

USC Viterbi School of Engineering

INF 558/CSCI 563: Building Knowledge Graphs

Units: 4

Term—Day—Time:

Fall 2017 – Friday – 2-5:20pm

Location: GFS 116

Instructor: Craig Knoblock

Office: GER 207

Office Hours: TBD

Contact Info: knoblock@usc.edu, 310-448-8786.

Instructor: Pedro Szekely

Office: GER 207

Office Hours: TBD

Contact Info: szekely@isi.edu, 310-448-8641.

Teaching Assistant: Minh Pham

Office: TBD

Office Hours: TBD

Contact Info: minhpham@usc.edu

Catalogue Course Description

Foundations, techniques, and algorithms for building knowledge graphs and doing so at scale. Topics include information extraction, data alignment, entity linking, and the Semantic Web.

Expanded Course Description

This course focuses on foundations, techniques, and algorithms for building knowledge graphs. Students will learn the theory and applications of the techniques needed to build and query massive knowledge graphs. Topics include crawling web sites, wrapper learning, information extraction, source alignment, string matching, entity linking, graph databases, querying knowledge graphs, data cleaning, Semantic Web, linked data, graph analytics, and intellectual property. The class will be run as a lecture course with lots of student participation and significant hands-on experience. As an integral part of the course each student will do a project using the research and tools covered in the class.

Learning Objectives

The learning objectives for this course are:

- Understand the algorithms and techniques for crawling web sites, structured data extraction, and information extraction from unstructured text.
- Understand the theory and techniques for cleaning, aligning, matching, and linking data.
- Understand the foundations and techniques of the Semantic Web, including RDF, ontologies, SPARKL, and linked data.
- Understand how to work with graph databases, including how to load massive datasets into such databases, how to organize the data for efficient access, and how to efficiently query the contents.
- Understand the entire process of how to design, construct, and query a knowledge graph to solve real-world problems.

- Understand how to apply the big data tools and infrastructure (e.g., Spark) to build and query knowledge graphs.

Required Preparation:

Prerequisite(s): INF 551 or CSCI 585
INF 552 or CSCI 567

Recommended Background: Experience programming in Python

Course Notes

The course will be run as a lecture class with student participation strongly encouraged. There are weekly readings and students are encouraged to do the readings prior to the discussion in class. All of the course materials, including the readings, lecture slides, home works will be posted online. The class project is a significant aspect of this course and at the end of the semester students will present their projects in class.

Required Readings and Supplementary Materials

Required Textbook: none

We use a set of technical papers and book chapters that are all available online. All of the required readings are listed in the course schedule.

Description and Assessment of Assignments

Homework Assignments

There will be weekly homework assignments for the first 10 weeks of class. The assignments must be done individually. The homework assignments are expected to take 8-10 hours per week. Each assignment is graded on a scale of 0-100 and the specific rubric for each assignment is given in the assignment. The homework topics are listed in the Course Schedule.

Course Project

An integral part of this course is the course project, which builds on the topics and techniques covered in the class. Students can work in teams of up to two people on this project. They will present their project proposals in class, conduct the project, and then create a video demonstration of the work and present the project in class.

Project Timeline:

- Week 2-4: Project proposals presented in class (team members, topic)
- Week 10: Project status update due (1 page status report)
- Week 15: Project presentation in class (short talk and video demonstration)

Project description: Each project team will build a knowledge graph for a topic of their choice. The knowledge graph must combine data from at least 3 different sources and at least 2 of those sites must be from online web sites. The best projects build on many of the topics covered in the class. The homeworks have been designed so that you can work on your projects in the process of doing your homework.

An example project would be to build a knowledge graph of used bicycles that could be purchased near the USC campus. This project would combine data from used sources, such as Craig's List, new bike sources such as BikeNashbar, and bicycle review sites, such as bicycling.com. The project would collect the data from each of these sources using wrapper techniques, extract the details of the used bicycle ads from Craig's List using information extraction techniques, align the data across these various sources to a domain ontology, link the entities across sources to combine the used data with the reviews from bicycling.com and prices from BikeNashbar, store all of the data into a graph database such as elasticsearch, and then build a simple user interface to show the results by executing queries against the graph database.

