

# Matthew Malensek

Associate Professor  
Department of Computer Science · University of San Francisco

☎ (415) 422-4756 | ✉ [mmalensek@usfca.edu](mailto:mmalensek@usfca.edu) | 🏠 [www.cs.usfca.edu/~mmalensek](http://www.cs.usfca.edu/~mmalensek)

## Research Interests

---

My research falls broadly within the area of big data and systems software, with a focus on performing analytics at scale, accounting for the storage and management of voluminous data streams, and transparently managing federation across cloud and fog domains.

## Professional Experience

---

### Academic Director, CS Graduate Programs

Department of Computer Science, University of San Francisco

*San Francisco, California*

*Nov 2023 — Present*

### Associate Professor

Department of Computer Science, University of San Francisco

*San Francisco, California*

*Aug 2023 — Present*

### Assistant Professor

Department of Computer Science, University of San Francisco

*San Francisco, California*

*Aug 2017 — Aug 2023*

### Graduate Research Assistant

Colorado State University

*Fort Collins, Colorado*

*Jan 2011 — Aug 2017*

### System Administrator

Colorado State University

*Fort Collins, Colorado*

*Jan 2010 — Jan 2011*

## Education

---

### Colorado State University

Ph.D., Computer Science

*Fort Collins, Colorado*

*Summer 2017*

### Colorado State University

M.S., Computer Science

*Fort Collins, Colorado*

*Fall 2012*

### Colorado State University

B.S., Applied Computing Technology  
Minor: Business Administration

*Fort Collins, Colorado*

*Fall 2009*

## Awards & Honors

---

2018	<b>IEEE TCSC Outstanding Ph.D. Dissertation Award</b> , Presented by the IEEE Technical Committee on Scalable Computing	Guangzhou, China
2014	<b>Best Paper Award</b> , IEEE/ACM Symposium on Big Data Computing (22% acceptance rate) with Jared Koontz and Sangmi Lee Pallickara	London, UK
2012	<b>Best Paper Award</b> , IEEE/ACM Conference on Utility and Cloud Computing (27% acceptance rate) with Sangmi Lee Pallickara and Shrideep Pallickara	Chicago, IL

## Funding

---

### Community Engaged Scholars in Computer Science

March 2019–February 2024

Sophie Engle (PI), Alark Joshi (Co-PI), **Matthew Malensek** (Co-PI),  
Xornam Apedoe (Co-PI), Chris Brooks (Co-PI), Gian Bruno (Key Contributor),  
Eileen Fung (Key Contributor), Sami Rollins (Former Co-PI, Key Contributor),  
Christina Tzagarakis-Foster (Former Co-PI), and Malik Henfield (Former Co-PI)  
National Science Foundation (NSF) Scholarships in  
Science, Technology, Engineering, and Mathematics (S-STEM) Grant DUE-1833718, \$650,000.  
<https://scholars.cs.usfca.edu>

### Conduit: Fast and Flexible Big Data Analytics Pipelines

2023

#### Matthew Malensek

University of San Francisco Faculty Development Fund (FDF) Award 2420, \$3,600

### Streaming Big Data Visualization Research

2022

#### Matthew Malensek

University of San Francisco Faculty Development Fund (FDF) Award 2253, \$3,000

### Travel for IEEE TCSC Outstanding PhD Dissertation Award

2018

#### Matthew Malensek

University of San Francisco Faculty Development Fund (FDF) Award, \$2,500

### MAGICS Lab Startup Funding

2018

P. Intrevado, D. G. Brizan, **Matthew Malensek**



University of San Francisco Faculty Development Fund (FDF) Award, \$5,000

## Publications

---

Publications are listed in reverse chronological order. Preprint versions in PDF format are available online at:  
<https://www.cs.usfca.edu/~mmalensek/publications>

### Books and Book Chapters

- [B2] Peter Pacheco and **Matthew Malensek**. *An Introduction to Parallel Programming, 2nd Edition*. 2nd ed. Morgan Kaufmann, Mar. 2022. ISBN: 978-0128046050. DOI: 10.1016/C2015-0-01650-1. [ [PDF](#) · <http://ippbook.com/>].
- [B1] Sangmi Lee Pallickara, **Matthew Malensek**, and Shrideep Pallickara. "On the Processing of Extreme Scale Datasets in the Geosciences". *Handbook of Data Intensive Computing*. Ed. by Borko Furht and Armando Escalante. New York, NY: Springer New York, 2011, pp. 521–537. ISBN: 978-1-4614-1415-5. DOI: 10.1007/978-1-4614-1415-5\_20. [ [PDF](#)].

