



# AI training catalog for modern enterprise teams

An executive overview of Vizulara's corporate AI training portfolio: practical modern AI adoption, deep AI systems capability, and broad AI fluency for non-technical teams.

**\$2M**

Annual recurring training revenue

**8,000+**

Corporate employees trained globally

**237**

Courses across the learning catalog

**201K**

YouTube subscriber learning community

**Built for enterprises that need more than inspiration.**

Every track is designed around clear outcomes, project work, capstone evidence, and evaluation reporting.

**Remote-first. Cohort-ready. Evaluation-backed.**

Programs can be delivered as leadership workshops, technical cohorts, or multi-bundle AI capability rollouts.

# About Vizudara AI Labs

Vizudara AI Labs is an AI education and product company headquartered in India, operating under Vizudara Technologies Private Limited. The company's mission is to make serious AI education accessible: from non-technical AI fluency to build-from-scratch model engineering, production LLM systems, robotics, vision-language-action models, and research-oriented bootcamps.

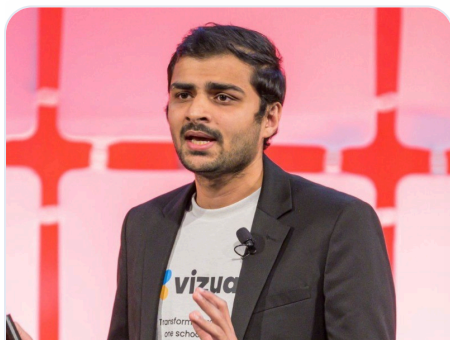
## What Vizudara offers

Vizudara combines live corporate training, pre-recorded courses, research bootcamps, academic minor programs, and AI-native products. For enterprises, the practical offering is simple: train teams, evaluate projects, and build a measurable AI capability pipeline.

## Why enterprises work with Vizudara

The team brings deep technical instruction from MIT, Purdue, and IIT Madras alumni, a large public learning footprint, and experience delivering AI programs for corporate teams across practical and advanced technical topics.

## Founding team



**CO-FOUNDER · MIT PHD  
· IIT MADRAS**

### Dr. Raj Dandekar

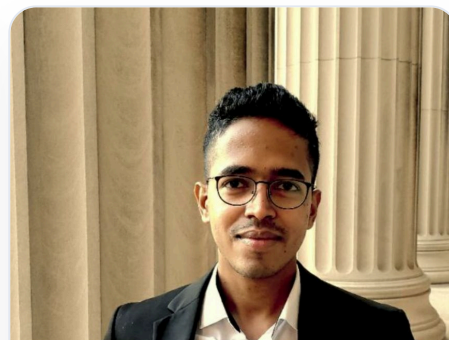
Focus areas include GenAI, LLMs, AI agents, RAG, small language models, scientific machine learning, inference engineering, and GPU systems.



**CO-FOUNDER · PURDUE  
PHD · IIT MADRAS**

### Dr. Rajat Dandekar

Focus areas include reinforcement learning, RLHF, reasoning models, agentic systems, robotics, and advanced AI training programs.



**CO-FOUNDER · MIT PHD  
· IIT MADRAS**

### Dr. Sreedath Panat

Focus areas include computer vision, ML foundations, scientific ML, robotics, VLA systems, modern software development, and AI teaching at scale.

**Founder-authored book:** The founders are co-authors of **Build DeepSeek from Scratch**, published by Manning.

## Overall portfolio

The enterprise catalog in this PDF is one part of a broader Vizulara portfolio.

### **Corporate training**

Remote-first workshops, executive sessions, technical cohorts, capstones, and evaluation reports.

### **Course library**

237 courses and 9,878 lessons across foundations, GenAI, ML, CV, robotics, LLMs, and production systems.

### **Research bootcamps**

Structured research programs across SciML, GenAI, CV, RL, ML/DL, GPU systems, and high school research.

### **AI products**

Vizz-AI for video tutoring, VizPods for text and notebook learning, ArcEval for AI collaboration evaluation, and DynaRoute for LLM routing.

# Table of contents

This document is structured for fast executive review first, then detailed curriculum inspection.

01

## About Vizudara AI Labs

Who Vizudara is, founder background, corporate training history, and broader product portfolio.

02

## Three bundles at a glance

Modern AI, Advanced Systems, and AI Fluency, with audience, pricing, and training intent.