Grading breakdown of the course project:

- Proposal: 10%
- Project video: 30%
- Presentation: 30%
- Overall project: 30%

Grading Breakdown

Quizzes: There will be weekly quizzes based on the material from the week before. There is no mid-term for this class.

Homework: There will be weekly homework based on the topics of the class each week.

Final Exam: There is a final exam at the end of the semester covering all of the material covered in the class.

Class Project: Each student will do a group class project based on the topics covered in the class. Students will propose their own project, do the research and build a proof-of-concept, create a video demonstration of the proof-of-concept, and present the project in class.

Grading Schema:

Quizzes	20%
Homework	20%
Final:	25%
Class Project	35%
<hr/>	
Total	100%

Grades will range from A through F. The following is the breakdown for grading:

94 - 100 = A	74 - 76.9 = C
90 - 93.9 = A-	70 - 73.9 = C-
87 - 89.9 = B+	67 - 69.9 = D+
84 - 86.9 = B	64 - 66.9 = D
80 - 83.9 = B-	60 - 63.9 = D-
77 - 79.9 = C+	Below 60 is an F

Assignment Submission Policy

Homework assignments are due at 11:59pm on the due date and should be submitted in Blackboard. You can submit homework up to one week late, but you will lose 20% of the possible points for the assignment. After one week, the assignment cannot be submitted.

Course Schedule: A Weekly Breakdown

	Topics/Daily Activities	Readings	Quizzes & Homeworks	Instructor
Week 1 Aug 25	Course Introduction & Use Case	Pedro Szekely, et al. Building and using a knowledge graph to combat human trafficking. In Proceedings of the 14th International Semantic Web Conference (ISWC 2015), 2015. http://iswc2015.semanticweb.org/sites/iswc2015.semanticweb.org/files/93670175.pdf	Homework 1: Project Ideas	Knoblock
	Crawling	The Anatomy of a Large Scale Hypertextual Web Search Engine Sergey Brin and Lawrence Page, Seventh		Szekely

		<p>International World Wide Web Conference, 1998. http://ilpubs.stanford.edu:8090/361/1/1998-8.pdf</p> <p>Searching the Web Arvind Arasu et al., ACM Transactions on Internet Technology, 2001 http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.25.5183</p>		
Week 1 Sep 1	Web Data Extraction	<p>Ion Muslea, Steve Minton, and Craig A. Knoblock. A hierarchical approach to wrapper induction. In Proceedings of the 3rd International Conference on Autonomous Agents, Seattle, WA, 1999. http://www.isi.edu/integration/papers/muslea99-agents.pdf</p> <p>AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapter 9. Morgan Kaufmann, 2012. http://www.sciencedirect.com/science/book/9780124160446</p> <p>W. Crescenzi, G. Mecca, and P. Merialdo. RoadRunner. Towards automatic data extraction from large web sites. 2001. http://www.vldb.org/conf/2001/P109.pdf</p> <p>B. Cenk Gazen and Steven Minton. Overview of autofeed: An unsupervised learning system for generating webfeeds. In Proceedings of AAAI, 2006. http://www.isi.edu/integration/courses/csci548/Papers/gazen06-aaai.pdf.</p>	<p>Quiz 1</p> <p>Homework 2: Crawling a website</p>	Knoblock
Week 3 Sep 8	Information Extraction	<p>Matthew Michelson and Craig A. Knoblock. Semantic Annotation of Unstructured and Ungrammatical Text. In Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-2005), Edinburgh, Scotland, 2005. http://www.isi.edu/integration/papers/michelson05-ijcai.pdf</p> <p>Andrew McCallum. Information Extraction: Distilling Structured Data from Unstructured Text . ACM Queue, volume 3, Number 9, November 2005. http://people.cs.umass.edu/~mccallum/papers/acm-queue-ie.pdf</p>	<p>Quiz 2</p> <p>Homework 3: Building wrappers</p>	Szekely
Week 4 Sep 15	Data Cleaning	<p>Wrangler: Interactive visual specification of data transformation scripts. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2011. http://vis.stanford.edu/papers/wrangler.</p> <p>Bo Wu, Pedro Szekely, and Craig A. Knoblock. Minimizing user effort in transforming data by example. In Proceedings of the International Conference on Intelligent User Interface, 2014. http://www.isi.edu/integration/papers/wu14-iui.pdf.</p> <p>Open Refine, Explore data.</p>	<p>Quiz 3</p> <p>Homework 4: Information Extraction</p>	Knoblock

