WHAT ARE WE TEACHING?

Automated Evaluation of CS Curricula Content Using Topic Modeling

Jean Michel Rouly

April 30, 2015

George Mason University

OVERVIEW

COMPUTER SCIENCE EDUCATION

- From 2007, number of new CS undergrads has increased 28.9%
- Fifth straight year of increase

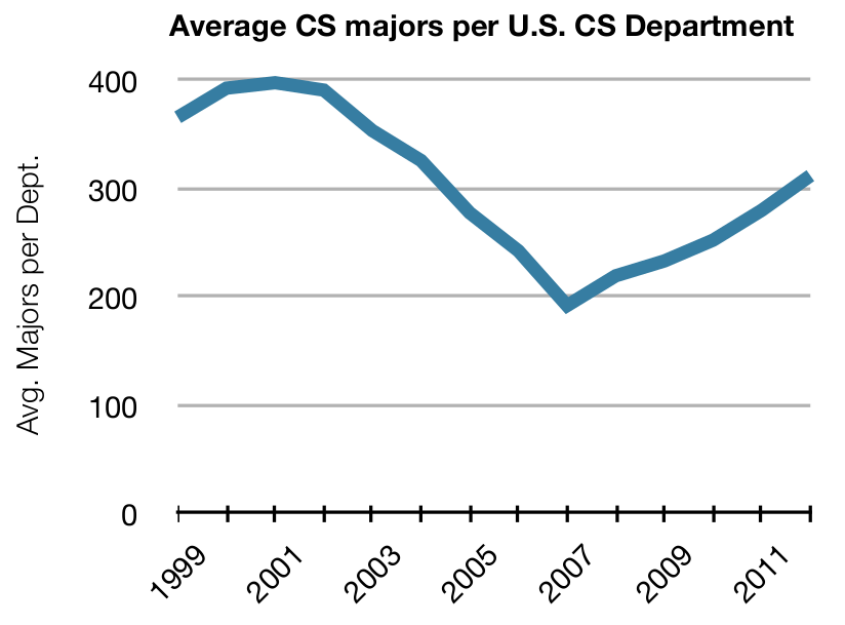

Figure 1: From Zweben, 2011.

DEPARTMENTAL EVALUATION

Evaluating a department on its conceptual coverage involves...

DEPARTMENTAL EVALUATION

Evaluating a department on its conceptual coverage involves...

Relative Standing comparing against other, similar departments

DEPARTMENTAL EVALUATION

Evaluating a department on its conceptual coverage involves...

Relative Standing comparing against other, similar departments

Absolute Performance benchmark against standardized expectations (published by the ACM)

DEPARTMENTAL EVALUATION

Evaluating a department on its conceptual coverage involves...

Relative Standing comparing against other, similar departments

Absolute Performance benchmark against standardized expectations (published by the ACM)

Both of these require data about the topics covered in a course.

AVAILABLE DATA

What we've got:

AVAILABLE DATA

What we've got:

1. Widely available university course description data.

AVAILABLE DATA

What we've got:

1. Widely available university course description data.
2. Descriptions detail what concepts are taught in a course.

AVAILABLE DATA

What we've got:

1. Widely available university course description data.
2. Descriptions detail what concepts are taught in a course.
3. **Human-readable** descriptions require manual inspection.

AVAILABLE DATA

What we've got:

1. Widely available university course description data.
2. Descriptions detail what concepts are taught in a course.
3. **Human-readable** descriptions require manual inspection.

So what to do?

RESEARCH GOAL

Through the application of probabilistic machine learning methods, specifically LDA topic modeling, a corpus of unstructured course descriptions can be digested and mined for topics. In this scenario, each topic represents a core concept covered by the courses.

A research framework will be constructed to read data from the Internet, digest into a common internal format, pipeline into an LDA topic model, and ultimately visualize in an interactive manner.

Ultimately the automatically discovered topics can be used in end-user university evaluation processes.

In other words

Build a tool that automatically...

RESEARCH GOAL

In other words

Build a tool that automatically...

- ingests large collection of course descriptions

RESEARCH GOAL

In other words

Build a tool that automatically...

- ingests large collection of course descriptions
- infers topics from course descriptions

RESEARCH GOAL

In other words

Build a tool that automatically...

- ingests large collection of course descriptions
- infers topics from course descriptions
- computes comparisons between departments

RESEARCH GOAL

In other words

Build a tool that automatically...

