

Goal-based Course Recommendation

Weijie Jiang^{1,2} (presenter)

Zachary A. Pardos¹ (presenter)

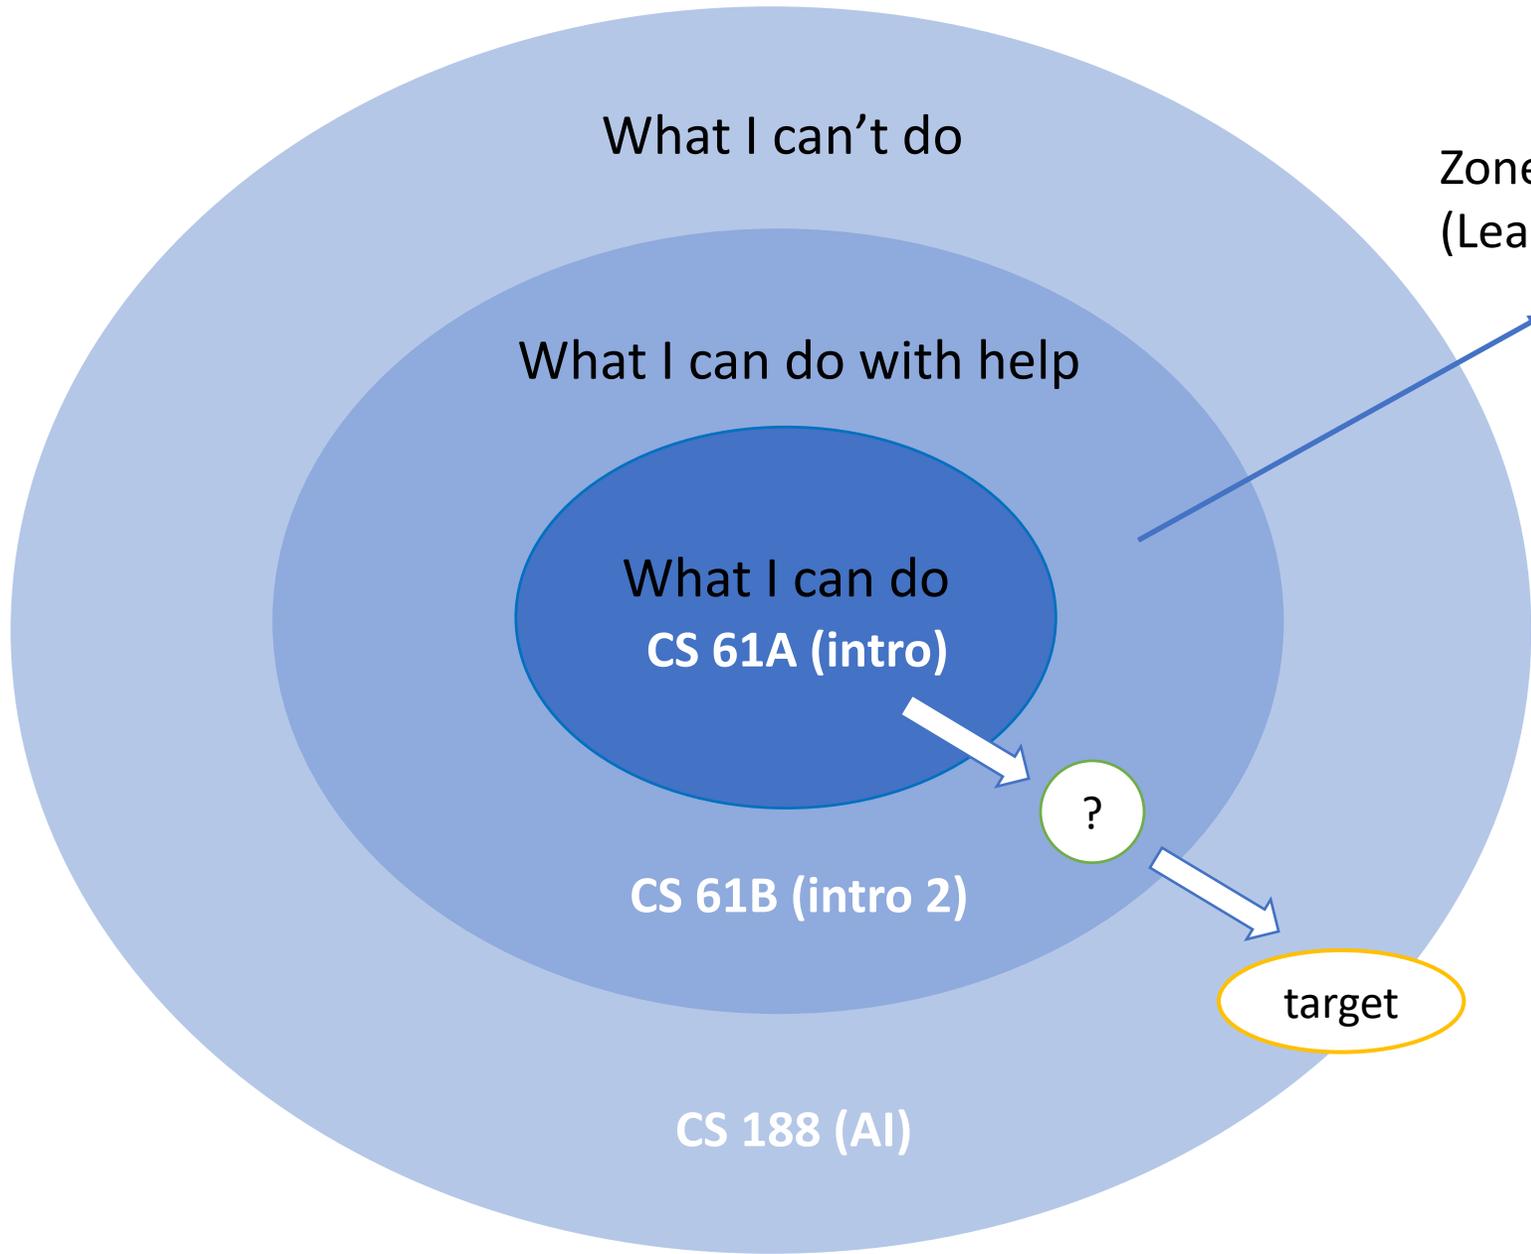
Qiang Wei²

1. *University of California at Berkeley*, 2. *Tsinghua University*



Computational Approaches to Human Learning (CAHL) research lab





What I can't do

What I can do with help

What I can do
CS 61A (intro)

CS 61B (intro 2)

