

DSCI 550: Data Science at Scale

Class Info

Spring Semester, 2022

Location : OHE 100D and online

Time: Th 3:30-6:50pm

Class number: 32413D

Class number: 32426D (DEN)

Website: <https://sites.google.com/view/dsci-550-spring2022/home>

Instructors



Dr. Keith Burghardt

<https://www.isi.edu/directory/keithab/>

E-Mail: keithab@isi.edu

Office Hours: 2:30pm-3:30pm (right before class)



Dr. Goran Muric

<https://www.isi.edu/directory/gmuric/>

E-Mail: gmuric@isi.edu

Office Hours: Thu 2:30pm-3:30pm (right before class) on Zoom:

<https://usc.zoom.us/j/91571066080>

Teaching Assistant



Bahareh Harandizadeh

<https://www.isi.edu/directory/harandiz>

E-Mail: harandiz@usc.edu

DSCI 550 Overview

This course is designed as an overview course to give students a broad understanding of Informatics topics for Big Data and to get practical experience with key Big Data informatics techniques. Topics include roadmap of informatics, the data lifecycle, the role of the data scientist, and analyzing and exploring Big Data with real world use cases in data analytics, and big data. Understanding Big Data involves understanding of digital file formats, their detection and data extraction from them. Emphasis areas include Document Type Detection; Parsing and extraction; Metadata understanding and analysis; Language Identification and detection from files and finally file formats and representation. The class also has a specific focus on Content Detection and Analysis from large data sets. Datasets used in the course are publicly collected by the instructor or his collaborators involved in national Big Data initiatives including DARPA, NASA and other projects. The course is designed to be accessible to students with experience programming in Python and Java at an intermediate level. The course will introduce the students to topical software frameworks that deal with Big Data including Tika, Solr, Elasticsearch™, TensorFlow, Nutch and Apache Hadoop™. The course will be a combination of lecture, in-class discussion, readings, group-based assignments and a final exam.

The objective of this course is to train students to be able to understand Big Data and Large Data Environments, e.g., file formats, their representation, and how to automatically extract information from large datasets of files. Specifically, students successfully completing this course will achieve three main objectives:

1. Develop sufficient proficiency in Big Data frameworks to write software capable of automatically extracting information from data including its text and metadata and language.
2. Develop sufficient proficiency in techniques with Large Data sets collected from the Web and other places (Intranet, Science Data Sets, Public Data Sets).
3. Develop sufficient proficiency in Python and Java to write and execute software that is “File Aware” and that automatically extracts text and metadata from large data sets.

The primary teaching methods will be discussion, case studies, and lectures. Students are expected to perform directed self learning outside of class which encompasses, among other things, a considerable amount of literature review. Leadership training in open source is provided and encouraged, and students leave with an experience in open source that makes them more marketable to companies and institutions looking to hire in Big Data, and Data Science.

In addition to foundations, and practical experience with Big Data and Data Science, the class will also introduce the student to the state-of-the-art in content detection research, future trends and state-of-the-practice. Students are expected to attend class regularly, and participate (as directed) in all class discussions, and most importantly, have fun!

USC ACADEMIC INTEGRITY

Statement on Academic Conduct and Support Systems

Academic Conduct Plagiarism - presenting someone else's ideas as your own, either verbatim or recast in your own words - is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in SCampus in [Section 11, Behavior Violating University Standards](#). Other forms of academic dishonesty are equally unacceptable. See additional information in [SCampus and university policies on scientific misconduct](#). Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the [Office for Equity, Equal Opportunity, and Title IX](#) or to the [Department of Public Safety](#). This is important for the safety of the USC community. Another member of the university community - such as a friend, classmate, advisor, or faculty member - can help initiate the report, or can initiate the report on behalf of another person. The [Sexual Violence Prevention & Services](#) provides 24/7 confidential support, and the sexual assault resource center webpage sarc@usc.edu describes reporting options and other resources.

Support Systems

A number of USC's schools provide support for students who need help with scholarly writing. Check with your advisor or program staff to find out more. Students whose primary language is not English should check with the [American Language Institute](#) which sponsors courses and workshops specifically for international graduate students. The [Office of Disability Services and Programs](#) provides certification for students with disabilities and helps arrange the relevant accommodations. If an officially declared emergency makes travel to campus infeasible, [USC Emergency Information](#) will provide safety and other updates, including ways in which instruction will be continued by means of blackboard, teleconferencing, and other technology.