03

## Modern AI Bundle

Agents, Claude Code software development, context engineering, RAG, and fine-tuning.

04

## Advanced Systems Bundle

Inference engineering, parallelism, robotics, world models, LLM internals, and VLA systems.

05

## AI Fluency Bundle

Shared AI language, use-case discovery, risk awareness, and leadership alignment.

06

## Evaluation with ArcEval

Why training needs evaluation, what ArcEval captures, and how cohort reporting works.

07

## Delivery and pricing notes

How Vizudara can package programs for enterprise rollout.

# Three bundles at a glance

The catalog is intentionally simple: one practical adoption bundle, one deep systems bundle, and one broad fluency bundle.

## 01 · BUNDLE 01

### Modern AI Bundle

Agents, Claude Code software development, context engineering, RAG, fine-tuning, and applied AI production workflows.

**Audience:** Software teams, product teams, analytics teams, managers, and transformation leaders who need to use modern AI tools in real workflows.

**\$250-\$300 per learner**

## 02 · BUNDLE 02

### Advanced Systems Bundle

Inference engineering, parallelism, robotics, world models, build-from-scratch LLMs, and VLA transformers.

**Audience:** Backend, ML, platform, AI infrastructure, robotics, vision, edge AI, and specialist R&D teams.

**\$250-\$500 per learner**

## 03 · BUNDLE 03

### AI Fluency Bundle

Executive AI fluency, AI vocabulary, use-case discovery, risk awareness, and sponsor-facing rollout reporting.

**Audience:** Non-technical staff, managers, analysts, operations teams, HR, support, compliance-adjacent teams, and senior leaders.

**\$250 per learner**

## Recommended rollout pattern

Start with the Modern AI Bundle, which is where most enterprises begin: agents, Claude Code software development, context engineering, RAG, and fine-tuning. Then customize with AI Fluency for broad non-technical alignment or Advanced Systems for specialist engineering teams.

## What makes Vizuara different

The training is project-backed and evaluation-backed. The outcome is not just attendance; it is evidence of what the cohort can understand, build, and explain.

# Modern AI Bundle

Turn AI curiosity into practical operating capability.

## WHO IT IS FOR

Software teams, product teams, analytics teams, managers, and transformation leaders who need to use modern AI tools in real workflows.

## WHY THIS BUNDLE EXISTS

This is the recommended starting point for most enterprises. It gives teams a shared understanding of agents, context engineering, AI-assisted software development, retrieval, fine-tuning, and evaluation. The goal is not to create isolated demos; the goal is to help teams build useful systems and measure whether they work.

## EXPECTED OUTCOME

Teams leave with working prototypes, better AI development discipline, and a clear mental model for where agents, RAG, context design, and Claude Code-style workflows fit inside enterprise operations.

## TRACKS INCLUDED

Context Engineering

Claude Code Software Development

Agents

RAG + Fine-Tuning

PRACTICAL, ROLE-AWARE · MIXED

\$300/person

## Context Engineering

A hands-on workshop for building production-grade AI agents through retrieval, tool orchestration, memory, context compression, and observability.

### DURATION

10 sessions  
over 2 weeks

### TIME

14 live  
hours

### SIGNAL

Renasas and Ramco  
training experience

### AUDIENCE

Engineers, analytics teams, product teams, and managers who need a practical vocabulary for AI systems.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Foundations of Context
- The instructional layer and system prompts

## WHAT PARTICIPANTS LEARN

## ENTERPRISE RELEVANCE

- Architect reliable agents that combine instructions, retrieval, tools, and memory.
- Understand when to use RAG, fine-tuning, prompt design, or structured tool calls.
- Ship a capstone agent with tracing, guardrails, and measurable quality.

Corporate training experience includes Renesas and Ramco; course participants include engineers and leaders from top AI and software companies.

Claude

OpenAI

Google Gemini

Ramco

Renesas

## CURRICULUM MAP

1. Context foundations: from prompts to context engineering
2. System prompts, instruction hierarchy, and dynamic prompt assembly
3. RAG architecture, chunking, embeddings, ChromaDB, and hybrid search
4. Advanced RAG: re-ranking, contextual compression, multi-index retrieval, and RAG evaluation
5. Function calling, MCP servers, tool orchestration, chaining, and fallbacks
6. Memory architectures: write, select, compress, isolate
7. Token optimization, observability, tracing, logging, cost control, and caching
8. Production capstone: RAG plus tools plus memory plus deployment patterns
9. Mega build session I with Dr. Raj Dandekar
10. Mega build session II with Dr. Rajat Dandekar

## CAPSTONE

Production-grade context-engineered agent for an enterprise workflow.

