

HENRIQUE GODOY

São Paulo, SP

2dgodoy01227@gmail.com [LinkedIn](#) [GitHub](#) [Website](#)

ABOUT

I build AI at CloudWalk, a \$1.2B ARR fintech, across JIM (US) and InfinitePay (Brazil), reporting directly to the [CEO and founder](#) and serving millions of daily interactions across 7M+ merchants. I previously led the first production deployment of a large language model inside one of the largest investment banks in Latin America, and I founded Doki, a personal-finance product for doctors in Brazil that recorded over R\$12 million on the platform. I have published two papers: [Extract-0](#) and [Alvorada-Bench](#).

EDUCATION

Inteli - Instituto de Tecnologia e Liderança Bachelor of Science, Computer Science R\$500k scholarship; valedictorian of the class. The program runs on ten-week project-based learning cycles delivered with partner companies, including Uber, Pirelli, and the Faculdade de Medicina da USP (full project list).	2022-2026	<i>São Paulo, SP</i>
University of São Paulo (USP) Admitted to the Mathematics program at 15; ranked top 200 in the OBMU (Brazilian Mathematical Olympiad for University Students).	2019	<i>São Paulo, SP</i>
University of Melbourne Discrete Optimization	2023	<i>Melbourne, Australia</i>

EXPERIENCE

AI Engineer CloudWalk (cloudwalk.io) After leaving Tela to take a “sabbatical” and get some situational awareness, I published some research, and Luis (CloudWalk’s CEO and founder) called me to work on AI reporting directly to him. Since then I’ve worked on post-training for language models serving 7M+ merchants, on the harness the models run inside, on the inference engine (vLLM/SGLang), on credit and risk models, and on feedback loops that dispatch thousands of agents to come up with emergent insights on closed-loop problems — all of it on the biggest GPU cluster in LatAm, along with the company’s strategy around AI.	Oct 2025 - Present <i>São Paulo, SP / Sunnyvale, CA</i>
Founding Engineer Tela (tela.com) I was the 10th person on the team, which is now over 200 people, and helped define a lot of the strategy and product from scratch. I delivered AI use-case implementations for enterprise clients, working as an FDE with big companies in Brazil.	Oct 2024 - Apr 2025 <i>São Paulo, SP</i>
Co-founder Doki During university, I developed an iOS/Android app, landing pages, dashboards, and a cloud-deployed backend for financial management aimed at doctors. I led the complete user journey, from interviews to development, which resulted in over R\$12 million registered in the app — transaction management, the mobile interface, and the backend systems, all full-stack on modern tooling and cloud infrastructure.	Jan 2024 - Sep 2024 <i>São Paulo, SP</i>
ML Research Intern BTG Pactual (btgpactual.com) I fine-tuned an LLM at the largest investment bank in Latin America, automating the analysis of fund communications for the investment funds division and scaling it across several divisions.	Jan 2023 - Dec 2023 <i>São Paulo, SP</i>
Research Engineer Higher Order Company (higherorderco.com) I worked on the first version of Kind, a modern proof language, and on Kindelia, a peer-to-peer functional computer, contributing to compiler development and functional programming language design in Rust. Part of the work was implementing Yves Lafont’s interaction combinators , an optimal way of computing.	Feb 2022 - Aug 2022 <i>Remote</i>
Freelance Developer <i>Remote</i> With 12yo started freelancing by making Minecraft with Forge (write-up), then started liking programming and taking freelance work around it, doing small and medium-scale projects in a lot of languages, from PHP to Java.	Until Feb 2022

PROJECTS

Extract-0: A Specialized Language Model for Document Information Extraction arXiv:2509.22906 During my work at Tela I identified that the biggest bottleneck was extracting data from text, so I decided to focus on that. I generated an asymmetric synthetic data pipeline on top of a lot of open source data (FDA, arXiv papers, ...), with the core idea that $X \rightarrow Y$ is hard — and $f(x)$ is what we actually want to map the intelligence onto — while $Y \rightarrow X$ is cheap (solving a sudoku is hard, but building one is not so much). Then I did SFT + GRPO (RLVR) on top of that data.	Sep 2025
---	----------

[arXiv:2508.15835](#)

Benchmarking AI models on Brazilian university entrance exams to address alignment and capability in a local context, and the concerns that come with the Global South being underrepresented in the data these models are trained on. I collected all the data from these exams, benchmarked the models on it, and ran the analysis.