## Journal Articles

- [J16] Daniel Rammer, Thilina Buddhika, **Matthew Malensek**, Shrideep Pallickara, and Sangmi Lee Pallickara. "Enabling Fast Exploratory Analyses Over Voluminous Spatiotemporal Data Using Analytical Engines". *IEEE Transactions on Big Data* 8.1 (Feb. 2022), pp. 213–228. DOI: 10.1109/tbdata.2019.2939834. [[PDF](#) · Impact Factor: 4.11].
- [J15] Thilina Buddhika, **Matthew Malensek**, Shrideep Pallickara, and Sangmi Lee Pallickara. "Living on the Edge: Data Transmission, Storage, and Analytics in Continuous Sensing Environments". *ACM Transactions on Internet of Things* 2.3 (Aug. 2021), pp. 1–31. DOI: 10.1145/3450767. [[PDF](#) · Impact Factor: 3.135].
- [J14] Naman Shah, **Matthew Malensek**, Harshil Shah, Shrideep Pallickara, and Sangmi Lee Pallickara. "Scalable Network Analytics for Characterization of Outbreak Influence in Voluminous Epidemiology Datasets". *Concurrency and Computation: Practice and Experience* 31.7 (Oct. 2019), e4998. DOI: 10.1002/cpe.4998. [[PDF](#) · Impact Factor: 1.55].
- [J13] **Matthew Malensek**, Walid Budgaga, Ryan Stern, Shrideep Pallickara, and Sangmi Lee Pallickara. "Trident: Distributed Storage, Analysis, and Exploration of Multidimensional Phenomena". *IEEE Transactions on Big Data* 5.2 (June 2019), pp. 252–265. DOI: 10.1109/TBDATA.2018.2817505. [[PDF](#) · Impact Factor: 4.11].
- [J12] Thilina Buddhika, **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Synopsis: A Distributed Sketch over Voluminous Spatiotemporal Observational Streams". *IEEE Transactions on Knowledge and Data Engineering* 29.11 (Nov. 2017), pp. 2552–2566. ISSN: 1041-4347. DOI: 10.1109/TKDE.2017.2734661. [[PDF](#) · Impact Factor: 4.561].
- [J11] Walid Budgaga, **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "A Framework for Scalable Real-Time Anomaly Detection over Voluminous, Geospatial Data Streams". *Concurrency and Computation: Practice and Experience* 29.12 (Mar. 2017). DOI: 10.1002/cpe.4106. [[PDF](#) · Impact Factor: 1.55].
- [J10] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Hermes: Federating Fog and Cloud Nodes to Support Query Evaluations in Continuous Sensing Environments". *IEEE Cloud Computing* 4.2 (Mar. 2017), pp. 54–62. DOI: 10.1109/MCC.2017.26. [[PDF](#) · Impact Factor: 5.59].
- [J9] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Fast, Ad Hoc Query Evaluations over Multidimensional Geospatial Datasets". *IEEE Transactions on Cloud Computing* 5.1 (Jan. 2017), pp. 28–42. ISSN: 2168-7161. DOI: 10.1109/TCC.2015.2398437. [[PDF](#) · Impact Factor: 5.93].
- [J8] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Analytic Queries over Geospatial Time-Series Data Using Distributed Hash Tables". *IEEE Transactions on Knowledge and Data Engineering* 28.6 (June 2016), pp. 1408–1422. ISSN: 1041-4347. DOI: 10.1109/TKDE.2016.2520475. [[PDF](#) · Impact Factor: 4.561].
- [J7] Cameron Tolooee, **Matthew Malensek**, and Sangmi Lee Pallickara. "A Scalable Framework for Continuous Query Evaluations over Multidimensional, Scientific Datasets". *Concurrency and Computation: Practice and Experience* 28.8 (June 2016), pp. 2546–2563. DOI: 10.1002/cpe.3651. [[PDF](#) · Impact Factor: 1.55].
- [J6] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Autonomous Cloud Federation for High-Throughput Queries over Voluminous Datasets". *IEEE Cloud Computing* 3.3 (May 2016), pp. 40–49. DOI: 10.1109/MCC.2016.65. [[PDF](#) · Impact Factor: 5.59].
- [J5] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Minerva: Proactive Disk Scheduling for QoS in Multitier, Multitenant Cloud Environments". *IEEE Internet Computing* 20.3 (May 2016), pp. 19–27. DOI: 10.1109/MIC.2016.48. [[PDF](#) · Impact Factor: 1.92].
- [J4] Walid Budgaga, **Matthew Malensek**, Sangmi Pallickara, Neil Harvey, F. Jay Breidt, and Shrideep Pallickara. "Predictive Analytics Using Statistical, Learning, and Ensemble Methods to Support Real-time Exploration of Discrete Event Simulations". *Future Generation Computer Systems* 56.C (Mar. 2016), pp. 360–374. ISSN: 0167-739X. DOI: 10.1016/j.future.2015.06.013. [[PDF](#) · Impact Factor: 7.187].
- [J3] Zhiqian Sui, **Matthew Malensek**, Neil Harvey, and Shrideep Pallickara. "Autonomous Orchestration of Distributed Discrete Event Simulations in the Presence of Resource Uncertainty". *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 10.3 (Sept. 2015), 18:1–18:20. ISSN: 1556-4665. DOI: 10.1145/2746345. [[PDF](#) · Impact Factor: 2.48].