CS 188 (AI)

Zone of Proximal Development (Vygotsky)
(Learners can do with guidance)

Goal of our course recommendation

- Given a specified target course
- Suggest course most likely to bring the target within ZPD

Model requirements for this task

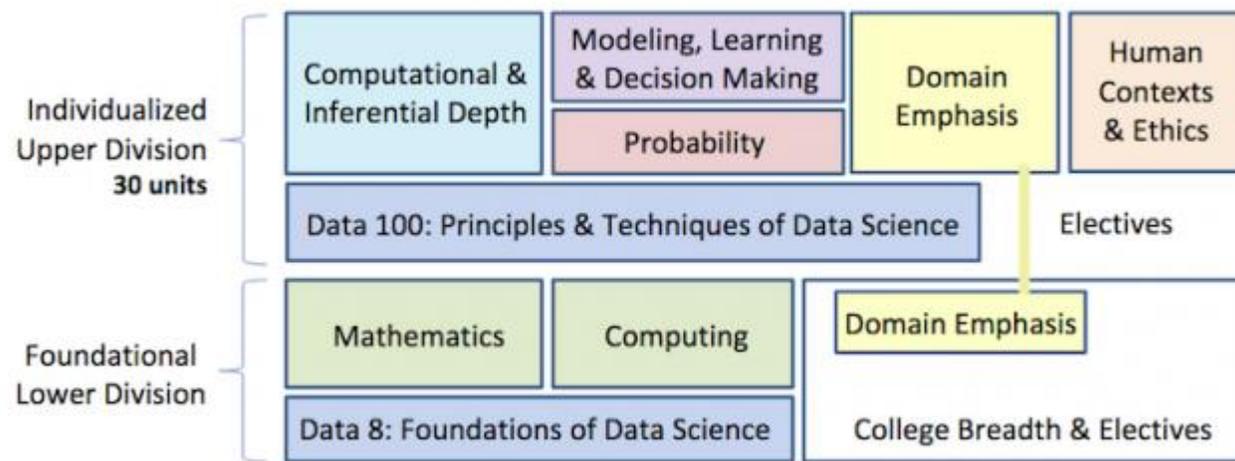
- Estimate a student's ZPD
- Encode prerequisite relationships

target

?

