

Jacob Valdez

AI systems & full-stack engineer

Agent infrastructure, full-stack AI products, and the runtime glue that turns research ideas into shipped software.

jacob@humanrobots.ai +1 (469) 968-9490 jvboid.dev github.com/JacobFV [@jvboid](https://twitter.com/jvboid) San Francisco, CA

HIGHLIGHTS

- Architected production integration surfaces for mobile and web clients at AGI, Inc.
- Building VibeStartup, an end-to-end platform for spinning up startups (planning, code, infra, growth) on top of agent workflows.
- Founded Limboid — a long-horizon robotics company; ship the software side end to end (control, sim, agent stack, dashboards).
- Authored a wide technical archive (ML systems, agents, multi-agent networks, consciousness, dev platforms) backing each project.

SKILLS

AI agents, agent runtime, tool use, schemas, evals, TypeScript, Next.js, React, Python, FastAPI, iOS, on-device LLMs, model quantization, multimodal pipelines, Postgres, Redis, Vercel, Cloudflare, Modal, research synthesis, technical writing, rapid prototyping

EXPERIENCE

- Jan 2026 – Apr 2026 **API / Integration Architect** · AGI, Inc.
Owned the integration surface across API, SDKs, and partner-facing edges — iOS, on-device LLMs, model quantization, agent control plane. Led schema + integration work through fast product pivots.
[agents](#) · [ios](#) · [on-device](#) · [schemas](#)
- Oct 2025 – Jan 2026 **Software Engineer** · AGI, Inc.
Infrastructure for AGI — agent runtime, mobile glue, eval tooling. Promoted into the integration architect role.
[agents](#) · [infra](#) · [mobile](#)
- Apr 2025 – Oct 2025 **Software Engineer** · Breezy
Voice-AI automation product — agent workflows, conversational interfaces, integration plumbing.
[voice-ai](#) · [agents](#)
- Sep 2024 – Dec 2024 **Applied ML Engineer (Intern)** · Deepshard
The Truffle computer — large-model experimentation, multi-agent research, on-device inference tooling.
[ml](#) · [on-device](#) · [research](#)
- Jun 2024 **Full-Stack Pipeline Engineer** · FLORA
Generative-AI creative pipeline — TypeScript, Python, Modal, fal.ai, ComfyUI, diffusion models, serverless Next.js front end.
[generative-ai](#) · [diffusion](#) · [full-stack](#)
- Jun 2022 – Jan 2023 **Software Engineer** · Motio
Built Soterre for Qlik Sense — full-stack engineering across product surfaces in Java/Hibernate. Started as an intern; converted to full-time in August.
[full-stack](#) · [java](#) · [product](#)
- 2020 – 2022 **Research Assistant** · UT Arlington research labs
Multi-agent learning, differentiable programming, world models, robotics — early prototypes that became the lineage in the project archive.
[research](#) · [ml](#) · [agents](#)

2020 – 2022

B.S., Computer Science · The University of Texas at Arlington

CS coursework alongside heavy lab work, independent ML/robotics prototypes, and a steady research output. GPA 3.6/4.0.

education · cs

2016 – 2018

A.A.S., Mathematics · Navarro College

Math associate degree taken dual-credit during high school. GPA 3.9/4.0.

education · math

PROJECTS

SOFTWARE & AI — FOCUS

- 2026 **SynthUX** — A low-compute synthetic visual computer-use dataset generator: goals expand into concrete Node Tree-style grammar trees, terminal affordances are replayed as low-level mouse and keyboard input in browser desktop simulators, and the resulting trajectories are recorded as frames, observations, videos, and alignment metadata.
- 2026 **Bonk** — A tiny AI life coach: a friendly desktop doge watches your active window, tracks goals, and gently bonks you back toward the work you meant to do.
- 2026 **yt2ctx** — A YouTube-to-context compiler that turns reference videos into transcripts, representative frames, style bibles, shot specs, agent prompts, and ZIP artifacts.
- 2026 **notion-vibestartup** — A Notion Developer Platform Hackathon project: OAuth into a founder workspace, generate the company operating system, publish product routes, and write real runtime events back into Notion.
- 2026 **macos-web-next** — macOS desktop simulation in the browser — companion to windows-web-next for computer-use agent training.
- 2026 **Chem-0** — A prototype autonomous chemistry lab exploring how AI agents can operate a low-cost robot arm, camera loop, lab workspace, and evented measurement stack.
- 2026 **Trash Sorter** — A robotic trash-sorting prototype built as an initial hands-on learning experience with a LeRobot arm at the Founders Inc Physical AI Hack World Tour SF hackathon.
- 2026 **personal site** — This site. Next.js + MDX, every artifact is a node in a graph, the navigation is the relationship.
- 2026 **Lunar Rover** — A lunar rover prototype built at RoboHacks at Y Combinator to collect regolith in a moon-surface scenario, with a four-node autonomous control architecture spanning remote viewing, ground control, satellite relay, and robot control.
- 2026 **brain-model / Resonance** — A cortical dynamics and affect-geometry project that became Resonance: algorithmic music targeted at a six-dimensional model of joy, flow, and self-model quieting.
- 2026 **belief-graph-orchestrator** — Event-sourced multi-scale cognitive control system for autonomous GUI agents.
- 2026 **Recursive Omnimodal Video Action Model** — ROAM studies looped attention and spatiotemporal canvas layouts for video-action diffusion models, grounding a cortical-canvas hypothesis in 26 experiments across toy models, MuJoCo behavior cloning, and a CogVideoX-2B robot-video graft.
- 2026 **precisionBOM** — AI-powered BOM optimization with live DigiKey inventory, price breaks, alternates, and export paths that move a hardware team from spreadsheet to sourcing decision.
- 2026 **The Tensor Computer** — A complete von Neumann computer — ALU, registers, cache, virtual memory, bus, GPU, peripherals — rebuilt as differentiable tensor operations in JAX, so that programs written in machine code can, in principle, be learned end-to-end by gradient descent.
- 2026 **canvas-engineering** — A type system for multimodal latent dynamics that turns video diffusion models into structured world-modeling substrates.
- 2025 **lifelogger** — Personal activity tracking daemon with LLM summarization.
- 2025 **browser-os** — Browser-native desktop shell — generic OS chrome that windows-web-next and macos-web-next inherit from.
- 2025 **imgpt** — A Python CLI for generating OpenAI images from one-off prompts or whole prompt directories, with persisted defaults, model controls, and cost-aware batch behavior.
- 2022 **TensorCode** — A framework for encoding and decoding arbitrary Python objects, differentiable programming, and runtime code generation — Programming 2.0 with the abstractions software engineers already know.
- 2022 **MPNets** — Multi-Paradigm Networks: a graph-based neural architecture experiment for training one policy across supervised, self-supervised, unsupervised, reinforcement, Hebbian, feedback, and structural learning signals.
- 2022 **Node Tree** — A Python toolkit for programming tree-shaped generative processes that can render natural-language descriptions and execute matching demonstrations in an environment.
- 2021 **The Multi-Agent Network (MAN)** — A modular framework for fusing pretrained foundation models with self-organising learning agents on a shared, typed workspace.
- 2021 **tf-som** — TensorFlow layers for locally competitive, self-organizing-map-style feature learning.
- 2021 **Predictive General Intelligence** — Theoretical framework an plan for developing artificial general intelligence