- [J2] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Evaluating Geospatial Geometry and Proximity Queries Using Distributed Hash Tables". *IEEE Computing in Science Engineering (CiSE)* 16.4 (July 2014), pp. 53–61. ISSN: 1521-9615. DOI: 10.1109/MCSE.2014.48. [[PDF](#) · Impact Factor: 1.71].
- [J1] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Exploiting Geospatial and Chronological Characteristics in Data Streams to Enable Efficient Storage and Retrievals". *Future Generation Computer Systems* 29.4 (June 2013), pp. 1049–1061. ISSN: 0167-739X. DOI: 10.1016/j.future.2012.05.024. [[PDF](#) · Impact Factor: 7.187].

## Refereed Conference Proceedings

- [C15] Sami N. Rollins, Alark Joshi, Xornam Apedoe, Sophie Engle, **Matthew Malensek**, and Gian Bruno. "Understanding Professional Identity Development Among Computer Science Students". *2021 ASEE Virtual Annual Conference Content Access*. <https://peer.asee.org/37962>. Virtual Conference: ASEE Conferences, July 2021. [[PDF](#)].
- [C14] Walid Budgaga, **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Concerto: Leveraging Ensembles for Timely, Accurate Model Training Over Voluminous Datasets". *2020 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)*. 2020, pp. 106–115. DOI: 10.1109/BDCAT50828.2020.00024. [[PDF](#)].
- [C13] Mingxin Lu, Edmund Wong, Daniel Barajas, Xiaochen Li, Mosopefoluwa Ogundipe, Nate Wilson, Pragya Garg, Alark Joshi, and **Matthew Malensek**. "AGAMI: Scalable Visual Analytics over Multidimensional Data Streams". *2020 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)*. 2020, pp. 57–66. DOI: 10.1109/BDCAT50828.2020.00020. [[PDF](#)].
- [C12] Alark Joshi, Gian Bruno, Xornam Apedoe, Sophie Engle, Sami Rollins, and **Matthew Malensek**. "Engendering Community to Computer Science Freshmen through an Early Arrival Program". *2020 ASEE Virtual Annual Conference Content Access*. <https://peer.asee.org/34545>. Virtual On line: ASEE Conferences, June 2020. DOI: 10.18260/1-2--34545. [[PDF](#)].
- [C11] Naman Shah, Harshil Shah, **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Network analysis for identifying and characterizing disease outbreak influence from voluminous epidemiology data". *Proceedings of the 2016 IEEE International Conference on Big Data*. Washington, D.C., USA, Dec. 2016, pp. 1222–1231. DOI: 10.1109/BigData.2016.7840726. [[PDF](#) · 18.68% Acceptance Rate].
- [C10] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Alleviation of Disk I/O Contention in Virtualized Settings for Data-Intensive Computing". *Proceedings of the 2015 IEEE/ACM 2nd International Symposium on Big Data Computing (BDC)*. Limassol, Cyprus, Dec. 2015, pp. 1–10. DOI: 10.1109/BDC.2015.32. [[PDF](#) · 16% Acceptance Rate].
- [C9] Jared Koontz, **Matthew Malensek**, and Sangmi Lee Pallickara. "GeoLens: Enabling Interactive Visual Analytics over Large-Scale, Multidimensional Geospatial Datasets". *Proceedings of the 2014 IEEE/ACM International Symposium on Big Data Computing (BDC)*. London, UK, Dec. 2014, pp. 35–44. DOI: 10.1109/BDC.2014.12. [[PDF](#) · 22% Acceptance Rate · **Best Paper Award**].
