

RESEARCH ARTICLE

Improving models for student retention and graduation using Markov chains

Mason N. Tedeschi^{1✉}, Tiana M. Hose², Emily K. Mehlman³, Scott Franklin⁴, Tony E. Wong^{1*}

1 New College of Florida, Sarasota, Florida, United States of America, **2** School of Mathematical Sciences, Rochester Institute of Technology, Rochester, New York, United States of America, **3** College of Science, Rochester Institute of Technology, Rochester, New York, United States of America, **4** School of Physics and Astronomy, Rochester Institute of Technology, Rochester, New York, United States of America

✉ Current address: Department of Biological Sciences, Texas Tech University, Lubbock, Texas, United States of America

* aewsma@rit.edu



OPEN ACCESS

Citation: Tedeschi MN, Hose TM, Mehlman EK, Franklin S, Wong TE (2023) Improving models for student retention and graduation using Markov chains. PLoS ONE 18(6): e0287775. <https://doi.org/10.1371/journal.pone.0287775>

Editor: Philipp Bitzenbauer, Leipzig University, Universität Leipzig, GERMANY

Received: January 31, 2023

Accepted: June 13, 2023

Published: June 26, 2023

Copyright: © 2023 Tedeschi et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The code accompanying this work is available from https://github.com/tonywong/LA_Markov_analysis. Raw data contains private student information and may not be shared. Processed summary data may be available directly from the authors, subject to Rochester Institute of Technology IRB approval. Raw data may be obtained via the RIT Human Subjects Research Office (hsro@rit.edu) for researchers who meet the criteria for access to confidential data.

Abstract

Graduation rates are a key measure of the long-term efficacy of academic interventions. However, challenges to using traditional estimates of graduation rates for underrepresented students include inherently small sample sizes and high data requirements. Here, we show that a Markov model increases confidence and reduces biases in estimated graduation rates for underrepresented minority and first-generation students. We use a Learning Assistant program to demonstrate the Markov model's strength for assessing program efficacy. We find that Learning Assistants in gateway science courses are associated with a 9% increase in the six-year graduation rate. These gains are larger for underrepresented minority (21%) and first-generation students (18%). Our results indicate that Learning Assistants can improve overall graduation rates and address inequalities in graduation rates for underrepresented students.

Introduction

University decision-makers invest in educational programs to support the success of diverse students. Assessing programmatic efficacy, however, is made complicated by complexities in program implementation and measurement of program outcomes. Administrators often measure programmatic efficacy using a single metric of student success, such as course failure rate (commonly rates of D, F, or withdrawal; DFW), and make decisions about funding or program longevity using this information. While this strategy provides an important view of one aspect of program efficacy, the analysis is limited to a single point in a student's academic career. This limitation risks missing other important student success outcomes across longer timescales, including retention and graduation rates, that the educational program may impact. Thus, there is a need for statistical methods to incorporate measures of longer-term, or “downstream”, student success into program evaluation metrics.

Additionally, if an educational program has a small scope (e.g., implementation in a single section of a course) or administrators are interested in the program's impacts on students

Funding: M.T., S.F., and T.W. were supported in part by a grant from the National Science Foundation, DUE #1757477 (https://www.nsf.gov/awardsearch/showAward?AWD_ID=1757477). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. There was no additional external funding received for this study.

Competing interests: The authors have declared that no competing interests exist.

from minoritized identities (e.g., women in science majors or underrepresented minority students), then program assessment is challenged by small sample sizes. Smaller sample sizes lead to lower confidence in measures of student success such as DFW or retention rates. This creates a need for methods to estimate more accurate statistics for program efficacy metrics with small sample sizes.

When program managers and administrators use graduation rates as a measure of program success, there is an inherent discrepancy between the students who are currently engaged in the particular program under study, and the students whose data are part of the graduation rate calculation. Specifically, very few of the students who are currently in an academic program are involved in a “current” estimate of the program’s graduation rate. Most of these students have departed the program, with or without a degree (or other marker of successful completion of the program). Thus, there is a benefit to developing an “online” estimate of graduation rates—that is, an estimate of graduation rates that can be updated regularly [1] and includes data from all cohorts of students, including those who are still enrolled in the program. Such an online estimate of graduation rate is arguably more appropriate for program evaluation in that it makes use of the most up-to-date information on student retention and persistence. A traditional estimate of six-year graduation rate (SYGR), on the other hand, relies entirely on data that is at least six years old. If, five years ago, there was a sizable improvement in first-year retention in the academic program under study, then the six-year graduation rate as estimated in the following year will be higher as a result (all other things being equal). This highlights the strength of the online estimate, in that it can incorporate all available information on year-to-year transitions.

