Data and Literature:
Introduction to Literary Text Mining

LLCU 255 – Fall 2017
T/Th 8:35 - 9:55, Rm. 361 688 Sherbrooke

Professor Andrew Piper
Office: Rm 484, 688 Sherbrooke
Phone: 514-398-4400 x094504
Email: andrew.piper@mcgill.ca
Office Hours: T 10 - 11 am, Th 4 - 5 pm

Course Description

This course will serve as an introduction to the new tools and techniques being developed to study literature and culture at a vastly greater scale. How does the ability to analyze several hundred to hundreds of thousands of texts give us new insights into the history of literature and culture? How might thinking about literature as data change our understanding of foundational categories like author, text, work, narrative, plot, character or even language? In order to address these questions, this course will introduce you to the basic concepts and practices of text mining (vector space models, distributional semantics, sentiment analysis, topic modeling, and social network analysis) and the ways in which they are being applied to the study of literature. Weekly assignments will introduce you to the R software environment and will culminate in a final project of your own choosing. No prior programming experience is required.

Reading List

All readings are available through myCourses.

Weekly Assignments

Wk. 1 09.05 *What is Text Mining?*

09.07 *What's it for? Some examples.*
- Lancashire, "Vocabulary Change in Agatha Christie."
- Piper/So, "Quantifying the Weepy Best-Seller." https://newrepublic.com/article/126123/quantifying-weepy-bestseller
- Matt Daniels, "The Largest Vocabulary in Hip Hop." Poster. http://poly-graph.co/vocabulary.html
- *Extra Credit*: Michel et al., "Quantitative Analysis of Culture Using Millions of Digitized Books."

Wk. 2 09.12 *What are literary features and how can we model them? (Feature Selection)*

09.14 *Preparing your data* -- the TM *package in R*

Wk. 3 09.19 *How can we understand the significance of literary features and what do we mean by significant?*

09.21 *Significance Testing*

Assignment 1 (Due one week from handout)

Wk. 4 09.26 *What are Vector Space Models and what can you do with them?*
- Turney et al.

09.28 *Clustering your documents to discover new relationships*

Wk. 5 10.03 *What kind of learning is machine learning?*

10.05 *The* kernlab *package in R*

Assignment 2

Wk. 6 10.10 *Applied Machine Learning 1: Thinking about Genre*
- Underwood, "The Life Cycles of Genres."
http://culturalanalytics.org/2016/05/the-life-cycles-of-genres/

10.12 *Applied Machine Learning 2: Studying Cultural Capital*
- Piper/Portelance, "How Cultural Capital Works."
http://post45.research.yale.edu/2016/05/how-cultural-capital-works-prizewinning-novels-bestsellers-and-the-time-of-reading/

Wk. 7 10.17 *Narrative Structure 1 (Emotional Arcs)*
- Mathew Jockers, "Revealing Sentiment and Plot Arcs with the Syuzhet Package."
http://www.matthewjockers.net/2015/02/02/syuzhet/
- Reagan et al., "The emotional arcs of stories are dominated by six basic shapes."
- Ben Schmidt, "Plot arceology: a vector space model of narrative structure."

10.19 *Sentiment Analysis*

Wk. 8 10.24 *Social Network Analysis*
- Healy, "Using Metadata to find Paul Revere."
https://kieranhealy.org/blog/archives/2013/06/09/using-metadata-to-find-paul-revere/
- Newmann, Networks, "Introduction" and Chap. 4

10.26 *The* igraph *package in R*

Assignment 3

Wk. 9 10.31 *Narrative Structure 2 (Social Networks)*
- Sack, "Character Networks for Narrative Generation."
- Kraicer/Piper, "Socializations: Gender, Genre and the Social Networks of Contemporary Fiction."

11.02 *Studying Character Networks - Using* BookNLP

Wk.10 11.07 *Topic Modeling*
- Mohr, "Topic Models: What they are and why they matter."
- Underwood, "Topic Modeling Made Just Simple Enough."
- Jockers, "Significant Themes in 19C Literature."

11.09 *The* topicmodels *package in R*

Assignment 4

Wk.11 11.14 *Spatial Analysis*
- Wilkens, "The Geographic Imagination of Civil War Era American Fiction."

11.16 *Spatial Analysis in R*

Wk.12 11.21 *Beyond Text: The Visual Page*
- Piper/Cheriet, "Footnote Detection." http://txtlab.org/?p=395

11.23 *How to build a literary model*

Wk.13 11.28 Review of Final Projects

11.30 Review of Final Projects

****Final Paper due as a hard copy on December 12 by 4 pm in Room 425**

Academic Integrity

McGill University values academic integrity. Therefore all students must understand the meaning and consequences of cheating, plagiarism and other academic offences under the Code of Student Conduct and Disciplinary Procedures (see http://www.mcgill.ca/integrity/ for more information).

Course Requirements

In accord with McGill University's Charter of Students' Rights, students in this course have the right to submit in English or in French any written work that is to be graded.

Class Participation	15%
Weekly Assignments (4x)	40%
Final Paper (6-8 pp.)	45%

Class Participation. You are expected to attend every class and actively participate in class discussions with observations and questions derived from close and thoughtful reading of each weeks' texts. Our aim is to engage critically with existing studies of literature and to think creatively about new ways of understanding texts.

Weekly Assignments. Weekly assignments are designed to introduce you to using the R software environment for text analysis. You will move from the straightforward implementation of existing scripts to the analysis of results in the form of a 1-2 pp paper. In each case you will be provided with a choice of data sets and a particular script which you will learn how to "tune." You are free, indeed encouraged, to construct your own data sets. The aim of these assignments is to give you a hands-on understanding of how computational analysis works and how to critically analyze your results.

Final Paper. The final paper will consist of the following steps: a) design an experimental study; b) choose your data; c) implement one or more R scripts for analysis; d) write a detailed and thoughtful engagement with your results. The aim of this paper is to have you work through the entire analytical process, from the choice of appropriate data, the relevance of your analytical techniques, to the potential significance of your findings. What data did you choose to work with and why? What has your method told you about your texts? What challenges did you encounter? What do you remain uncertain about? Why is this an important question to be asking in the first place? As with the weekly assignments you may choose an existing data set or create one of your own.

Weekly Tutorial (Optional). A trouble-shooting tutorial for using R will be held every week at a specified time. This is an opportunity for you to improve your R programming skills.

Late papers will lose a half-grade for every class late. Students who receive a grade of D,F, or J will not be allowed to do supplemental work. All papers will be submitted to the text-matching software per university policy. Three or more missed classes will result in a lowering of the student's overall grade. According to Senate regulations, instructors are not permitted to make special arrangements for final exams. Please consult the Calendar, section 4.7.2.1, General University Information and Regulations at www.mcgill.ca. In the event of extraordinary circumstances beyond the University's control, the content and/or evaluation scheme in this course is subject to change.