, aerial
- Perception modality: LiDAR, vision, tactile sensors
- Algorithmic family: SLAM, reinforcement learning, model predictive control
- Application domain: manufacturing, medical, service robotics

Example of a well-scoped question:

> **Example Research Question**
> "What deep learning approaches have been proposed for sim-to-real transfer in contact-rich robot manipulation tasks, and what are their primary failure modes?"

2.2 Establish Inclusion and Exclusion Criteria

Decide in advance which papers are eligible for your review. Document these criteria explicitly.

Include	Exclude
Peer-reviewed journal articles and conference papers (see venue list below)	Blog posts, press releases, marketing material
Publications from the last 5 years	Papers older than 15 years unless foundational

Papers addressing your specific sub-domain	Tangentially related papers outside your scope
Papers with reproducible or publicly described methods	Abstracts-only with no accessible full text

2.3 Define a Search Protocol

A reproducible search protocol strengthens the credibility of your review. It should specify:

- Which databases you searched (see Section 3)
- Which keywords and Boolean operators you used
- Date range and language filters applied
- How many results were retrieved and how many were selected

3. Finding Relevant Literature

3.1 Primary Search Databases

For robotics research, the following databases are recommended:

- IEEE Xplore (ieeexplore.ieee.org) — Essential for IEEE conference and journal papers; covers the majority of robotics publications
- Google Scholar — Broad coverage including preprints; use for discovery, not final citation lists
- arXiv (specifically cs.RO and cs.CV, cv.LG section) — Preprint server; check for latest work, but verify if published elsewhere
- ACM Digital Library — Relevant for HRI, cognitive robotics, and robotics-adjacent CS topics
- Scopus / Web of Science — Broad scientific databases useful for citation tracking and impact metrics
- Semantic Scholar — AI-powered search with citation graphs; useful for finding related work

> **Practical Tip**
> Start with IEEE Xplore and Semantic Scholar for most robotics topics. Use Google Scholar for a wide first pass, but confirm sources are peer-reviewed before citing them.

3.2 Keyword Strategy

Effective search requires both precision and breadth. Use a combination of:

- Core terms: exact domain vocabulary (e.g., "SLAM", "reinforcement learning", "impedance control")
- Synonyms: alternative phrasings used in the community (e.g., "Imitation Learning" vs. "Learning from demonstration")
- Boolean operators: AND to narrow, OR to broaden (e.g., "manipulation" AND ("deep learning" OR "neural network"))
- Wildcards: to capture variants (e.g., "localiz*" captures localization, localizing, localized)

Iterate your keyword strategy: review the titles and abstracts of the first 20 results to discover new relevant terms used by the community.

3.3 Snowballing

Once you have a core set of papers, expand your coverage through:

- Backward snowballing: follow the references cited within papers you have already found
- Forward snowballing: find papers that cite your core papers (use Google Scholar's "Cited by" feature)
- Review papers: identify recent survey papers in your area; their reference lists are valuable maps of the field

4. Assessing Relevance and Quality

4.1 Evaluating Publication Venues

Not all publications carry equal scientific weight. In robotics, venue prestige is a meaningful quality signal. The table below provides an overview of high-quality robotics venues:

Venue	Publisher	Tier / Notes
IEEE Transactions on Robotics (T-RO)	IEEE	A* – Flagship robotics journal
International Journal of Robotics Research (IJRR)	SAGE	A* – Oldest, highly cited
Robotics and Autonomous Systems (RAS)	Elsevier	A – Broad robotics scope
IEEE Robotics & Automation Letters (RA-L)	IEEE	A – Fast publication letters
Science Robotics	AAAS	A* – High impact, interdisciplinary
ICRA (Conference)	IEEE	A – Premier annual conference
IROS (Conference)	IEEE	A – Large international conference
RSS (Conference)	RSS	A* – Selective, high quality
CoRL (Conference)	CoRL	A* – Robot learning focus
International Journal of Computer Vision (IJCV)	Springer	A – classic CV journal; many perception papers applied to robotics
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	IEEE	A – leading journal for vision and learning methods widely adopted in robotics
NeurIPS (Conference)	NeurIPS Foundation	A* – Robot learning, RL, perception, sim-to-real are major tracks
ICML (Conference)	International Machine Learning Society	A – ICML traditionally features more content on reinforcement learning, robotics, and optimization theory than its peers
ICLR (Conference)	ICLR	A – explicitly lists robotics as an application area; covers deep learning, reinforcement learning, and representation learning

CVPR (Conference)	IEEE CS / CVF	A* – perception, 3D understanding, autonomous driving, manipulation
ICCV (Conference)	IEEE CS / CVF	A – broad computer vision; strong overlap with object detection, tracking, scene understanding for robots
ECCV (Conference)	Springer	A – alternates with ICCV; similar scope

4.2 Assessing Author Reputation

Author credibility is a complementary signal to venue quality. Established researchers in robotics tend to publish consistently in the top venues listed above.

