

# USC Viterbi School of Engineering

## **CSCI 548: Information Integration on the Web**

**Units: 3**

**Term—Day—Time:**

**Spring 2015 – MW – 3:30-4:50pm**

**Location:** ZHS 163

**Instructor: Jose Luis Ambite**

**Office:** Outside classroom

**Office Hours:** Immediately after class

**Contact Info:** [ambite@isi.edu](mailto:ambite@isi.edu), 310-448-8472.

**Instructor: Craig Knoblock**

**Office:** AFH B55a

**Office Hours:** Wednesdays 10-11am

**Contact Info:** [knoblock@usc.edu](mailto:knoblock@usc.edu), 310-448-8786.

**Teaching Assistant: Bo Wu**

**Office:** plaza between RTH cafe and EEB

**Office Hours: Tuesday 10am-12pm**

**Contact Info:** [wubo@usc.edu](mailto:wubo@usc.edu)

### **Catalogue Course Description**

Foundations, techniques, and algorithms for information integration. Topics include Semantic Web, linked data, data integration, entity linkage, source modeling, and information extraction.

### **Expanded Course Description**

This course focuses on foundations, techniques, and algorithms for information extraction, modeling and integration. Topics covered include semantic web (RDF, OWL, SPARQL), linked data and services, mash-ups, theory of data integration, schema mappings, record/entity linkage, data cleaning, source modeling, and information extraction. The class will be run as a lecture course with lots of student participation and significant hands-on experience.

### **Learning Objectives**

The learning objectives for this course are:

- Understand the foundations and techniques of the Semantic Web, including RDF, OWL, SPARKL, linked data, and linked services
- Understand the theory and techniques of traditional data integration, including view integration, schema mapping, record linkage
- Understand the algorithms and techniques for data cleaning, source modeling, building mashups, semi-structured extraction, and information extraction
- Understand the theory and application of the state-of-the-art software and tools for information extraction
- For any given integration problem, be able to select and apply the most relevant information integration techniques to solve that problem

**Prerequisite(s):** CSCI 561

**Co-Requisite (s):** none

**Concurrent Enrollment:** none

**Recommended Preparation:** CSCI 585 and some programming experience

### 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, homeworks will be posted online on Blackboard.

### Technological Proficiency and Hardware/Software Required

Students are expected to know how to program in a language such as Java, C++, or Python. Students are also expected to have their own laptop or desktop computer where they can install and run software to do the weekly homework assignments.

### Required Readings and Supplementary Materials

Required Textbook: Principles of Data Integration by Doan, Halevy, & Ives, Morgan Kaufmann, 2012

The book is available online from the USC library and is available for purchase.

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 12 weeks of class. The assignments must be done individually. The homework assignments are expected to take 6-8 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.

#### 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.

Grading Schema:

Quizzes	25%
Homework	50%
Final:	25%
<hr/>	
Total	100%

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

94 - 100 = A	74 - 76 = C
90 - 93 = A-	70 - 73 = C-
87 - 89 = B+	67 - 69 = D+
84 - 86 = B	64 - 66 = D
83 - 83 = B-	60 - 63 = D-
77 - 79 = 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
<b>Week 1</b> Jan 12	<b>Course Introduction</b>	Frank Manola and Eric Miller. Rdf primer. Technical report, W3C, February 2004. <a href="http://www.w3.org/TR/2004/REC-rdf-primer-20040210/">http://www.w3.org/TR/2004/REC-rdf-primer-20040210/</a> .	Homework 0: Academic Integrity (due on Friday)	Professor Ambite
Jan 14	<b>RDF, Graph data model</b>	Tim Berners-Lee. Why rdf model is different from the xml model. Technical report, W3C, 1998. <a href="http://www.w3.org/DesignIssues/RDF-XML.html">http://www.w3.org/DesignIssues/RDF-XML.html</a> .		Professor Ambite
<b>Week 2</b> Jan 21	<b>RDF Schema</b>	Rdf vocabulary description language 1.0: Rdf schema. Technical report, W3C, February 2004. <a href="http://www.w3.org/TR/2004/REC-rdf-schema-20040210/">http://www.w3.org/TR/2004/REC-rdf-schema-20040210/</a> .  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. <a href="http://www.w3.org/TR/2012/NOTE-rdfa-primer-20120607/">http://www.w3.org/TR/2012/NOTE-rdfa-primer-20120607/</a> .	Quiz 1 (on Wednesday) Homework 1: Creating a Wrapper (due on Friday)	Professor Knoblock
<b>Week 3</b> Jan 26	<b>SPARQL query language</b>	Steve Harris and Andy Seaborne. Sparql 1.1 query language. Technical report, W3C, January 2012. <a href="http://www.w3.org/TR/2012/PR-sparql11-query-20121108/">http://www.w3.org/TR/2012/PR-sparql11-query-20121108/</a> .	Quiz 2 (on Monday)  Homework 2: RDF (due on Friday)	Professor Ambite
Jan 28	<b>OWL 2</b>	Krtzsch Markus, Simancik Frantisek, and Horrocks Ian. A description logic primer. 2012. <a href="http://arxiv.org/pdf/1201.4089.pdf">http://arxiv.org/pdf/1201.4089.pdf</a> .		Professor Ambite
<b>Week 4</b> Feb 2	<b>Linked Data</b>	Aduna B.V. Http communication protocol for sesame 2. In System documentation for Sesame 2.x, chapter 8. October 2013. <a href="http://www.csee.umbc.edu/courses/graduate/691/spring14/01/examples/sesame/openrdf-sesame-2.6.10/docs/system/ch08.html">http://www.csee.umbc.edu/courses/graduate/691/spring14/01/examples/sesame/openrdf-sesame-2.6.10/docs/system/ch08.html</a> .  Chimezie Ogbuji. Sparql 1.1 graph store http protocol. Technical report, W3C, May 2012. <a href="http://www.w3.org/TR/sparql11-http-rdf-update/">http://www.w3.org/TR/sparql11-http-rdf-update/</a> .	Quiz 3 (on Monday)  Homework 3 SPARQL (due on Friday)	Professor Knoblock