Statement on Diversity

The diversity of the participants in this course is a valuable source of ideas, problem solving strategies, and engineering creativity. We encourage and support the efforts of all of our students to contribute freely and enthusiastically. We are members of an academic community where it is our shared responsibility to cultivate a climate where all students and individuals are valued and where both they and their ideas are treated with respect, regardless of their differences, visible or invisible.

TEXTBOOK

Chris A. Mattmann, and Jukka Zitting. *Tika in Action*, 256 pages. New York: Manning Publications, November 2011. ISBN: 9781935182856.

C. Mattmann. *Machine Learning with TensorFlow: 2nd Edition*. 456 pages. New York: Manning Publications, December 2020. ISBN 9781617297717

ASSIGNMENTS and EXAMINATIONS

Name	Description	Weight
Exam	An exam testing your understanding of the lecture materials	25%
Assignments	Assignments where you will build on the Big Data and Data Science topics in course and make a contribution to one of the existing open source frameworks (Tika, Solr, Nutch, TensorFlow, OODT, etc.).	45%
Individual Presentation	An individual presentation demonstrating the student's understanding of one of the required paper readings in the course.	25%
Participation	Participation in lectures, by asking questions and contributing to the conversation. Attending lectures (physically and remotely) and positively contributing to the class experience.	5%

Project Submission Guidelines

Submission guidelines will be specified in each assignment.

Schedule (subject to change; check regularly)

Week	Lecture Topic	Assigned Readings	Assignments & Deadlines
1	Course Introduction	Tika in Action, Chapter 1	Resources:
Jan 13th, 2022	Introduction to Big Data	Mattmann, Chris. A vision for data science. Nature, Vol. 493, No. 7433, pp. 473-475, January 24, 2013.	DARPA I2O (DARPA Dan) Video
	DARPA XDATA Program - Overview Slides	Lynch, Clifford. "Big data: How do your data grow?." Nature 455.7209 (2008): 28-29.	Zero Dark Thirty Workbench Video
	Breakout Groups on Big Data	Howe, Doug, et al. "Big data: The future of biocuration." Nature 455.7209 (2008): 47-50. Wigan, Marcus R., and Roger Clarke. "Big data's big unintended consequences." Computer 46.6 (2013): 46-53. Schwartz, J. A. N. A., et al. "Measuring the value of Big Data exploitation systems: Quantitative, non-subjective metrics with the user as a key component." Parsons Journal for Information Mapping 6 (2014): 1-12. Sotera Defense Solutions. A Survey of Big Data Methods, Assessments, and Approaches. November 2012 De Mauro, Andrea, Marco Greco, and Michele Grimaldi. "What is big data? A consensual definition and a review of key research topics." AIP conference proceedings. Vol. 1644. No. 1. AIP, 2015.	
2	Report out from Big Data Breakouts	Tika in Action, Chapter 2	
Jan 20th, 2022	A Taxonomy of File Formats	Crocker, David. RFC 822 "Standard for the format of ARPA Internet text messages." (1982).	
	Content Detection Libraries	Freed, Ned and Nathaniel Borenstein. RFC 1341. MIME (Multipurpose Internet Mail Extensions). Mechanisms for Specifying and Describing the Format of Internet Message Bodies. June 1992.	
	Language Bindings for Apache Tika	Freed, Ned, and Nathaniel Borenstein. RFC 2045. Multipurpose internet mail extensions (MIME) part one: Format of internet message bodies. 1996.	
	Individual Student Presentations - Week 1 Papers	Freed, Ned, and Nathaniel Borenstein. RFC 2046 Multipurpose internet mail extensions (MIME) part two: Media types, November, 1996. Freed, Ned. RFC 2048 "Multipurpose internet mail extensions (MIME) part four: Registration procedures." ISI (1996). Hicks, Ben J., et al. "Organizing and managing personal electronic files: A mechanical engineer's perspective." ACM Transactions on Information Systems (TOIS) 26.4 (2008): 23. Shim, Jungwon Roy. "Arium: Beyond the Desktop Metaphor: A new way of navigating, searching, and organizing personal digital data." Masters Thesis, Carnegie Mellon University (2012). Crowder, Jerome, Jonathan Marion, and Michele Reilly. "File Naming in Digital Media Research: Examples from the Humanities and Social Sciences." Journal of Librarianship and Scholarly Communication 3.3 (2015). Jackson, Andrew N. "Formats over time: Exploring UK web history." arXiv preprint arXiv:1210.1714 (2012).	
3	Report outs from the in class discussion around classifying files and the MIME taxonomy	Tika in Action, Chapter 3	Resources:
Jan 27th, 2022		Bik, Elisabeth M., Casadevall, Arturo, Fang, Ferrie C. The Prevalence of Inappropriate Image Duplication in Biomedical Research Publications. Manku, Gurmeet Singh, Arvind Jain, and Anish Das Sarma. "Detecting near-duplicates	

