

Vlad-Stefan Harbuz

Peony Software
86-90 Paul Street
EC2A 4NE
London, UK

Email vlad@vlad.website
Website vlad.website
GitHub [@vladh](https://github.com/~vladh)
Sourcehut [~vladh](https://sourcehut/~vladh)
Codeberg [@vladh](https://codeberg.org/~vladh)
Lives in Edinburgh, Scotland

I'm a researcher working on software and philosophy that contributes to the public good. I work as a software developer and frequent Open Source contributor, with 17 years of experience spread across systems programming, data science (particularly life sciences), performance optimisation, 3D graphics, web technologies of all types, compilers and C. I have also played a thought leadership role in the Open Source community, and have considerable experience in teaching programming and technical writing. Most importantly, I am good at quickly and thoroughly learning difficult technologies and doing deep research on unfamiliar topics. I am also a PhD researcher in the philosophy department at the University of Edinburgh, where I work on exploitation and ideology. I love cats.

PROFESSIONAL EXPERIENCE

Peony Software

2022–present

Founder

[Peony Software](#) is a software studio that I run, focused on exclusively developing software that benefits the public good. I am the only permanent member, but frequently collaborate with around a dozen prolific Open Source contributors who also have extensive experience in building technically complex projects used by millions, as well as researchers and artists. Here is a selection of my work.

- **Open Source Pledge:** I co-maintain the Open Source Pledge, which I helped develop in collaboration with [Sentry](#). The Pledge, which [launched internationally](#) on 8 Oct 2024, is an initiative which aims to get companies to pay the maintainers of the Open Source software they depend on. This would help correct the dramatic underfunding of Open Source software, which leads to [maintainer burnout](#) and serious security issues on a [global scale](#).

I developed most of the software the Pledge relies on, liaised with decision makers such as CEOs at prominent tech companies, and developed the philosophy of the Pledge together with my collaborators as part of our thought leadership role. I also had to create innovative Open Source supply chain analysis software, produce long-form video content, create photorealistic 3D graphics and supervise psychological research. Our members, some of which I recruited, have paid \$4,853,685 to Open Source maintainers since our launch.

My work on the Pledge has been publicly praised by the [Open Source Initiative](#), the [PHP Foundation](#), [Open Source Collective](#), [Ruby Central](#), the [Perl and Raku Foundation](#), the [Django Software Foundation](#), venture capital firm [Accel](#), the creators of [Open Collective](#), [ESLint](#), [Vue.js](#), and others. The Pledge has been featured on the Nasdaq screen in Times Square three times (1, 2, 3). I have been interviewed in the [SF Gazetteer](#), [spoke at FOSDEM 2025](#) (the largest Open Source event in the world), and my work has been covered by [TechCrunch](#), [GitHub](#), [The New Stack](#), [InfoQ](#), [IEEE Spectrum](#), [Open Source Ready](#), [Theo - t3.gg](#), [The Register](#) (1, 2), [InfoWorld](#) (1, 2), [It's FOSS](#) (1, 2), [Forkable](#) (1, 2), [The Repository](#), [Technical.ly](#), and taught in classes at [Columbia University](#).

- [thanks.dev](#): I am a core developer of [thanks.dev](#), a service that enables companies to analyse their codebases in order to determine which Open Source maintainers they depend on, then pay those maintainers for their work. This service, which is making a big impact on the state of Open Source funding, is unique in its transparency and ease of use. As part of my work on [thanks.dev](#), I've had to write code that processes millions of lines of code worth of private codebases, while ensuring that code analysis and payment flows are always maximally intuitive for users. My work is regularly relied on by large companies such as Twitter founder Jack Dorsey's [Cash App](#), [Sentry](#), and others.

Here is a selection of my Open Source contributions:

- I am one of the maintainers of the [Hare](#) programming language. My contributions include a [regular expression engine](#) written from scratch that offers better performance than those of programming languages such as Python or Javascript, [neural networks](#), comprehensive [mathematical programming](#) tools, [date/time](#) arithmetic, [2D and 3D graphics](#) libraries, a [linear algebra](#) library, as well as other contributions to the language such as tuple unpacking. My work has enabled programmers to use the Hare language for 2D/3D graphics and numerical/scientific programming. I also serve a technical leadership role in my areas of technical expertise, and previously served a cultural leadership role, guiding over 100 contributors. Hare has been covered in [The Register](#).
- I have made significant contributions to [Sciris](#), a scientific Python library used by hundreds of researchers, including in one of the most widely adopted [models](#) of COVID-19. I have also contributed to [Atomica](#), a simulation engine for compartmental models which can be used to simulate disease epidemics and health care cascades. This work is currently in use by the [Gates Foundation](#) as part of their internal portfolio management tools.
- I am the primary maintainer of [mpld3](#), a scientific data visualisation tool originally developed by Google researcher Jake Vanderplas. [mpld3](#) is among the top 0.6% of the most downloaded Python packages on PyPI, with over [13 million](#) downloads.
- I have contributed to [SourceHut](#), the largest Open Source hacker forge, providing code hosting and other services to over 40,000 users. My contributions have included new features such as license support in Git repositories.

