Marina Danilevsky, PhD

San Jose, CA | (847) 533-5669 | marina.danilevsky@gmail.com | https://www.marinadanilevsky.com

Senior Research Scientist and **AI Model Specialist** with 10+ years of results, impacting strategic vision to guide executive management, and leading teams of distributed and co-located researchers and developers. Experienced in spearheading projects and initiatives that focus on aligning downstream business objectives and outcomes with technical innovation though scoping and connecting dots across teams. A proven track record of fostering cultures focused on solving real-world problems. Adept at translating complex concepts into compelling and intuitive strategy for consumption up and down leadership chains. Author of over 40 papers, 6 granted patents and media appearances with over 700k views. Seeking to leverage multi-disciplinary background to make valuable contributions as a technical, strategic and people leader in language understanding, model evaluation and a human-in-the-loop perspective on artificial intelligence.

Core Competencies

- Large Language Models
- Lifecycle of AI Data and AI Evaluation
- Human in the Loop and AI
- People & Engineering Management
- Strategic Thinking and Leadership
- Project Management & Scoping
- Cross-division Collaboration
- Technical Communication
- Executive Leadership Support

PROFESSIONAL EXPERIENCE

IBM Research, Almaden, CA — Feb 2022 – Current

Senior Research Scientist & People Manager

Senior research scientist with a portfolio of both technical and people leadership responsibilities focused on large language model data collection, evaluation and quality control to meet the demand of IBM LLM model performance

- Drove scoping, technical and business agendas as the lead of IBM Research AI subtheme “human in the loop (HITL) data collection and evaluation.” Spearheaded the conception, design, and execution of HITL evaluation methodology for Gen AI.
- Built and scaled IBM’s investment in data annotation and evaluations, owning the $500k engagement with Defined.AI, one of the largest in IBM research, leveraging the methodology defined and implemented by my team.
- Significantly altered and improved the in-house LLM evaluation methodology of a large international bank, as the evaluation lead during a $1M engagement. Processes were taken on as the foundational best practice by their entire data analysis team for LLM usage going forward.
- Supervised and directed a portfolio of 14 AI and ML research scientists & research engineers across Almaden and Yorktown research labs.

IBM Research, Almaden, CA — Jan 2021 – Jan 2022

Chief of Staff to Vice President of AI at IBM Research

Responsible for coordinating the AI division’s interactions and initiatives within and across senior & executive management, impacting over 500 team members. Transformed the Head of AI’s strategic business goals into actionable technical strategies and operating plans impacting 1/3 of IBM Research footprint.

- Supported the VP of AI Research in developing world-wide research strategies and transforming them into operating plans; Tracked progress towards meeting quarterly and annual performance goals by coordinating information across 8 departments
- Ensured that divisional technical workstreams remained aligned to corporate business objectives via sustained active communications and metrics tracking with junior and senior AI technical leads; Gathered and synthesized complex technical material from across the division to support the AI executive leadership team in their internal and external communication needs
- Spearheaded the creation of a demo illustrating the research and product roadmap of IBM Conversational AI; presented the demo as part of the spotlight talk of the VP of AI Research at THINK, the flagship IBM conference for its clients and partners. Functionalities defined in demo use case (centering on RAG for Conversational Assistants) were taken up and implemented by a cross-departmental team over the following several months.
- On the leadership team for the 2021 GTO initiative on Foundation Models (GTO is the annual research-wide initiative to envision the technical and business future in the 5–10-year time frame, with implications for IBM’s current and potential lines of business); co-created and co-presented to the CEO and executive technical leadership team of IBM.

IBM Research, Almaden, CA — Jun 2014 – Dec 2020

Research Scientist

Created multiple award winning research systems & technologies at the intersection of semi-structured text data, information extraction, language modeling and human-in-the-loop techniques for AI. Acted as point of contact for multiple external client engagements (including Elf Cosmetics) and acted as technical lead on multiple cross division initiatives.

