

# Ves Grantcharov

## Education

**University of New Mexico** (2024 - present) PhD in Nanoscience and Microsystems Engineering

**Georgia Institute of Technology** (2014 - 2019) BSc in Discrete Mathematics, Minor in Computer Science (Intelligence)

## Publications

- Ryan C Barron\*, **Vesselin Grantcharov\***, Selma Wanna, Maksim E Eren, Manish Bhattarai, Nicholas Solovyev, George Tompkins, Charles Nicholas, Kim Ø Rasmussen, Cynthia Matuszek, Boian S Alexandrov. “Domain-Specific Retrieval-Augmented Generation Using Vector Stores, Knowledge Graphs, and Tensor Factorization.” ICMLA 2024. (35 citations, \*equal contribution)
- Manish Bhattarai, Ryan Barron, Maksim E. Eren, Minh N. Vu, **Vesselin Grantcharov**, Ismael Ismael, Valentin Stanev, Cynthia Matuszek, Vladimir I Valtchinov, Kim Rasmussen, Boian S. Alexandrov. “HEAL: Hierarchical Embedding Alignment Loss for Improved Retrieval and Representation Learning.” KnowledgeNLP Workshop @ COLING 2025.
- [Manuscript in preparation] Blind Source Separation Model for Identification of Sub-populations in Dynamic Light Scattering. Target: Journal of Chemical Physics.

## Research Experience

### University of New Mexico / Los Alamos National Laboratory

*PhD Researcher, [Smart Tensors Group](#) (2024 - present)*

*Advisor: Dimitar Petsev (UNM), Mentor: Boian Alexandrov (LANL)*

*Funded by LANL through Spring 2025*

#### Blind source separation for dynamic light scattering spectroscopy

Developing method for identifying sub-populations in polydisperse (multimodal) samples from DLS data, addressing limitations of existing techniques that practitioners distrust for mixture analysis. Formulated problem as Laplace inversion with Gaussian mixture priors and working toward a separability criterion analogous to Rayleigh’s criterion in optics. Implemented box-constrained second-order trust-region optimizer in JAX to handle the ill-posed inverse problem. Manuscript in preparation targeting Journal of Chemical Physics.

#### Domain-specific retrieval-augmented generation

Contributing to information retrieval system for materials science research, currently focused on perovskite literature. Designed hybrid architecture routing queries between graph database (for bibliometric/trend queries) and vector search (for semantic retrieval over research text) based on query classification. Published at ICMLA 2024 (35 citations); follow-up work on hierarchical embedding alignment published at KnowledgeNLP 2025.

### University of New Mexico

*Research Assistant, advised by Amanda Bienz (2025 - present)*

#### Distributed second-order optimization for neural networks

Investigating whether curvature-informed optimization can be wall-time competitive for deep learning. Building on Hessian-free methods (Martens & Sutskever, 2012), replacing conjugate gradient with LSQR for solving the Gauss-Newton subproblem, avoiding explicit formation of Jacobian or Hessian matrices. Implemented distributed dense LSQR using MPI all-reduce; extending to GPU/NCCL with JAX.

## Georgia Tech

*Undergraduate Researcher, advised by Ernest Croot (2018 - 2019)*

### Additive combinatorics and the cap set problem

Explored lower bounds for cap sets (maximal sets of vectors containing no three-term arithmetic progressions). Developed software for representing, validating, and constructing cap sets via graph search; reduced state-space through isomorphism identification.

## Los Alamos National Laboratory

*Intern (Summer 2016)*

Developed Julia packages for nonnegative matrix factorization and SVM, integrated into MADS.jl library.

## Industry Experience

### Fitch Ratings

*Senior Machine Learning Engineer (2022 - 2024), New York*

Built retrieval-augmented generation system over 270k research documents, from initial prototype to production feature for premium subscribers. Diagnosed relevance failures in naive RAG approach and developed NER-based query preprocessing to extract entities (company, industry, year) for metadata filtering, improving retrieval from demo-able to production-ready. Explored agentic RAG architectures using LangGraph. Feature announced in [public press release](#) after departure.

Led ML for ESG data extraction project automating metric extraction from company documents using instruction-tuned LLMs. Developed custom heuristics for paragraph segmentation using spatial and textual features of document elements.

### Lifion by ADP

*Machine Learning Engineer / Applied Scientist (2020 - 2022), New York*

Proposed and implemented learning-to-rank model to personalize search results based on user history, drawing on approaches from Yelp's search ranking. Implemented hybrid search combining HNSW vector indices with n-gram matching before vector databases were widely available.

## Georgia Pacific

*Data Engineer (2019 - 2020), Atlanta*

Developed LSTM model for predicting plywood quality grades from time series data of dryer temperatures and line velocity, supporting manufacturing optimization.

## Awards

Graduate Student Excellence Award, University of New Mexico School of Engineering (2025)

Putnam Mathematical Competition - 70th percentile (2016), 60th percentile (2015)

ACM ICPC Southeast Regionals - 4th place (2018)

## Skills

**Graduate Coursework:** Deep Reinforcement Learning, Parallel Processing, Numerical Linear Algebra, Numerical ODEs, Stochastics, Material Characterization Methods, Chemistry and Physics at the Nanoscale, Nanofabrication

**Languages & Tools:** Python (JAX, PyTorch, NumPy, SciPy, Pandas, LangChain/LangGraph), C, MPI/NCCL, CUDA, Slurm, SQL, Spark/DeltaLake, Elastic/OpenSearch, Neo4J, AWS, Azure, GCP, Jenkins, Docker, Kubernetes, LaTeX, Git, Bash