- [C8] **Matthew Malensek**, Walid Budgaga, Sangmi Pallickara, Neil Harvey, F. Jay Breidt, and Shrideep Pallickara. "Using Distributed Analytics to Enable Real-Time Exploration of Discrete Event Simulations". *Proceedings of the 2014 IEEE/ACM 7th International Conference on Utility and Cloud Computing*. London, UK, Dec. 2014, pp. 49–58. DOI: 10.1109/UCC.2014.13. [[PDF](#) · 19% Acceptance Rate].
- [C7] Cameron Tolooee, **Matthew Malensek**, and Sangmi Lee Pallickara. "A Framework for Managing Continuous Query Evaluations over Voluminous, Multidimensional Datasets". *Proceedings of the 2014 IEEE International Cloud and Autonomic Computing Conference (ICCAC)*. London, UK, Sept. 2014, pp. 73–82. DOI: 10.1109/ICCAC.2014.25. [[PDF](#)].
- [C6] **Matthew Malensek**, Sangmi Pallickara, and Shrideep Pallickara. "Polygon-Based Query Evaluation over Geospatial Data Using Distributed Hash Tables". *Proceedings of the 2013 IEEE/ACM 6th International Conference on Utility and Cloud Computing (UCC)*. Dresden, Germany, Dec. 2013, pp. 219–226. DOI: 10.1109/UCC.2013.46. [[PDF](#) · 24% Acceptance Rate].
- [C5] **Matthew Malensek**, Sangmi Pallickara, and Shrideep Pallickara. "Autonomously Improving Query Evaluations over Multidimensional Data in Distributed Hash Tables". *Proceedings of the 2013 ACM Cloud and*

*Autonomic Computing Conference (CAC)*. Miami, Florida, USA, Sept. 2013, 15:1–15:10. DOI: 10 . 1145 / 2494621 . 2494638. [[PDF](#) · 35% Acceptance Rate ].

- [C4] **Matthew Malensek**, Zhiquan Sui, Neil Harvey, and Shrideep Pallickara. "Autonomous, Failure-resilient Orchestration of Distributed Discrete Event Simulations". *Proceedings of the 2013 ACM Cloud and Autonomic Computing Conference (CAC)*. Miami, Florida, USA, Sept. 2013, 3:1–3:10. DOI: 10 . 1145 / 2494621 . 2494625. [[PDF](#) · 35% Acceptance Rate ].
- [C3] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Expressive Query Support for Multidimensional Data in Distributed Hash Tables". *Proceedings of the 2012 IEEE/ACM 5th International Conference on Utility and Cloud Computing (UCC)*. Chicago, Illinois, USA, Nov. 2012, pp. 31–38. DOI: 10 . 1109 / UCC . 2012 . 41. [[PDF](#) · 27% Acceptance Rate · **Best Paper Award** ].
- [C2] **Matthew Malensek**, Sangmi Lee Pallickara, and Shrideep Pallickara. "Galileo: A Framework for Distributed Storage of High-Throughput Data Streams". *Proceedings of the 2011 IEEE/ACM 4th International Conference on Utility and Cloud Computing (UCC)*. Melbourne, Australia, Dec. 2011, pp. 17–24. DOI: 10 . 1109 / UCC . 2011 . 13. [[PDF](#) · 26.7% Acceptance Rate ].
- [C1] Sangmi Lee Pallickara, **Matthew Malensek**, and Shrideep Pallickara. "Enabling access to timeseries, geospatial data for on-demand visualization". *IEEE Symposium on Large Data Analysis and Visualization, (LDAV)*. Providence, Rhode Island, USA, Oct. 2011, pp. 141–142. DOI: 10 . 1109 / LDAV . 2011 . 6092339. [[PDF](#) ].