- 2026 **ascii-art** — Multi-layer ASCII art generator with dithering, edge detection, and compositing — browser-native.
- 2026 **standup-ai** — AI scrum-lead assistant — automates standup task management across Notion and Linear.
- 2026 **windows-web-next** — A browser-based Windows 11 simulation built for human annotation and computer-use agent research.
- 2026 **general-unified-world-modeling** — A typed causal ontology of civilization that projects economics, politics, firms, people, and events onto one structured latent world model.
- 2020 **multi-graph-former** — A TensorFlow/Keras experiment showing that transformer attention can be read as graph neural network message passing, then generalized from seq2seq to graph2graph.
- 2019 **Computatrum** — It's time for AI grow up
- 2017 **Polonius as a Fool** — Animation of two scenes from Hamlet for 11th grade English class in Spring 2017.
- 2015 **Space Pong** — A two-player Pong-style game built in Blender during Summer 2015, with a space-themed arena and keyboard controls for playing against a friend or yourself.

ADJACENT WORK

- 2026 **ESP32-S3 USB Webcam** — Firmware that makes a Seeed XIAO ESP32S3 Sense enumerate as a real USB Video Class webcam, with Linux seeing it as a normal /dev/video camera.
- 2025 **eggroll-trainer** — Python implementation of EGGROLL (Evolution Guided General Optimization via Low-rank Learning) for scalable evolution strategies training
- 2025 **bsbr** — novel attention mechanism for efficient processing of long sequences in transformer architectures
- 2024 **Tiles** — Free and open source tiles matching game!
- 2023 **Jacob's Hits 2023 (album)** — Some of my favorite personal music creations
- 2022 **AI Proverbs** — The treasure trove of AI ponderances. Pictures and text by AI.
- 2021 **MLN-Dashboard** — Modular server for data visualization
- 2021 ~~The Fertile Cresent~~ **The Fertile Cresent** — Ecosystem of projects, repositories, and libraries to promote the advancement of artificial intelligence
- 2021 **Broadening and Building Beyond Classical Reinforcement Learning** — Affective psychology + multi-agent reinforcement learning
- 2021 **jumpy** — Another autograd / deep learning / RL library
- 2021 **Summer Break 2021! (album)** — A collection of various songs capturing various emotions I experienced during the 2021 summer break
- 2021 **Dash** — Covid19 Analysis Visualization Dashboard
- 2021 **20Q** — 20 Questions Android App
- 2021 **Stanford Open Datathon Group Project** — Data scientist in a 5-student international group
- 2021 **DesparadosAEye** — Project manager, general app developer, and ML engineer for a 5-person group developing open-ended Android chatbot
- 2020 **Home Internet Factory** — CNC router + 3D printer + Raspberry Pi 4 + camera + Internet
- 2019 **Workplace Surveillance System** — When you need to catch employees on their phone, this will do the trick!
- 2019 **Full Stack Artificial Intelligence** — Computer interaction + deep reinforcement learning
- 2019 **Limboid** — Hydraulic musculoskeletal humanoid robot
- 2016 **SQTEST** — Flexible Spanish – English translation system
- 2016 **Labatron** — Automated stoichiometry solver for CHEM 1411/12
- 2015 **Sale** — Basic Spanish – English translation system
- 2015 **Cookie Cutter CNC** — Full stack engineering project
- 2014 **Dolphin Rocket** — Bottle-rocket project and mini documentary from 9th grade: a low-budget MythBusters x October Sky crossover inspired by Apollo and early spaceflight history.
- 2014 **Copyright Calculator** — Scripted conversation bot to help students stay safe
- 2014 **The Right Night Light** — Blender animation about a monkey searching for the right night light so he can cross the ground, completed in May 2014.
- 2013 **Looking for Princess Suzzane** — Blender animation made to teach conjugation of -AR verbs in Spanish, published on December 3, 2013.
- 2013 **Cookie Baker 3D Printer** — Programmable 12x12 2D gantry