	Document Similarity and Deduplication	for web crawling." Proceedings of the 16th international conference on World Wide Web. ACM, 2007.	The MIME guys: How two Internet gurus changed e-mail forever
	Individual Presentations - Week 2 Papers	Henzinger, Monika. "Finding near-duplicate web pages: a large-scale evaluation of algorithms." Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2006.	IEEE Computer Interview with Nathaniel Bornstein
		Cooper, Matthew, Jonathan Foote, and Andreas Girgensohn. "Automatically organizing digital photographs using time and content." Image Processing, 2003. ICIP 2003. Proceedings. 2003 International Conference on. Vol. 3. IEEE, 2003.	
		Manber, Udi. "Finding similar files in a large file system." Usenix Winter. Vol. 94. 1994.	
		Chim, Hung, and Xiaotie Deng. "Efficient phrase-based document similarity for clustering." IEEE Transactions on Knowledge and Data Engineering 20.9 (2008): 1217-1229.	
4	Remaining presentations from last week (if any)	Tika in Action, Chapter 4	Resources:
Feb 3rd, 2022	Document Type Detection	Amirani, Mehdi Chehel, Mohsen Toorani, and A. Beheshti. A new approach to content-based file type detection. Computers and Communications, 2008. ISCC 2008. IEEE Symposium on. IEEE, 2008.	The Economist - Digital Bit Rot
	Individual Presentations - Week 3 papers	McDaniel, Mason, and M. Hossain Heydari. Content based file type detection algorithms. System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on. IEEE, 2003.	Bit Rot - on Digital Vellum (by Vint Cerf, Google) TEDx Talk
	Advanced File System Statistics and Understanding	Alamri, Nasser S., and William H. Allen. "A comparative study of file-type identification techniques." SoutheastCon 2015. IEEE, 2015.	
		Li, Wei-Jen, et al. "Fileprints: Identifying file types by n-gram analysis." Information Assurance Workshop, 2005. IAW'05. Proceedings from the Sixth Annual IEEE SMC. IEEE, 2005.	
		Shahi, Ashim. "Classifying the classifiers for file fragment classification." Masters Thesis, Universiteit van Amsterdam (2012).	
		Ahmed, Irfan, et al. "Fast file-type identification." Proceedings of the 2010 ACM Symposium on Applied Computing. ACM, 2010.	
		Pierris, Georgios, and Stilianos Vidalis. "Forensically classifying files using HSOM algorithms." Emerging Intelligent Data and Web Technologies (EIDWT), 2012 Third International Conference on. IEEE, 2012.	
		Harris, Ryan M. "Using artificial neural networks for forensic file type identification." Master's Thesis, Purdue University (2007).	
		Douceur, John R., and William J. Bolosky. A large-scale study of file-system contents. ACM SIGMETRICS Performance Evaluation Review 27.1 (1999): 59-70.	
5	Introduction to Assignment 1	Tika in Action, Chapter 5	Resources:
Feb 10th, 2022	Content Extraction	Kilicoglu, Halil, et al. "Semantic MEDLINE: a web application for managing the results of PubMed Searches." Proceedings of the third international symposium for semantic mining in biomedicine. Vol. 2008. 2008.	Introduction to Information Retrieval - Chris Mattmann
	Individual Presentations - Week 4 papers	Kobayashi, Mei, and Koichi Takeda. "Information retrieval on the web." ACM Computing Surveys (CSUR) 32.2 (2000): 144-173.	
		Voorhees, Ellen M., and Donna Harman. "Overview of the sixth text retrieval conference (TREC-6)." Information Processing & Management 36.1 (2000): 3-35.	
		Arasu, Arvind, and Hector Garcia-Molina. Extracting structured data from web pages. Proceedings of the 2003 ACM SIGMOD international conference on Management of data. ACM, 2003.	

