

Amar Viswanathan Kannan

✉ amarviswanathan@gmail.com • 🌐 amarviswanathan.github.io • in amarviswanathan
📄 amarviswanathan

I am a hands on ML research manager at Verisk AI, where I am involved in the conceptualization and construction of large scale multimodal knowledge graphs. As a researcher I work on a gamut of problems that fall at the intersection of Information Extraction (NLP), Knowledge graph representation, construction, summarization, optimization, intersection of vision and language. Prior to this I was a Lead ML Scientist at Best Buy. I was the **Co-Principal Investigator** of the highly successful \$1M Siemens DARPA-ASKE project (2018-2020). I completed my PhD under Dr. James A.Hendler at the Tetherless World Constellation where I focused on large scale knowledge graph query (SPARQL) reformulation.

Employment

- **Verisk AI** New York City, NY
Machine Learning Research Manager June 2021 – Present
- **BestBuy** Minneapolis/Remote
Lead Machine Learning Scientist March 2021 – June 2021
- **Siemens Corporate Technology** Princeton
Acting Senior Key Expert - Knowledge Graphs September 2019 – Feb 2021
- **Siemens Corporate Technology** Princeton
Staff Research Scientist August 2018 – Feb 2021
- **Infosys Technologies Limited** Bengaluru
Senior Systems Engineer, Infosys Labs October 2007– July 2011
- **IBM T.J. Watson Research Center** Yorktown Heights
Summer Research Intern, Mentor: Dr.Geeth De Mel Summer 2015
- **IBM T.J. Watson Research Center** Yorktown Heights
Summer Research Intern, Mentor: Dr.Kaoutar el Maghraoui Summer 2014

Education

- **Rensselaer Polytechnic Institute** Troy, New York
PhD Computer Science , CGPA 3.6 2011–2018(Spring)
- **Rensselaer Polytechnic Institute** Troy, New York
M.S. Computer Science , CGPA 3.9 2014–2016
- **Anna University** Chennai, India
B.E. Computer Science, 84% with Distinction 2003-2007

Projects

Industry.....

Paraphrastic Summarization: (2021-PRESENT): Abstractive Summarization system of Large scale multimodal event news data utilizing SOTA Transformer based NLG systems (T5, RoBERTa). End to end pipeline(Kubeflow) and has extracted 10k events. ROUGE-F1 scores of > 35% (higher than industry standards) along with SOTA models to correct hallucination. Tested on SOTA datasets such as CNNDM and XSum **Keywords:** T5, RoBERTa, OpenIE, Graph Correction

Event Knowledge graphs: (2021-PRESENT): Leading the effort in constructing an event knowledge graph for 'The Machine'. Involves a pipeline that includes user annotation practices, NER, Relation Extraction and Schema mapping. Knowledge graph of currently a million entities along with 80 unique property edges deployed as both RDF graph and Property graph representation. **Keywords:** Amazon Neptune, Neo4j, Virtuoso, NER, Relation Extraction, spaCy, PyTorch Lightning

Product Knowledge graphs: (2021): Led the effort in constructing a product knowledge graph that improves the personalization and recommendation experience of the Best Buy search platform. Mapped the entire product catalog of best buy to build a Best Buy Product Knowledge Graph to create 603 unique product categories and 200,000 entities along with 63 property edges. Pipeline created using Kubeflow. Worked on multimodal representation, entity knowledge augmentation and entity disambiguation. Led a team of 5 scientists and 3 software engineers. **Keywords:** LSTM, Encoder-Decoder, Transformers, Embeddings, Knowledge graphs

KG enabled caption generation: (2019-2020): Built a caption generation system that utilized Knowledge graphs (DBpedia, YAGO) and entity linking along with a encoder-decoder network for contextual caption generation. Image and Caption training Dataset was the 330K images and 80 object category COCO dataset from the VQA challenge. Led a team of 2 research scientists and 2 software engineers for this task.**Keywords:** LSTM, Encoder-Decoder, Transformers, Embeddings, Knowledge graphs

Visual Question Answering: (2019-2020): Built a Visual Question answering system that utilized Knowledge graphs (DBpedia, YAGO), entity linking along with a GRU based network for question answering. Led a team of 2 research scientists and 2 software engineers for this task. Datasets were the COCO datasets from the VQA challenge and Pascal VOC dataset. **Keywords:** GRU, Embeddings, NLP, Knowledge graphs