- Technical lead for a 7 person development team for creation of language model for parsing, extracting and understanding contracts, our contributions integrated with output from a cross functional, multi-national team of research scientists and 50+ developers ; our model, which was productized by the IBM cloud product team as a part of Watson Compare and Comply, was GA released in 9 months, and within 12 months was available on both public cloud and on-premises.
- Conceptualized, designed and managed Label Sleuth, an open-source no-code system for text annotation and building text classifiers. With over 50 releases, and nearly 250 stars on GitHub, project was recognized and awarded by IBM senior management for significant scientific impact, teaming and reusability.
- Restructured and led the development of the Financial Understanding Knowledgebase, a financial domain instantiation of a multi-year IBM Research cross-lab initiative on content knowledgebases. Presented a GTO proposal based on this project to IBM Research leadership.
- Primary internal and external (e.g. Wells Fargo) facing technical point-of-contact for demonstrations and tutorials on SystemT, a declarative information extraction system, as well as technical contributor to the core SRL system. My team and I were recognized by IBM senior leadership for our scientific and business impact, teaming and reusability.

TECHNICAL PROFICIENCIES

LLMs | Data Analysis | Database Schema Architecture | Git | Agile | Trello | Python | SQL | R | Java | C++ | Angular | CSS | HTML

EDUCATION

University of Illinois at Urbana-Champaign, Urbana, IL

Doctor of Philosophy in Computer Science in Data Mining| Advisor Jiawei Han

Dissertation: Discovering latent topical phrases in document collections and networks with text components: Leveraging text mining and information network analysis for human oriented applications

University of Illinois at Urbana-Champaign, Urbana, IL

Master of Science in Computer Science in Data Mining | Advisor Jiawei Han

Master Thesis: SCENE: Structural Conversation Evolution Network

The University of Chicago, Chicago, IL

Bachelor of Science in Mathematics with General Honors

PUBLICATIONS AND MEDIA

Mixture of Experts [link]
Marina Danilevsky is a frequent guest on this weekly IBM podcast that dives into the ever-evolving landscape of AI.
IBM Technology 2024.

IBM experts break down LLM benchmarks and best practices [link]
IBM Blog, 2024.

What Is Retrieval-Augmented Generation (RAG)? [link]
Senior Research Scientist **Marina Danilevsky** explains the LLM/RAG framework in this video with **700k+** views
IBM Technology 2023.

Language Models in Plain English: Humans Understanding How Machines Understand Language. [link]
Austin Eovito and **Marina Danilevsky**.
O'Reilly 2021.

PATENTS GRANTED

Producing Explainable Rules via Deep Learning
Pritviraj Sen, Siddhartha Brahma, Yunyao Li, Laura Chiticariu, Rajasekar Krishnamurthy, Shivakumar Vaithyanathan, **Marina Danilevsky Hailpern**
US-11900070-B2 2024.

Producing Explainable Rules via Deep Learning
Xilun Chen, Laura Chiticariu, Alexandre Evfimieski, **Marina Danilevsky Hailpern**, Prithviraj Sen
US-111650970-B2 2023

Table recognition in portable document format documents
Douglas Burdick, Wei Cheng, Alexandre Evfimieski, **Marina Danilevsky Hailpern**, Rajasekar Krishnamurthy, Shajith Ikbal Mohamed, Prithviraj Sen, Shivakumar Vaithyanathan
US-11200413-B2 2021.

Constraint tracking and inference generation
Huaiyu Zhu, Michael Wehar, **Marina Danilevsky Hailpern**, Mauricio Hernandez-Sherrington, Yunyao Li
US-11074508-B2 2021.

Cross-lingual information extraction program
Alan Akbik, Laura Chiticariu, **Marina Danilevsky Hailpern**, Yunyao Li, Huaiyu Zhu
US-10042846-B2 2018.

Generation of a natural language resource using a parallel corpus
Alan Akbik, Laura Chiticariu, **Marina Danilevsky Hailpern**, Yunyao Li, Huaiyu Zhu
US-9898460-B2 2018.