As Markov chains model transitions of a system between states, they provide a natural mathematical structure to represent the year-to-year transitions of students between academic states (e.g., freshman year, sophomore year). A (first-order) Markov model is one in which the probability of the system transitioning from one state to another is conditionally independent of all previous states of the system, given the current state. Previous work has leveraged Markov chains to construct estimates of college students’ retention and persistence [1–3]. Boumi and Vela [1] further allowed for reflexive transitions from a state to itself; that is, students can remain in the same academic level at multiple points in time. This “multi-level absorbing Markov chain” allowed multiple visits to a state either in sequence (e.g., a student remaining in the sophomore state for more than one academic year) or non-contiguous (a student withdrawing for multiple, non-contiguous semesters then re-enrolling). Here, the term “absorbing” refers to the fact that once a student enters either the Graduation state or the Drop-Out state, they are assumed to remain in that state. Thus, in the parlance of Markov chains, these system states are said to be absorbing states. By contrast, each academic year is a transient state because students pass through these states en route to a permanent outcome (i.e., one of the two absorbing states). However, the introduction of reflexive transitions by Boumi and Vela [1] leads to an underestimation bias in Markov-based estimates of SYGR. Those authors develop a correction to account for this bias, but to avoid this issue, we focus on graduation within six calendar years of students’ matriculation.

Here, we assess the ability of a Markov chain model to constrain estimates of SYGR by comparing its performance to that of a traditional graduation rate calculation. We use eight years of data from “UNI”, a large, private university in the northeast of the United States, to compute SYGR estimates. SYGR metrics are typically used as descriptive statistics as opposed to a prediction, although there are instances where one might take a predictive view of SYGR. One example is a prospective student who is using SYGR to compare the relative strengths of two candidate academic programs. Such a student is assuming that the past performance (SYGR) is an indicator of future performance (their anticipated academic experience and likelihood of

successful completion of a degree). By contrast, the typical descriptive interpretation of SYGR measures the academic performance of a specific cohort from (at least) six years ago. In practice, a traditionally-calculated SYGR is used as a point estimate, with no quantification of uncertainty. However, this metric is only exact for the cohort on which it is based. These students necessarily entered the university at least six years ago, and may have departed the university a year ago or more. Mapping this traditionally-calculated SYGR to the present day leads to the interrelated issues of uncertainty stemming from variations in SYGR over time and sensitivity of the calculated SYGR to potentially small cohort sizes. While decision-makers typically do not quantify uncertainty in SYGR, for the reasons outlined above, we characterize these sensitivities and uncertainties by constructing bootstrap confidence intervals.

We hypothesize that the use of more data in the Markov estimates will lead to more confident estimates of SYGR than a traditional calculation. This ability to use more data makes an online estimate of graduation rates suitable for small groups of students. We demonstrate the power of the analysis by using it on a localized academic program, the Learning Assistant (LA) program, and smaller demographic sub-groups within the program. However, we stress that an online estimate of graduation rates has no drawback relative to a traditional calculation. The online estimate will make use of all of the same information as a traditional estimate, as well as the year-to-year transition information from not-yet-departed students. This makes the online estimate appropriate for other settings as well. We assess, through our real-world data from UNI, the benefits of using the Markov model to estimate SYGR for smaller, minoritized groups of students (underrepresented minorities and first-generation college students). Finally, using the LA program at UNI as a test case, we evaluate the impacts of students' exposure to LAs in terms of improvements to SYGR for students who have experienced LA support in a high-attrition science course.

Methods

Data and context

We analyze the academic progress of students at UNI, the second largest producer of STEM graduates among all private universities in the United States. UNI is a primarily undergraduate institution with approximately 15,000 undergraduate and 3,000 graduate students. The student body is rich in diversity, with more than 2,500 international students and 2,900 African American, Latino American, and Native American (AALANA) students. UNI services a substantial deaf/hard-of-hearing student population and has developed extensive accommodation resources, with over 1,500 deaf/hard-of-hearing students taking courses alongside their hearing peers. Data for this study includes all UNI students, and comes from the UNI Office of Institutional Data. The data include previous and subsequent academic records, program (major) at the time of each class, and, if the student graduates, the graduation year and degree obtained. All data are fully anonymized prior to use in our analysis and the requirement for informed consent was waived for this retrospective study.