How to assess the author's reputation:

- Google Scholar Profile: check h-index, total citations, and the journals/conferences they publish in regularly. An h-index above 20 indicates a well-established researcher.
- Semantic Scholar Profile: provides a cleaner citation graph and co-author network
- DBLP (dblp.org): comprehensive computer science bibliography; excellent for tracking an author's publication history over time
- ResearchGate: less authoritative, but useful for finding full texts and preprints
- Reputation of the first author is usually not most relevant one. In many publications first authors are (PhD) students while last authors are the supervising professors.

Key questions to ask about the authors:

- Are they affiliated with a recognized research institution or lab?
- Do they publish in top-tier venues consistently, or predominantly in lower-tier outlets?
- Are their papers regularly cited by others in the field (not just self-citations)?

4.3 Critically Reading a Paper

Do not simply accept what a paper claims. Apply the following critical reading checklist:

- Is the problem statement clearly defined and motivated?
- Is the proposed method technically sound?
- Are experiments conducted on standard benchmarks or real hardware?
- Are baselines fairly chosen and up to date?
- Are limitations discussed honestly?
- Are results statistically significant (error bars, ablation studies)?
- Is the code or dataset publicly available for reproducibility?

5. Organizing and Synthesizing the Literature

5.1 Use a Reference Manager

Manual management of citations is error-prone. Use a dedicated tool from the start:

- Zotero (zotero.org) — Free, open-source, excellent browser integration; highly recommended
- Mendeley — Free with cloud storage; integrates with Elsevier
- JabRef — Ideal for LaTeX users; works directly with .bib files

Immediately annotate each paper you read with a short note summarizing: the core contribution, the method, the dataset/hardware used, key results, and relevance to your topic.

5.2 Build a Synthesis Matrix

A synthesis matrix is a table that maps papers (rows) against key dimensions (columns). Example dimensions for a robotics review:

- Task: grasping, navigation, planning, etc.
- Approach: learning-based, model-based, hybrid, etc.
- Evaluation: simulation only, real robot, public benchmark, etc.

This matrix allows you to see patterns, identify overrepresented approaches, and spot gaps systematically rather than impressionistically.

5.3 Organize Thematically, Not Chronologically

A good literature review is not a list of paper summaries. It synthesizes the field into coherent themes. Common organizational strategies in robotics:

- By methodology: group papers by algorithmic family (e.g., model-based vs. data-driven approaches)
- By sub-problem: group by the aspect of the problem addressed (perception, planning, actuation, learning)
- By evaluation setting: group by deployment context (simulation, structured lab, unstructured real world)

Within each theme, compare and contrast: what do papers agree on? Where do they conflict? Which assumptions are made?

6. Using AI Tools: Benefits and Risks

Important Note: The suggestions bellow do not replace the university and faculty guidelines on usage of AI tools in teaching

Please check for official guidelines provided by the university or faculty about usage of AI tools in exam-relevant work. This is specifically relevant for any written text/report/survey you submit as part of the evaluation.

AI language models such as ChatGPT and Claude can provide meaningful assistance at various stages of the literature review process. However, they introduce specific risks that you should understand and actively manage.

6.1 Overview of Benefits and Risks

AI Use Case	Potential Benefit	Risk / Caution
Scoping & brainstorming	Generate initial keyword lists, identify sub-fields, map related topics	May miss niche or highly recent areas; verify against databases
Explaining concepts	Clarify unfamiliar terminology, methods, or mathematical foundations	Explanations may be oversimplified or occasionally wrong; cross-check

Summarizing papers	Assist digesting long papers quickly, identify key contributions	AI may hallucinate details; always read the original abstract and method
Identifying gaps	Prompt to suggest potential research gaps from a list of summaries	AI suggestions are generic; validate with domain expert knowledge
Writing & paraphrasing	Improve clarity of your own writing, check grammar and structure	Risk of plagiarism if AI text used verbatim; always rewrite in own words
Citation management	NOT recommended – AI invents non-existent references frequently	NEVER trust AI-generated citations; use Zotero, Mendeley, or Semantic Scholar

6.2 Recommended AI-Assisted Workflows

Keyword generation

Prompt an AI tool with your research question and ask it to suggest relevant keywords, related sub-fields, and synonyms. Then validate these against your own search results.

