AI Forward-Deployed Engineering Sprint

Full Week Schedule

Monday – Friday

Day 1 Curriculum

Themes: Kickoff, Agentic AI Prototyping, Evals I, RAG

🎯Goals

- 🧑‍💻Engineering: Dev env setup, E2E stack, simple evals, Prompting, RAG
- 💼Product: “Concrete” ideas, product visions
- 🧑‍🤝‍🧑Group Project: Pitches w/ problem & proposed solution

📚Learning Outcomes

- 🗺️Overview of the AI Product-Engineering Program and what you will learn
- 🤖Edge: Understand how Product & Engineering are evolving in the age of AI
- 💼Product: Learn how leading VCs make concrete problems worth solving into products
- 🧑‍💻Learn key components of an AI Engineer's interactive development environments
- 🧑‍💻Understand the end-to-end prototyping stack we will use throughout the program
- 🧑‍💻Understand how to perform simple evals
- 🧑‍💻Understand core constructs of Pythonic RAG applications

💻Activities

- 💼Articulate a problem worth solving with your group
- 💼Determine if your solution requires RAG
- 💼Understand how to articulate a product vision
- 💼Pitches (end of day)
- 🧑‍💻Set up interactive development environments
- 🧑‍💻Build an end-to-end prototype in your dev environment
- 🧑‍💻Run simple LLM-as-a-judge evals and vibe checks
- 🧑‍💻Build a RAG application from scratch in Python

📜Recommended Reading

- Word2Vec (Jan 2013) (Topic: Embeddings)

- Chain of Thought (Jan 2022) (Topic: Prompting/Evals)
- Self-Refine (Mar 2023) (Topic: Prompting/Evals)
- RAG (May 2020) (Topic: RAG)
- The LLM Application Stack (June 2023) (Topic: Prototyping)
- 12-Factor Agents (June 2025) (Topic: Context Engineering)

Day 1 Detailed Schedule

AM

- 9:00 AM: 👋Dev Environment Set Up/Meet & Greet Teams
- 10:00 AM: 🏈Kickoff
- 10:45 AM: 🤖(Lecture - 15 min) Introduction to Agentic AI Prototyping
- 11:00 AM: 💼(Lecture - 5 min) Intro to "Concrete" problems worth solving
- 11:05 AM: 👫(Groups - 30 min) Articulate a "concrete" problem worth solving
- 11:35 AM: 🤩 (Lecture + Demo - 10 min) Example Demo Day Presentation
- 12:00 PM: Lunch

PM

- 1:00 PM: 👩‍🏫(Lecture - 10 min) Intro to deve env setup & end-to-end prototype stack
- 1:10 PM: 👩‍💻(Demo - 15 min) dev env setup & E2E prototype build
- 1:25 PM: 👩‍🏫(Lecture - 15 min) Intro to Evals
- 1:40 PM: 👩‍💻(Demo - 10 min) Evals / LLM-as-a-judge notebook demo
- 1:50 PM: 👫(Groups - 30 min) How will you leverage vibe checks & evals? **
- 2:20 PM: 👩‍🏫(Lecture - 25 min) Intro to RAG & Dense Vector Retrieval
- 2:45 PM: 👩‍💻(Demo - 15 min) Pythonic RAG from scratch notebook
- 3:00 PM: 👫(Groups - 30 min) Where will you get data? How will you chunk it? **
- 3:30 PM: 👩‍🏫(Lecture - 5 min) Defining a Product Vision
- 3:35 PM: 👫(Groups - 30 min) Defining your product vision & preparing pitches
- 4:05 PM: 🎤(Pitches - 5 min * 8)
- 4:45 PM: 👩‍🏫(Lecture - 15 min) Conclusions & Day 1 Wrap

** = [Optional: Run Notebook]

Day 2 Curriculum: Tuesday

Themes: Agents, Memory, Harnesses, Skills, MCP, and LangChain

🎯Goals

- 👩‍💻Engineering: Agents, Memory, Harnesses, Skills, MCP, LangChain
- 💼Product: product sense, user empathy
- 👫Group Project: Share vibe-coded frontend & backend infrastructure diagram