		Lewandowski, Dirk. "Web searching, search engines and Information Retrieval." <i>Information Services & Use</i> 25.3, 4 (2005): 137-147.	
		Weninger, Tim, William H. Hsu, and Jiawei Han. "CETR: content extraction via tag ratios." <i>Proceedings of the 19th international conference on World wide web.</i> ACM, 2010.	
		Karpathy, Andrej, and Li Fei-Fei. "Deep visual-semantic alignments for generating image descriptions." <i>Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.</i> 2015.	
6	Assignment 1 questions	Tika in Action, Chapter 6	Resources:
Feb 17th, 2022	Last Week Tonight: Robocalls	Gowda, Thamme, and Chris A. Mattmann. "Clustering Web Pages Based on Structure and Style Similarity (Application Paper)." <i>Information Reuse and Integration (IRI), 2016 IEEE 17th International Conference on.</i> IEEE, 2016.	Last Week Tonight - John Oliver (HBO) - Robocalls
	RoboKiller Video		New App to stop Robocalls - RoboKiller
	Understanding Metadata	Anquetil, Nicolas, and Timothy Lethbridge. <i>File clustering using naming conventions for legacy systems.</i> Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research. IBM Press, 1997.	
	Individual Presentations - week 5 papers	Swierk, Edward, et al. "The Roma personal metadata service." <i>Mobile Networks and Applications</i> 7.5 (2002): 407-418.	
		Karypis, Michael Steinbach George, Vipin Kumar, and Michael Steinbach. "A comparison of document clustering techniques." <i>KDD workshop on Text Mining.</i> 2000.	
		Marchionini, Gary. "Exploratory search: from finding to understanding." <i>Communications of the ACM</i> 49.4 (2006): 41-46.	
7	Video - Understanding Metadata TedX talk	Tika in Action, Chapter 7	Resources:
Feb 24th, 2022	Information Clustering	Koehn, Philipp, et al. "Moses: Open source toolkit for statistical machine translation." <i>Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions.</i> Association for Computational Linguistics, 2007.	The Power of Metadata: Deepak Jagdish and Daniel Smilkov at TEDxCambridge 2013
	Example - Machine Learning with TensorFlow - Activity Clustering (Android)	Post, Matt, et al. "Joshua 5.0: Sparser, better, faster, server." <i>Proceedings of the Eighth Workshop on Statistical Machine Translation.</i> 2013.	
	Week 6 Individual Presentations	Lins, Rafael Dueire, and Paulo Gonçalves. <i>Automatic language identification of written texts.</i> Proceedings of the 2004 ACM symposium on Applied computing. ACM, 2004.	
		Papineni, Kishore, et al. "BLEU: a method for automatic evaluation of machine translation." <i>Proceedings of the 40th annual meeting on association for computational linguistics.</i> Association for Computational Linguistics, 2002.	
		Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. "Neural machine translation by jointly learning to align and translate." <i>arXiv preprint arXiv:1409.0473</i> (2014).	
		Tromp, Erik, and Mykola Pechenizkiy. "Graph-based n-gram language identification on short texts." <i>Proc. 20th Machine Learning conference of Belgium and The Netherlands.</i> 2011.	
		Lopez-Moreno, Ignacio, et al. "Automatic language identification using deep neural networks." <i>Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on.</i> IEEE, 2014.	
		Bertoldi, Nicola, et al. "MMT: New open source MT for the translation industry." <i>Proceedings of The 20th Annual Conference of the European Association for Machine Translation (EAMT).</i> 2017.	
8	Exam Review	Tika in Action, Chapter 8	Resources:
March 3rd, 2022	Assignment 2	Tjong Kim Sang, Erik F., and Fien De Meulder. "Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition." <i>Proceedings</i>	Forensic Linguistic Profiling & What Your Language Reveals About