## Ph.D. Dissertation and Master's Thesis

- [T2] **Matthew Malensek**. "Low-latency, Query-Driven Analytics over Voluminous Multidimensional, Spatiotemporal Datasets". PhD Dissertation. Colorado State University, July 2017. Advisors: Shrideep Pallickara and Sangmi Lee Pallickara. [[IEEE TCSC Outstanding Ph.D. Dissertation Award](#) ].
- [T1] **Matthew Malensek**. "On the Evaluation of Exact-match and Range Queries over Multidimensional Data in Distributed Hash Tables". MS Thesis. Colorado State University, Aug. 2012. Advisor: Shrideep Pallickara.

## Research

---

My research has produced software in the areas of big data, cloud computing, and distributed systems. These projects are released as open source software under permissive licenses to encourage community involvement and collaboration. A complete overview of my projects is available at: <https://www.cs.usfca.edu/~mmalensek/research>.

## MAGICS Lab

I am a co-founder and co-director of the USF MAGICS Lab: **M**achine Learning, **A**rtificial Intelligence, **G**eneral Intelligence, and **C**omputing at **S**cale. The MAGICS lab brings together USF and Bay Area students from a broad variety of backgrounds and meets weekly for student research presentations, paper reviews, discussions, and tutorials. Current lab members, projects, and publications can be viewed at: <http://magics.cs.usfca.edu>.

## Current Research Projects

### Agami: Scalable Streaming Visualization

Data exploration often involves producing visualizations, but if the dataset in question is too large, traditional visualization techniques and storage systems may be unable to cope with the extreme data volumes. Agami provides live-updating, interactive visualizations over streaming data.

<https://github.com/agami-viz>

### Minerva: Dynamic Disk Scheduling in Cloud Environments

Cloud environments typically provide resource isolation between virtual machines (VMs), but certain usage patterns by *noisy neighbors* can monopolize resources and have a negative impact on other VMs. Minerva leverages containerization functionality in the Linux kernel to balance and enforce fair scheduling policies for the cloud.

<https://www.cs.usfca.edu/~mmalensek/projects/minerva>

## Galileo: Management and Exploration of Extreme-Scale Spatiotemporal Data

Galileo is a scalable distributed storage framework for managing multidimensional datasets. I have continued development on the project at USF, with a recent focus on Internet-of-Things (IoT) and mobile devices.

<http://galileo.cs.colostate.edu>

## Forager: Distributed Execution Engine

Distributed execution engines often manage long-running, CPU-intensive computations. However, large quantities of short-term, stateless computations require a different set of considerations. Forager is part of the *Granules* cloud runtime, with a focus on scheduling and managing small tasks.

<http://granules.cs.colostate.edu>

## Teaching

---

I primarily teach systems courses in the USF CS department. I created the department's *CS 677: Big Data* course offering, helped develop the curriculum for and teach the first iterations of our *CS 521: Systems Programming* bridge course, and am the primary instructor for *CS 326: Operating Systems*.

A complete listing of my courses is available at: <https://www.cs.usfca.edu/~mmalensek/teaching>. All my course materials, including slides, notes, videos, projects, lab assignments, and test cases are publicly available.

