

# Preface

## 0.1 Getting Started: About the Project and How to Navigate It

### 0.1.1 Why This Document?

This document is intended as a structured, accessible on-ramp to deep learning and computer vision—both for students and early-stage practitioners, and for experienced engineers who want a quick way to enter new subfields or refresh core ideas. The aim is to replace resource overload with a coherent path: concise explanations, consistent notation, and curated references that connect lectures, key papers, and practical recipes into a navigable whole.

This volume originated in the summer of 2024 as a structured companion to the University of Michigan’s EECS498 curriculum, taught by Justin Johnson. By revisiting each lecture and distilling its core ideas, the text organizes the material into a practical, pedagogically coherent resource that extends the course with clarified explanations, curated references, and contextual notes designed to support durable understanding.

The material is tailored for undergraduate and early graduate students in computer science, electrical engineering, software engineering, and related fields. Readers are expected to have foundational knowledge in Python programming, calculus, and linear algebra. If you feel uncertain about these prerequisites, I recommend reviewing the following:

- Learn Python for programming fundamentals,
- Khan Academy Calculus for introductory calculus, and
- 3Blue1Brown’s Linear Algebra series for geometric intuition and matrix operations.

While prior coursework in machine learning is not strictly required, it can be highly beneficial. For a beginner-friendly option, I recommend Andrew Ng’s Machine Learning course on Coursera. For those seeking a more theoretical treatment, Stanford’s CS229: Machine Learning offers a rigorous mathematical perspective.

**A Living Resource.** This document is intended to evolve into a living resource—open, modular, and continuously improved by the community. In the near future, I plan to release the full source files and compilation scripts (e.g., LaTeX + Markdown + BibTeX) to support versioning, feedback, and collaborative contributions. Readers will be able to:

- Suggest clarifications, corrections, and additions,
- Contribute new sections, figures, or examples,
- Help maintain up-to-date coverage of evolving topics in the field.

The goal is to build a shared, reliable foundation for learning deep learning and computer vision—one that reflects the generosity and rigor of the research community it draws from.

### 0.1.2 Your Feedback Matters

This project was created independently during my personal time and at the time of the first publication, has been reviewed by only a small number of readers. As such, your feedback—whether in the form of corrections, clarifications, or suggestions—is highly appreciated and vital to improving the clarity, accuracy, and usefulness of the text.

As previously mentioned, looking ahead, I plan to make this project fully open-source to support broader collaboration. If you are interested in contributing—whether by expanding coverage, improving explanations, suggesting exercises, or proposing structural changes—I would love to hear from you. Please don't hesitate to reach out via email, suggest pull requests via the project GitHub and thank you in advance for helping shape this resource into a living, community-driven guide for deep learning and computer vision.

### 0.1.3 How to Use This Document Effectively

This document is designed to support a wide range of learners—from students discovering deep learning for the first time to practitioners revisiting foundational concepts. It is not intended to be read linearly from start to finish. Instead, treat it as a flexible reference that can adapt to your goals and level of experience. Here are some recommended strategies:

1. **If You're New to the Field:** Begin by following the EECS498 lecture series on YouTube, and read the corresponding sections in parallel. Use this document to:

- Clarify ideas presented quickly in the video.
- Reference equations or diagrams while watching.
- Revisit chapters after the lecture to solidify your understanding.

Try pausing to reflect on key definitions or derivations before reading through the explanations—they're designed to help you reason things out.

2. **If You're a Practitioner or Advanced Learner:** Use the table of contents to jump directly to the topics that interest you. This document is structured to support:

- Quick lookups and topic refreshers.
- In-depth conceptual reviews when needed.
- Bridging gaps between foundational material and modern techniques.

3. **Use It Selectively:** There is no expectation that you read the entire document. Each section is written to stand on its own (as much as possible), so feel free to explore topics non-linearly. If you're preparing for a specific project, paper, or interview, jump directly to relevant chapters—such as optimization, normalization, transformers, or generative models.