PEER-REVIEWED PUBLICATIONS

Evaluating Robustness of Dialogue Summarization Models in the Presence of Naturally Occurring Variations.
Ankita Gupta, Chulaka Gunasekara, Hui Wan, Jatin Ganhotra, Sachindra Joshi and **Marina Danilevsky**.
NLP4ConvAI@ACL 2024.

How Can Personalized Context Help? Exploring Joint Retrieval of Passage and Personalized Context
Hui Wan, Hongkang Li, Songtao Lu, Xiadong Cui and **Marina Danilevsky**.
ICASSP 2024.

Semi-Structured Object Sequence Encoders
Rudra Murthy, Riyaz Bhat, Chulaka Gunasekara, Siva Patel, Hui Wan, Tejas Dhamecha, Danish Contractor, and **Marina Danilevsky**. EMNLP 2023 Findings.

PriMeSRL-Eval: A Practical Quality Metric for Semantic Role Labeling Systems Evaluation
Ishan Jindal, Alexandre Rademaker, Khoi-Nguyen Tran, Huaiyu Zhu, Hiroshi Kanayama, **Marina Danilevsky**, and Yunyao Li.
EACL 2023 Findings.

Label Sleuth: From Unlabeled Text to a Classifier in a Few Hours
Eyal Shnarch, Alon Halfon, Ariel Gera, **Marina Danilevsky**, Yannis Katsis, Leshem Choshen, Martin Santillan Cooper, Dina Epelboim, Zheng Zhang, and Dakuo Wang. *https://www.label-sleuth.org/*
EMNLP 2022.

Learning to Robustly Aggregate Labeling Functions for Semi-supervised Data Programming
Ayush Maheshwari, Krishnateja Killamsetty, Ganesh Ramakrishnan, Rishabh Iyer, **Marina Danilevsky**, and Lucian Popa.
ACL 2022 Findings.

Explainability for Natural Language Processing
Marina Danilevsky, Shipi Dhanorkar, Yunyao Li, Lucian Popa, Kun Qian, and Anbang Xu.
KDD 2021.

XNLP: A Living Survey for XAI Research in Natural Language Processing
Kun Qian, **Marina Danilevsky,** Yannis Katsis, Ban Kawas, Erick Oduor, Lucian Popa, and Yunyao Li.
IUI 2021 Companion.

SemEval-2021 Task 9: Fact Verification and Evidence Finding for Tabular Data in Scientific Documents (SEM-TAB-FACTS)
Nancy Wang, Diwakar Mahajan, **Marina Danilevsky**, and Sara Rosenthal.
SemEval@ACL 2021.

Facilitating Knowledge Sharing from Domain Experts to Data Scientists for Building NLP Models

Soya Park, April Yi Wang, Ban Kawas, Q. Vera Liao, David Piorkowski, and **Marina Danilevsky**.
IUI 2021.

Development of an Enterprise-Grade Contract Understanding System
Arvind Agarwal, Laura Chiticariu, Poornima Chozhiyath Raman, **Marina Danilevsky**, Diman Ghazi, Ankush Gupta, Shanmukha Guttula, Yannis Katsis, Rajasekar Krishnamurthy, Yunyao Li, Shubham Mudgal, Vitobha Munigala, Nicholas Phan, Dhaval Sonawane, Sneha Srinivasan, Sudarshan R. Thitte, Mitesh Vasa, Ramiya Venkatachalam, Vinitha Yaski, Huaiyu Zhu.
NAACL 2021.

A Survey of the State of Explainable AI for Natural Language Processing
Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen.
AACL-IJCNLP 2020.

Learning Explainable Linguistic Expressions with Neural Inductive Logic Programming for Sentence Classification
Prithviraj Sen, **Marina Danilevsky**, Yunyao Li, Siddhartha Brahma, Matthias Boehm, Laura Chiticariu, Rajasekar Krishnamurthy.
EMNLP 2020.