- ingests large collection of course descriptions
- infers topics from course descriptions
- computes comparisons between departments
- evaluates departments on their concepts

BACKGROUND

Topic Modeling

Attempts to discover the abstract topics of a dataset.

Topics

A topic is a frequency distribution over terms, roughly representing a concept taught in a course.

Overview

Latent Dirichlet Allocation (LDA) is a form of *Latent Variable Modeling* that can infer topics from within a document.

LDA takes a generative approach to latent variable modeling, assuming the topics occur in some proportion within each document.

DATA

UNIVERSITY DATASET

University	Course Count	Web
American University	32	american.edu
George Mason University	145	gmu.edu
Kansas State University	83	ksu.edu
Louisiana State University	59	lsu.edu
Portland State University	190	pdx.edu
Rensselaer Polytechnic Institute	61	rpi.edu
University of South Carolina	64	sc.edu
Stanford University	69	stanford.edu
University of Utah	142	utah.edu
University of Tennessee, Knoxville	29	utk.edu
ACM Exemplar Courses	68	—

Table 1: University course descriptions

SAMPLE COURSE DESCRIPTIONS

Raw Course Description

Capstone course focusing on design and successful implementation of major software project, encompassing broad spectrum of knowledge and skills, developed by team of students. Requires final exhibition to faculty-industry panel.

Cleaned course description

capston focus design success implement major softwar project encompass broad spectrum knowledg skill develop team student requir final exhibit faculti industri panel

TRAJECTORY

Trajectory is a tool that automatically ingests course description data from the Internet and presents an accessible interface for departmental evaluation.

Lines of Python	3193
Lines of Java	631
Lines of HTML/CSS/JS	1828
Lines of JSON	3219
Lines of Bash	165
Size on disk	6.7M

Table 2: Code statistics

TRAJECTORY

Four primary modules:

Scrape	web-scrape online university catalogs
Import/Export	pass structured data between Learn and Scrape
Learn	estimate LDA topic model on data
Web	visualization tool

Underneath the entire system is a structured relational database layer.

TRAJECTORY / WEB

Browse	collected data by university or department
Understand	courses through inferred topics
Analyze	conceptual overlap in prerequisite chains
Compare	departments based on conceptual composition
Evaluate	departments against ACM benchmarks

(Not So) Live Demo

DISCLAIMER

Images have been modified to fit in this presentation.

Visit `trajectory.rouly.net` on your smart device to follow along.

BROWSE

#	Department	Course Count	Web
1	Applied Information Technology	39	-
2	Computer Science	106	-

Figure 2: Browse a university's departments.