Motivation

DATA SCIENCE

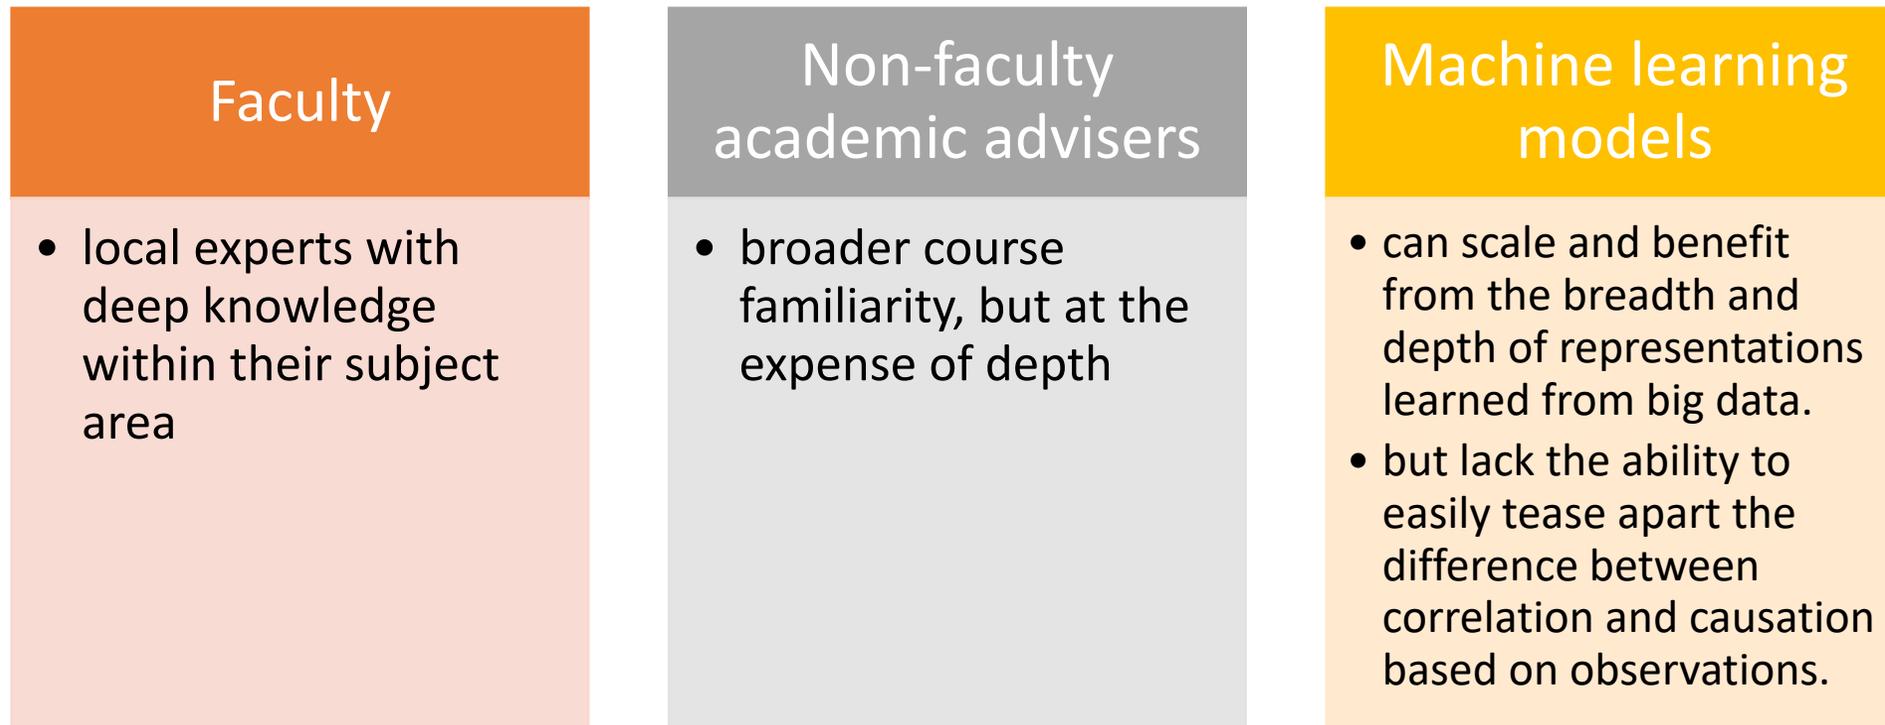


The future of work



Motivation

- The task of recommending an appropriate course personalized to any student's course history and any arbitrary target course is arguably of unreasonable human difficulty.



Both resources are scarce in higher education compared to the number of students enrolled.

 We explore if, given enough constraints, reasonable suggestions can be reliably extracted from such a model.

Motivation

Why not just use course prerequisite listing?

- They may not be comprehensive
- Neglect to include combinations of courses from different departments that together would cover the requirement material.
- Do not take into account what an individual student already knows, and are thus often bypassed by students if not enforced.

Related Work

Student performance prediction

- Models capturing co-enrollment information surpass those only using features of students or of courses in capturing variation in student performance (Gardner and Brooks, 2018)
- Capturing the interaction effect that can occur when a student takes two difficult courses at once (Brown et al., 2018)
- Better discrimination between grade labels could be achieved by binarizing the ordinal grade label (Polyzou and Karypis, 2018)
- Classical collaborative filtering was used to predict course grades, which helps to infer the prerequisite graph meant to potentially help students towards higher rates of on-time graduation (Backenköhler et al., 2018)
 - High accuracy but limited context, only 72 courses from a single department

Related Work

RNN for assessment in education contexts

- Game context, to predict outcomes based on game activity (Akram et al., 2018)
- Math tutoring systems, to predict responses to questions of various skills given response histories. (Piech et al., 2015)
- Produce prerequisite graphs in the same math tutoring contexts as well as in MOOCs. (Chen et al., 2018)
- A deployed campus system to predict the next courses a student is likely to enroll in, given their course taking history. (Pardos, Fan, & Jiang, 2019)

Related Work

- Operationalization of grade prediction models
 - Early-warning type systems meant to signal to students or advisers when struggle is occurring or is imminent. (Hlosta et al., 2017, Zhang and Rangwala et al., 2018)
 - unintended consequences, such as greater course drop-out (Jayaprakash et al., 2014)
 - Showing grade distributions of courses to students and common course sequences for each course before enrollment.
 - Course selection behaviors were not affected by the system but GPA was, leading to an unexpected quarter of a grade point decrease in GPA. (Chaturapruek et al., 2018)

It is important for analytic-based interventions to encourage learners to set goals and for these interventions to strive to be more prescriptive in scaffolding personalized avenues to achieving them.

