



Adam Gradzki

Email: hi@adamgradzki.com

Telephone (Global): +1 301 785 1297

Work Authorization: EU Citizenship (Poland) & US Citizenship

Country of Residence: United States of America

References available upon request

SUMMARY

Adam Gradzki is a software developer and data engineer with experience in implementing financial analytics and computer vision models in both cloud and handheld devices deployed in the field. His work revolves around implementation and automation of business processes in the cloud to bring lasting business value to clients including finance and business analysts.

Highlights (chronological; oldest to newest):

- Led a team to develop and deploy a **self-service dashboard for radiology** chest x-ray and biopsy pathology slide classification system in the cloud
- Reduced **model training cost** and **faster model testing** with a custom hybrid cloud topology, a blend of on-premises (custom) and cloud server (AWS) technology for model training
- **Implemented and deployed** to production externally audited & essential regulatory consumer credit risk management **loss forecasting models**: IFRS9 and CECL
- Implemented a macroeconomically sensitive **bank stress testing model search** engine

Adam worked for the University of Pittsburgh Drug Discovery Institute to link clinical trial outcomes with performance of Lab-on-a-Chip (LOC) devices which made testing the effects of drugs in human organs without requiring human involvement possible after completing Advanced Individual Training (AIT) at the U.S. Army Intelligence Center of Excellence at Ft. Huachuca. Adam led a team of 4 developers at SemanticMD to develop cloud computer vision dashboards (HTML and REST APIs) for chest x-ray and pathology slide self-service image upload, parsing pixel data, and display systems for high resolution medical image formats such as DICOM and SVS and disease signatures within. Adam worked in a team of ten developers and business analysts to launch international loss forecasting models into production at Capital One Financial Corporation after leaving SemanticMD. Adam is currently a consultant supporting a TOP 40 US BANK implement loan metric alerting systems to automate regulatory audits with auto loans, credit cards, and deposit account performance. Finally, Adam deployed a macroeconomically sensitive bank stress testing model search engine.

SKILLS

- **Linux**: 10 years of professional experience building heterogenous Linux systems including Arch Linux, Alpine Linux, LFS, Ubuntu, Debian, CentOS, and Amazon Linux AMIs on EC2.
- **Data Acquisition, Data Refinement, Qualification Requirement Specification, Software Architecture**: Worked with business analysts and finance teams to deploy new business logic in a rapid and flexible manner.
- **Software Testing**: Introduced Jenkins **CI/CD with pytest** integration and tests for production software and machine learning model deployment.
- **Big Data and Distributed Systems**: Deployed code to **distribute image processing** spanning multiple servers on the cloud.
- **Python**: 10 years of professional experience. Experienced in optimized numerical routines, low-level web APIs, high-level numerical APIs such as Pandas and NumPy, **cloud VM/container deployments**
- **Tableau and SQL**: Implemented **loan metric alerting systems** to automate regulatory audits with auto loans, credit cards, and deposit account performance
- **Javascript**: 4 years of professional experience. Developed dimensionality reduction techniques for **2D and 3D visualization** of semantically related data clusters to allow users to easily see hidden relationships within the data using Three.js with WebGL and native pixel coordinate based DOM manipulation.

HARD SKILLS

Arch Linux, Python, SQL, Pandas, Tableau Desktop, Tableau Server, PostgreSQL, SnowSQL, SQLAlchemy, FastAPI, Parallel Processing, Cloud Engineering and Resource Optimization, QEMU, SSH, Vim, PyCharm, Visual Studio Code, WebStorm, NumPy, Numba, React 18, React Xstate, Ubuntu Linux, CentOS, AWS, Amazon Linux, AWS EC2, AWS RDS, AWS Route53, Keras, Tensorflow, NumPy, SciPy, scikit-learn, scikit-image, Cython, OpenCV 3, Redis, Django, Mailgun, Flask, IPython, Jupyter, unicorn, yarn, gunicorn, uwsgi, pytest, requests, toolz, psycpg2, asyncpg, Python 3 asyncio, systemd, supervisord, zsh, bash, sh, matplotlib, VTK, Point Cloud Library (PCL), OpenSlide, BeautifulSoup 4, lxml, KVM VPS management, VirtualBox, VMware Workstation, VMware Fusion, Golang, Rust, DataGrip

EXPERIENCE

Organization: **Flying Phase**

Role: **Data Engineer**

Dates: July 2020 - Present

Tasks, Goals, and Accomplishments:

- Consistently brought in consulting revenue and successfully completed multiple reports on three **TOP 40 US BANK** data modernization projects
 - Chief Credit Office
 - Reengineered the Board of Directors' Credit Card Risk Layering report - A convoluted 1000 line SAS program was converted to **straightforward, clean, and maintainable Snowflake SQL** fixing multiple logic errors and inconsistencies along the way
 - First Line of Defense
 - Deposit Monthly Credit Risk - Monitoring solution developed from scratch for **application volume, credit risk approval rates, money out exception rates, and dispute losses**
 - Consumer Loan Originations - Monitoring solution developed from scratch for **loan application, approval and booking trends**
 - Consumer Loan Originations Performance - Monitoring solution developed from scratch for application and **loan origination performance metrics** by application data and origination date
 - Consumer Loan Portfolio Performance Metrics - Upgraded an existing report to include **additional balance segmentations** by FICO, APR, Origination FICO, Product Mix, Delinquencies, Losses, and Recovery Rate
 - Credit Risk Review
 - Aligned extremely messy data sources containing numerous logic errors across multiple lines of business and the IT team to **achieve a consistent view** of demand deposit account (DDA) balances across first and second line of defense
- Developed **customer commitment extraction tool** for SCG to automatically identify, extract, and verify written customer commitments from image scans and text documents using **natural language processing and regular expressions**.
- Developed methodology for **cloud migration strategy** cloud migration proposals
- Developed **customer complaint management prototype** with a focus on analyzing tweets from Twitter using machine learning methods