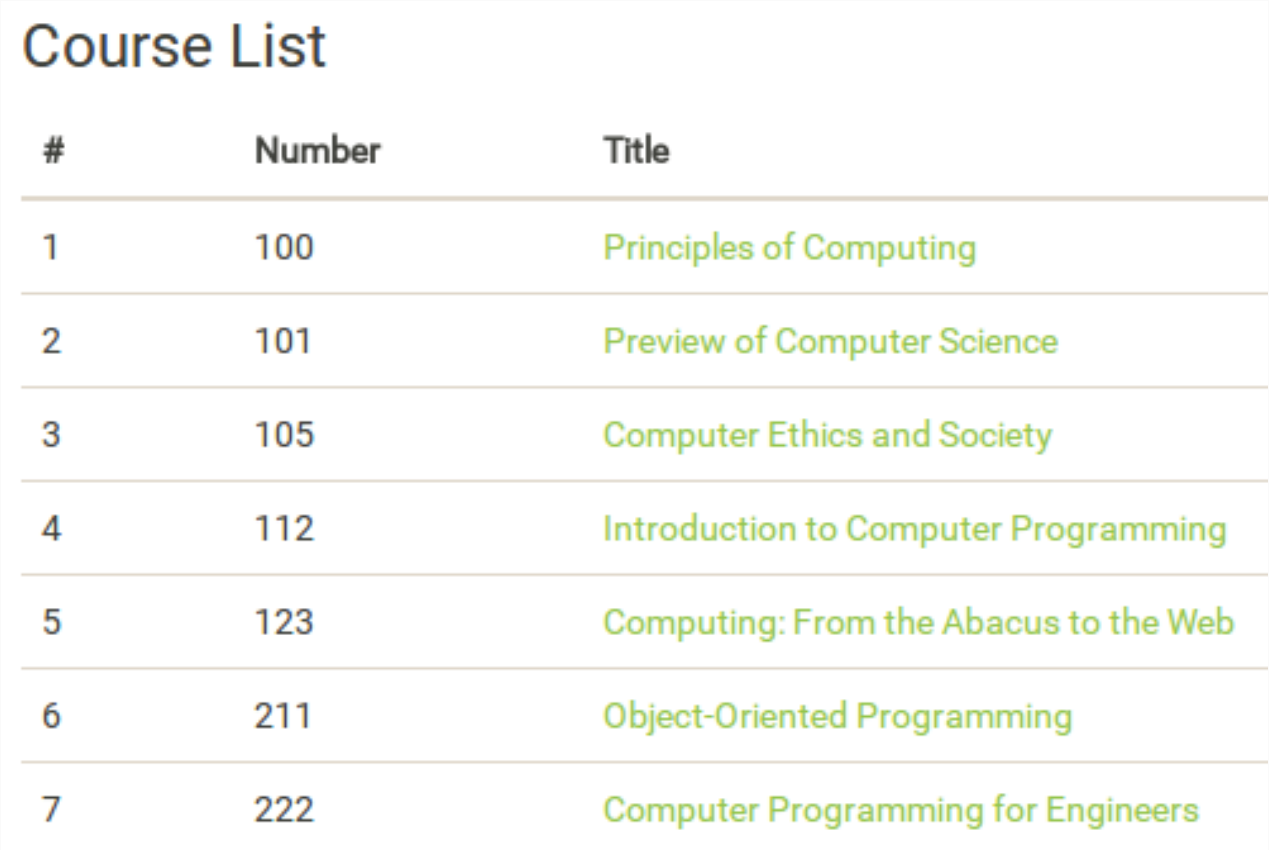

Course List

#	Number	Title
1	100	Principles of Computing
2	101	Preview of Computer Science
3	105	Computer Ethics and Society
4	112	Introduction to Computer Programming
5	123	Computing: From the Abacus to the Web
6	211	Object-Oriented Programming
7	222	Computer Programming for Engineers

Figure 3: Browse a department's courses.