		<p>http://youtu.be/B70J_H_zAWM.</p> <p>Open Refine, Clean and transform data. http://youtu.be/cO8NVCs_Ba0.</p> <p>Open Refine, Reconcile and match data. http://youtu.be/5tsyz3ibYzk.</p>		
Week 5 Sep 22	Graph Representation RDF, RDF Schema, JSON-LD	<p>Frank Manola and Eric Miller. Rdf primer. Technical report, W3C, February 2004. http://www.w3.org/TR/2004/REC-rdf-primer-20040210/</p> <p>Tim Berners-Lee. Why rdf model is different from the xml model. Technical report, W3C, 1998. http://www.w3.org/DesignIssues/RDF-XML.html.</p> <p>Rdf vocabulary description language 1.0: Rdf schema. Technical report, W3C, February 2004. http://www.w3.org/TR/2004/REC-rdf-schema-20040210/</p> <p>Ben Adida, Ivan Herman, Manu Sporny, and Mark Birbeck. Rdfa 1.1 primer rich structured data markup for web documents. Technical report, W3C, June 2012. http://www.w3.org/TR/2012/NOTE-rdfa-primer-20120607/</p>	<p>Quiz 4</p> <p>Homework 5: Data Cleaning</p>	Szekely
Week 6 Sep 29	Ontologies / RDF Mapping Tools	<p>Krtzsch Markus, Simancik Frantisek, and Horrocks Ian. A description logic primer. 2012. http://arxiv.org/pdf/1201.4089.pdf.</p> <p>R2rml: Rdb to rdf mapping language. http://www.w3.org/TR/r2rml/</p>	<p>Quiz 5</p> <p>Homework 6: RDF</p>	Knoblock
Week 7 Oct 6	String Matching	<p>AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapters 4. Morgan Kaufmann, 2012. http://www.sciencedirect.com/science/book/9780124160446</p>	<p>Quiz 6</p> <p>Homework 7: Ontologies/R2 RML</p>	Szekely
Week 8 Oct 13	Entity Linking	<p>AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapters 7. Morgan Kaufmann, 2012. http://www.sciencedirect.com/science/book/9780124160446</p>	<p>Quiz 7</p> <p>Homework 8: String Matching</p>	Knoblock

<p>Week 9 Oct 20</p>	<p>Semantic Labeling / Table Extraction</p>	<p>Ramnandan, S.; Mittal, A.; Knoblock, C. A.; and Szekely, P. Assigning Semantic Labels to Data Sources. In <i>Proceedings of the 12th ESWC</i>, 2015. http://usc-isi-i2.github.io/papers/ramnandan15-eswc.pdf</p> <p>Venetis, P., Halevy, A., Madhavan, J., Pasca, M., Shen, W., Wu, F., Miao, G., Wu, C.: Recovering semantics of tables on the web. <i>Proc. VLDB Endow.</i> 4(9), 528–538 (2011) http://www.vldb.org/pvldb/vol4/p528-venetis.pdf</p> <p>Syed, Z., Finin, T., Mulwad, V., Joshi, A.: Exploiting a web of semantic data for interpreting tables. In: <i>Proceedings of the Second Web Science Conference</i> (2010) http://ebiquity.umbc.edu/_file_directory_/papers/478.pdf</p>	<p>Quiz 8</p> <p>Homework 9: Entity Linking</p>	<p>Pham/ Ghasemi Gol</p>
<p>Week 10 Oct 27</p>	<p>Source Modeling</p>	<p>Mark James Carman and Craig A. Knoblock. Learning semantic descriptions of web information sources. In <i>Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)</i>, January 2007. http://www.isi.edu/integration/papers/carman07-ijcai.pdf</p> <p>José Luis Ambite, Sirish Darbha, Aman Goel, Craig A. Knoblock, Kristina Lerman, Rahul Parundekar, and Thomas Russ. Automatically constructing semantic web services from online sources. In <i>Proceedings of the 8th International Semantic Web Conference (ISWC 2009)</i>, 2009. http://www.isi.edu/integration/papers/ambite09-iswc.pdf</p> <p>Craig A. Knoblock and Pedro Szekely. Exploiting semantics for big data integration. <i>AI Magazine</i>, 2015. http://usc-isi-i2.github.io/papers/knoblock15-aimagazine.pdf</p>	<p>Quiz 9</p> <p>Homework 10: Source Modeling</p>	<p>Knoblock</p>
<p>Week 11 Nov 3</p>	<p>Querying Knowledge Graphs / ElasticSearch / SPARQL</p>	<p>Oren, E., Delbru, R., Catasta, M., Cyganiak, R., Stenzhorn, H., & Tummarello, G. (2008). Sindice. com: a document-oriented lookup index for open linked data. <i>International Journal of Metadata, Semantics and Ontologies</i>, 3(1), 37-52. http://wtlab.um.ac.ir/images/e-library/linked_data/other/Sindice.pdf</p> <p>Freitas, A., Oliveira, J. G., Curry, E., O’Riain, S., & da Silva, J. C. P. (2011, June). Treo: combining entity-search, spreading activation and semantic relatedness for querying linked data. In <i>Proc. of 1st Workshop on Question Answering over Linked Data (QALD-1) at the 8th Extended Semantic Web Conference (ESWC 2011)</i>.</p>	<p>Quiz 10</p> <p>Homework 11: SPARQL/ElasticSearch</p>	<p>Szekely/Slepicka</p>