## EVALUATION

ArcEval-style project review for architecture quality, tool-use discipline, retrieval quality, prompt clarity, and iteration logs.

HANDS-ON ENGINEERING · TECHNICAL

\$300/person

# Claude Code Software Development

A five-day practical bootcamp on AI-assisted software development, coding agents, MCP, secure vibe coding, testing, deployment, and DevOps.

### DURATION

5 days

### TIME

10 live hours

### SIGNAL

Ramco corporate cohort signal

### AUDIENCE

Engineering teams, tech leads, QA teams, platform teams, and product engineers.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Modern software development with Claude Code

## WHAT PARTICIPANTS LEARN

- Move from prompt-based coding to a disciplined plan-generate-modify-verify workflow.

## ENTERPRISE RELEVANCE

A Ramco corporate cohort has used this track; it is designed for software teams adopting Claude Code and

- Use Claude Code, Cursor, Copilot, MCP servers, and agentic terminals in real projects.
- Create test, security, review, and deployment workflows around AI-generated code.

agentic development.

Ramco

Claude Code

Cursor

GitHub Copilot

Vercel

## CURRICULUM MAP

1. Day 1: LLM intuition, prompting, developer workflow, and AI development environment
2. Project: CLI app from scratch
3. Project: REST API from a napkin spec
4. Day 2: coding agent architecture, tool use, MCP, Warp, and autonomy levels
5. Project: build a coding agent from scratch

6. Project: custom MCP server connected to GitHub, Notion, or a database
7. Day 3: PRD-driven development, multi-file codebases, UI generation, v0, Bolt, Lovable
8. Project: full-stack SaaS app and Chrome extension
9. Day 4: secure vibe coding, OWASP risks, Semgrep, Snyk, tests, reviews, and documentation
10. Day 5: CI/CD, Docker, Terraform, monitoring, incident response, and final deploy

## CAPSTONE

Ship a secure full-stack application with auth, database, tests, CI/CD, and live URL.

## EVALUATION

ArcEval is the strongest fit here: it can inspect prompt history, code edits, terminal commands, tests, security fixes, and final implementation quality.

APPLIED · MIXED

\$250/person

# Agents

A practical agent track for teams that want LLM systems to take actions, call tools, reason over state, and automate repeatable workflows.

### DURATION

10-day agent bootcamp or 8-week voice-agent track

### TIME

10-16 estimated hours

### SIGNAL

Enterprise automation rollout track

### AUDIENCE

Business teams, operations teams, software teams, and AI transformation groups.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Build Mini ClawdBot from Scratch
- Build OpenClaw-RL and VoiceAgents using Claude Code

## WHAT PARTICIPANTS LEARN

## ENTERPRISE RELEVANCE

- Understand agent patterns without drowning non-AI teams in research terminology.
- Build multi-step agents with tools, memory, and fallback behavior.
- Design automation opportunities for operations, analytics, support, and software teams.

Learner network includes professionals from Microsoft, Google, NVIDIA, AWS, Anthropic, and OpenAI ecosystems.

Claude

OpenAI

Google

AWS

Enterprise operations

## CURRICULUM MAP

1. Agent fundamentals and the difference between chatbots, workflows, and autonomous agents
2. Tool calling and function-calling patterns
3. Multi-agent workflows and orchestration
4. Agent memory and state management

5. Voice agents: ASR to LLM to TTS pipeline
6. Streaming, latency, and production voice-agent architecture
7. Enterprise agent safety: approvals, permissions, and audit trails
8. Capstone: internal workflow agent

## CAPSTONE

Workflow automation agent for a corporate process such as ticket triage, document routing, or analytics Q&A.

## EVALUATION

Custom agent evaluator: task completion, tool-call correctness, retrieval grounding, escalation behavior, and auditability.

## PRODUCTION BUILDER · TECHNICAL

\$250/person

# RAG + Fine-Tuning

A production-oriented track covering document ingestion, retrieval architecture, embedding quality, instruction fine-tuning, and evaluation.