📚Learning Outcomes

- 🤖Edge: The Latest from the Lang-X ecosystem
- 💼Learn why product sense and user empathy are important
- 👩‍💻Understand what an agent is, and how it uses tools
- 👩‍💻Understand agent memory and the different types
- 👩‍💻Understand what an agent harness is, and its role in agentic systems
- 👩‍💻Understand Agent Skills vs. MCP servers

💻Activities

- 💼Vibe-code frontend UI options
- 💼Draw your solution's infrastructure diagram with agents, tools, and harness
- 💼Share your best frontend and infrastructure diagrams with the group
- 👩‍💻Build Pythonic agent from scratch
- 👩‍💻Build Pythonic agent harness from scratch
- 👩‍💻Build Pythonic agent memory from scratch
- 👩‍💻Build Pythonic spec-driven-development from scratch
- 👩‍💻Build Pythonic agent skills & MCP server from scratch
- 👩‍💻Build an agent and harness with RAG tool, memory, skills, MCP server in Langchain

📜Recommended Reading

- Philosophy, by LangChain (Topic: Agent Orchestration)

- Context engineering in agents, by LangChain (Topic: Context Engineering)
- Thinking in LangGraph, by LangGraph (Topic: Product meets Engineering)
- Model Context Protocol (Nov 2024) (Topic: MCP)
- Agent Skills (Oct 2025) (Topic: Skills)
- Deep Agents Overview, by LangChain (Topic: Agent Harnesses)
- Effective Harnesses for Long-Running Agents (Topic: Agent Harnesses)
- Harness Design for Long-Running Agents (Topic: Agent Harnesses)
- Why product sense is the only skill that will matter in the AI age

Day 2 Detailed Schedule

AM

- 9:00 AM: 👩‍🏫(Lecture - 20 minutes) Introduction to Agents
- 9:20 AM: 🧑‍💻(Demo - 10 minutes) Building an Agent from Scratch
- 9:30 AM: 👫(Groups - 30 min) What agents and tools will you use in your app? **
- 10:00 AM: 👩‍🏫(Lecture - 15 minutes) Introduction to Agent Harnesses
- 10:15 AM: 🧑‍💻(Demo - 10 minutes) Building an Agent Harness from Scratch
- 10:25 AM: 👫(Groups - 35 min) Draw the agents, harnesses, and tools you'll need **

10:30 AM: BREAK (10 minutes)

- 11:00 AM: 👩‍🏫(Lecture - 10 minutes) Introduction to product sense & user empathy
- 11:10 AM: 👫(Groups - 20 min) Draw the user interface that your group agrees on
- 11:30 AM: 👫(Groups - 20 min) Vibe code user interfaces individually
- 11:50 AM: 👫(Groups - 10 min) Choose the best UI; share with entire group*
- 12:00 PM: Lunch

PM

- 1:00 PM: 👩‍🏫(Lecture - 10 minutes) Introduction to Agent Memory
- 1:10 PM: 🧑‍💻(Demo - 10 minutes) Building Agent Memory from Scratch
- 1:20 PM: 💼(Lecture - 10 min) Intro to product sense & user empathy
- 1:30 PM: 👫(Groups - 30 min) Draw how you will leverage agent memory **
- 2:00 PM: 👩‍🏫(Lecture - 15 minutes) Introduction to Agent Skills & MCP
- 2:15 PM: 🧑‍💻(Demo - 10 minutes) MCP vs. Agent Skills

2:30 PM: BREAK (10 minutes)

- 2:25 PM: 👫(Groups - 35 min) Determine what skills & servers you will use **
- 3:00 PM: 👩‍🏫(Lecture - 15 minutes) Introduction to LangChain, LangGraph, Deep Agents