- I am working on the [peony game engine](#), a research project in building a full-featured 3D video game engine entirely from scratch in C++. My work on this game engine has been used for educational purposes by students at top USA [universities](#), and featured in a [PhD thesis](#).
- I have made contributions, some large and some very small, to numerous other Open Source projects, such as [SDL](#) and [Glad](#) (widely-used graphics technologies), [npm](#) (the most widely-used package manager in the world), [tldraw](#), [iced](#) (a popular Rust GUI library), [Wiktionary](#) and [moti](#) via [digerō](#), [uxn](#), [senpai](#), [dezoomify-rs](#), [REUSE](#) and others.

Saffron

2012–2022

Co-founder, Head of Software

In 2012, I founded [Saffron](#), a software development company, together with my business partner. Despite starting with no funding, I played an instrumental part in growing Saffron to a company with over £400,000 in yearly revenue, offices in Romania, Switzerland and the Netherlands, and 12 employees. Our main focus was data visualisation software for scientific data, as well as web applications. Across a decade, I was responsible for the technical architecture of all of our projects, and worked closely with our team on the projects' development. This required me to have expert-level knowledge of web and scientific programming, Python, Javascript, and many other technologies. I also mentored around 25 employees in total, teaching programmers of all experience levels from first principles. I led work on over 30 projects, including:

- **50five & Ajusto:** I coordinated all technical aspects of a 6-year-long effort funded by multinational energy company [Engie](#) to build a home services application for the Netherlands, Belgium, the UK and Romania. I built large swaths of backend, frontend and mobile software to be used across all deployments of the service, architected time-critical systems such as scheduling and invoicing together with the client and rolled out the necessary infrastructure. My leadership role involved coordinating around 50 international collaborators and stakeholders. This project alone attracted over £1m in revenue, and the software I architected is still active in Belgium and the Netherlands, having been used by over 10,000 customers.
- **ICGenealogy:** In collaboration with the [Centre for Neural Circuits and Behaviour](#) at the [University of Oxford](#) and the [Laboratory of Computational Neuroscience](#) at the [EPFL](#), I developed a piece of software for visualising ion channels, a kind of neuroscience research data. Because this project was to be used in cutting-edge research at dozens of universities worldwide, its development required me to combine scientific programming expertise with user-friendly architecture and interfaces. A [professor](#) in Computational Neuroscience at the University of Cambridge described my work as having “a level of UI sophistication and user-friendliness that is simply unprecedented in our field”.
- **Optima:** I worked together with researchers from the [Optima Consortium for Decision Science](#) to help them build Open Source data science tools that help decision-makers choose the best public health investments. My work allowed researchers to adapt code to different diseases roughly 10 times faster than before. The Optima model has been used in around 40 countries, most significantly by the [Gates Foundation](#), and has been used to make decisions on how

to allocate roughly US\$1 billion of funding across different health priorities, including HIV, tuberculosis, nutrition, and malaria.

- Other clients I've worked with include the [Swiss Post](#), the [City of Basel, Switzerland](#), the [University of Copenhagen](#), [GlaxoSmithKline](#), the [Burnet Institute](#), and the [Romanian Government](#).

Submodule

2018–2022

Co-founder, Head of Software

At [Submodule](#), my co-founder and I built the GB01, a device that allows users to manage, Nintendo Game Boy video game cartridges using a modern interface. This project was technically demanding as it required me to develop both a cross-platform desktop application and embedded C software on the device itself. The GB01 has been purchased by over a thousand customers and positively reviewed online.

clumsy computer

2020–present

Host

Due to my frustration with the shortcomings of current programming resources, I started a Twitch and YouTube channel called [clumsy computer](#) in 2020. On this channel, I teach programming from first principles based on a “from scratch” approach, which enables students to understand fundamentals of software more deeply. *clumsy computer* has received strong praise from viewers despite the relatively small amount of content I have produced so far.