		https://www.deri.ie/sites/default/files/publications/freitas_qald_2011_0.pdf Steve Harris and Andy Seaborne. Sparql 1.1 query language. Technical report, W3C, January 2012. http://www.w3.org/TR/2012/PR-sparql11-query-20121108		
Week 12 Nov 10	Linked Data/ DBpedia	Tom Heath and Christian Bizer Linked Data: Evolving the Web into a Global Data Space Synthesis Lectures on the Semantic Web: Theory and Technology Chapters 1 to and including section 4.2. https://west.uni-koblenz.de/files/ws1213/seminar-web-science/linked-data.pdf	Quiz 11 Homework 12: Linked Data	Knoblock
Week 13 Nov 17	Graph Analytics Intellectual Property	Rajaraman, J. Leskovec and J. D. Ullman, <i>Mining of Massive Datasets</i> , Cambridge University Press, 2012. http://infolab.stanford.edu/~ullman/mmds/ch10.pdf Thomas P. Vartanian and Robert H. Ledig. Scrape it, scrub it and show it: The battle over data aggregation. http://web.archive.org/web/20070818130311/http://www.ffhsj.com/bancmail/bmarts/aba_art.html . Kembrew McLeod. Intellectual property law, freedom of expression, and the web, 2003. http://www.electronicbookreview.com/thread/technocapitalism/proprietary . Electronic frontier foundation. http://www.eff.org/issues/intellectual-property .	Quiz 12	Szekely
Week 14 Nov 24	Thanksgiving break! No class.		Quiz 13	
Week 15 Dec 1	Student Presentations			Szekely Knoblock Pham
FINAL Dec 8 2-4pm	Final Exam		During assigned time in the <i>Schedule of Classes</i> at www.usc.edu/soc .	

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* <https://scampus.usc.edu/1100-behavior-violating-university-standards-and-appropriate-sanctions>. Other forms of academic dishonesty are equally unacceptable. See additional information in *SCampus* and university policies on scientific misconduct, <http://policy.usc.edu/scientific-misconduct>.

Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the *Office of Equity and Diversity* <http://equity.usc.edu> or to the *Department of Public Safety* <http://capsnet.usc.edu/department/department-public-safety/online-forms/contact-us>. This is important for the safety of the whole 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 Center for Women and Men* <http://www.usc.edu/student-affairs/cwm/> provides 24/7 confidential support, and the sexual assault resource center webpage <http://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* <http://dornsife.usc.edu/ali>, which sponsors courses and workshops specifically for international graduate students. *The Office of Disability Services and Programs* http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html 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* <http://emergency.usc.edu> will provide safety and other updates, including ways in which instruction will be continued by means of blackboard, teleconferencing, and other technology.