	Video - Linguistic Forensics	of the seventh conference on Natural language learning at HLT-NAACL 2003-Volume 4. Association for Computational Linguistics, 2003.	You Harry Bradford TEDxStoke
	Language Identification	Nadeau, David, and Satoshi Sekine. "A survey of named entity recognition and classification." <i>Linguisticae Investigationes</i> 30.1 (2007): 3-26.	The Shallowness of Google Translate - The Atlantic - Douglas Hofstadter
	Week 7 Individual Presentations	Ritter, Alan, Sam Clark, and Oren Etzioni. "Named entity recognition in tweets: an experimental study." <i>Proceedings of the Conference on Empirical Methods in Natural Language Processing</i> . Association for Computational Linguistics, 2011.	Apache cTAKES
		Mattmann, Chris A., and Madhav Sharan. "An automatic approach for discovering and geocoding locations in domain-specific web data." <i>Proceedings of the 2016 IEEE 17th International Conference on Information Reuse and Integration (IRI'16)</i> . 2016.	Apache UIMA Apache OpenNLP
		Khodak, Mikhail, Nikunj Saunshi, and Kiran Vodrahalli. "A Large Self-Annotated Corpus for Sarcasm." <i>arXiv preprint arXiv:1704.05579</i> (2017).	Stanford Core NER/NLP NLTK
		Hutto, Clayton J., and Eric Gilbert. "Vader: A parsimonious rule-based model for sentiment analysis of social media text." <i>Eighth international AAAI conference on weblogs and social media</i> . 2014.	
		Geyer, Kelly, et al. "Named Entity Recognition in 140 Characters or Less." <i># Microposts</i> . 2016.	
9	Exam		Resources:
March 10th, 2022	Spillover lecture - Language Identification		How to Save the World Wayne Moses Burke TEDxBeaconStreet
	Individual Presentations		
10	No classes today. Spring recess.	No classes today. Spring recess.	
March 17th, 2022			
11	Individual Presentations	Tika in Action, Chapter 9	Resources:
March 24th, 2022	Discussion on Named Entity Recognition	Dean, Jeffrey, and Sanjay Ghemawat. <i>MapReduce: simplified data processing on large clusters</i> . <i>Communications of the ACM</i> 51.1 (2008): 107-113.	Apache Distributed Release Audit Tool (DRAT)
	Hadoop Spark and Tika: Large Scale Content Detection and Analysis	Zaharia, Matei, et al. <i>Spark: cluster computing with working sets</i> . <i>Proceedings of the 2nd USENIX conference on Hot topics in cloud computing</i> . Vol. 10. 2010.	DRAT Video
		Elsayed, Tamer, Jimmy Lin, and Douglas W. Oard. "Pairwise document similarity in large collections with MapReduce." <i>Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers</i> . Association for Computational Linguistics, 2008.	
		M. Bernaschi, M. Cianfriglia, A. Di Marco, A. Sabellico, G. Me, G. Carbone, G. Totaro. <i>Forensic Disk Image Indexing and Search in an HPC environment</i> . <i>IEEE International Conference on High Performance Computing & Simulation (HPCS)</i> , 2014.	
		Meusel, Robert, Peter Mika, and Roi Blanco. "Focused crawling for structured data." <i>Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management</i> . ACM, 2014.	
		Niu, Feng, et al. "DeepDive: Web-scale Knowledge-base Construction using Statistical Learning and Inference." <i>VLDS 12</i> (2012): 25-28.	
		Mattmann, C. A., Oh, J. H., Palsulich, T., McGibbney, L. J., Gil, Y., & Ratnakar, V. (2015, November). <i>DRAT: An Unobtrusive, Scalable Approach to Large Scale Software License Analysis</i> . In <i>Automated Software Engineering Workshop (ASEW), 2015 30th IEEE/ACM International Conference on</i> (pp. 97-101). IEEE.	
12	Assignment 3 - Introduction	Tika in Action, Chapter 10	Resources:
March 31st, 2022			TedXTalk Analyzing and modeling complex and big