Feb 4	<b>Data Cleaning</b>	<p>Wrangler: Interactive visual specification of data transformation scripts. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2011. <a href="http://vis.stanford.edu/papers/wrangler">http://vis.stanford.edu/papers/wrangler</a>.</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. <a href="http://www.isi.edu/integration/papers/wu14-iui.pdf">http://www.isi.edu/integration/papers/wu14-iui.pdf</a>.</p> <p>Open Refine, Explore data. <a href="http://youtu.be/B70J_H_zAWM">http://youtu.be/B70J_H_zAWM</a>.</p> <p>Open Refine, Clean and transform data. <a href="http://youtu.be/cO8NVCs_Ba0">http://youtu.be/cO8NVCs_Ba0</a>.</p> <p>Open Refine, Reconcile and match data. <a href="http://youtu.be/5tsyz3ibYzk">http://youtu.be/5tsyz3ibYzk</a>.</p>		Mr. Wu
<b>Week 5</b> Feb 9	<b>Database theory basics</b>	AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapter 2.1, 2.2, 2.3 and 2.4. Morgan Kaufmann, 2012.	Quiz 4 (on Monday)	Professor Ambite
Feb 11	<b>Logical Data Integration</b>	<a href="http://www.sciencedirect.com/science/book/9780124160446">http://www.sciencedirect.com/science/book/9780124160446</a>	Homework 4: OWL (due on Friday)	Professor Ambite
<b>Week 6</b> Feb 18	<b>Scalable data integration</b>	<p>Alon Halevy and Rachel Pottinger. A scalable algorithm for answering queries using views. The VLDB Journal The International Journal on Very Large Data Bases, 2001. <a href="http://www.vldb.org/conf/2000/P484.pdf">http://www.vldb.org/conf/2000/P484.pdf</a>.</p> <p>Scalable query rewriting: a graph-based approach, 2001. <a href="http://www.isi.edu/ambite/konstantinidis2011-sigmod.pdf">http://www.isi.edu/ambite/konstantinidis2011-sigmod.pdf</a>.</p>	<p>Quiz 5 (on Monday)</p> <p>Homework 5: Data Cleaning (due on Friday)</p>	Mr. Konstantinidis
<b>Week 7</b> Feb 23	<b>Linked Services</b>	<p>Barry Norton and Reto Kruppenacher. Consuming dynamic linked data. In Proceedings of the 1st International Workshop on Consuming Linked Data, 2010. <a href="http://ceur-ws.org/Vol-665/NortonEtAlCOLD2010.pdf">http://ceur-ws.org/Vol-665/NortonEtAlCOLD2010.pdf</a>.</p> <p>Mohsen Taheriyani, Craig A. Knoblock, Pedro Szekely, and Jose Luis Ambite. Rapidly integrating services into the linked data cloud. In Proceedings of the 11th International Semantic Web Conference (ISWC 2012), 2012. <a href="http://www.isi.edu/integration/papers/taheriyani12-iswc.pdf">http://www.isi.edu/integration/papers/taheriyani12-iswc.pdf</a>.</p>	<p>Quiz 6 (on Monday)</p> <p>Homework 6: Logical Data Integration (due on Friday)</p>	Mr. Taheriyani
Feb 25	<b>Semi-Automatic Source Modeling</b>	<p>Craig A. Knoblock, Pedro Szekely, Jose Luis Ambite, Aman Goel, Shubham Gupta, Kristina Lerman, Maria Muslea, Mohsen Taheriyani, and Parag Mallick. Semi-automatically mapping structured sources into the semantic web. In Proceedings of the Extended Semantic Web Conference, Crete, Greece, 2012. <a href="http://www.isi.edu/integration/papers/knoblock12-eswc.pdf">http://www.isi.edu/integration/papers/knoblock12-eswc.pdf</a>.</p>		Professor Knoblock

<b>Week 8</b> Mar 2	<b>RDF Mapping Tools</b>	[2] R2rml: Rdb to rdf mapping language. <a href="http://www.w3.org/TR/r2rml/">http://www.w3.org/TR/r2rml/</a> .	Quiz 7 (on Monday)	Professor Ambite
Mar 4	<b>Schema Mapping</b>	Jayant Madhavan, Shawn R. Jeffery, Shirley Cohen, Xin (Luna) Dong, David Ko, Cong Yu, and Alon Halevy. Web-scale data integration: You can only afford to pay as you go, 2007. <a href="http://www.docin.com/p-47000224.html">http://www.docin.com/p-47000224.html</a> .  AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapter 5. Morgan Kaufmann, 2012. <a href="http://www.sciencedirect.com/science/book/9780124160446">http://www.sciencedirect.com/science/book/9780124160446</a>	Homework 7: Triple Stores (due on Friday)	Professor Ambite
<b>Week 9</b> Mar 9	<b>Source Modeling</b>	Mark James Carman and Craig A. Knoblock. Learning semantic descriptions of web information sources. In Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI), January 2007. <a href="http://www.isi.edu/integration/papers/carman07-ijcai.pdf">http://www.isi.edu/integration/papers/carman07-ijcai.pdf</a> .  Jos'e 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 Proceedings of the 8th International Semantic Web Conference (ISWC 2009), 2009. <a href="http://www.isi.edu/integration/papers/ambite09-iswc.pdf">http://www.isi.edu/integration/papers/ambite09-iswc.pdf</a> .	Quiz 8 (on Monday)	Professor Ambite
Mar 11	<b>String Matching</b>	[1] AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapters 4. Morgan Kaufmann, 2012. <a href="http://www.sciencedirect.com/science/book/9780124160446">http://www.sciencedirect.com/science/book/9780124160446</a>	Homework 8: Karma (due on Friday)	Professor Knoblock
<b>Week 10</b> Mar 23	<b>Record Linkage</b>	AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapters 7. Morgan Kaufmann, 2012. <a href="http://www.sciencedirect.com/science/book/9780124160446">http://www.sciencedirect.com/science/book/9780124160446</a>	Quiz 9 (on Monday)	Professor Knoblock
Mar 25	<b>Data Warehousing</b>	AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapter 10. Morgan Kaufmann, 2012. <a href="http://www.sciencedirect.com/science/book/9780124160446">http://www.sciencedirect.com/science/book/9780124160446</a>  ]	Homework 9: String Similarity (due on Friday)	Professor Ambite