- 3:15 PM: 🧑‍💻(Demo - 15 minutes) Build deep agent w/ RAG, memory, skills, MCP server
- 3:30 PM: 👫(Groups - 30 min) Draw a final, presentable, infrastructure diagram **
- 4:05 PM: 🎤(Front end and infrastructure diagram presentations - 5 min * 8)
- 4:45 PM: 👩‍🏫(Lecture - 15 min) Latest from LangChain, Conclusions, Day 2 Wrap

* = recommend trying different tools (e.g., Claude Code, Codex, Cursor, v0, etc.)
** = [Optional: Run Notebook]

Day 3 Curriculum: Wednesday

Themes: Multi-Agent Systems, SDG, Evals, Advanced Retrieval, Deep Research

🎯Goals

- 👩‍💻Engineering: multi-agent systems, synthetic data generation, evals, advanced retrieval strategies, deep research
- 💼Product: Customer validation & demo readiness
- 👫Group Project: Map evals & vibe checks; identify and email customers

📚Learning Outcomes

- 💼Product: Understand customer validation and build-measure-learn loop
- 👩‍💻Understand multi-agent systems and common design patterns
- 👩‍💻Learn to leverage Synthetic Data Generation (SDG) for your evals
- 👩‍💻Understand best-practice metrics for evaluating retrieval pipelines
- 👩‍💻Understand advanced retrieval techniques and the importance of metadata

💻Activities

- 💼Map where your solution will require systematic evals vs. vibe checks
- 💼Identify users and email them a demo of your app, asking for feedback
- 👩‍💻Build a multi-agent report generator with LangGraph
- 👩‍💻Implement SDG and RAG evals using RAGAS
- 👩‍💻Build a RAG app, use evals and advanced retrieval techniques to improve it
- 👩‍💻Build a multi-agent report generator in LangGraph
- 👩‍💻Unroll OSS Deep Research from LangGraph

📜Recommended Reading

- BM25 (2009) (Topic: Classic Information Retrieval)
- Reciprocal Rank Fusion (2009) (Topic: Classic Information Retrieval)
- Lessons from Improving AI Applications (May 2025) (Topic: Evals)
- In Defense of Evals (Sep 2025) (Topic: Evals)
- Read How we built our multi-agent research system, (June 2025) (Topic: Deep Research)
- Learning the Bitter Lesson (July 2025) (Topic: Deep Research)
- The Customer Development Methodology

Day 3 Detailed Schedule

AM

- 9:00 AM: 👩‍🏫(Lecture - 15 min) Introduction to Multi-Agent Systems
- 9:15 AM: 👩‍💻(Demo - 15 min) Building a Multi-Agent report generator in LangGraph
- 9:30 AM: 👫(Groups - 30 min) Will your solution benefit from multi-agents? **
- 10:00 AM: 👩‍🏫(Lecture - 15 min) Introduction to Synthetic Data Generation (SDG)
- 10:15 AM: 👩‍💻(Demo - 15 min) SDG for RAG Evals
- 10:30 AM: 👫(Groups - 30 min) Do you have the data you need? Will you use SDG?**
- 11:00 AM: 👩‍🏫(Lecture - 20 min) Introduction to Systematic Evals
- 11:20 AM: 👩‍💻(Demo - 15 min) Building systematic RAG eval pipelines
- 11:35 AM: 👫(Groups - 25 min) Map all evals & vibe checks across your solution **
- 12:00 PM: Lunch

PM

- 1:00 PM: 👩‍🏫(Lecture - 15 min) Advanced Retrieval: Hybrid Search, Re-ranking, Query Decomposition
- 1:15 PM: 👩‍💻(Demo - 15 min) Advanced retrieval pipeline notebook
- 1:30 PM: 👫(Groups - 30 min) Does your app need advanced retrieval? Design it. **
- 2:00 PM: 👩‍🏫(Lecture - 15 min) Lessons from Deep Research
- 2:15 PM: 👩‍💻(Demo - 15 min) Building a deep research agent
- 2:30 PM: 👩‍🏫(Lecture - 15 min) Introduction to Customer Validation & Lean Startup
- 2:45 PM: 👫(Groups - 30 min) Revisit harness design / known users / prepare demos**
- 3:15 PM: 🎤(Functional demos required (frontends not required) - 5 min * 8)
- 3:55 PM: 👩‍🏫(Lecture - 5 min) Conclusions & Day 3 Wrap
- 4:00 PM: 🎉Leave for party at AI Aspire