Example Prompt

"I am writing a literature review on sim-to-real transfer for contact-rich robot manipulation. Suggest 15 relevant search keywords and related sub-fields I should explore."

Conceptual clarification

If you encounter a method or concept you do not understand (e.g., "Riemannian manifold optimization in trajectory planning"), ask an AI to explain it in plain language. This accelerates comprehension. Always verify technical claims against textbooks or primary sources.

Writing assistance

Use AI to improve the clarity, structure, and grammar of text you have already written. Do not ask AI to write sections for you from scratch as the results may be too generic and occasionally wrong. It further deprives you of the intellectual engagement central to academic work, and in most institutions constitutes academic misconduct if undisclosed.

6.3 Critical Risks to Manage

Hallucinated references — the most serious risk

AI language models frequently generate plausible-sounding but entirely fictional paper titles, authors, journal names, and DOIs. This is called hallucination. If you copy such a reference into your work without verification, it constitutes a citation to a non-existent paper — a serious academic integrity violation.

Rule: Never Cite AI-Generated References

Every reference must be verified independently in a trusted database (IEEE Xplore, ACM DL, Scopus, Google Scholar) before inclusion. There are NO exceptions to this rule. If you cannot locate the paper, do not cite it.

Outdated knowledge

AI models have a training data cutoff and are unaware of papers published after that date. In robotics, where the field advances very rapidly, this means significant recent work will be invisible to AI. Always supplement AI-assisted searching with direct database queries using date filters.

Superficial synthesis

AI tools can produce fluent-sounding summaries that are actually shallow, incorrect, or missing nuance. A literature review written primarily from AI summaries will lack the depth expected at the graduate level and may misrepresent papers. Read primary sources yourself.

Academic integrity and disclosure

Check your institution's policy on AI use before employing these tools. Many universities now require explicit disclosure of AI assistance in submitted work. Using AI to write portions of your literature review without disclosure, where prohibited, may constitute academic misconduct. When in doubt, ask your supervisor.

6.4 Useful AI-Powered Research Tools

These tools are specifically designed for academic research and are generally safer than general-purpose chatbots:

- Semantic Scholar (semanticscholar.org) — AI-powered paper search with citation graphs and research summaries
- Connected Papers (connectedpapers.com) — Visualizes citation networks to find related work
- Elicit (elicit.com) — Designed for systematic reviews; searches and extracts structured information from papers (account required, free version)
- Research Rabbit (researchrabbit.ai) — Finds clusters of related papers from a seed set (account required, free version)

Recommendation

Use these domain-specific research AI tools in preference to general-purpose chatbots for the search and discovery phases. They are grounded in real literature and carry a much lower risk of hallucination.

7. Writing the Literature Review

7.1 Structure

A standard literature review in a robotics thesis or paper typically contains:

- Introduction to the field: situate the review within the broader research context
- Taxonomy of the field: present the major sub-categories or approaches
- Thematic discussion: synthesize and critically compare papers within each theme
- Identification of gaps: articulate what is missing, understudied, or unresolved
- Positioning statement: briefly explain how your own work addresses one or more of these gaps

7.2 What Makes a Good Literature Review

- Critical, not merely descriptive: go beyond summarizing — compare, evaluate, and draw conclusions
- Well-scoped: demonstrates mastery of a focused area rather than superficial coverage of everything
- Up to date: includes recent publications, particularly from the last 3–5 years
- Balanced: represents multiple research groups and perspectives, not only one lab's work
- Concise and precise: no padding; every sentence contributes
- Properly cited: all claims are attributed; citation format is consistent

7.3 Common Mistakes to Avoid

- Writing a list of summaries without synthesis or comparison
- Ignoring contradictory evidence or inconvenient findings
- Over-relying on a small number of papers from a single research group
- Citing secondary sources (textbooks, Wikipedia, review blogs) instead of primary papers
- Neglecting to discuss evaluation methodology and reproducibility
- Failing to acknowledge limitations of cited methods