

ALEXANDER J KALINOWSKI

(609) · 405 · 9238 ◇ kalinowski.aj@gmail.com

github.com/akalino

EDUCATION

Drexel University Ph.D. in Information Science	<i>Expected 2023</i>
DePaul University M.S. in Predictive Analytics	<i>May 2015</i>
The George Washington University B.S. in Mathematics	<i>May 2012</i>

WORK EXPERIENCE

Zoetis <i>Senior Data Scientist</i>	September 2021 - Present <i>Cherry Hill, NJ</i>
<ul style="list-style-type: none">· Assess the viability and impact of artificial intelligence and machine learning tasks by implementing proof-of-concepts· Support multiple use cases across Veterinary Medicine Research and Development (VMRD), specifically in the areas of video and text analytics· Build cloud automation for training, testing and deploying machine learning pipelines· Mentor associate data scientists by coaching on best practices and reviewing code for completeness· Lead the data science reading group, a forum for data scientists to share and explain state-of-the-art research papers	
Wells Fargo <i>Knowledge Graph Engineer, Innovation Group</i>	September 2019 - September 2021 <i>Philadelphia, PA</i>
<ul style="list-style-type: none">· Served as the data science specialist within the innovation group, a research and development group dedicated to emerging technologies· Built and deployed semantic search capabilities to interface with enterprise metadata repositories· Developed machine learning capabilities for entity resolution, automating a previously manual task· Worked with legal teams to assess technical innovations, leading to two patents· Researched and provided opinion pieces on new approaches for knowledge graph machine learning	
Jornaya <i>Senior Data Scientist</i>	August 2018 - May 2019 <i>Philadelphia, PA</i>
<ul style="list-style-type: none">· Developed and deployed a URL classification model using PySpark that served as the companies flagship “Consumer Journey” model· Utilized AWS SageMaker to deploy predictive models in a more sustainable, API friendly way· Provided mentoring and support for junior team members, aiding their transition from analysts to data scientists	
Wells Fargo <i>NLP Data Scientist</i>	July 2017 - August 2018 <i>Philadelphia, PA</i>
<ul style="list-style-type: none">· Built complaint detection and classification models used in efforts to address customer concerns· Championed team effectiveness tools, moving team to Jira and enforcing proper reusability and versioning of code via Git	

- Partnered with technology teams to build enterprise-wide distributed data science environment, now supporting 200+ scientists

Sidecar

Data Scientist

January 2016 - July 2017

Philadelphia, PA

- Built and deployed machine learning algorithms to optimize advertising within Google Shopping
- Mined search query data, including using named entity recognition and sentiment analysis, to evaluate consumer intent and extract product-level features
- Converted existing R scripts to Python/PySpark in order to meet scaling needs, including re-writing packages (see Github page)

DoubleDown Interactive

Data Scientist

April 2014 - November 2015

Seattle, WA

- Defined and modernized the company data warehouse, moving from offline Hadoop processing to cloud-based stream processing
- Build dashboards to assess customer churn and product stickiness
- Executed ad hoc analyses to answer root cause questions and build insightful reports

KPMG

Economic Valuation Associate

July 2011 - September 2013

Philadelphia, PA

- Developed statistical models for employee stock option pricing and valuation
- Wrote SQL queries to assist with financial remediation projects

PUBLICATIONS

An, Y., **Kalinowski, A.**, & Greenberg, J. (2022) Exploring Wasserstein Distance across Concept Embeddings for Ontology Matching. ISWC'22 Workshop on Ontology Matching (OM-2022).

Kalinowski, A., & An, Y. (2022). Repurposing Knowledge Graph Embeddings for Triple Representation via Weak Supervision. 2022 Third International Conference on Intelligent Data Science Technologies and Applications (IDSTA). **Best Paper Award.**

Kalinowski, A., & An, Y. (2022). Quantifying the Impact of Predicate Similarities on Knowledge Graph Triple Embeddings. AIC 2022, 8th International Workshop on Artificial Intelligence and Cognition.

Kalinowski, A. (2021). Optimal Transport Methods for Aligning Knowledge Graph Triples with Natural Language in Unsupervised Settings. International Semantic Web Conference.

Kalinowski, A., & An, Y. (2021). Exploring Sentence Embedding Structures for Semantic Relation Extraction. 2021 International Joint Conference on Neural Networks (IJCNN), 1-7.

An, Y., **Kalinowski, A.**, & Greenberg, J. (2021). OTMapOnto: optimal transport-based ontology matching. OM@ISWC.

An, Y., **Kalinowski, A.**, & Greenberg, J. (2021). Clustering and Network Analysis for the Embedding Spaces of Sentences and Sub-Sentences. 2021 Second International Conference on Intelligent Data Science Technologies and Applications (IDSTA), 138-145.

Kalinowski, A., & An, Y. (2020). A Survey of Embedding Space Alignment Methods for Language and Knowledge Graphs. ArXiv, abs/2010.13688.

Kalinowski, A., & An, Y. (2020). A Comparative Study on Structural and Semantic Properties of Sentence Embeddings. ArXiv, abs/2009.11226.