*need to 3:45 stop

** = [Optional: Run Notebook]

Day 4 Curriculum: Thursday

Themes: Prompt Optimization, OSS Models, Fine-Tuning, Inference, Caching, Guardrails

🎯Goals

- 🧑‍💻Engineering: Prompt optimization techniques, OSS model landscape (fine-tuning & inference), guardrails for agent safety
- 💼Product: product-market fit
- 🧑‍🤝‍🧑Group Project: optimize prompts; choose models, design guardrails, full demos

📚Learning Outcomes

- 🤖Edge: What are the best OSS models available today?
- 💼Product: How to define product-market fit for your agent
- 🧑‍💻Understand prompt optimization techniques
- 🧑‍💻Understand the OSS model landscape and how to choose models
- 🧑‍💻Understand when to choose fine-tuning of OSS models in 2026
- 🧑‍💻Understand how to think about guardrails and alignment with users
- 🧑‍💻Learn a few production-minded changes: endpoints, caching, guardrails

💻Activities

- 💼Evaluate your product against PMF frameworks
- 💼Identify which prompts in your solution can be optimized
- 💼Demo your current prototype to the group for feedback
- 🧑‍💻Run prompt optimization with DSPy
- 🧑‍💻Fine-tune a small OSS embedding model for your RAG pipeline
- 🧑‍💻Deploy an OSS model to a production endpoint for inference
- 🧑‍💻Implement a semantic cache for your RAG application
- 🧑‍💻Implement off-the-shelf guardrails

📜Recommended Reading

- The AI Guardrails Index (Feb 2025) (Topic: Guardrails)
- Caching (July 2025) (Topic: Semantic Caching)
- DSPy: Compiling Declarative Language Model Calls (Oct 2023) (Topic: Prompt Optimization)
- The Product-Market Fit Framework (Topic: PMF)

Day 4 Detailed Schedule

AM

- 9:00 AM: 👩‍🏫(Lecture - 15 min) Introduction to Prompt Optimization
- 9:15 AM: 👩‍💻(Demo - 15 min) Systematic prompt optimization with DSPy
- 9:30 AM: 👫(Groups - 25 min) Identify fixed prompts vs. optimizable prompts **
- 10:00 AM: 👩‍🏫(Lecture - 10 min) Introduction to OSS Models; when to use/how to choose
- 10:10 AM: 👩‍🏫(Lecture - 10 min) Introduction to LLM training & fine-tuning
- 10:20 AM: 👩‍💻(Demo - 10 min) Fine-tuning OSS embeddings for improved retrieval
- 10:30 AM: 👩‍🏫(Lecture - 10 min) Introduction to LLM inference & optimization
- 10:40 AM: 👩‍💻(Demo - 10 min) Advanced LLM inference optimization techniques
- 10:50 AM: 👫(Discussion - 30 min) When will you benefit from OSS in production? **
- 11:20 AM: 👩‍🏫(Lecture - 15 min) Intro to production & caching
- 11:35 AM: 👩‍💻(Demo - 10 min) Setting up semantic caches for prompts and embeddings
- 11:45 AM: 👫(Groups - 15 min) What prompts that you would be likely to cache first? **
- 12:00 PM: Lunch