## Courses Taught at the University of San Francisco

<b>CS 326</b> · Operating Systems	Fall 2024 · Enrollment: 32
<b>CS 490</b> · Senior Team Project	Fall 2024 · Enrollment: 15
<b>CS 326</b> · Operating Systems	Spring 2024 · Enrollment: 30
<b>CS 521</b> · Systems Programming	Spring 2023 · Enrollment: 20
<b>CS 677</b> · Big Data	Spring 2023 · Enrollment: 21
<b>CS 286</b> · Special Topics: Peer-Led Seminar in Tech ( <i>co-teacher</i> )	Fall 2022 · Enrollment: 09
<b>CS 326</b> · Operating Systems	Fall 2022 · Enrollment: 34
<b>CS 677</b> · Big Data	Fall 2022 · Enrollment: 17
<b>CS 326</b> · Operating Systems	Spring 2022 · Enrollment: 33
<b>CS 521</b> · Systems Programming	Spring 2022 · Enrollment: 21
<b>CS 326</b> · Operating Systems	Fall 2021 · Enrollment: 30
<b>CS 386</b> · Special Topics: Mentoring and Professional Identity Development	Fall 2021 · Enrollment: 07
<b>CS 677</b> · Big Data	Fall 2021 · Enrollment: 26
<b>CS 326</b> · Operating Systems ( <b>two sections</b> )	Spring 2021 · Enrollment: 51
<b>CS 521</b> · Systems Programming	Spring 2021 · Enrollment: 19
<b>CS 326</b> · Operating Systems	Fall 2020 · Enrollment: 28
<b>CS 677</b> · Big Data	Fall 2020 · Enrollment: 31
<b>CS 326</b> · Operating Systems ( <b>two sections</b> )	Spring 2020 · Enrollment: 64
<b>CS 326</b> · Operating Systems	Fall 2019 · Enrollment: 32
<b>CS 677</b> · Big Data	Fall 2019 · Enrollment: 35
<b>CS 326</b> · Operating Systems ( <b>two sections</b> )	Spring 2019 · Enrollment: 51
<b>CS 326</b> · Operating Systems	Fall 2018 · Enrollment: 34
<b>CS 677</b> · Big Data	Fall 2018 · Enrollment: 31
<b>CS 220</b> · Introduction to Parallel Computing	Spring 2018 · Enrollment: 33
<b>CS 326</b> · Operating Systems	Spring 2018 · Enrollment: 34
<b>CS 220</b> · Introduction to Parallel Computing	Fall 2017 · Enrollment: 33
<b>CS 686</b> · Special Topics in Big Data	Fall 2017 · Enrollment: 31

## Open Source Teaching Tools

I aim to automate as much of the administrative tasks in my courses as possible so that I can focus on what matters: one-on-one help, developing students' skills, and providing tailored feedback. To this end, I have developed the following teaching-oriented software:

### Automated Testing and Grading Framework

I provide test cases for projects that can be run by students to evaluate their work and automatically submit grades to Canvas. Tests are run in a secure, sanitized container on a remote server. The framework also provides support for *checkpoints* to encourage students to start working early and keep making progress through the project's duration.

<https://github.com/malensek/cowtest>

### Course Site Generator

The websites for my courses are generated by software that produces an assignment schedule, syllabus, list of upcoming deadlines, etc. from text files. Since many students use Canvas in their other classes, all the materials are synced to Canvas so students can access them in either location.

<https://github.com/malensek/wake>

### Github Classroom Utility Scripts

My courses use the industry-standard `git` version control system extensively, and student source code is hosted on *Github*. I have developed a set of scripts that automatically back up project repositories, enable guided reviewing and grading of submissions, and uploading feedback.

<https://github.com/malensek/classroom-utils>

## Directed Reading and Research

The USF CS department gives students the opportunity to participate in directed studies with the *CS 498 / CS 698: Directed Reading and Research* course. Students can work on a variety of semester-long research projects for course credit. Often these courses are spun off to form the basis for fully-fledged research projects and publications.

I have supervised the following instances of CS 498 / CS 698.