Scientific Knowledge graph construction: (2018 - 2020): As co-PI for DARPA's ASKE Project, I built an end to end text BiLSTM-CRF extraction system that converted 960 scientific papers to queryable knowledge graphs. I contributed heavily to the NLP pipeline creation was the owner of the Knowledge base construction and oversaw the image extraction architecture. NER and RE Annotation was done using Brat. Knowledge graph ontology was created using Protege. Led a team of 15 members (industry and academic partner) and we were successful in taking this to two phases and also publishing this extensively. More details at <https://github.com/deepcurator/DCC>. **Keywords:** BiLSTM, Named Entity Recognition (NER), Relationship Extraction (RE), KG construction, RDF Knowledge graph

Infosys Technologies Limited: (2007-2011): I worked on Semantic Wikis and Business Process Execution Language (BPEL). In addition, I also worked on Sentiment analysis to build iSEE (Infosys Sentiment Extraction Engine). This also resulted in a patent,

which was awarded in 2013. **Keywords:** Sentiment Analysis, Text Analysis, Suggestion Mining

Academic

Knowledge graph querying (Thesis): (FALL 2014 - SPRING 2018): Built a framework for graph query reformulation, which utilized a specificity and MMR diversity based system for suggesting queries. The work was selected at the AAAI Doctoral Consortium in 2016. <http://tinyurl.com/aaaidc16>. **Keywords:** Flexible Querying, SPARQL Query Relaxation, Query Reformulation, Maximal Marginal Relevance (MMR), RDF, RDFS, Optimization

Question Answering System: (FALL 2013 - SUMMER 2014): I developed an extension for the Watson QA pipeline to answer Knowledge Graph queries. We used a corpus of ≈ 8 billion triple statements to train our system. More details on the system at <http://tinyurl.com/watsonrdf>. **Keywords:** Question Answering, RDF, RDFS, OWL

Large Scale Text Analysis of International Open Government Metadata: (SUMMER 2013): I worked on the textual understanding of the metadata collected by the IOGDS project and analyzed the textual metadata using traditional Named Entity Recognition and Information Retrieval measures. The results were also presented in different visualizations developed using d3.js. The details can be found at <http://tinyurl.com/logd2>. **Keywords:** TF-IDF, Clustering, Word Clouds

Awards

DARPA: Co-PI for the \$1M project **Deep Code Curator : A system to represent scientific papers as knowledge graphs**. This was part of the DARPA ASKE program in 2018. (Siemens share is \$760,000)

US2TS 2018: NSF Award for Student presenters at the U.S. Semantic Technologies Symposium 2018

AAAI 2016: SIGAI Doctoral Consortium Award

Finalist: 3MT Three Minute Thesis at RPI's Graduate research symposium 2016

Professional Service

Reviewer: CVPR 2021, CVPR 2022, CogSci 2016, CogSci 2017, CogSci 2018, CogSci 2019, CogSci 2020, AAAI 2019, AAAI 2020, AAAI 2022, Big Data Journal

Organizer: WebSci 2017

Organizer: Multimodal Forensics (CVPR 2020)

PC Member: Semantic Web for Social Good (SWSG 2018), SWSG 2019

Conference Chair: CIKM 2020, Demos and Posters Session Co-Chair

Skills

NLP: Sentiment Analysis, Named entity extraction, Relation extraction, Question Answering, Seq2seq models, contextual models, embeddings

Semantic Web: RDF, RDFS, OWL, SPARQL

Graphs(KG): Entity Summarization, Knowledge Graph(KG) construction, KG evaluation, RDF2Vec

Querying: Query reformulation

Vision: Object Detection, Object Tracking, Visual Question Answering

Tools and Libraries

Languages: Java, Python

Web: XHTML, CSS, Bootstrap, jQuery, d3.js

Graph DB: Protégé, Virtuoso, Apache Jena (Java based), Blazegraph (Java based), Neo4j

Libraries: PyTorch, NLTK, spaCy, scikit-learn, opencv, brat, pandas, Stanford NLTK (java based), Lucene, Solr

Select Talks and Posters

2020: Amar Viswanathan. A multi modal knowledge graph for deep learning papers and code, **The 29th ACM Conference on information and knowledge management, CIKM 2020**

2019: Amar Viswanathan. Deep Code Curation - A project to build multi modal knowledge graphs from deep learning papers, **The Scientific Literature Knowledge Bases workshop at AKBC 2019**

2016: Amar Viswanathan. Pragmatics Aware Querying in Heterogeneous Knowledge Graphs, **Thirtieth AAAI Conference on Artificial Intelligence, Doctoral Consortium, Phoenix, AZ (USA) 02/12/2016**.

2016: Amar Viswanathan, Geeth de Mel, James A. Hendler. Pragmatics and Discourse in Knowledge Graphs, **Workshop on Symbiotic Cognitive Systems, Thirtieth AAAI Conference on Artificial Intelligence, Phoenix, AZ(USA) 02/12/2016**.