### DURATION

Modular 2-8 week track

### TIME

8-24 estimated hours

### SIGNAL

Used by working data and engineering teams

### AUDIENCE

Teams building document Q&A, research assistants, support copilots, policy search, and internal knowledge systems.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Production level RAG Workshop: Part 1
- Introduction to LLM Instruction Fine-tuning

## WHAT PARTICIPANTS LEARN

- Build retrieval systems that work with messy enterprise documents.
- Choose between RAG, fine-tuning, and hybrid approaches for real business constraints.

## ENTERPRISE RELEVANCE

Built from Vizulara's production RAG and fine-tuning course library used by working engineers and data teams.

Data teams

Search teams

Knowledge ops

- Measure faithfulness, relevance, coverage, and latency before rollout.

#### Compliance workflows

## CURRICULUM MAP

1. Document ingestion with PyMuPDF and preprocessing pipelines
2. Chunking, embeddings, vector databases, and metadata design
3. Hybrid retrieval, query expansion, and re-ranking
4. RAG evaluation: faithfulness, relevance, coverage, and error analysis
5. Instruction fine-tuning: behavioral adaptation rather than information dumping
6. Fine-tuning data design, prompt-response pairs, and validation splits
7. Deployment patterns for enterprise knowledge assistants
8. Capstone: production RAG assistant with eval harness

## CAPSTONE

Enterprise knowledge assistant with retrieval metrics, failure cases, and improvement plan.

## EVALUATION

Custom RAG and fine-tuning evaluator: answer faithfulness, source precision, retrieval recall, regression tests, and latency profile.

# Advanced Systems Bundle

Build deep technical capability for teams working near AI infrastructure and frontier systems.

## WHO IT IS FOR

Backend, ML, platform, AI infrastructure, robotics, vision, edge AI, and specialist R&D teams.

## WHY THIS BUNDLE EXISTS

Some teams need to go deeper than application-layer AI. This bundle is for engineers who must understand latency, memory, serving, distributed training, GPU communication, transformers, VLA systems, robotics policies, and world models from first principles.

## EXPECTED OUTCOME

Teams leave with the vocabulary and implementation intuition needed to reason about production LLM systems, inference costs, parallelism choices, model internals, and embodied AI systems.

## TRACKS INCLUDED

Inference Engineering

Parallelism

Robotics and World Models

Build LLM from Scratch

VLA Transformers

## ADVANCED PRACTICAL SYSTEMS · TECHNICAL

\$500/person

# Inference Engineering

A systems course on the full LLM inference stack: tokenization, KV cache, quantization, speculative decoding, FlashAttention, serving engines, edge deployment, and evals.

### DURATION

4 weeks

### TIME

14 lectures,  
2-3 hours  
per day

### SIGNAL

Guest bench from  
frontier AI and infra  
teams

### AUDIENCE

Platform, backend, ML engineering, AI infrastructure, and architecture teams.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- How the vLLM inference engine works
- Multi-Query Attention and KV cache memory

## WHAT PARTICIPANTS LEARN

- Understand what makes LLM applications slow, expensive, or memory-hungry.

## ENTERPRISE RELEVANCE

Guest speaker bench includes Anthropic, NVIDIA, Apple, Microsoft, AnyScale, Red Hat, Amazon, and Mastercard.

- Choose the right optimization across KV cache, quantization, batching, kernels, and serving engines.
- Build mental models for secure, low-latency enterprise deployment.

Anthropic

NVIDIA

Apple

Microsoft

AnyScale

Red Hat

Amazon

Mastercard

## CURRICULUM MAP

1. The inference stack: tokenization, forward pass, autoregressive loop
2. Transformer deep dive: MHA, MQA, GQA, embeddings, positional encoding
3. Prefill, decode, KV cache, chunked prefill, prefix caching, prompt caching
4. GPU architecture, tensor cores, roofline, memory-bound vs compute-bound work
5. Quantization: FP16, BF16, INT8, INT4, GPTQ, AWQ, GGUF
6. Speculative decoding: draft-target, n-gram, EAGLE, Medusa
7. FlashAttention, kernel fusion, PagedAttention, batching, vLLM, SGLang, TensorRT-LLM
8. MoE, model parallelism, edge deployment, voice, multimodal inference, production systems, structured outputs, evals, fine-tuning, and distillation

## CAPSTONE

Latency and cost optimization project for an LLM serving pipeline.

## EVALUATION

Custom inference evaluator: benchmark reproducibility, latency/cost profile, memory analysis, throughput, and production design review.