Our Methodology

- Three assumptions and their respective prediction validations to, collectively, suggest if the approach may warrant testing in the wild.

Assumption 1: Students have a zone of proximal development with respect to course material and that course recommendations should be limited to courses they are expected to be able to succeed in.

- **Validation 1: An RNN-based deep knowledge tracing model** adapted to course grades prediction.

Assumption 2: Such a model of course performance is capable of inferring prerequisite information that can subsequently be used to recommend courses anticipated to be appropriate preparation for a target course.

- **Validation 2: A technique based on course grades prediction model** to infer the university's existing prerequisite courses list, which tests the grade prediction model's ability to infer these existing dependencies.

Assumption 3: The generated recommendations ought to be followed more frequently by students who succeed in a target course than students who underachieve.

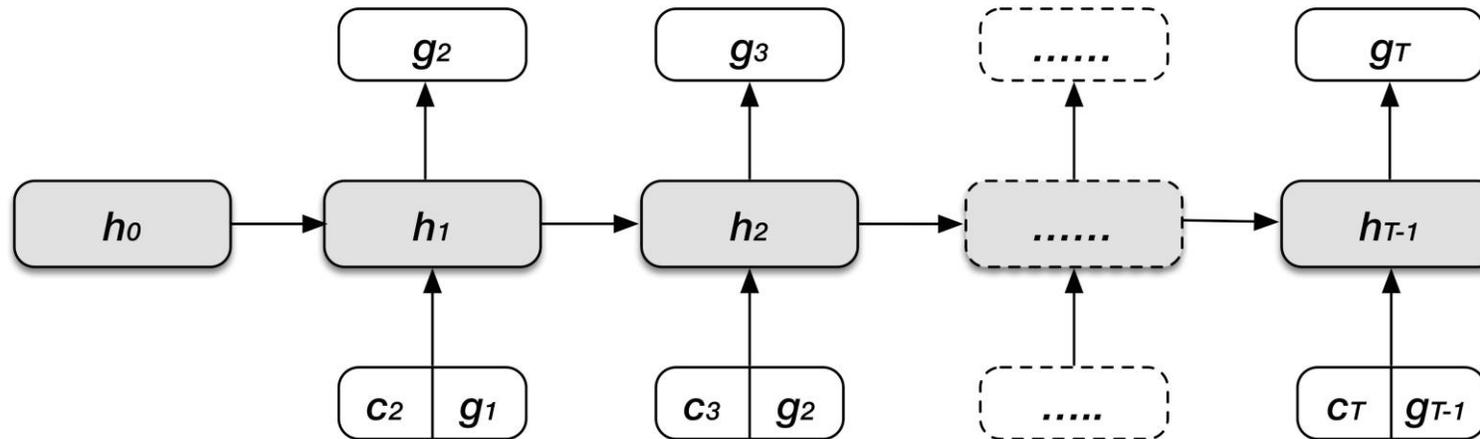
- **Validation 3: Another technique based on course grades prediction model** to predict the previous semester course enrollments before a historically difficult course in the next semester.

Course Grade Prediction Model

- Student Information: enrolled courses, grades and majors in the past semesters.
- Just one model for all the courses
- Input features of three kinds of proposed models
 - Model 1: Course grades
 - Model 2: Course grades + enrolled courses
 - Model 3: Course grades + enrolled courses + student major(s)

Course Grade Prediction Model

- **Students' grades in previous semesters** and **course co-enrollment composition in the current semester** will both impact their performance (Brown et al. 2018).
- Interaction effects among courses enrolled in together
 - The zero sum of student available time and the time demands of each enrolled courses
 - Positive synergistic effective among courses
 - Learning Data Structure and Discrete Math together may reinforce learning between courses because they share similar content.



g_i : grades in semester i , c_i : enrolled courses in semester i

Course Grade Prediction Model

- Training Objective

- **The standard cross entropy loss:** maximize the similarity between the distributions of the model output and the label.