UNDERSTAND

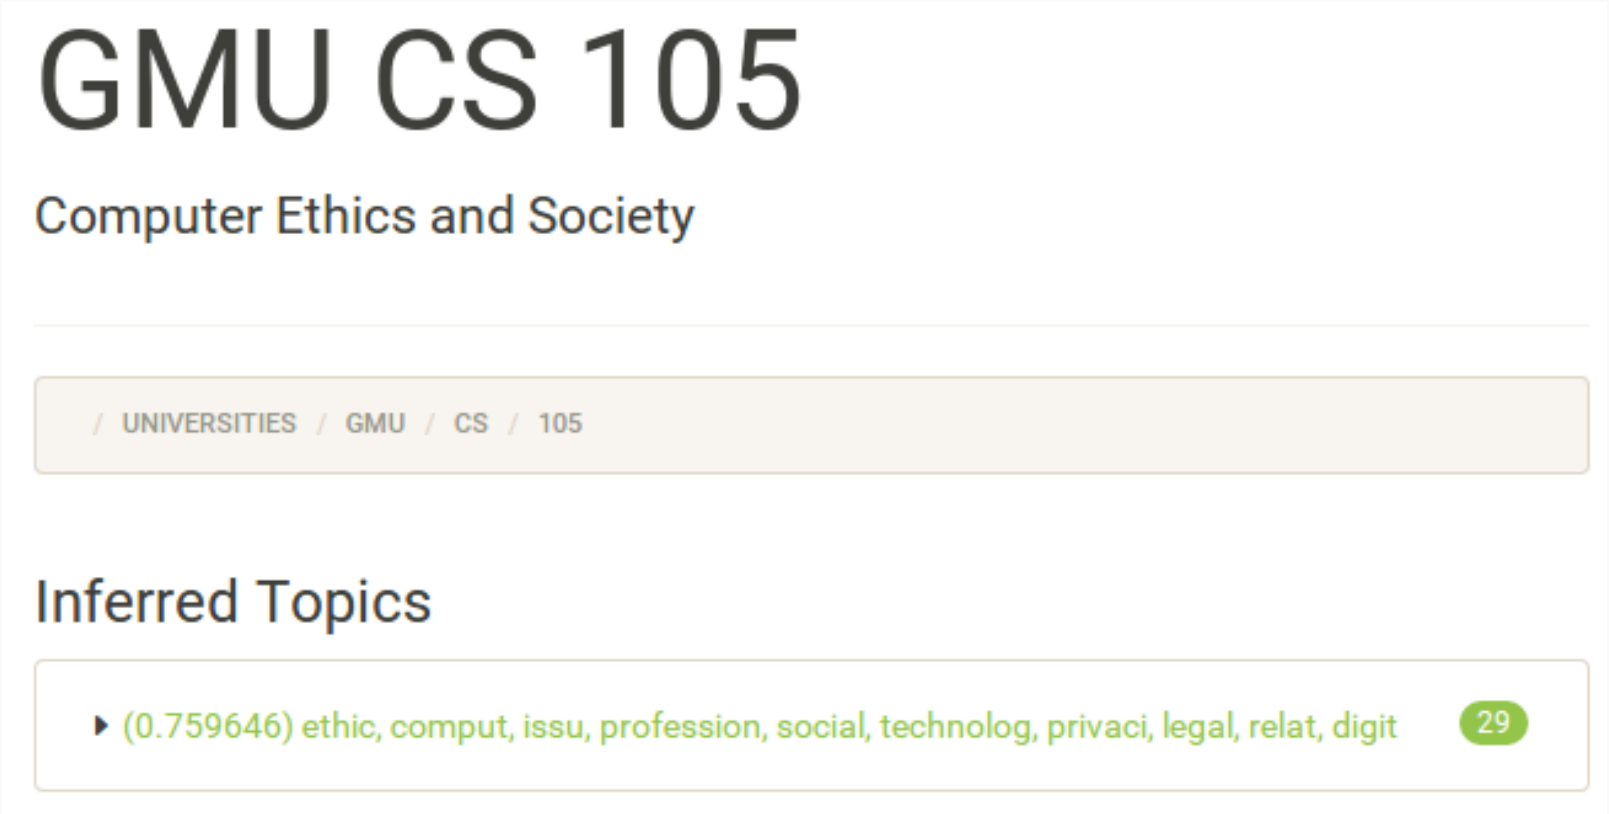

Figure 4: View a course's inferred topics.

Figure 5: View related courses.

Figure 6: View related courses.