2015: Amar Viswanathan, Geeth de Mel, James A. Hendler. Pragmatic Query Reformulation and Answer Generation in Knowledge Graphs, **IBM Cognitive Computing Symposium, RPI, Troy, NY 11/09/2015**. (Poster Session)

2015: Amar Viswanathan. Pragmatic Query Reformulation in Heterogeneous Knowledge Graphs, **Intern Talk and Poster Session, IBM, Yorktown Heights, NY, 08/22/2015**.

2014: Amar Viswanathan. "Not Elementary, My dear Watson.."- Extending Watson for Question Answering on Linked Open Data, **IBM Cognitive Computing Symposium, IBM, Yorktown Heights, NY 10/30/2014**. (Poster Session)

2014: Amar Viswanathan. Semi Supervised Pattern Summarization of Client Resolution Data, **IBM Summer Intern Poster Event IBM, Yorktown Heights, NY, 08/22/2014**. (Poster Session), Public talk on 08/06/2014

2013: John Erickson, Amar Viswanathan, Josh Shinavier, Yongmei Shi, James A. Hendler. Text Analysis of International Open Government Data, **NY State Health Data Codeathon, RPI, Troy, NY, 12/20/2013**.

Publications

Aditi Roy, Ioannis Akrotirianakis, Amar Kannan Viswanathan, Dmitriy Fradkin, Arquimedes Canedo, and Kaushik Koneripalli. Diag2graph : Representing deep learning diagrams in research papers as knowledge graphs. In *International Conference in Image Processing (ICIP 2020)*. **ICIP 2020**, 2020.

Amar Viswanathan, Dmitriy Fradkin, Aditi Roy, Ioannis Akrotirianakis, and Arquimedes Canedo. A multi modal knowledge graph for deep learning papers and code. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (To appear in CIKM 2020)*. **CIKM 2020**, 2020.

Amar Viswanathan Kannan. *Schema-and Data-aware Query Reformulation in Knowledge Graphs*. PhD thesis, Rensselaer Polytechnic Institute, 2018.

Sabbir Rashid, Amar Viswanathan, Ian Gross, Elisa Kendall, and Deborah McGuinness. Leveraging Semantics for Large-Scale Knowledge Evaluation. In *WebSci-17 Workshop on Industrial Knowledge Graphs*. **WebScience-17**, 2017.

Amar Viswanathan, James R Michaelis, Geeth R de Mel, and James Hendler. In context query reformulation for failing sparql queries. In *Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR VIII*. **SPIE-17**, 2017.

Amar Viswanathan. Pragmatic reformulation in knowledge graphs. In *AAAI-16 Doctoral Consortium*. **AAAI**, 2016.

Amar Viswanathan, Geeth De Mel, and James Hendler. Pragmatics and Discourse Knowledge Graphs. In *AAAI-16 Workshop on Symbiotic Cognitive Systems*. **AAAI**, 2016.

Rajesh Balakrishnan, Bintu G Vasudevan, Amar Viswanathan, Prasanna Venkatesh Raghunathan, and Umadas Ravindran. Methods for analyzing user opinions and devices thereof, July 19 2013. **US Patent App. 13/946,832**.

John S Erickson, Amar Viswanathan, Joshua Shinavier, Yongmei Shi, and James A Hendler. Open Government Data: A Data Analytics Approach. *IEEE Intelligent Systems*, pages 19–23, 2013.

Amar Viswanathan, Prasanna Venkatesh, Bintu G Vasudevan, Rajesh Balakrishnan, and Lokendra Shastri. Suggestion Mining from Customer Reviews. 2011.

Toufeeq Hussain, Rajesh Balakrishnan, and Amar Viswanathan. Semantic wiki aided business process specification. In *Proceedings of the 18th international conference on World Wide Web*, pages 1135–1136. ACM, 2009.

Working Papers

2022: Amar Viswanathan, Shikha Bordia, Arjun Manoharan, Manish Shrivastava, Maneesh Singh, "A Knowledge graph guided factuality preserving summarization system"

2021: Amar Viswanathan, "Knowledge graph summarization utilizing language models"

Relevant Graduate Courses

Advanced Semantic Technologies, Ontology Engineering, Advanced Web Science, Semantic E-Science Foundations of Data Science, Data Science, Data Mining, NLP with Watson, Knowledge Graphs from IE Text, Analysis of Algorithms, Foundations of Network Science, Linear Algebra, Numerical Computing

Relevant Certifications

Udacity: Deep Learning Nanodegree

Udacity: Machine Learning Nanodegree

Udacity: Computer Vision Nanodegree

Udacity: Natural Language Processing Nanodegree

Coursera: Machine Learning Andrew Ng

Coursera: Deep Learning Specialization

Coursera: Text Analysis and Search Engines