

Charles Leifer

Topeka, KS • coleifer@gmail.com • github.com/coleifer • charlesleifer.com

SUMMARY

Principal engineer and open source author with 17+ years building production Python systems. Creator of widely-used libraries (*peewee*, *huey*, *sqlite-web*, 22K+ combined GitHub stars) and co-creator of *Read the Docs*. Currently shipping production AI features in a real-estate SaaS, including evaluation infrastructure for reliable prompt iteration, structured extraction from complex documents, and automated contract review.

OPEN SOURCE

- *peewee* (12K stars): full-featured Python ORM in a compact, expressive package.
- *huey* (6K stars): Python task queue with comprehensive scheduling, retry, and workflow primitives.
- *sqlite-web* (4.1K stars): web-based database browser for SQLite.
- Co-creator of *readthedocs.org* (Djangodash 2010, 3rd place); now hosts documentation for thousands of open-source projects, serving millions of pageviews monthly.
- *charlesleifer.com*: actively maintained technical blog on Python, databases, and software internals.

EXPERIENCE

Paperless Pipeline, *Principal Engineer*

November 2019 – Present

Responsible for the full application lifecycle of a SaaS platform serving 1,800+ real estate brokerages and processing roughly 500K transactions annually. Owns core backend services, architecture, infrastructure, AI integration, and production operations.

- Designed and built evaluation harnesses for AI feature work: maintains control corpora with expected outputs, runs candidate prompts at scale, synthesizes diffs across runs, and uses an LLM feedback loop to propose prompt and tooling refinements.
- Shipping AI features: structured contract extraction, contract review to surface deal risks, financial data synthesis, workflow automations.
- Migrated large production codebase from Python 2.7 to 3.9, modernized platform and overhauled AWS infrastructure and server architecture, reducing costs across the board.
- Rewrote entire PDF processing stack, vastly improving throughput and reliability (10M+ PDFs/yr).
- Built a complete Document Signing platform, opening a new revenue stream.
- Overhauled commission disbursement tooling, solving complex graph problem of distributing commissions at close among multiple parties with deductions and contributions flowing through diverse channels.
- Designed and shipped a comprehensive reporting suite covering all major broker and agent financials, built on the structured data layer from the commission system.

Propylon, *Senior Software Engineer*

August 2016 – November 2019

- Authored *TimeArc*, a change-management tool used by legal and regulatory officials to track changes across large corpora of documents.
- Built efficient ingestion tooling to parse and normalize diverse relational data formats at scale.
- Worked closely with clients to integrate historical data from enterprise CMS platforms.
- Mentored junior developers via one-on-one coaching and code review.

BATS Global Markets, *Software Engineer*

August 2015 – August 2016

- Developed web-based tools to monitor and operate the third-largest U.S. equities market.
- Built a critical-alert dispatch interface integrating physical alert beacons, desktop software, and SMS.

Counsyl, *Senior Software Engineer*

March 2013 – August 2015

- Automated and optimized billing infrastructure for a medical genomic testing company, increasing revenue and operational efficiency.
- Built a trouble-ticketing framework adopted across the company that improved cross-team workflow.

- **Mediaphormedia** (2009–2012): Built core features for the award-winning *Ellington CMS*; reduced backend latency for high-traffic news sites; led complex production migrations.
- **Nerdery / Sierra Bravo** (2012–2013): Performance work for a high-traffic site (Python, PostgreSQL, Memcached, Apache).
- **Lazerpro Digital Media** (2008–2009): Full-stack work on a community website for a college town (food ordering, classifieds, MLS, news).

SKILLS

Python ecosystem: deep contributor experience including ORMs, task queues, database tooling, Django, Flask; C extensions and performance work

Production AI: prompt engineering, structured-output extraction from complex documents, evaluation infrastructure, multi-model workflows (Claude, Gemini, GPT)

Data & systems: PostgreSQL, MySQL, SQLite (extensively), Redis; performance profiling and optimization

Infrastructure: AWS, Linux administration, Nginx, deployment automation, CI/CD

Other languages: Cython, C, JavaScript, SQL, shell

EDUCATION

Bachelor of Arts in English, Honors

December 2006

University of Kansas, Lawrence, KS • GPA: 3.7 / 4.0

HONORS & PUBLICATIONS

- Co-author, *Learning Flask Framework*, Packt Publishing, 2015.
- Article on clustering algorithms published in *Hacker Monthly*, February 2013.
- College of Liberal Arts and Sciences Honors; Department of English Honors.
- Summerfield Scholarship recipient (four years).