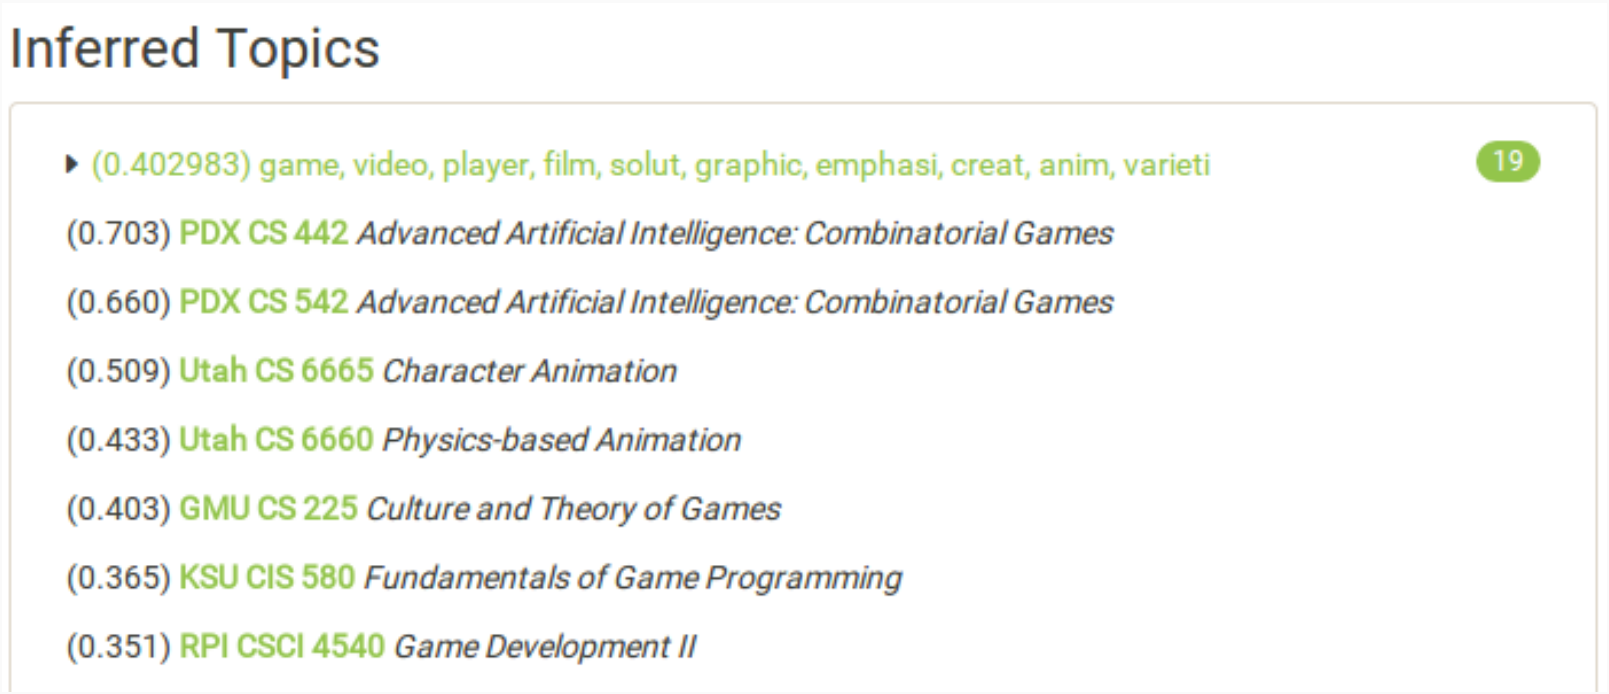

Figure 7: View related courses.

ANALYZE

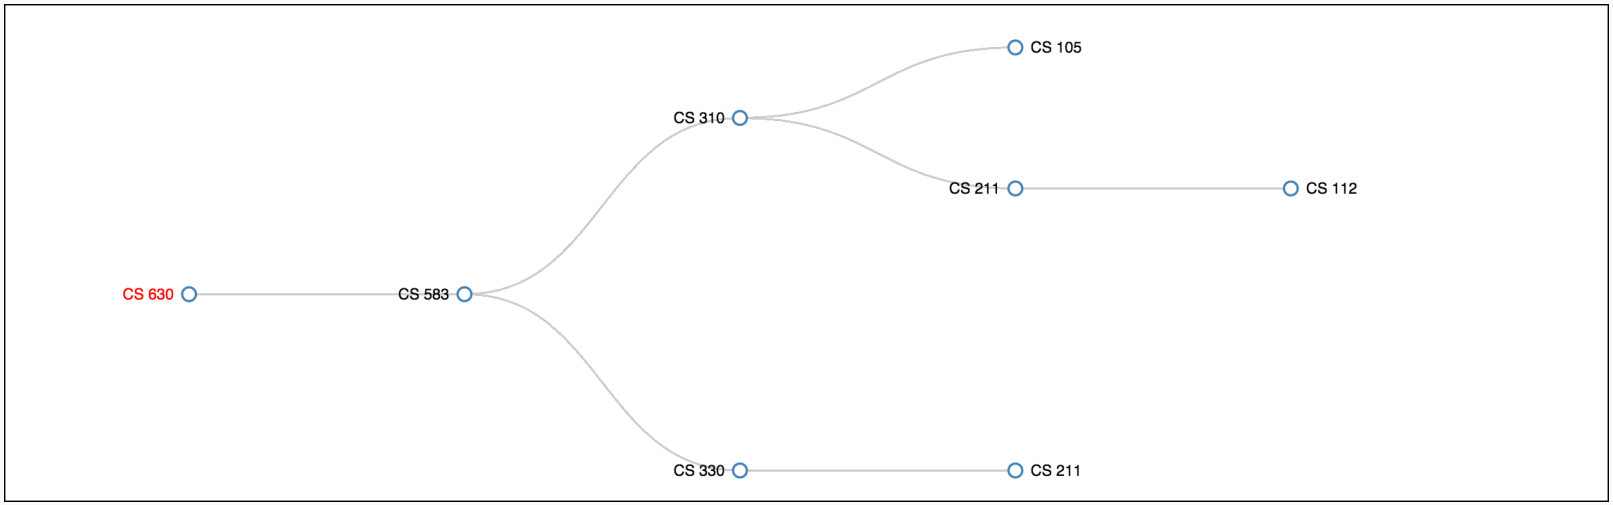

Figure 8: Course prerequisite tree.