WhyFlow: Explaining Errors in Data Flows Interactively
Maeda F. Hanafi, Azza Abouzeid, **Marina Danilevsky** and Yunyao Li.
DaSH@KDD 2020.

Active Learning for BERT: An Empirical Study
Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, **Marina Danilevsky,** Ranit Aharonov, Yoav Katz, Noam Slonim.
EMNLP 2020.

Table extraction and understanding for scientific and enterprise applications
Douglas Burdick, **Marina Danilevsky**, Alexandre V Evfimievski, Yannis Katsis, and Nancy Wang.
VLDB 2020 Tutorial.

Making AI Machines Work for Humans in FoW
Sihem Amer-Yahia, Senjuti Basu Roy, Lei Chen, Atsuyuki Morishima, James Abello Monedero, Pierre Bourhis, François Charoy, **Marina Danilevsky**, Gautam Das, Gianluca Demartini, Shady Elbassuoni, David Gross-Amblard, Emilie Hoareau, Munenari Inoguchi, Jared Kenworthy, Itaru Kitahara, Dongwon Lee, Yunyao Li, Ria Mae Borromeo, Paolo Papotti, Raghav Rao, Sudeepa Roy, Pierre Senellart, Keishi Tajima, Saravanan Thirumuruganathan, Marion Tommasi, Kazutoshi Umemoto, Andrea Wiggins, and Koichiro Yoshida.
SIGMOD Record

Declarative Text Understanding with SystemT
Huaiyu Zhu, Yunyao Li, Laura Chiticariu, Marina Danilevsky, Sanjana Sahayaraj.
KDD 2019.

DIMSIM: An Accurate Chinese Phonetic Similarity Algorithm Based on Learned High Dimensional Encoding
Min Li, **Marina Danilevsky**, Sara Noeman, Yunyao Li.
CoNLL 2018.

Web Information Extraction
Laura Chiticariu, **Marina Danilevsky**, Howard Ho, Rajasekar Krishnamurthy, Yunyao Li, Sriram Raghavan, Frederick Reiss, Shivakumar Vaithyanathan, Huaiyu Zhu.
Encyclopedia of Database Systems (2nd ed), *pp. 1–9*, Springer New York, 2018.

Creation and Interaction with Large-scale Domain-specific Knowledge Bases
S. Bharadwaj, L. Chiticariu, M. Danilevsky, S. Dhingra, S. Divekar, A. Carreno-Fuentes, H. Gupta, N. Gupta, S.-D. Han, M. Hernandez, H. Ho, P. Jain, S. Joshi, H. Karanam, S. Krishnan, R. Krishnamurthy, Y. Li, S. Manivannan, A. Mittal, F.
VLDB 2017.

A Method to Accelerate Human in the Loop Clustering
Anni Coden, **Marina Danilevsky**, Daniel Gruhl, Linda Kato, and Meena Nagarajan.
SDM 2017.

A Rectangle Mining Method for Understanding the Semantics of Financial Tables
Xilun Chen, Laura Chiticariu, Marina Danilevsky, Alexandre Evfimievski and Prithviraj Sen
ICDAR 2017.

Multilingual Information Extraction with POLYGLOTIE
Alan Akbik, Laura Chiticariu, **Marina Danilevsky**, Yonas Kbrom, Yunyao Li, Huaiyu Zhu.
COLING 2016.

Generating High Quality Proposition Banks for Multilingual Semantic Role Labeling
Alan Akbik, Laura Chiticariu, **Marina Danilevsky**, Yunyao Li, Shivakumar Vaithyanathan, Huaiyu Zhu.
ACL 2015.

Automatic Construction and Ranking of Topical Keyphrases on Collections of Short Documents.
Marina Danilevsky, Chi Wang, Nihit Desai, Xiang Ren, Jingyi Guo, and Jiawei Han.
SDM 2014.

