

PETER SOBOT 🙌

petersobot.com
hire@petersobot.com
Brooklyn, NY

Staff software engineer with a passion for scaling, optimization, teaching, mentorship, and reverse engineering.

Languages of Choice: Python, Scala, C++, Java, C, Objective-C, Ruby, Go, JavaScript (TypeScript), Swift, SQL

Areas of Focus: Audio Processing, Machine Learning, Debugging, Performance Optimization, Open Source, Public Speaking

📌 Work Experience

Staff Machine Learning Engineer at [Spotify](#) in [New York, New York](#) (February 2020–Present)

- Designed and built 🧩 🛠️ **Pedalboard**, a high-performance Python library for working with audio with nearly 5,000 stars on GitHub
- Provided technical leadership to Spotify's Audio Intelligence (MIQ) research lab, which focuses on applying machine learning to audio
 - Built machine learning tooling and frameworks to increase research productivity by **40x** (experiments/quarter)
 - Trained, optimized, and deployed machine learning models for music information retrieval and content identification
 - Built, deployed, and patented Spotify's first on-device machine learning for audio processing
 - Drove cost-optimization efforts to optimize data pipelines, reducing cloud spend by millions of dollars per year
 - Helped improve open source libraries including [TensorFlow](#), [Apache Beam](#), [hnsplib](#), and [TensorFlow Datasets](#).
 - Acted as interim manager for 3 engineers and 5 research scientists *and didn't get scared away from people management*
- Member of Spotify's company-wide incident response team, serving as the first line of defense for severe incidents and outages
- Re-designed, re-implemented, and open-sourced Spotify's approximate nearest-neighbour search library: 🛩️ [Voyager](#)

Staff Software Engineer at [Spotify](#) in [New York, New York](#) (July 2019–February 2020)

- Provided technical leadership, software design, and code to the Personalization Platform team (~40 people)
 - Owned and maintained the core machine learning systems for music recommendation across all of Spotify
 - Powered features including Home, *Discover Weekly*, *Daily Mix*, Radio, Suggested Songs, and *Fans Also Like*
 - Made contributions to open source libraries crucial to Spotify's recommender systems, including [Annoy](#) and [Scio](#)
 - Rewrote Spotify's core recommendation models to increase training speed by **5x**
- Focused on teaching, mentorship, and education
 - Led engineer onboarding bootcamp for all new US-based engineering hires
 - Redesigned introductory engineering curriculum for all global engineering hires
 - Mentored 7 engineers, resulting in 7 promotions
- Gave dozens of internal and external talks to other teams, companies, and conferences across three continents

Senior Backend Engineer at [Spotify](#) in [New York, New York](#) (May 2017–July 2019)

- Designed and implemented scalable backend services in Java for music recommendation and personalization
- Built data pipelines in Scala and Python, processing hundreds of terabytes per day to serve 200m+ Spotify users
- Won Spotify's internal company-wide hack week two years in a row (with projects sadly too secretive to list)
- Led the Spotify NYC Musicians' club and ran bi-monthly *Friday Night Live* internal employee concert series

Senior Software Engineer at [PagerDuty](#) in [Toronto, Ontario](#) (November 2016–April 2017)

- Provided technical direction, design input, mentorship and engineering work to Incident Management team
- Authored and deployed code contributions to 34 projects in 10 languages (primarily Ruby, Javascript, and Swift)
- Won 9 of 12 monthly engineering hack days in 2016 with projects including:
 - [#oncallselfie](#), native Twitter integration for PagerDuty's iOS and Android apps

- Rich HTML Email support for hundreds of thousands of PagerDuty users
- Internal real-time display of outgoing notifications with Go, Kafka, WebSockets & WebGL
- PagerDuty's watchOS app for displaying on-call status and scheduling on your wrist
- Versioning and undo for PagerDuty's on-call scheduling service
- Fastest progression from entry-level to senior engineer in company history (21 months)

Software Engineer II at PagerDuty in Toronto, Ontario (October 2015–November 2016)

- Core contributor for product improvements on Workflow and Incident Management: Systems teams
- Spearheaded efforts to increase overall data throughput by 100x, enabling future customer growth
- Mentored and onboarded 2 interns and 3 new full-time employees

Software Engineer I at PagerDuty in Toronto, Ontario (February 2015–October 2015)

- Helped ship core product enhancements including Incident Snooze and Incident Urgencies
- Ensured reliability and uptime of PagerDuty by participating in primary on-call rotations
- Recorded, sang, and released the BarberDuty ringtones, PagerDuty's on-call barbershop quartet that wakes you up at night

Entrepreneur in Residence at The Working Group in Toronto, Ontario (June 2014–February 2015)

- Directed TWG Labs, a company division focused on experiments, learning and tools
- Managed a team of 3 developers on internet-of-things projects (Raspberry Pi, iBeacons, NodeJS and Heroku)
- Built bleeding-edge product prototypes for national Canadian brands including Tim Hortons