## DEEP SYSTEMS · TECHNICAL

**\$300/person**

# Parallelism

A deep systems course covering data, tensor, pipeline, sequence/context, and expert parallelism for training and serving frontier AI models.

### DURATION

12 lectures  
plus optional  
research  
phase

### TIME

30+  
content  
hours

### SIGNAL

Indexed 12-lecture  
specialist course wiki

### AUDIENCE

AI infrastructure teams, systems engineers, GPU engineers, and high-performing ML engineers.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Matrix multiplication on GPUs
- How the vLLM inference engine works

## WHAT PARTICIPANTS LEARN

- Build a first-principles understanding of distributed LLM training and serving.

## ENTERPRISE RELEVANCE

This track has a complete indexed lecture wiki and is designed for teams working near GPU systems, distributed

- Interpret GPU memory, communication, NCCL, all-reduce, and profiling behavior.
- Prepare advanced engineers for GPU systems research and production architecture.

training, and inference infrastructure.

NVIDIA

OpenAI Triton

GPU systems

MLSys

## CURRICULUM MAP

1. GPU architecture and CUDA fundamentals
2. Triton and custom GPU kernels
3. NCCL and distributed communication
4. Data parallelism, DDP, FSDP, and ZeRO
5. Tensor parallelism
6. Pipeline parallelism

7. Context and sequence parallelism
8. Expert parallelism and MoE
9. 5D parallelism integration
10. Diffusion language models
11. Speculative decoding and inference
12. Research methods and paper writing

## CAPSTONE

Distributed training or inference research/project artifact with profiling and design documentation.

## EVALUATION

Custom GPU evaluator: code correctness, profiler interpretation, memory accounting, communication analysis, and research proposal review.

## SPECIALIST · TECHNICAL

\$300/person

# Robotics and World Models

A specialist track covering robot learning, Diffusion Policy, VLA models, world models, SO-101 robot-arm deployment, and autonomous systems.

### DURATION

8+ lectures or  
6-month  
robotics minor

### TIME

16-60  
estimated  
hours

### SIGNAL

Modern robotics minor  
and VLA track signal

### AUDIENCE

R&D teams, innovation labs, edge AI groups,  
robotics groups, and advanced technical  
cohorts.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Robot Imitation Learning
- The Transformer Architecture for Robotics

## WHAT PARTICIPANTS LEARN

- Understand how modern robotics is moving from control stacks to learned policies.
- Build VLA and world-model intuition with notebooks and visualizations.

## ENTERPRISE RELEVANCE

Robotics and VLA tracks are connected to Vizura's modern robotics minor and advanced robot learning cohorts.

Renasas

Robotics

Autonomous systems

Edge AI

- Deploy model components on real robot hardware when the corporate track includes labs.

## CURRICULUM MAP

1. Diffusion Policy for visuo-motor policy learning
2. ResNet18, Spatial Softmax, 1D U-Net, FiLM, and DDPM components
3. Vision-Language-Action models: SigLIP, ViT, transformer, VLM, flow matching, pi0
4. SmolVLA, OpenVLA, SO-101 robot deployment, and real-robot data handling
5. World Models: IRIS, DIAMOND, DINO-WM, DreamZero, and LeWorld
6. Action chunking with transformers
7. Hugging Face model management
8. Demo day and applied robotics capstone

## CAPSTONE

Robotics policy or VLA deployment project with evaluation artifacts.

## EVALUATION

Custom robotics evaluator: simulation metrics, policy rollout analysis, hardware-readiness checklist, and capstone demo scoring.

## DEEP FOUNDATIONS · TECHNICAL

\$250/person

# Build LLM from Scratch

A first-principles course for engineers who need to understand tokenization, embeddings, self-attention, transformers, training, and fine-tuning.

### DURATION

5-day guided cohort or 3-4 week self-paced track

### TIME

20-30 estimated hours

### SIGNAL

Long-running LLM foundations library

### AUDIENCE

Engineers who will own LLM systems, architecture reviews, model selection, and technical strategy.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Large Language Models basics
- The GPT tokenizer: Byte Pair Encoding

## WHAT PARTICIPANTS LEARN

- Demystify how LLMs work internally rather than treating them as magic APIs.
- Build transformer components from scratch in code.
- Give engineering teams the mental model needed for production architecture decisions.