<p><b>Week 11</b> Mar 30</p> <p>Apr 1</p>	<p><b>Mashup principles</b></p> <p><b>Mashup tools</b></p>	<p>Shubham Gupta and Craig A. Knoblock. Building geospatial mashups to visualize information for crisis management. In Proceedings of the 7th International Conference on Information Systems for Crisis Response and Management, 2010. <a href="http://www.isi.edu/integration/papers/gupta10-iscram.pdf">http://www.isi.edu/integration/papers/gupta10-iscram.pdf</a>.</p> <p>Jeffrey Wong and Jason I. Hong. Making mashups with marmite: towards end-user programming for the web. In ACM SIGMOD Record, 2007. <a href="http://repository.cmu.edu/cgi/viewcontent.cgi?article=1063&amp;context=hcii">http://repository.cmu.edu/cgi/viewcontent.cgi?article=1063&amp;context=hcii</a>.</p> <p>Rob Ennals, Eric Brewer, Minos Garofalakis, Michael Shadle, and Prashant Gandhi. Intel mash maker: join the web. 2007. <a href="http://23.30.224.201/publications/intel-mash-maker-join-web">http://23.30.224.201/publications/intel-mash-maker-join-web</a>.</p> <p>Huynh David, Mazzocchi Stefano, and Karger David. Piggy bank: Experience the semantic web inside your web browser. 2007. <a href="http://simile.mit.edu/papers/iswc05.pdf">http://simile.mit.edu/papers/iswc05.pdf</a>.</p> <p>Leo Sauermann and Richard Cyganiak. Cool uris for the semantic web. Technical report, 2008. <a href="http://www.w3.org/TR/cooluris/">http://www.w3.org/TR/cooluris/</a>.</p>	<p>Quiz 10 (on Monday)</p> <p>Homework 10: Record Linkage (due on Friday)</p>	<p>Professor Knoblock</p> <p>Professor Knoblock</p>
<p><b>Week 12</b> Apr 6</p> <p>Apr 8</p>	<p><b>Information Extraction</b></p>	<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. <a href="http://www.isi.edu/integration/papers/michelson05-ijcai.pdf">http://www.isi.edu/integration/papers/michelson05-ijcai.pdf</a></p> <p>Andrew McCallum. Information Extraction: Distilling Structured Data from Unstructured Text . ACM Queue, volume 3, Number 9, November 2005. <a href="http://people.cs.umass.edu/~mccallum/papers/acm-queue-ie.pdf">http://people.cs.umass.edu/~mccallum/papers/acm-queue-ie.pdf</a></p> <p>Charles Elkan, Tutorial on Log-linear Models and Conditional Random Fields. <a href="http://videolectures.net/cikm08_elkan_llmacrf/">http://videolectures.net/cikm08_elkan_llmacrf/</a></p>	<p>Quiz 11 (on Monday)</p> <p>Homework 11: Mashups (due on Friday)</p>	<p>Professor Knoblock</p> <p>Professor Knoblock</p>