Founder & Lead Developer at Appstruments in Toronto, Ontario (2011–2014)

- Crafted viral music apps for the web, Android and iOS:
 - *the Wub Machine*, an automatic music remixer used to make nearly 5 million remixes
 - *forever.fm*, an endless, automatic & beatmatched radio station

Infrastructure Software Engineering Intern at Inkling in San Francisco, California (2013)

- Rewrote Inkling's aging deployment system
- Created a new onboarding process for engineers
- Replicated Inkling's entire backend infrastructure for a business-critical, Apple-style product launch

Software Engineering Intern at Capcom Game Studio Vancouver in Burnaby, British Columbia (2012)

- Finished Capcom's internal testing infrastructure for Dead Rising 3, a launch title for the Xbox One
- Created analytics tools using real-time gameplay data to aid game designers
- Got my first (and only, so far) video game credit

Software Developer Co-op at The Working Group (now Deloitte) in Toronto, Ontario (2011)

- Rails and iOS development for a variety of clients and internal projects.

Web Developer Co-op at Imagination Plus (now Human_Code) in Hamilton, Ontario (2009)

- PHP and JavaScript development. Led development of a custom PHP-driven learning management system.

Education

Bachelor of Software Engineering (with distinction) from the University of Waterloo (2009–2014)




- Resident of VeloCity, Waterloo's startup incubator residence
- Built MixBox, an iPad music remixing app based on the Echo Nest's (now Spotify's) Remix API (2012–2014, team of 3)
- Wrote JoosBox, a compiler for a large subset of Java in Scala (Winter 2014, team of 3)

- Developed "Colour of the Web," a research project to find the average colour of the internet (Fall 2010, group of 2)

Patents (14)

- Systems and methods for importing audio files in a digital audio workstation (2021)
United States Patent #US20230139415A1
European Patent #EP3962093A1
- Access control for on-device machine learning models (2021)
United States Patent #US20220343004A1
- Systems and methods for providing play context recognition using ultrasound codes (2020)
United States Patent #US20220060791A1
European Patent #EP3962093A1
- Method, System, and Computer-readable Medium for Creating Song Mashups (2019)
United States Patent #US20230075074A1
United States Patent #US20210201863A1
European Patent #EP3843083A1
- Systems and methods for embedding data in media content (2019)
United States Patent #US20230186883A1
United States Patent #US11545122B2
United States Patent #US10777177B1
European Patent #EP3799045A1
- Systems and Methods for Dynamic and Interactive Visualizations for Navigating Media Content (2018)
United States Patent #US20200037019A1
United States Patent #US11449195B2

Selected Open-Source Projects

-   **Pedalboard**, a Python library for high-performance audio work (including VST/AU hosting)
 - Designed and built at Spotify to provide faster data augmentation for training machine learning models
 - Open-sourced to critical acclaim from the Python audio community, gaining 3,000 stars on GitHub in four months
-  **Voyager**, a C++, Python, and Java library for cross-language, in-process nearest-neighbor search
- `keynote-parser`, a Python library for decoding (and re-encoding) Apple Keynote presentation files
- **MusicBoy**, a hardware project to release an album on a custom Nintendo Game Boy® cartridge
 - Built with Toronto-based electro-pop artist Pusher, who provided music and artwork
 - Based on *Super Mario Land*, plus an ARM-based coprocessor in the Game Pak to play back CD-quality audio
 - Presented at !!con 2019 in New York City
- **SampleScanner**, a command-line app to convert hardware synthesizers into software instruments (VSTs)

Selected Talks and Presentations

- It's All About Cost: How to Think about Machine Learning Products
presented at _____ on December 15th, 2022
- Working with Audio in Python (feat. Pedalboard)
presented at _____ in Dublin on July 13th, 2022
- Even more vintage: releasing music on a custom-built Game Boy cartridge!
presented at _____ in NYC on May 12th, 2019

- Music Recommendations at Scale with Cloud Bigtable
presented at _____ in San Francisco on April 9, 2019
- The Recommendations Engine
presented at _____ in Toronto on October 12th, 2018
- Words Are Hard: Talking Good via Computer
presented at _____ in Montréal on January 13th, 2017
- How the Shazam audio fingerprinting algorithm works
presented at _____ on September 2nd, 2015

Activities & Interests

Music Composition, Performance, and Production (Drums, Bass, Guitar, Piano) · Electronic Music Production · Photography · Twitter · Painstakingly Emulating Obsolete Digital Synthesizers from the 90s as Software Instruments · Video Production for Concerts · Home Maintenance of Brooklyn Condos · Cycling around Prospect Park · Cycling around Central Park · Cycling across the George Washington Bridge to 9W · Excessive Usage of Comic Sans