4. **Refresh your knowledge over time** Use this as a long-term companion. When you return after months to backpropagation, attention, or KL divergence, jump directly to the relevant section. Short definitions, worked mini-examples, and “enrichment” boxes are designed to support both first-pass learning and quick second-pass review
5. **Turn the notes into an interactive guide (RAG + podcast workflows)** Upload this document—together with a short reading list of papers on your target topic—into retrieval-augmented tools such as **NotebookLM** or **ChatGPT/Gemini with file support**. You can then:
  - Generate *podcast-style briefings* or episode outlines that walk through the papers in a narrative order.
  - Ask follow-up questions in natural language and get citations back to specific sections, figures, or equations.
  - Summarize long derivations, contrast methods head-to-head, and extract recurring themes or assumptions.
  - Auto-create quizzes, flashcards, and timelines to reinforce retention.

This workflow turns a static reading list into an interactive study companion, letting you survey a subfield, revisit concepts, or spin up a quick refresher with minimal friction.

#### 0.1.4 Staying Updated in the Field

While this document aims to provide a reliable and up-to-date foundation, deep learning continues to evolve at a remarkable pace—often with new methods, architectures, and theoretical insights emerging daily. As such, readers are encouraged to treat this material as a starting point rather than an endpoint.

To stay current with developments:

- Use tools like **Connected Papers** to explore the citation graph and discover related work.
- Follow **Trending Papers** to track trending papers, benchmark results, and open-source implementations.
- Subscribe to **DailyArxiv** to receive daily email digests of new arXiv submissions in machine learning, computer vision, and related fields.
- Watch high-level paper reviews and discussions on YouTube channels such as Yannic Kilcher and Andrej Karpathy.
- Follow key researchers and labs on **X** (formerly Twitter), such as @ylecun, @karpathy, @smerity, and @david\_ha, to stay informed about emerging ideas and discussions that may precede formal publication.
- Review accepted papers and tutorials from top-tier conferences such as **CVPR**, **ICCV**, **ECCV**, **NeurIPS**, **ICLR**, and **ICML**.
- Bookmark curated repositories like the Image Matching Papers List to explore focused subfields efficiently.

By combining this foundational document with regular engagement in the broader research ecosystem, you can maintain both conceptual depth and up-to-date knowledge as the field progresses.