## ENTERPRISE RELEVANCE

Vizuara's build-from-scratch library has reached learners from Microsoft, Google, NVIDIA, OpenAI, Anthropic, AWS, and similar frontier teams.

Microsoft

Google

NVIDIA

OpenAI

Anthropic

## CURRICULUM MAP

1. Token prediction and probability distributions
2. Tokenization and byte pair encoding
3. Token embeddings and positional encoding
4. Self-attention and multi-head attention from scratch

5. Transformer feed-forward networks and residual architecture
6. Loss functions and optimization
7. Pre-training and scaling intuition
8. Instruction fine-tuning and classification fine-tuning
9. Production-level language model concepts

## CAPSTONE

Small GPT-style model built from scratch and fine-tuned on a task.

## EVALUATION

Custom notebook evaluator for implementation correctness, conceptual quizzes, training curves, and model behavior analysis.

BEGINNER TO INTERMEDIATE · TECHNICAL

\$300/person

# VLA Transformers

A specialist track on transformer-based vision-language-action systems: perception, language grounding, multimodal fusion, and action/control pipelines.

### DURATION

8-week applied systems track

### TIME

16-24 estimated hours

### SIGNAL

Advanced vision and robotics portfolio signal

### AUDIENCE

Vision AI, robotics, autonomous systems, edge AI, and applied transformer teams.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Introduction to ACT for autonomous driving
- Connect, collect data, train, and infer on a robot car

## WHAT PARTICIPANTS LEARN

- Understand how vision transformers, VLMs, and VLA systems connect.
- Build multimodal systems that map perception and language into action.
- Connect transformer-based vision systems to robotics, autonomous driving, and edge AI.

## CURRICULUM MAP

1. Foundations of autonomous systems and the sense-plan-act paradigm

## ENTERPRISE RELEVANCE

VLA and robotics courses are part of Vizuar's advanced AI systems portfolio, with strong interest from embedded AI and robotics audiences.

Renesas

Vision AI

Edge AI

Autonomous systems

2. Vision for driving: cameras, image representation, lane detection, and obstacle detection

3. Transformer-based language understanding for driving instructions
4. Vision-language fusion: early fusion, late fusion, and attention-based fusion
5. Action spaces, control systems, safety constraints, and fail-safe mechanisms
6. Imitation learning and image-language-action datasets
7. Real-world deployment constraints: latency, noise, edge optimization, and sim-to-real gaps
8. Final VLA transformer project: end-to-end perception-language-action system

## CAPSTONE

---

End-to-end VLA transformer system with perception, language understanding, control policy, and evaluation report.

## EVALUATION

---

Custom VLA evaluator: model architecture review, perception/action metrics, safety checks, latency profile, and final demo scoring.

# AI Fluency Bundle

Create a shared AI language across business, operations, and leadership teams.

## WHO IT IS FOR

Non-technical staff, managers, analysts, operations teams, HR, support, compliance-adjacent teams, and senior leaders.

## WHY THIS BUNDLE EXISTS

Enterprise AI adoption fails when only a small technical group understands the vocabulary. This bundle gives broader teams a realistic understanding of what AI can do, what it cannot do, where risks sit, and how to identify good use cases.

## EXPECTED OUTCOME

Teams leave with a common AI vocabulary, an actionable use-case map, and a grounded way to discuss value, risk, human approval, privacy, and evaluation.

## TRACKS INCLUDED

AI Fluency

EXECUTIVE AND OPERATIONS FRIENDLY · NON-TECHNICAL

\$250/person

## AI Fluency

A broad modern AI orientation for non-technical and semi-technical teams: what AI can do, what it cannot do, where risk sits, and how to identify good use cases.

### DURATION

3-5 day applied fluency sprint

### TIME

6-10 guided hours plus self-paced practice

### SIGNAL

Broad corporate fluency training experience

### AUDIENCE

Non-technical staff, managers, support teams, analysts, HR, operations, and compliance-adjacent teams.

**Delivery model:** Blended format: self-paced prep and practice, one live founder-led session, and graded assignments/projects with sponsor-ready cohort reporting.

## TEACHING PREVIEWS IN THE WEB CATALOG

- Generative AI for Absolute Beginners: Introduction
- Generative AI vs Conventional AI

## WHAT PARTICIPANTS LEARN

- Create a shared AI vocabulary across business, compliance, and technical teams.
- Help teams identify realistic automation and knowledge-work opportunities.