Constructing Topical Hierarchies in Heterogeneous Information Networks
Chi Wang, **Marina Danilevsky**, Jialu Liu, Nihit Desai, Heng Ji, and Jiawei Han.
ICDM 2013.

Information Graph Model and Application to Online Advertising
Marina Danilevsky and Eunyee Koh.
UEO 2013.

AMETHYST: A System for Mining and Exploring Topical Hierarchies in Information Networks
Marina Danilevsky, Chi Wang, Fangbo Tao, Son Nguyen, Gong Chen, Nihit Desai, and Jiawei Han.
KDD 2013.

A Phrase Mining Framework for Recursive Construction of a Topical Hierarchy
Chi Wang, **Marina Danilevsky,** Nihit Desai, Yinan Zhang, Phuong Nguyen, Thrivikrama Taula, and Jiawei Han.
KDD 2013.

Entity Role Discovery in Hierarchical Topical Communities
Marina Danilevsky, Chi Wang, Nihit Desai, and Jiawei Han.
MDS@KDD 2013.

Large-Scale Spectral Clustering on Graphs
Jialu Liu, Chi Wang, **Marina Danilevsky** and Jiawei Han.
IJCAI 2013.

ACES: A Cross-Discipline Platform and Method for Communication and Language Research
Joshua Hailpern, **Marina Danilevsky**, Andrew Harris, Sunah Suh, Reed LaBotz and Karrie Karahalios.
CSCW 2013.

Efficient Mining of Correlated Sequential Patterns Based on Null Hypothesis
Cindy Xide Lin, Ming Ji, **Marina Danilevsky**, and Jiawei Han.
Web-KR@CIKM 2012.

Locality Preserving Feature Learning
Quanquan Gu, **Marina Danilevsky**, Zhenhui Li, and Jiawei Han.
AISTAT 2012.

The Joint Inference of Topic Diffusion and Evolution in Social Communities
Xide Lin, Qiaozhu Mei, JIawei Han, Yunliang Jiang, and **Marina Danilevsky**.
ICDM 2011.

Ranking-Based Classification of Heterogeneous Information Networks
Ming Ji, Jiawei Han, and **Marina Danilevsky**.
KDD 2011.

ACES: Aphasia Emulation, Realism, and the Turing Test
Joshua Hailpern, **Marina Danilevsky**, and Karrie Karahalios.
ASSETS 2011.

SCENE: Structural Conversation Evolution NEtwork
Marina Danilevsky, Joshua Hailpern, and Jiawei Han.
ASONAM 2011.

WinaCS: Construction and Analysis of Web-Based Computer Science Information Networks
Timothy Weninger, **Marina Danilevsky**, Fabio Fumarola, Joshua Hailpern, Jiawei Han, Ming Ji, Thomas J. Johnston, Surya Kallumadi, Hyungsul Kim, Zhijin Li, David McCloskey, Yizhou Sun, Nathan E. TeGrotenhuis, Chi Wang, and Xiao Yu.
SIGMOD 2011.

ACES: Promoting Empathy Towards Aphasia Through Language Distortion Emulation Software
Joshua Hailpern, **Marina Danilevsky**, Andrew Harris, Karrie Karahalios, Gary Dell, and Julie Hengst.
CHI 2011.

Graph Regularized Transductive Classification on Heterogeneous Information Networks
Ming Ji, Yizhou Sun, **Marina Danilevsky**, Jiawei Han, and Jing Gao.
ECMLPKDD 2010.

Clustering in the Creative Industries: Insights from the Origins of Computer Software
Martin Campbell-Kelly, **Marina Danilevsky**, Daniel D. Garcia-Swartz, and Shane Pederson.
Industry and Innovation 17(4), 2010.

The Economics of Market Coordination for the pre-Fed Check Clearing System: A Peek into the Bloomington (IL) Node
Howard H. Chang, **Marina Danilevsky**, David S. Evans, and Daniel D. Garcia-Swartz.
Explorations in Economic History 45(4): 445-461, 2008.