	Readout - Named Entity Recognition Group Presentations	Bialecki, Andrzej, et al. "Apache lucene 4." SIGIR 2012 workshop on open source information retrieval. 2012.	data Professor Maria Fasli TEDxUniversityofEssex
	Open Source Content Detection Technologies	Turtle, Howard, Yatish Hegde, and S. Rowe. "Yet another comparison of lucene and indri performance." SIGIR 2012 Workshop on Open Source Information Retrieval. 2012.	Textract Scrapy
	Individual Presentations	Bontcheva, Kalina, et al. "TwitIE: An Open-Source Information Extraction Pipeline for Microblog Text." RANLP. 2013. Cunningham, Hamish. "GATE, a general architecture for text engineering." Computers and the Humanities 36.2 (2002): 223-254. Atserias, Jordi, et al. "FreeLing 1.3: Syntactic and semantic services in an open-source NLP library." Proceedings of LREC. Vol. 6. 2006. Manning, Christopher D., et al. "The stanford corenlp natural language processing toolkit." ACL (System Demonstrations). 2014. Savova, Guergana K., et al. "Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): architecture, component evaluation and applications." Journal of the American Medical Informatics Association 17.5 (2010): 507-513.	Mattmann ApacheCon 2015 Content Talk
13	Ted Talk	Tika in Action, Chapter 11	Resources:
April 7th, 2022	Evaluating Content Detection	Nowell, Lucy Terry, et al. "Visualizing search results: some alternatives to query-document similarity." Proceedings of the 19th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 1996.	This is What happens when you reply to SPAM e-Mail. - James Veitch - Ted Talk
	Walkthrough of Polar Deep Insights		
	Individual Presentations	Shneiderman, Ben. "The eyes have it: A task by data type taxonomy for information visualizations." Visual Languages, 1996. Proceedings., IEEE Symposium on. IEEE, 1996. Gottron, Thomas. "Evaluating content extraction on HTML documents." Proceedings of the 2nd International Conference on Internet Technologies and Applications (ITA'07). 2007. Leuski, Anton. "Evaluating document clustering for interactive information retrieval." Proceedings of the tenth international conference on Information and knowledge management. ACM, 2001. Bailey, Peter, et al. "Evaluating search systems using result page context." Proceedings of the third symposium on Information interaction in context. ACM, 2010.	USC IRDS Polar Data Science website
14	Group Readouts on Evaluating Content Detection and Analysis	Palamuttam, Rahul, et al. "SciSpark: Applying in-memory distributed computing to weather event detection and tracking." Big Data (Big Data), 2015 IEEE International Conference on. IEEE, 2015.	Resources:
April 14th, 2022	Lecture on NoSQL	Leavitt, Neal. "Will NoSQL databases live up to their promise?." Computer 43.2 (2010).	2019 MIT Citi - 3 - Top Ten Big Data Blunders - Michael Stonebraker (Only watch first 15m)
	Lecture on SciSpark		
	Individual Presentations	Stonebraker, Michael. "SQL databases v. NoSQL databases." Communications of the ACM 53.4 (2010): 10-11. Stonebraker, Michael. "Stonebraker on NoSQL and enterprises." Communications of the ACM 54.8 (2011): 10-11. Rafique, Ansar, et al. "On the performance impact of data access middleware for nosql data stores." IEEE Transactions on Cloud Computing (2015). Moniruzzaman, A. B. M., and Syed Akhter Hossain. "Nosql database: New era of databases for big data analytics-classification, characteristics and comparison." arXiv preprint arXiv:1307.0191 (2013).	Data Manipulation at Scale - NoSQL rebuttal
15	Video - Scientific Data: Water and Snow in the Western US	Tika in Action, Chapter 12 - 14	Resources:
April		C. Mattmann, D. Freeborn, D. Crichton, B. Foster, A. Hart, D. Woollard, S. Hardman, P. Ramirez, S. Kelly, A. Y. Chang, C. E. Miller. A Reusable Process	

21st, 2022	Searching Scientific Datasets	Control System Framework for the Orbiting Carbon Observatory and NPP Sounder PEATE missions. In Proceedings of the 3rd IEEE Intl Conference on Space Mission Challenges for Information Technology (SMC-IT 2009), pp. 165- 172, July 19 - 23, 2009.	Thomas Painter of NASA JPL speaks at TEDxIS on a thought provoking take
	Scientific Data Processing (Airborne Snow Observatory)	Wilkinson, Mark D., et al. "The FAIR Guiding Principles for scientific data management and stewardship." Scientific data 3 (2016): 160018.	
	Individual Presentations		
	Big Data with an Eye Towards the Future: Discussion	Buneman, Peter, et al. "Archiving scientific data." ACM Transactions on Database Systems (TODS) 29.1 (2004): 2-42. Fox, Peter, and James Hendler. "Changing the equation on scientific data visualization." Science 331.6018 (2011): 705-708. Plale, Beth, et al. "Active management of scientific data." IEEE Internet Computing 9.1 (2005): 27-34. Gray, Jim, et al. "Scientific data management in the coming decade." ACM SIGMOD Record 34.4 (2005): 34-41. Ailamaki, Anastasia, Verena Kantere, and Debabrata Dash. "Managing scientific data." Communications of the ACM 53.6 (2010): 68-78.	