SolaRaft  
 Distributed Geopresence  
 Streaming Data Visualization  
 Geospatial DHT in Go  
 Efficient Geopresence Evaluation II  
 Efficient Geopresence Evaluation  
 Implementing Containers in C  
 Exploring MapReduce  
 Network Security Research  
 Geospatial Query Processing II  
 File Transfer Client  
 Geospatial Query Processing  
 Streaming Big Data Visualization  
 Non-Blocking TCP Server Performance Evaluation  
 Machine Learning Principles  
 Stock Portfolio Robo Advisor  
 GPU Geospatial Queries  
 Exploring MapReduce  
 Streaming Summaries with MapReduce  
 Solar File System  
 TensorFlow News  
 Smart City IoT Infrastructure  
 Disease Outbreak Simulation  
 Blockchain Implementation  
 Analytics in Meteorological Domains

Andrew Molis, Anh Nguyen · Spring 2023  
 Jason McGowan · Spring 2023  
 Ziyang Liu, Chuxi Wang, Yiqi Wei · Spring 2022  
 Jie Liu · Spring 2021  
 Ryan Dielhenn · Fall 2020  
 Ryan Dielhenn · Summer 2020  
 Chris Smith · Spring 2020  
 Chris Smith · Fall 2019  
 Daiya Masuda · Summer 2019  
 Davin Jimenez · Summer 2019  
 Pedram Namiranian · Spring 2019  
 Gabriel Cisneros · Spring 2019  
 Pragya Garg, Nate Wilson · Spring 2019  
 Chirag Jain · Spring 2019  
 Olivia Kumar · Spring 2019  
 Jinyue Song · Fall 2018  
 Michelle Dong · Fall 2018  
 Rugved Mavidipalli · Summer 2018  
 Chirag Jain · Summer 2018  
 Shengcai Cheng · Spring 2018  
 Xue Kang, Liang Wang · Spring 2018  
 Bhargavi Kommineni, Rushabh Shah · Spring 2018  
 Jiali Ding · Spring 2018  
 Priyam Jemesh Patel · Spring 2018  
 Chirag Jain · Fall 2017

## Professional & Academic Service

I maintain active involvement in big data, cloud computing, and distributed computing conferences and serve as a reviewer for related journals.

### IEEE/ACM International Symposium on Cluster, Cloud, and Internet Computing (CCGRID)

2024

Program Committee Member

Philadelphia, USA

### IEEE International Conference on Cloud Computing in Emerging Markets

2022

Student Project Showcase Chair

Zurich, Switzerland

### IEEE International Conference on Cloud Computing in Emerging Markets

2021

Technical Program Committee Member

Somerset, New Jersey, USA

### IEEE/ACM International Conference on Utility and Cloud Computing

2018

Tutorials Co-chair

Zurich, Switzerland

### IEEE/ACM International Conference on Utility and Cloud Computing

2016

Tutorials Co-chair

Shanghai, China



## IEEE/ACM International Conference on Utility and Cloud Computing

Session Chair

2013

Dresden, Germany

### Journal Reviewer

I regularly review manuscripts for journals, including:

- *Nature Scientific Data*
- International Journal of Geographical Information Science
- IEEE Transactions on Systems, Man, and Cybernetics: Systems
- IEEE Transactions on Cloud Computing
- IEEE Cloud Computing Magazine
- Future Generation Computer Systems

### Professional Memberships

- Association for Computing Machinery (ACM)
- Institute of Electrical and Electronics Engineers (IEEE)

## University Service

### Educational Technology Services Advisory Board

University-wide board that works with the ETS Team to advise on technologies that support teaching and learning at the University. The board evaluates everything from classroom layouts and technology to software such as learning management systems.

Fall 2018 – Spring 2019

### ETS Discussion Software Working Group

Evaluated several online discussion tools for classroom use. Built a rubric for evaluation and selected software based on usability, support for sharing code and math equations, mobile apps, polling, etc.

Fall 2021

### Faculty Adviser, USF Indian Student Organization (ISO)

Serving as an adviser and helping with organization logistics. The ISO is one of the larger and more active student organizations at USF.

