

John DeSanto

415-420-6742

jdesanto@gmail.com

Seattle, WA

Qualifications

- Experienced developer, lead, and manager
- Data Science, Machine Learning, and NLP
- Ontology, KB development, and Generative AI
- Graduate degrees in math and philosophy
- Python, Go, Lisp, Scala
- SQL, Redshift, AllegroGraph, GraphDB
- OWL, RDF/S, SPARQL, SHACL
- PyTorch, scikit-learn, Pandas, NLTK

Professional Experience

Ontology Developer (contract)

1/2024 – Present

Hyperthesis - Los Angeles, California

- Ontology development for an academic research search engine. Developed methods to utilize general and subject specific ontologies to improve search performance
- Developed methods to improve section identification and extraction from PDFs

Principal Engineer

3/2020 – 11/2023

Oracle – Seattle, WA

- Developed and maintained Oracle's Accelerated Data Science SDK, an open-source Python library available on GitHub and used within Oracle's Data Science service. Engineering lead for hyper-parameter tuning work.
- Built and deployed custom Python computing environments using industry standard libraries (e.g., PyTorch, TensorFlow, ONNX) and infrastructure for deploying models. Maintained supporting documentation, Jupyter notebooks, and automation tools.
- Lead for Oracle's [Guardian AI](#) open-source effort, a unified "Responsible AI" Python library for detecting bias and privacy risks in machine learning models and LLM output

Staff Engineer

11/2018 – 2/2020

Medium - San Francisco, CA

- Member of the data engineering team responsible for Medium's Redshift/Snowflake data infrastructure. Rebuilt scheduling infrastructure to use Apache Airflow.
- Developed data pipelines and infrastructure for PySpark and machine learning jobs.

Ontologist

11/2009 – 3/2020

Multiple Clients

- Ontology consultant on multiple projects. Created and updated ontologies in the following subject areas: food/climate, resumes, Alzheimer's disease, department store retail, printers and accessories (e.g. ink, cables, paper)

Principal Developer, Foundation

5/2016 – 9/2018

EvengX (TrustPipe) - San Francisco, CA

- Created the foundation of EvengX's network intelligence platform that provided real-time packet analysis and threat detection. Using Python and Go, I created and refined modified K-means clustering methods to identify malware in network packet data.
- Developed and maintained processing/testing pipelines for supervised training. Training data consisted of packet captures of labelled malicious/non-malicious traffic

Staff Software Engineer, Architecture

3/2015 – 5/2016

Vungle - San Francisco, CA

- Developed and implemented strategies in Scala for capturing, storing, and utilizing data generated by Vungle's platform and SDK. Improved ad customization and reconciliation via optimizations to Vungle's pub/sub infrastructure
- Team lead for designing and implementing Vungle's advertising exchange. We built an advertising exchange from the ground up in two quarters in Go.

Software Architect

5/2011 – 4/2015

AdPredictive - San Francisco, CA

- Developed AdPredictive's programmatic advertising platform in Python. Wrote code in Python for bidding engine, ad reconciliation, and analytical tools
- Constantly monitored and improved Python code for distributed real-time bidding (RTB) and large-scale ad arbitrage. Implemented Bayesian methods for adjusting bid prices.

Senior Ontologist

8/2007 – 7/2008

NASA / Intrinsic - Mountain View, California

- Created and tested OWL/RDFS ontologies for NASA data systems using TopBraid Composer and XML Schema tools. Worked with SMEs to establish model validity. Created and presented a common data model for integrating multiple ontology formats.

Director of Research and Development

8/2004 – 10/2006

Kozoru - Overland Park, Kansas

- Led Kozoru's research and development efforts and contributed to the newLISP open-source project. Managed junior developers. Developed answer ranking algorithms, distributed source acquisition code, and methods for testing query response changes
- Created and released a customizable chat-based search product written in Lisp, BYOMS (build your own mobile search), using an in-memory data store generated in real time. Implemented a distributed statistical parser using a customized map/reduce framework

Ontologist

11/2002 – 8/2004

Unicru - San Francisco, California

- Led development of the SmartMatch ontology. Supervised knowledge base content editors, developed ontology content and KB functional requirements
- Wrote Python code to scrape and post-process job role data from public sources and resumes. Created methods to unify job titles into roles.

Education

- **The Johns Hopkins University**
M.A. Philosophy
- **The University of Colorado**
M.A. Mathematics
- **The University of Michigan**
B.S. Mathematics and Philosophy

Publications

- Dan Crow and John DeSanto, [*A Hybrid Approach to Concept Extraction and Recognition-based Matching in the Domain of Human Resources*](#), 16th. IEEE International Conference on Tools with AI (ICTAI-04), Boca Raton 2004.
- Martin Quiroga, John DeSanto and John Flowers, [*"Natural language based search engine and methods of use therefore"*](#), U.S Patent #7,447,683.