John DeSanto

415-420-6742 jdesanto@gmail.com Seattle, WA

Qualifications

- Experienced software developer and ontologist
- Data Science, ML, NLP, and Generative AI
- Ontology and knowledge graph development
- Graduate degrees in math and philosophy
- Python, Go, Lisp, Scala
- SQL, Postgres, StarDog, GraphDB
- RDF, RDFS, SKOS, OWL, SPARQL
- PyTorch, scikit-learn, Pandas, NLTK

Professional Experience

Ontologist / Developer
1/2024–Present

Hyperthesis - Los Angeles, California

- Developed ontologies and software to improve search capabilities for academic papers
- Improved reasoning capabilities about diseases by extracting semantic data from PDFs to OWL/RDFS.
- Improved sharing capabilities by creating tools to export RDFS to Excel and back.

Principal Engineer
3/2020–11/2023

Oracle – Seattle, WA

- Developed and maintained Oracle's Accelerated Data Science SDK used in Oracles's hosted notebook service.
- Built and deployed custom Python computing environments for customer applications using industry standard libraries (e.g., PyTorch, TensorFlow, ONNX). Maintained supporting documentation, Jupyter notebooks, and automation tools.
- Led 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. Maintained pipelines for machine learning and analytics.
- Rebuilt scheduling infrastructure to use Apache Airflow to run more efficiently.
- Developed data pipelines and infrastructure for PySpark and machine learning jobs.

Principal Developer, Foundation
5/2016–9/2018

EvengX (TrustPipe) - San Francisco, CA

- Created the foundation of EvengX's network intelligence platform providing real-time packet analysis and threat detection.
- Created and refined modified K-means clustering methods to identify malware in network packet data in Python and Go. This significantly increased both detection speed and accuracy.
- Developed and maintained processing/testing pipelines for supervised training. Training data consisted of packet captures and fingerprinting 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. This allowed Vungle to improve valuations on ad inventory.

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.

Ontologist 2009–2020
Multiple Clients

- Permanent and contract ontology roles:
    - Electric Capital (6/2024 - 9/2024) - cryptocurrency taxonomy
    - Gro Technology (7/2019 - 11/2020) - food ontology for climate risk
    - Self-employed (2008-2010) - department store retail, technology retail
    - NASA (8/2007 - 8/2008) - ontologies for Constellation program

  (further history on request)
- Responsibilities:
    - Created, maintained, and evaluated application specific ontologies using Protege, TopBraid Composer, and custom ontology management tools.
    - Set assessment standards with stakeholders and subject matter experts.
    - Built extensive tooling for export to spreadsheets, visualization, A/B comparison, and querying.

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.