Organization: **Swift Pace Solutions**

Client Site: **Capital One**

Role: **Python Software Developer, Data Engineering**

Dates: February 2019 - April 2020

Tasks, Goals, and Accomplishments:

- **Automated loss forecasting model execution** processes resulting in a reduction in analyst labor in setting up experiments. Removed analyst obstacles by obtaining an order of magnitude time of execution improvement. Made testing and exploration possible of many different economic scenarios when model execution, modification of underlying model assumptions, and the ability to share a reproducible run configuration with other analysts reduced the time needed to launch to only one hour from a cold start instead of one week per analyst using the prior Excel based workflow.
- **Decreased model execution key person risk** with documentation and software process automation. Many

essential tasks used to be fully executable with one or two employees using a manual process.

- **Implemented a configuration sharing model** linking commonly used business analyst data sources to assets on the cloud. Many business processes are driven by ad-hoc Excel spreadsheets or CSV files scattered on the local file systems of business analyst laptops. My configuration system enabled production loss forecasting models to ingest input data and produce a quarterly assessment in an idempotent manner. This reduced the possibility of introducing model errors by eliminating manual alignment of an ad-hoc mixture of input data files from many different sources by each business analyst.
- **Automated model backtesting** to help analysts optimize and improve their modeling strategies before applying them to real world markets. Current Expected Credit Loss (CECL) backtesting analyst experiment cold start time reduced by 89% from 26 hours to 3 hours. Hot start time was reduced by 96% from 26 hours to 1 hour.
- **Unified international and domestic allowance calculations** (an estimate linked to expected credit losses on a financial asset that is applied to reduce the carrying amount of the financial asset in the statement of financial position) within a single model to reduce quarterly assessment execution cold start time from two weeks to one day.

Organization: **AGSK LLC**

Role: **Owner, Computer Vision Software Developer**

Dates: October 2017 - February 2019

Tasks, Goals, and Accomplishments:

- Managed a team of machine learning engineers to develop **skin cancer detection algorithms**.
- Worked with stakeholders to identify project scope, budget, performance measurements on real world test data, and reporting findings.
- **Deployed neural networks** on a low resource mobile device.
- **Optimized algorithm performance** using genetic algorithm hyperparameter search, neural architecture search, synthetic image generation using rule-based image augmentation and generative adversarial neural networks, training set undersampling using neighborhood cleaning rule, training set oversampling techniques, saliency analysis with generated heat map overlay, and 2D and 3D PCA for in-browser feature space visualization. Performed lesion segmentation with systems built using **Keras, PyTorch**, and **Tensorflow**, and **scikit-learn** machine learning frameworks written in the **Python** programming language and the **NumPy, SciPy, Pandas**, and **OpenCV** numerical and image processing libraries.

Organization: **SemanticMD**

Role: **Computer Vision Software Developer**

Dates: September 2014 - October 2017

Tasks, Goals, and Accomplishments:

Managed a team of data scientists in the development of:

- **Cellular video annotation tool** built with proven technologies including HTML5, ffmpeg, and Leaflet.js to provide long lasting value to the customer in terms of compatibility and extensibility.
- **Neural network optimal brain damage** resulting in doubling of inference performance at a cost of 0.02% increased error rate.
- **Cardiac ultrasound** ventricle ejection fraction measurement using video analysis.
- **Facial alignment** from unorganized facial images by performing image registration and elastic deformation.
- **Visualization of similar images in 3D space** using a Three.js/WebGL canvas running in a web browser.
- **PACS systems x-ray image extraction** with a self-hosted DICOM server.

Wrote software in the development of:

- **Automation scripts** for on-premises and cloud server infrastructure with Ansible, sh/bash/sh, and Python scripts
- **Accelerated image processing** using cluster computing methods
- **Distributed task queues** using Python/Celery
- **Identification of regions of interest** within an image with command line and REST API
- **X-ray tuberculosis detector** made available as a REST API

- **Bounding box annotation tool** designed for maximum annotator throughput
- **Automatic face crop** tool using a neural network for detection
- **Skin patches analysis** using Google Cloud Vision (GCV) API
- **Profile guided optimization** using Numba JIT and Cython
- **Ultra high resolution image viewer** for 60000x60000 pixel Aperio SVS images in a web browser as tiles of highly compressed JPEG images generated on the fly as the user pans across the canvas
- **Automated lion facial image extraction** in coordination with 3rd party database developers for LINC
- **Identified lion classification algorithms** using probabilistic regression techniques
- **Developed lion-specific image preprocessing algorithms** including a whisker spot pattern detector
- **Priced cloud resources optimally** for algorithm deployment
- **Automated process supervision scripts** to maintain high algorithm API uptime
- **Audited lion detection algorithm performance** using stratified k-folds cross validation

Organization: **University of Pittsburgh Drug Discovery Institute**

Role: **Python Web Software Developer**

Dates: December 2013 - September 2014

References and Point of Contact: Human Resources Department, tel. 412-624-7000

EDUCATION

U.S. Army Intelligence Center of Excellence (2012)

Advanced Individual Training (AIT)

PERSONAL INTERESTS

The cognitive science of human memory decay and the “spacing effect”