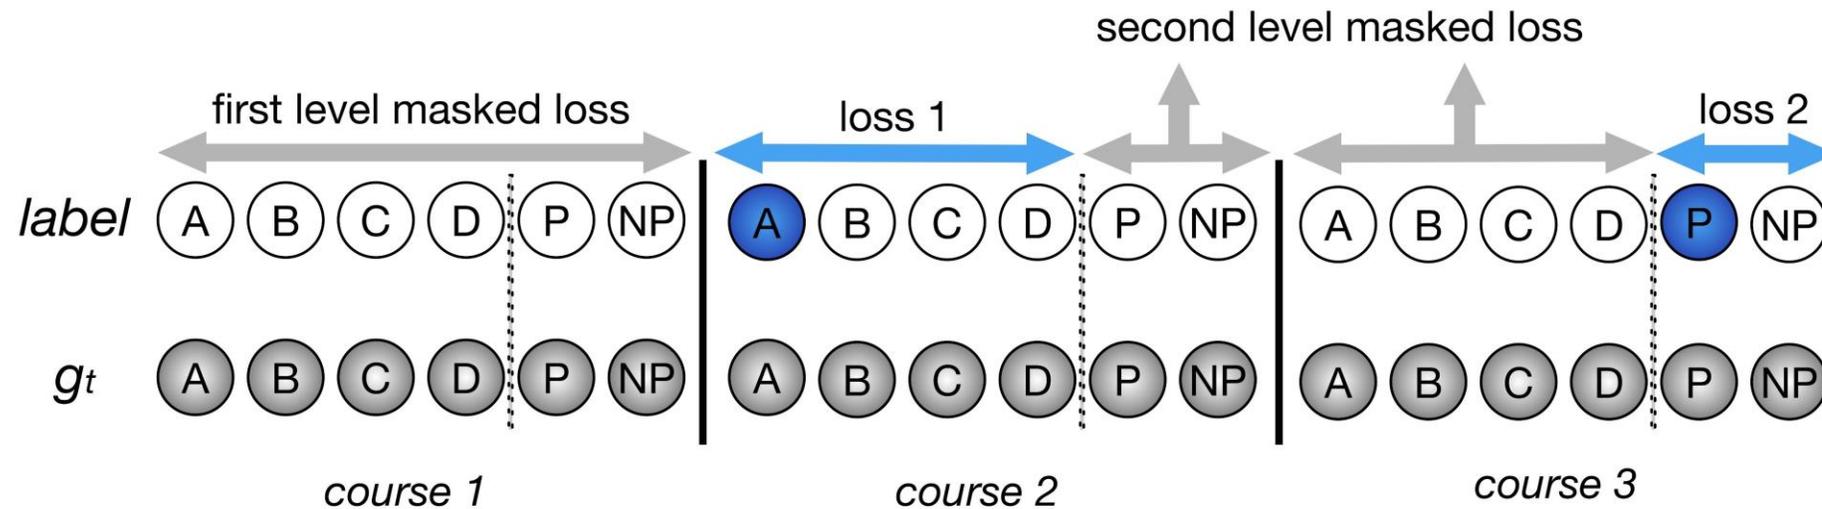
$$loss = - \sum_t \hat{g}_{t+1}^T \log g_{t+1}$$

- Two inappropriateness to use standard cross entropy
 - The predictions of courses not taken in the next semester is meaningless in training.
 - Students can only have a grade either in letter (e.g., ABCD) or in Pass/No-pass grade type

Course Grade Prediction Model

- **Our Masked cross entropy loss**

- Mask out courses not taken
- Mask out grade types not chosen



$$\text{loss} = - \sum_t \sum_{i, \hat{g}_{t+1}^i \neq 0} (g_{t+1}^{i1T} \log g_{t+1}^{i1} + g_{t+1}^{i2T} \log g_{t+1}^{i2})$$

Data Set

- Anonymized student course enrollments from Fall 2008 through Fall 2017 at UC Berkeley.
- 164,196 undergraduates with a total of **4.8 million** course enrollment records.
- 10,430 unique courses, including 9,714 unique primary lecture courses from 197 subjects in 124 different departments hosted in 17 different divisions of 6 different colleges.
- Training set: **Fall 2008 to Fall 2015**, validation set: Spring 2016
Testing set: **Spring 2017**

semester year	student ID (anon)	major	department	course number	grade
Spring 2014	x137905	Law	Law	178	B
Summer 2014	x137905	Law	Law	165	C
Fall 2014	x282243	Math	Math	140	D
Fall 2014	x282243	Math	Math	121	A

