

OZAN OZBEKER

Pittsburgh, PA · 304-651-6856 · contact@ozanozbeker.com · linkedin.com/in/ozanozbeker · ozanozbeker.com/resume

Industrial Engineer turned Analytics Engineer running solo on Fortune 500 client engagements.

Experience

OneMagnify (formerly Splash Analytics, acquired 2023)

Data Science Consultant | Remote | January 2023 to Present

- **Run the day-to-day client relationship** for two flagship Fortune 500 engagements. Lead the standing cadence (bi-weekly and monthly), scope new work directly with manager-level stakeholders on the call, and own delivery end to end across data engineering, analytics engineering, modeling, and reporting.
- **Own the recurring reporting cadence** for both engagements: monthly data refreshes, weekly and monthly Excel deliverables, Quarto reporting, and ad-hoc analysis. Pipelines are automated, idempotent, and reproducible via `uv run`; production code is scripts (never notebooks), orchestrated on Windows Task Scheduler.
- **Rebuilt a flagship customer-engagement scoring pipeline** (26M US and 8M Canada customers, weekly) from legacy R/dbplyr to Python/Polars, then native DB2 SQL under production pressure. Cut runtime from 16 hours to 4 by eliminating unnecessary data movement; added explicit data quality checks at every stage.
- **Designed and built oxy**, a Typer-based Python CLI for async market data ingestion via the Oxylabs HTTP API (over `httpx`), replacing brittle Selenium browser automation. Reduced multi-day scrape workflows to ~5 minutes; runs ~20,000 calls monthly across five internal projects. Two CLI namespaces (`oxy scrape`, `oxy build`); raw JSON preserved verbatim, processed layer materialized to Parquet.
- **Built omutil**, an internal Python package distributed via GitLab with optional dependency groups (`omutil[database]`, `omutil[email]`, `omutil[brand]`). Standardized SQLAlchemy connections per client database, branded Quarto reporting, and structured logging. Adopted ADBC as the default database transport for substantial speedups over ODBC on large extracts.
- **Replaced a third-party sales-estimate dependency** with `amazon-atlas`, an in-house LightGBM model that predicts rolling 30-day Amazon unit sales for any ASIN. Distributed as both an installable Python package and a FastAPI service; eliminated a recurring SaaS subscription and kept client reporting continuous.

West Virginia University

Adjunct Instructor | Morgantown, WV | January 2025 to Present

- **Sole instructor of record** for ~150 students across two academic years for IENG 331 (Computer Applications in Industrial Engineering).
- **Redesigned the curriculum** to cover Python, modern Python tooling (`uv`, `Ruff`, `basedpyright`), terminals and CLIs, SQL with DuckDB, and reproducible reporting with Quarto.
- **Author and maintain** a public Quarto textbook at ozanozbeker.com/cw4e.

Xylem, Inc.

Operations Leadership Development Program (Business Data Analyst, Manufacturing Engineer) | Uniontown, PA and Morton Grove, IL | July 2021 to December 2022

- **Standardized the annual physical-inventory-counting process** across three manufacturing facilities, improving auditability and recovering 0.5 to 1 day of production capacity per site annually.
- **Automated quarterly obsolete-and-excess inventory reporting** for three site materials managers, replacing three days of manual ERP work per manager per quarter with SQL pipelines and Tableau dashboards.

Education

Bachelor of Science, Industrial Engineering | West Virginia University | August 2017 to May 2021

Skills

- **Daily:** Python (Polars, DuckDB, Typer, `marimo`, `httpx`), SQL (DB2, Postgres, SQL Server, DuckDB), Quarto, `uv`, `Ruff`, `basedpyright`, Git, JSON, Parquet, NDJSON, ETL/ELT, REST APIs
- **Comfortable:** SQLAlchemy, FastAPI, `pydantic`, ADBC, Selenium, LightGBM, HTML/CSS, R (Tidyverse), Excel, Power BI, Shiny, Linux, Bash, Tableau, Windows Task Scheduler
- **Learning:** Dagster, dbt, dlt