### 0.1.5 Dependency Tree

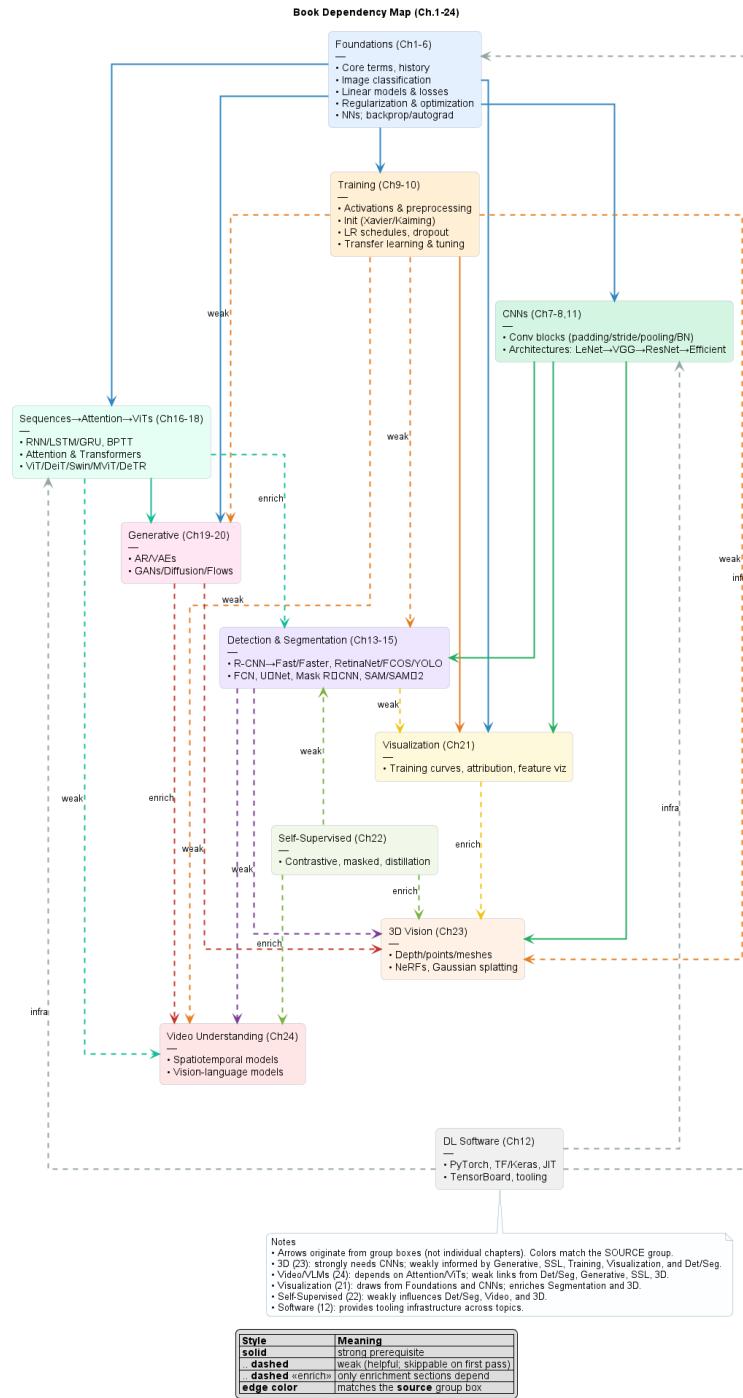


Figure 1: **Book chapter dependencies** Arrows indicate prerequisite relationships between chapters. Earlier foundations support later topics such as self-supervised learning, segmentation, video understanding, and 3D vision.

### 0.1.6 Contributors

- **Lead author and editor** — Ron Korine: compilation, writing, and curation across all sections.
- **Curriculum source** — Justin Johnson's EECS498 (University of Michigan): structure, topic sequencing, and foundational materials.
- **Figure and asset creators** — authors of publicly available figures, tutorials, and visualizations cited in-context throughout.

#### *Named contributors*

- **Adi Arbel** — contributions to the *ReLIC* family and *Plenoxels* sections; broad manuscript feedback on clarity, organization, and emphasis.

#### *Community feedback*

Thanks to colleagues and readers who suggested clarifications, spotted errors, and recommended references. If you have contributed content or identify material that should be credited differently or removed, please write to [eeecs498summary@gmail.com](mailto:eeecs498summary@gmail.com). Updates will be reflected in subsequent revisions.

### 0.1.7 The Importance of Practice

“What I cannot build, I do not understand”. Hands-on work is essential for mastering vision. Pair the lectures with the EECS498 assignments ([link](#)) to solidify core ideas through implementation. Use Google Colab for experiments and Kaggle for applied challenges, and track runs with WandB, ClearML, or TensorBoard. Building as you learn deepens theory and yields practical intuition for real-world tasks.

### 0.1.8 Final Remarks

Thank you for engaging with this material. Whether you are a student, researcher, or practitioner, I hope this document supports your efforts to develop a deeper and more structured understanding of deep learning and computer vision.

This resource is intended not only as a learning companion but also as a reference you can revisit over time. The field continues to evolve rapidly, and while the foundational ideas remain central, new architectures, techniques, and theoretical tools are introduced regularly. A solid conceptual base will enable you to adapt confidently to these developments.

I am also sincerely grateful to readers who have taken the time to share suggestions, point out errors, or recommend improvements. Your feedback helps refine and expand this work, and ensures that it remains a relevant and valuable resource for future learners.

*“Live as if you were to die tomorrow. Learn as if you were to live forever.”*

— Mahatma Gandhi