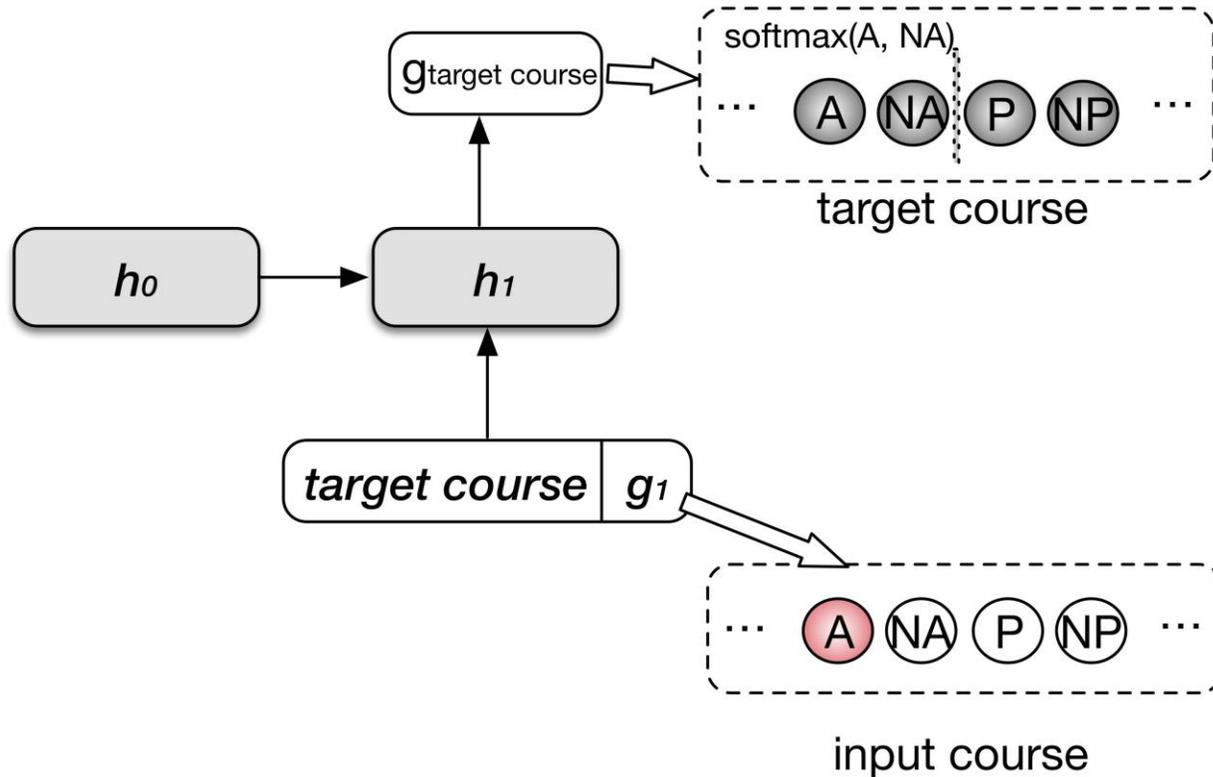
Other attributes: instructor, subgrade, credit ...

Student Grade Prediction

- Students can set the achievement level for the target course to either the A or B level. We refined the input and output of our model by setting this threshold (A or B), and then converted the categorical grades to binary classes, i.e., 'above or equal to grade threshold' and 'below threshold'.

model settings		letter grade		pass/no-pass
threshold	model	accuracy (%)	F-score	accuracy
B	Baseline-B	85.46	None	90.97
B	Model 1	87.75	39.05	90.71
B	Model 2	88.05	42.01	91.78
B	Model 3	87.78	40.21	91.76
A	Baseline-A	50.31	None	83.69
A	Model 1	74.61	55.05	85.42
A	Model 2	75.23	60.24	85.81
A	Model 3	75.19	58.86	86.05

Prerequisite Courses Prediction



Set the target course

Enumerate the prerequisite courses to find out the best to boost the target course to get a high grade.

Prerequisite Courses Prediction

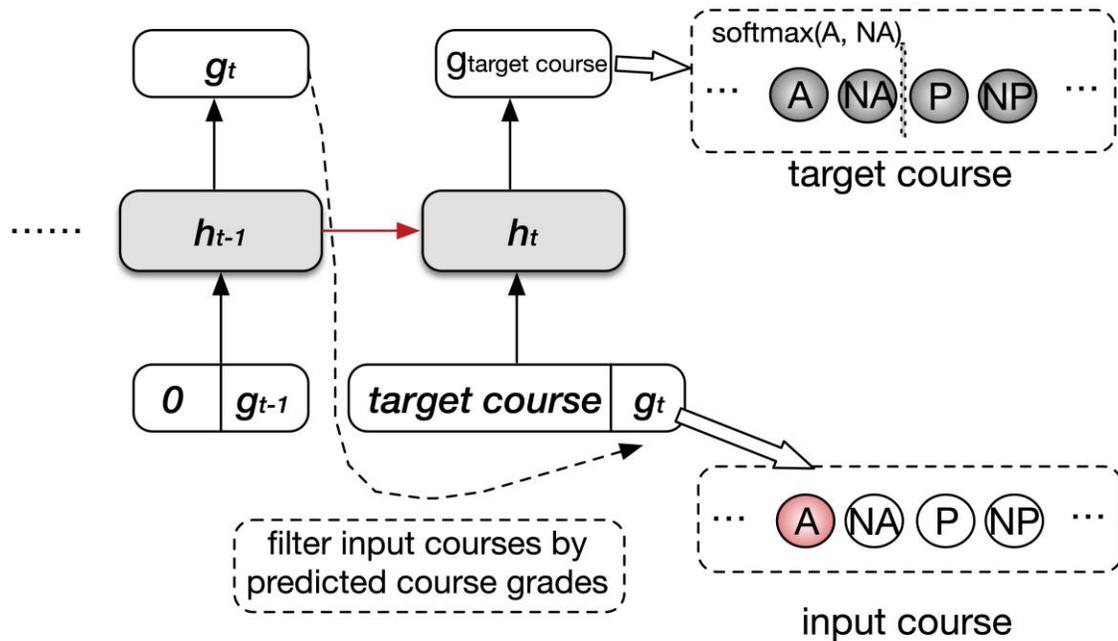
- Evaluation: the university provided a set of 2,300 prerequisite course pairs which contains 1,215 target courses

prerequisite course	target course
188 Computer Science	189 Computer Science
53 Mathematics	189 Computer Science
5 Statistics	266 Computer Science
3A Chemistry	3B Chemistry
1 Economics	100B Economics

model		accuracy (%)	
threshold	model	pairs	target courses
B	Model 2	30.48	44.86
B	Model 3	29.61	43.54
A	Model 2	29.72	43.77
A	Model 3	30.08	45.61