PM

- 1:00 PM: 👩‍🏫(Lecture - 10 min) Introduction to Guardrails
- 1:10 PM: 👩‍💻(Demo - 10 min) Building with off-the-shelf guardrails
- 1:20 PM: 👫(Groups - 25 min) Redraw infra diagram/agent harness with prod in mind**
- 1:45 PM: 👩‍🏫(Lecture - 10 min) Demo Day best-practices / Storytelling
- 1:55 PM: 👫(Groups - 60 min) Prepare full group demos (with slides)
- 2:55 PM: 🎤(Front end and infrastructure diagram presentations - 15 min * 8)
- 4:55 PM: 👩‍🏫(Lecture - 5 min) Conclusions & Day 4 Wrap

** = [Optional: Run Notebook]

Day 5 Curriculum: Friday

Themes: Spec-Driven Development, OpenClaw, SDG for Agent Evals, Demo Day

🎯Goals

- 🧑‍💻Engineering: Spec-Driven Development (SDD) in practice, agentic operations & observability, SDG for agent evals
- 💼Product: Final PMF refinement, Demo Day preparation
- 👫Group Project: Final demos & presentations

📚Learning Outcomes

- 💼Product: Deliver compelling Demo Day presentation with MVP-to-prod path
- 🧑‍💻Understand Spec-Driven Development (SDD) and where it's useful today
- 🧑‍💻Understand OpenClaw and how it differs from previous agent harnesses
- 🧑‍💻Understand how SDG can be used for long-horizon agent eval in 2026
- 🤩Present live for Demo Day!

💻Activities

- 💼Finalize your PMF narrative and product positioning
- 💼Prepare and rehearse Demo Day presentations
- 💼Final demos & presentations
- 🧑‍💻Implement SDD from scratch using our custom harness, including Ralph Loops
- 🧑‍💻Build Local AI with OpenClaw
- 🧑‍💻Use SDG to generate agent evals for your project
- 🧑‍💻Final prototype polish and integration

📜Recommended Reading

- Understanding Spec-Driven Development (Topic: SDD)
- Ralph Wiggum as SWE (Topic: Ralph Loops)

Day 5 Detailed Schedule

AM

- 9:00 AM: 👩‍🏫(Lecture - 10 min) Spec-Driven Development and Ralph Loops
- 9:10 AM: 🧑‍💻(Demo - 10 min) Implementing SDD from scratch
- 9:30 AM: 🧑‍🤝‍🧑(Group Discussion - 10 min) Can you use SDD in engagements today? **
- 9:40 AM: 👩‍🏫(Lecture - 10 min) Introduction to OpenClaw
- 9:50 AM: 🧑‍💻(Demo - 10 min) Building Local AI with OpenClaw
- 10:00 AM: 🧑‍🤝‍🧑(Group Discussion - 10 min) Where to recommend OpenClaw today? **
- 10:10 AM: 👩‍🏫(Lecture - 10 min) SDG for Agent Evals
- 10:20 AM: 🧑‍💻(Demo - 15 min) Using SDG to generate agent evals
- 10:30 AM: 🧑‍🤝‍🧑(Group Discussion - 10 min) Where to recommend OpenClaw today? **
- 10:40 AM: 👩‍🏫(Group Discussion - 20 min) Ironing out the Demo Day Run of Show
- 11:00 AM: 🧑‍🤝‍🧑(Groups - 60 min) Rehearse presentations with peer feedback (x4)
- 12:00 PM: Lunch

PM

- 1:00 PM: 🧑‍🤝‍🧑(Groups - 60 min) Rehearse presentations with peer feedback (x4)
- 2:00 PM: ☕Coffee Break & Final Polish (30 minutes)
- 2:30 PM: 🎤Demo Day Presentations (10 min * 8 — 5 min demo + 5 min Q&A)
- 4:30 PM: 🏆(15 min) Judge deliberation / awards & recognitionair
- 4:45 PM: 👩‍🏫(Lecture - 10 min) Conclusions / looking ahead / what we didn't cover
- 4:55 PM: 🎉(Leadership comments - 5 min) Program Wrap & Celebration

** = [Optional: Run Notebook]