PUBLICATIONS AND TALKS

Peer-Reviewed Papers

- Kerr CC, Sanz-Leon P, Abeyesuriya R, Chadderdon GL, **Harbuz VS**, Saidi P, del Mar Quiroga M, Martin-Hughes R, Kelly SL, Cohen JA, Stuart RM, Nachesa A. [Sciris: Simplifying scientific software in Python](#). The Journal of Open-Source Software 2023, 8(88), 5076.
- Kedziora DJ, Abeyesuriya R, Kerr CC, Chadderdon GL, **Harbuz VS**, Metzger S, Wilson DP, Stuart RM. [The Cascade Analysis Tool: software to analyze and optimize care cascades](#). Gates Open Research 2019 3:1488.

Contributions to Peer-Reviewed Papers

- Podlaski WF, Seeholzer A, Groschner LN, Miesenböck G, Ranjan R, Vogels TP. [Mapping the function of neuronal ion channels in model and experiment](#). eLife 2017 6:e22152.

Other Publications

- Anticipations Summer 2023 17 Jul 2023
Long-Term Members of our Communities Should Be Able to Vote

Talks

- [PPE Society](#) London, 15 Jul 2025
The Hidden Non-Market Workers That Keep Markets Going: Why Companies Should Pay Open Source Maintainers
- [The Association for Social and Political Philosophy](#) Glasgow, 15 Jul 2025
The Hidden Non-Market Workers That Keep Markets Going: Why Companies Should Pay Open Source Maintainers
- [AltCtrlOrg](#) — Invited, >200 attendees Basel, 06 Feb 2025
Open Source: Deceptive Power or Collective Governance?
- [FOSDEM 2025](#) — ~100 attendees Brussels, 02 Feb 2025
Why and How Companies Should Pay Open Source Maintainers
- [Statistical Office of Basel, Switzerland](#) — Invited, ~50 attendees Basel, 19 Apr 2019
Ways to Tell a Story With Data Visualisation

Media Appearances

- [Gazetteer SF](#) 24 Oct 2024
Wild billboards around the Bay call attention to tech companies free-loading off open source projects by Megan Rose Dickey

Education

- [PhD Philosophy](#), University of Edinburgh 2024-2028
As part of my PhD, I am researching ideology and the philosophy of work, to determine what parts of our culturally inherited ideas lead us to epistemic practices which cause us to eg accept our exploitation. My committee chair is [Barry Maguire](#), and the other members of my committee are [Berislav Marušić](#) and [Matthew Chrisman](#).
Areas of specialisation: Epistemology, Political and Social Philosophy.
Areas of competence: Philosophy of Language.
- [MScR Philosophy \(Distinction\)](#), University of Edinburgh 2023-2024
In my thesis, *Practical Belief and Epistemic Alienation*, I argued that we can believe not only for factual reasons, but also for practical reasons, ie because believing would bring about a favourable outcome. I explained that this can lead to epistemic alienation. My thesis received a Distinction, the highest mark available, and was examined by [Berislav Marušić](#) and [Robin McKenna](#). My supervisor was [Barry Maguire](#).

- **MA Philosophy** (Merit), University College London 2022-2023
In writing my thesis, titled “Epistemic Alienation”, I researched the philosophy of work and how self-alienation makes us epistemically blind to choices we would otherwise be able to make, and described this using Fricker’s idea of epistemic injustice. My supervisor was [Han van Wietmarschen](#).
- **MSc Computational Biology and Bioinformatics** (not completed) 2019–2020
ETH Zürich & Universität Basel
Relevant courses: Advanced Machine Learning, Biofluidynamics, Cell and Molecular Biology, Computational Biology, Data Mining, Spatio-Temporal Modelling in Biology. I left this two-year course after one year to pursue other work.
- **BSc Computer Science** (First Class), The University of Manchester 2013-2016
Relevant courses: Machine Learning and Optimisation, AI and Games, Symbolic AI, Advanced Computer Graphics, Computer Vision, Compilers, Operating Systems, Natural Language Systems. My Final Year Project was entitled “[Modelling Classical Music with Machine Learning](#)”, supervised by [Allan Ramsay](#).

I’ve also studied linguistics for a term at the University of Basel with [Heike Behrens](#), and classical composition and music theory for a year as part of unmatriculated studies at the Musik-Akademie Basel. I speak native English and Romanian, fluent German and intermediate Japanese.