Student Course Selection Prediction

- “To perform well on a target course of interest, which course shall I take the semester before, given my course enrollment and grade history?”



Feed histories into the model.

Set the target course

Knowledge-based filters

Enumerate the prerequisite courses to find out the best to boost the target course to get a high grade.

Student Course Selection Prediction

Filter out

the target course.

the higher level courses by the course number.

courses which are unavailable in the target semester.

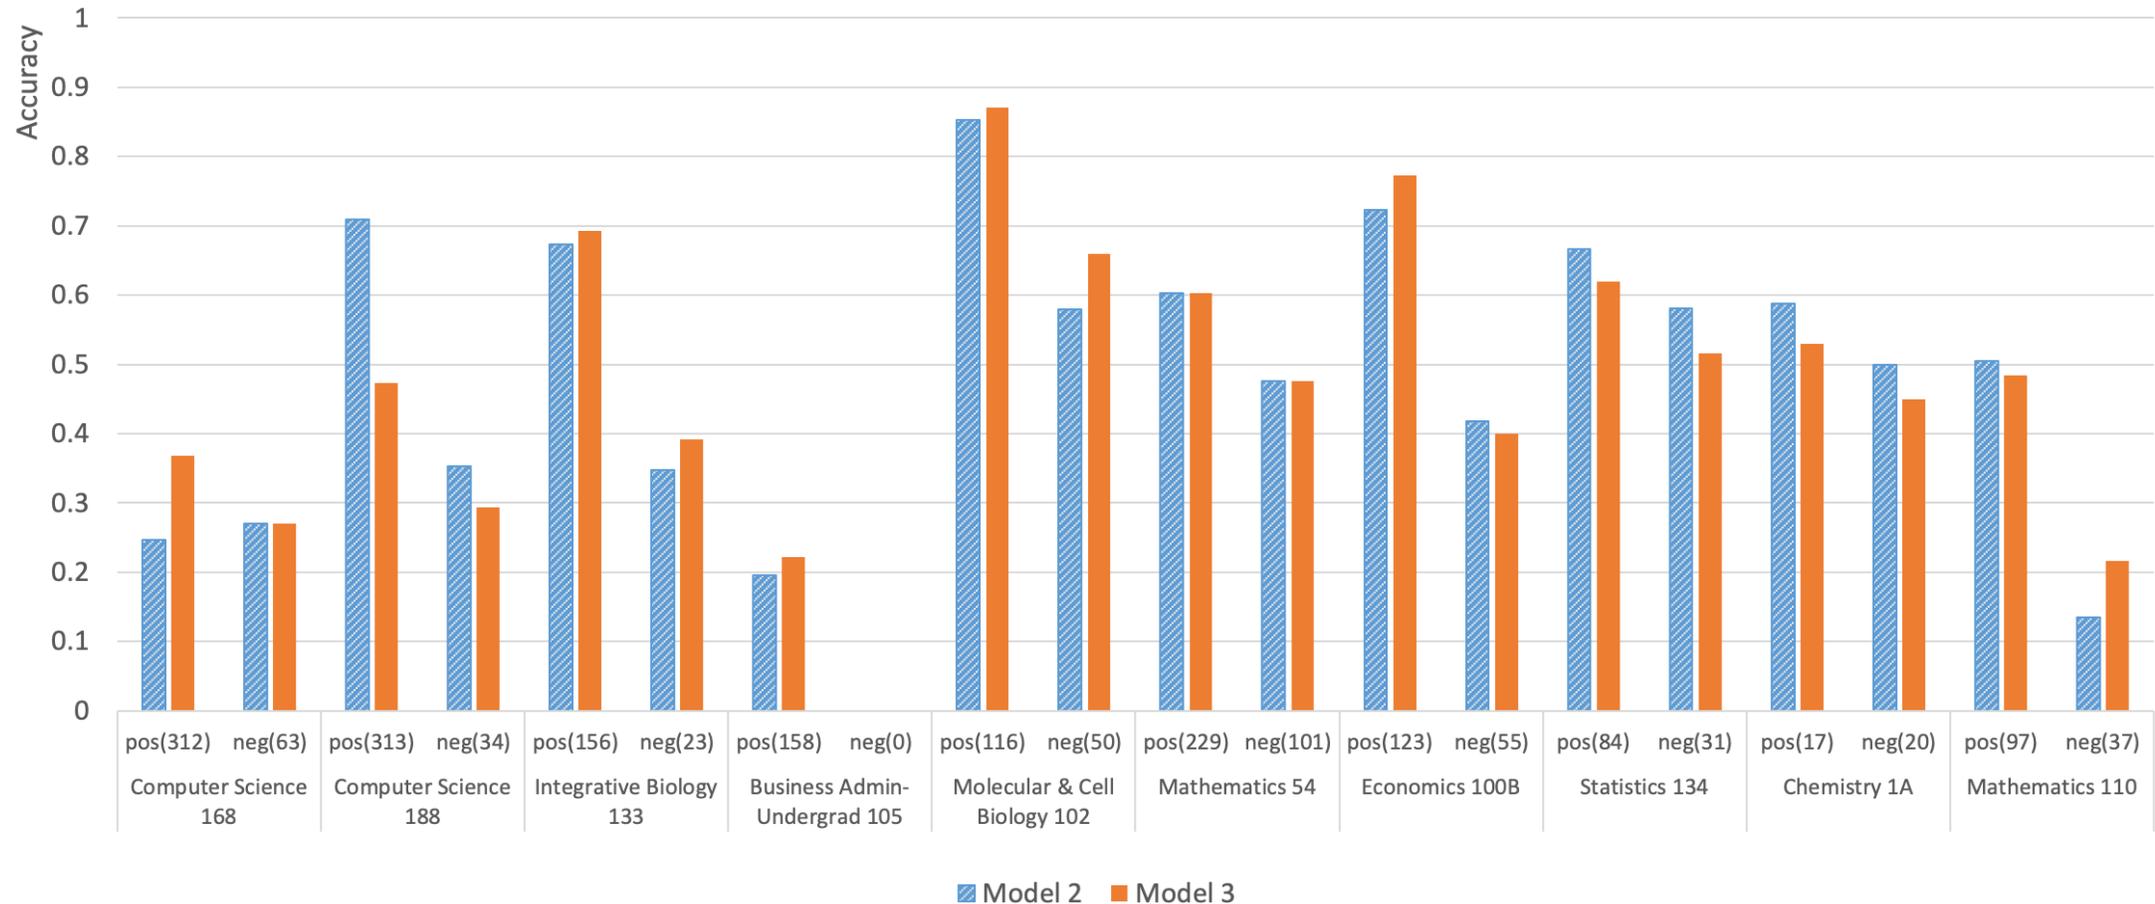
courses the student has already taken.