COMPARE

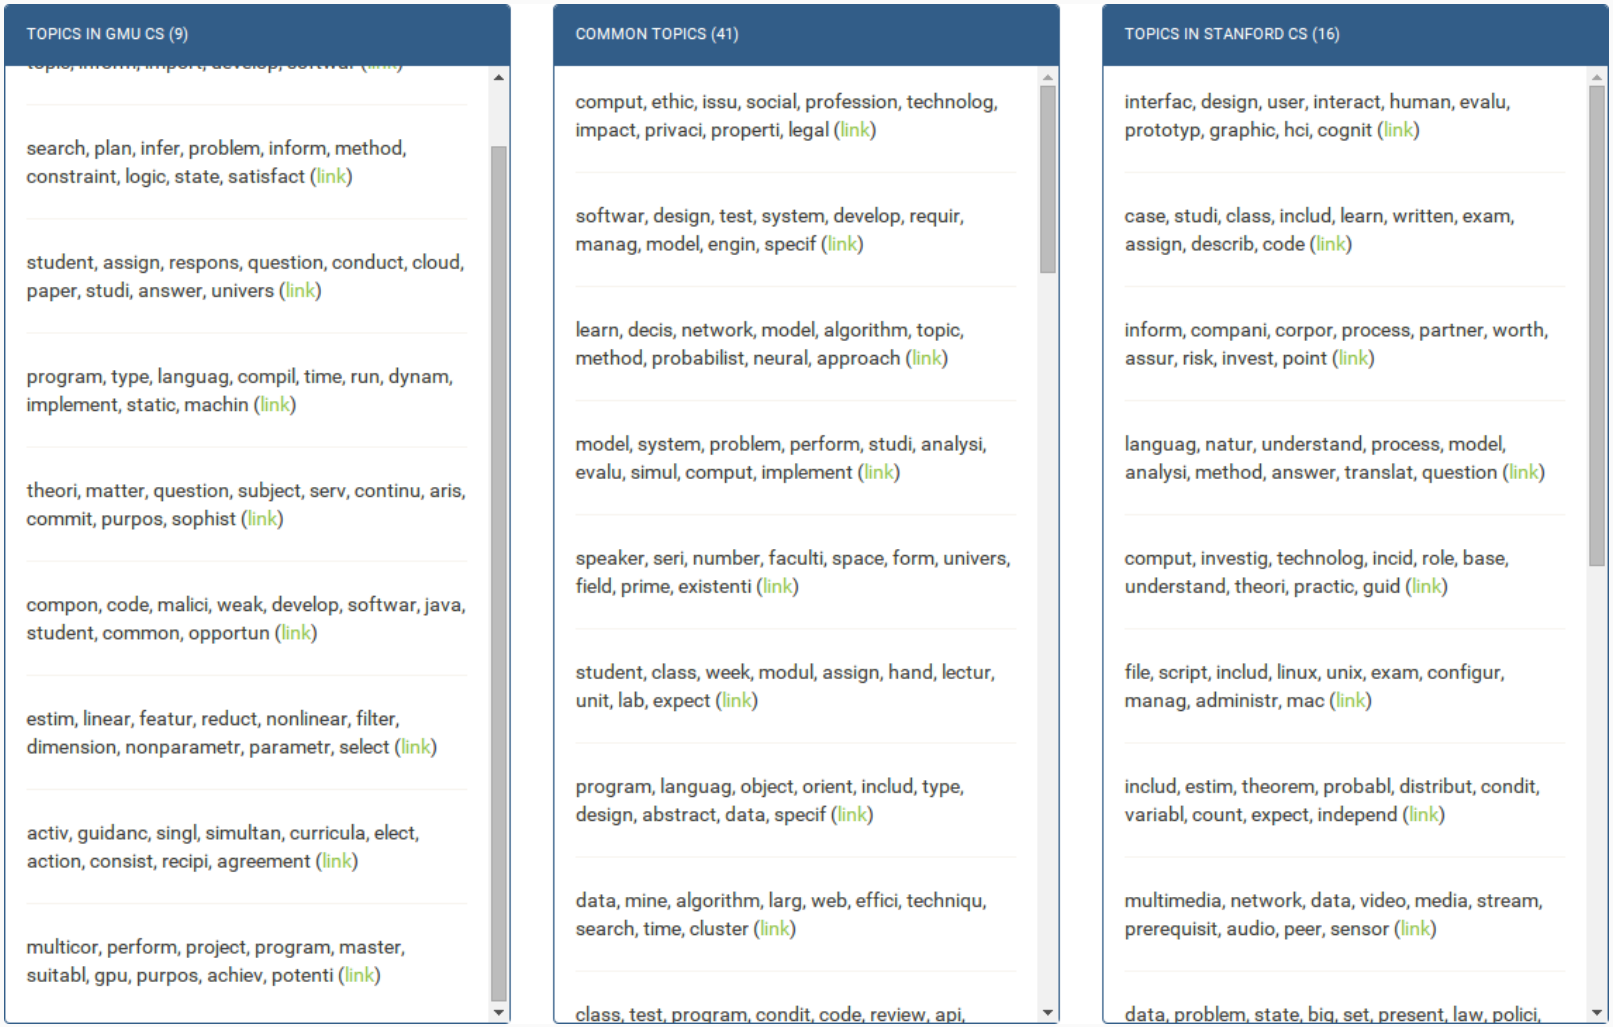

Figure 9: Compare departmental concept coverage.