## ENTERPRISE RELEVANCE

Built from First Principles Labs' corporate training programs and Vizura's experience making AI accessible to non-AI audiences.

- Reduce hype-driven decisions by teaching model limitations, risk, and eval thinking.

Leadership

Operations

Finance

Healthcare

Compliance

## CURRICULUM MAP

1. Predictive AI, generative AI, and agentic AI in plain language
2. Journey of a token through an LLM
3. What RAG, fine-tuning, agents, and copilots actually mean

## CAPSTONE

Department-level AI use-case map with risk, effort, and value scoring.

4. AI opportunities in financial services, healthcare operations, analytics, and support
5. Security, privacy, auditability, and human approval loops
6. Prompt and context patterns for knowledge workers
7. Use-case discovery workshop
8. Executive synthesis and rollout roadmap

## EVALUATION

Custom fluency evaluation: pre/post assessment, scenario-based prompts, use-case rubric, and manager-facing cohort report.

# Evaluation with ArcEval

The hardest question in corporate upskilling is not “Did people attend?” It is “Can they now perform better work?” Vizuara built ArcEval and custom course evaluators to answer that question with evidence.

## Why ArcEval exists

Most enterprise training programs measure the easiest signals: attendance, quiz completion, certificates, and survey satisfaction. Those signals are useful, but they do not prove that an engineer can use AI responsibly, debug with it, ship a working system, or explain trade-offs.

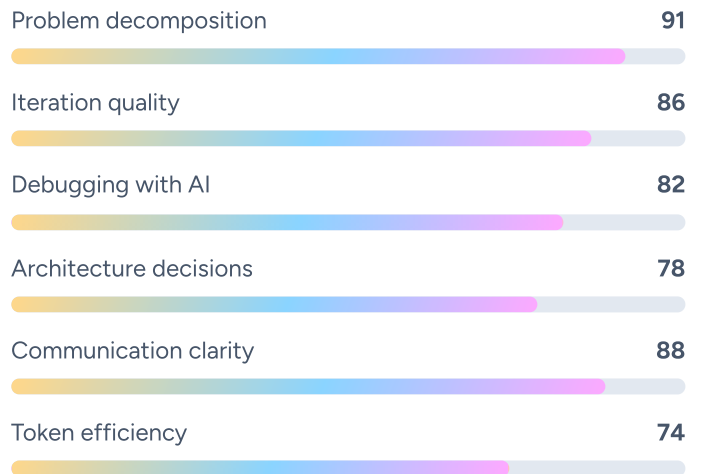
- Attendance does not measure applied skill.
- Multiple-choice quizzes miss real project behavior.
- Managers need cohort-level evidence, not just certificates.
- AI-assisted work needs process evaluation: prompts, edits, tests, decisions, and debugging loops.

### ARCEVAL REPORT PREVIEW

84 / 100

## Capstone: Build a production AI workflow

Evidence-based scoring across AI collaboration behavior.



Prompt history

Code edits

Terminal commands

Test runs

Debug loops

Final project output

1

Learner builds a project

2

Environment captures process

3

Rubric scores evidence

4

Sponsor receives report

### ArcEval tracks

Best suited for Modern Software Development and Context Engineering, where prompt history, code edits, terminal commands,

### Custom AI system evaluators

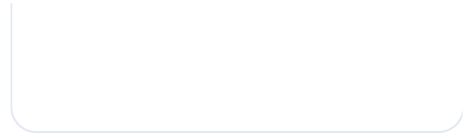
RAG, agents, inference, robotics, VLA, and parallelism tracks use specialized rubrics for grounding, latency, memory, task completion,

### Enterprise reporting

For sponsors, Vizuara can provide pre/post assessments, capstone pass rates, quality bands, learner evidence, team-level insights, and recommended next steps.

tests, and final implementation  
quality matter.

profiling, and deployment  
readiness.



## How this can be packaged for an enterprise rollout

Vizuara can structure the catalog as a single focused workshop, a multi-week technical cohort, or a broader AI transformation program across multiple audiences. Pricing in this catalog is shown per learner for each track, with capstones, evaluation, and blended delivery built into the training model.

**Suggested rollout:** Begin with the Modern AI Bundle as the core enterprise track. Add AI Fluency when non-technical teams and leaders need a shared AI language, and add Advanced Systems when specialist teams need deeper model, serving, infrastructure, robotics, or VLA capability.