courses with predicted grades lower than the student's goal. (outside ZPD)

courses both outside target course's own department and never host the prerequisite courses for the courses in the same department as the target course.

Student Course Selection Prediction

- 10 historically difficult courses from different departments across STEM and non-STEM disciplines.
- Evaluate on successful students and underachieving students



Student Course Selection Prediction

model		goal-based		grade prediction	prereq eval
threshold	model	successful students	Successful students – underachieving students	F-score	average
B	Model 2	57.65	21.21	42.01	37.67
B	Model 3	56.34	19.33	40.21	36.58
A	Model 2	36.09	22.93	60.24	36.74
A	Model 3	41.36	22.61	58.86	37.84

- Much higher recommendation accuracy on successful students than underachieving students.
- Models with a B threshold scored considerably higher than the same model with an A threshold (57.65% vs. 36.09%).
- A models are superior in the F-score of grade prediction but lower recommendation accuracy.
 - Maybe A threshold for a preparatory class is too high.

Assumption 1: Students have a zone of proximal development with respect to course material and that course recommendations should be limited to courses they are expected to be able to succeed in.

Contributions

- A novel approach to personalize prerequisite courses for goal-based recommendation based on adaptations of a recurrent neural network.
- The target course can be an arbitrary course offered at the university along with the level of achievement they wish to attain (a grade of A or B)
- Validated three model variants regarding to three assumptions:
 - grade prediction (ZPD)
 - prerequisite inference (prerequisite information)
 - preparation semester course recommendation (better followed by successful students)

Limitations and Future Work

- Infer causal relationships
 - Real-world evaluation to see the efficacy of the recommendations is expensive.
- The high stakes of such a recommendation
 - Need to be highly confident in the algorithm's performance in order to ethnically justify such a real-world evaluation.
 - Additional sources of validation:
 - Instructor ratings of suggested prerequisite courses for their own course given to students with different course histories
 - Academic adviser ratings of such recommendations
 - Ratings from students themselves
- A sequence of curricular preparation may be more desirable.
- Many other goals
 - intended career path
 - graduate on time

Goal-based Course Recommendation

Thank You!

Weijie Jiang (jiangwj@berkeley.edu)
Zachary A. Pardos (zp@berkeley.edu)
Qiang Wei
University of California at Berkeley



Questions?

Paper link: <http://tiny.cc/goal-based>

Code: (see paper)

Acknowledgement: Thank you to the National Science Foundation (award #1446641) and to UC Berkeley staff Andrew Eppig, Tom O'Brien, Rebecca Sablo, and Walter Wong for their assistance with the anonymized data and for sharing their wisdom on the complexities of the enrollment context.