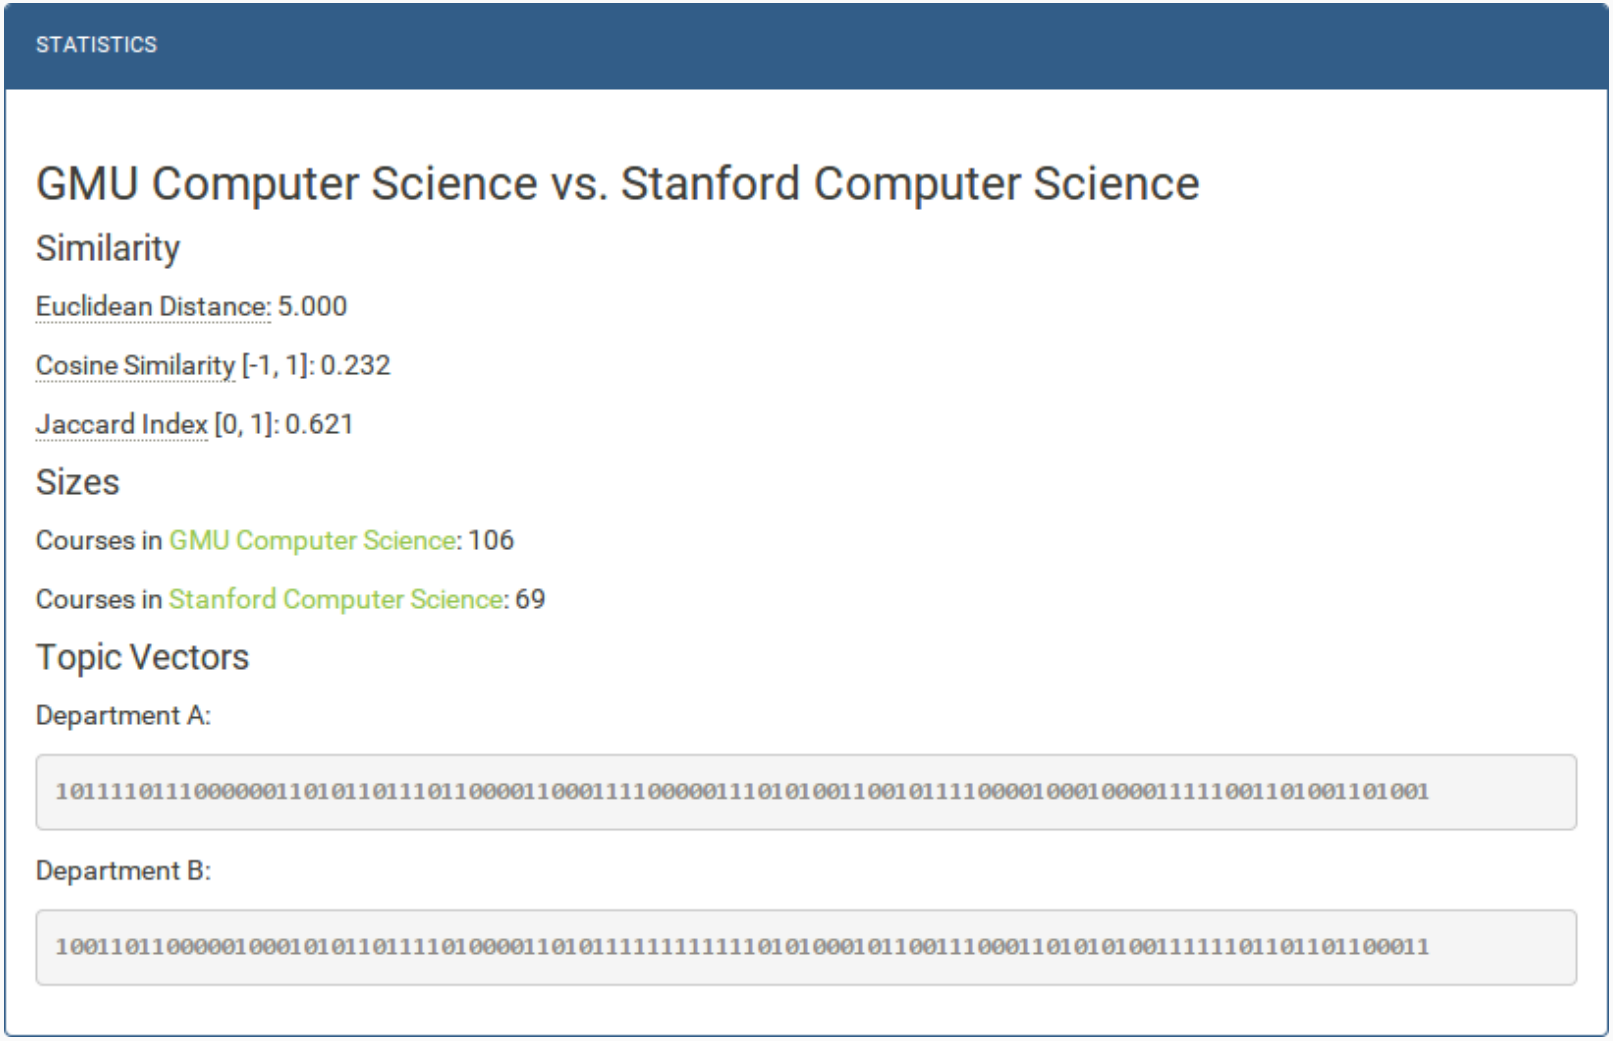

Figure 10: Quick departmental comparison.

EVALUATE

#	Number	Title	Knowledge Areas	Ground Truth
1	100	Principles of Computing	AL, CN, SP, IM	None.
2	101	Preview of Computer Science	CN	None.
3	105	Computer Ethics and Society	SE, SP	IAS, NC, SP
4	112	Introduction to Computer Programming	AL, CN, SP, IM	AL, AR, CN, PL, SDF
5	123	Computing: From the Abacus to the Web	SP	None.
6	211	Object-Oriented Programming	AL, CN, IM	AL, AR, CN, OS, PL
7	222	Computer Programming for Engineers	AL, CN, IM	PL
8	225	Culture and Theory of Games	SP	None.

Figure 11: Evaluate department against Knowledge Areas.

WRAP UP

FUTURE GOALS

- Learning Outcomes meta-analysis
- Alternative methods
- Topic summarization
- Extension to workforce preparedness
- Smart web scraping

SUMMARY

Try out **Trajectory** online at

`trajectory.rouly.net`

Get the source of this presentation and the **Trajectory** project at

`github.com/jrouly/trajectory`

CREDITS

Co-authors:

- Huzefa Rangwala
- Aditya Johri

Presentation theme:

- `github.com/matze/mtheme`

PRIMARY SOURCES

- D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet Allocation. *J. Mach. Learn. Res.*, 3:993–1022, Mar. 2003.
- ACM/IEEE-CS Joint Task Force on Computing Curricula. Computer science curricula 2013. Technical report, ACM Press and IEEE Computer Society Press, December 2013.
- S. Zweben. Computing degree and enrollment trends. Technical report, Computing Research Association, 2011.

Questions?