<p><b>Week 13</b> Apr 13</p> <p>Apr 15</p>	<p><b>OWL Profiles</b></p> <p><b>Wrapper Learning</b></p>	<p>Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. DI-lite: tractable description logics for ontologies. In Proc. of the 20th National Conference on Artificial Intelligence, 2005. <a href="http://www.aaai.org/Papers/AAAI/2005/AAAI05-094.pdf">http://www.aaai.org/Papers/AAAI/2005/AAAI05-094.pdf</a>.</p> <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. <a href="http://www.isi.edu/integration/papers/muslea99-agents.pdf">http://www.isi.edu/integration/papers/muslea99-agents.pdf</a>.</p> <p>AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. Principles of Data Integration, chapter 9. Morgan Kaufmann, 2012. <a href="http://www.sciencedirect.com/science/book/9780124160446">http://www.sciencedirect.com/science/book/9780124160446</a></p>	<p>Quiz 12 (on Monday)</p> <p>Homework 12: Information Extraction (due on Friday)</p>	<p>Professor Ambite</p> <p>Professor Knoblock</p>
<p><b>Week 14</b> Apr 20</p> <p>Apr 22</p>	<p><b>Ontology-based Data Integration</b></p> <p><b>Wrapper Generation</b></p>	<p>Hector Prez-Urbina, Ian Horrocks, and Boris Motik. Efficient query answering for owl 2. In International Semantic Web Conference, 2009. Efficient Query Answering for OWL 2. <a href="https://www.cs.ox.ac.uk/boris.motik/pubs/puhm09query-OWL2.pdf">https://www.cs.ox.ac.uk/boris.motik/pubs/puhm09query-OWL2.pdf</a>.</p> <p>W. Crescenzi, G. Mecca, and P. Merialdo. RoadRunner. Towards automatic data extraction from large web sites. 2001. <a href="http://www.vldb.org/conf/2001/P109.pdf">http://www.vldb.org/conf/2001/P109.pdf</a>.</p> <p>B. Cenk Gazen and Steven Minton. Overview of autofeed: An unsupervised learning system for generating webfeeds. In Proceedings of AAAI, 2006. <a href="http://www.isi.edu/integration/courses/csci548/Papers/gazen06-aaai.pdf">http://www.isi.edu/integration/courses/csci548/Papers/gazen06-aaai.pdf</a>.</p>	<p>Quiz 13 (on Monday)</p>	<p>Professor Ambite</p> <p>Professor Knoblock</p>
<p><b>Week 15</b> Apr 27</p> <p>Apr 29</p>	<p><b>Intellectual Property</b></p> <p><b>Course Review</b></p>	<p>Thomas P. Vartanian and Robert H. Ledig. Scrape it, scrub it and show it: The battle over data aggregation. <a href="http://web.archive.org/web/20070818130311/http://www.ffhsj.com/bancmail/bmarts/aba_art.html">http://web.archive.org/web/20070818130311/http://www.ffhsj.com/bancmail/bmarts/aba_art.html</a>.</p> <p>Kembrew McLeod. Intellectual property law, freedom of expression, and the web, 2003. <a href="http://www.electronicbookreview.com/thread/technoapitalism/proprietary">http://www.electronicbookreview.com/thread/technoapitalism/proprietary</a>.</p> <p>Electronic frontier foundation. <a href="http://www.eff.org/issues/intellectual-property">http://www.eff.org/issues/intellectual-property</a>.</p>	<p>Quiz 14</p>	<p>Professor Knoblock</p> <p>Professors Knoblock &amp; Ambite</p>

<b>FINAL</b> May 8 2-4pm	<b>Final Exam</b>		During assigned time in the <i>Schedule of Classes</i> at <a href="http://www.usc.edu/soc">www.usc.edu/soc</a> .	
--------------------------------	-------------------	--	------------------------------------------------------------------------------------------------------------------	--

## 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/academicssupport/centerprograms/dsp/home\\_index.html](http://sait.usc.edu/academicssupport/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.