2020 – 2023

## College Service

### Engineering Curriculum Committee

This committee was tasked with developing the curriculum for a new engineering program at USF. One of the most crucial aspects of this endeavor was building a program that would be well-aligned with USF's core strengths. As part of the committee, I traveled to the Franklin W. Olin College of Engineering in Needham, Massachusetts to observe the project-based and student-centered learning techniques employed there. My efforts within the committee were focused on the Computer Engineering and Electrical Engineering concentrations for the program, particularly building a nationally-competitive curriculum that would also leverage existing USF courses.

Spring 2018 – Spring 2019

The College of Arts and Sciences Engineering Program is now live, here:

<https://www.usfca.edu/arts-sciences/undergraduate-programs/engineering>

## Departmental Service

## CS Major Adviser

Each CS student is assigned a major adviser who helps with creating a degree plan, registering for classes, and providing general advice. Summary of students advised per semester:

- Spring 2024 **125**
- Fall 2023 **26**
- Spring 2023 **37**
- Fall 2022 **35**
- Spring 2022 **38**
- Fall 2021 **43**
- Spring 2021 **48**
- Fall 2020 **47**
- Spring 2020 **48**
- Fall 2019 **51**
- Spring 2019 **35**
- Fall 2018 **27**

2018 – Present

## Webtrack Advising

Participated in *Webtrack* during the summer to welcome and advise incoming USF CS students.

Summer 2022

## Co-Investigator, Community Engaged Scholars in Computer Science

Funded via S-STEM NSF grant DUE-1833718, the **CES|CS** program combines scholarships with community-building activities such as field trips to tech companies, guest speakers, faculty and alumni mentoring, social events, and coursework. As a faculty mentor in the program, I met with my advisees weekly. I also managed the CES|CS Alumni Mentoring Program, which paired participating students with USF CS alumni mentors from industry.

March 2019–February 2024

CES|CS program website: <https://scholars.cs.usfca.edu>

## Director, MAGICS Lab

While the USF CS MAGICS lab (Machine Learning, Artificial Intelligence, Game Intelligence, and Computing at Scale) is primarily a research lab, our weekly meetings also serve as a student organization where new CS students can find guidance from faculty and peers, join projects to collaborate on, and share ideas. Many of our talks provide tutorials for using a variety of machine learning or big data tools that students may not have used in their courses. Lab website: <http://magics.cs.usfca.edu>.

2018 – Present

## Faculty Adviser, USF ACM Chapter

The USF CS ACM Chapter hosts invited talks and helps coordinate student events such as CS Night and the annual CS picnic.

2018 – Present

## Organizer, Jeremy Kerby Hackathon

With a generous donation by USF CS alumni Jeremy Kerby, the USF CS ACM club hosted a competitive hackathon for junior and senior students that focused on interview-style coding problems. Beyond co-organizing this event, I also developed coding challenges and served as one of the judges.

Spring 2019

### **2019 Tenure-Track Faculty Search Committee**

Served as a member of our department's committee to hire a new tenure track faculty member. Helped review over 100 applications, interview prospective candidates, coordinate in person visits, and submit final decision to dean's office.

*Fall 2018 – Spring 2019*

### **Newmark Funding Committee**

Served as a member of the committee for allocating scholarship funds for students to attend the Grace Hopper Celebration, Tapia, and other conferences donated by Newmark Philanthropies.

*2018, 2020, 2021*

### **Newmark Tuition Scholarship Committee**

Served as a member of the committee for allocating tuition scholarship funds donated by Newmark Philanthropies.

*2022*

### **Faculty Representative, Grace Hopper Celebration**

Traveled to Houston with USF students for GHC '18. With a generous gift from Craig Newmark Philanthropies, our department was able to fully fund the travel and conference registration of 32 students.

*[Houston, Texas](#)*

*2018*

### **Assessment Reports**

I have provided data from my courses for CS assessment reports, which included syllabi, project descriptions and rubrics, as well as grade breakdowns.

*2018, 2021*

- CS 521 – Systems Programming (Fall 2021)
- CS 677 – Big Data (Fall 2018)

### **MSCS Application Review**

I review around 5-10 MSCS applications (transcripts, test scores, personal statements, etc.) per academic year.

*2017 – Present*