We demonstrate the power of Markov analysis by using it to evaluate the effectiveness of a localized program, UNI's Learning Assistant (LA) program. In this program, UNI hires undergraduate Learning Assistants (LAs) to support instructors using evidence-based, student-centered learning practices in their courses. LAs receive training in pedagogical knowledge and pedagogical content knowledge through planning sessions with faculty and a pedagogy course [4, 5] to support them as effective near-peer educators in the classroom. STEM faculty use LAs to implement evidence-based teaching strategies known to increase learning [6, 7] and lower DFW rates [7–10]. A scoping review [11] of 39 peer-reviewed studies found that LA-supported courses were associated with improvements in higher-order cognitive skills [12] and improved

conceptual learning [6], a result that did not depend upon a single set of curricular materials (e.g., [13–15]). DFW rates in LA-supported courses are lower than in non-supported courses [7, 9, 10], with increased impacts on students from underrepresented demographic groups [8]. Importantly, previous work has not closely examined how LA support relates to longer-term metrics such as graduation and retention rates, nor has a direct comparison between concurrent supported and unsupported classes been possible. At UNI, the LA program is implemented in a subset of sections of particular courses, and Markov analysis enables a comparison between the performance of the LA-supported students and the “control” group of students in unsupported sections. In this way, UNI’s LA program naturally gives rise to contexts in which improved program efficacy metrics for small data situations could provide decision-makers with more confident estimates of student success.

Traditional estimate of six-year graduation rate

Traditional estimates of graduation rates require first identifying a student cohort of interest. Suppose N_{start} is the initial number of students who matriculate in that cohort. Among those students, let N_{deg} represent the number of students who obtain a degree within six years. Then $\text{GR}_{\text{tr}} = N_{\text{deg}}/N_{\text{start}}$ is the estimate of the SYGR for this cohort of interest. A key limitation of using this traditional approach is that N_{start} is frequently quite small when the cohort of interest is an underrepresented group of students, or a relatively small subgroup of students who experience an academic intervention.

Markov chain-based estimate of six-year graduation rate

Model. In contrast to the traditional SYGR calculation, we employ a Markov model that estimates the SYGR by leveraging the probabilities of students persisting from year-to-year, conditioned on their initial class status (e.g., freshman, sophomore).

Define $S = \{1, 2, 3, 4, 5, 6, D, G\}$ as the state space for our Markov chain model, where the integers 1, 2, …, 6 represent the numbers of years that a student has been enrolled at UNI (e.g., students in state 2 are generally sophomores), D represents students who have dropped out and no longer are pursuing a degree, and G represents students who have earned a degree. States D and G are both absorbing states, while the others are transient.

Suppose that N_i represents the number of students in state i at the beginning of a year, and let N_j represent the number of those students who are in state j by the end of the year. Then the transition probability from state i to state j is

$$P_{ij} = P(X_{t+1} = j | X_t = i) = \frac{N_j}{N_i}. \quad (1)$$

This probability is, by definition, 0 if j is not in $\{i+1, D, G\}$. Additionally, by definition, $P(X_{t+1} = j | X_t = D) = 0$ if j not equal to D , and similarly $P(X_{t+1} = j | X_t = G) = 0$ if j not equal to G . Thus, the overall structure of the transition probability matrix is sparse (Table 1). The transition model is depicted visually in S1 Fig in S1 File.

The six-step transition probability of a student arriving in the graduation state (G) within six years of matriculation (originating in state 1) are $P^6_{1G} = P(X_{t+6} = G | X_t = 1)$. These provide an estimate of SYGR: $\text{GR}_{\text{ma}} = P^6_{1G}$, the row 1, column G (eighth column) of the sixth power of the transition probability matrix P .

A key strength of the Markov modeling approach is the ability to combine information from multiple cohorts of students. For example, let A and B represent two cohorts of interest. Let $N_{A,i}$ and $N_{B,i}$ represent the numbers of students from cohorts A and B who are in state i at the beginning of a year. Similarly, let $N_{A,j}$ and $N_{B,j}$ represent the numbers of students who are

Table 1. Markov model transition probability matrix.

	1	2	3	4	5	6	D	G
1	0	P_{12}	0	0	0	0	P_{1D}	P_{1G}
2	0	0	P_{23}	0	0	0	P_{2D}	P_{2G}
3	0	0	0	P_{34}	0	0	P_{3D}	P_{3G}
4	0	0	0	0	P_{45}	0	P_{4D}	P_{4G}
5	0	0	0	0	0	P_{56}	P_{5D}	P_{5G}
6	0	0	0	0	0	0	P_{6D}	P_{6G}
D	0	0	0	0	0	0	1	0
G	0	0	0	0	0	0	0	1

General structure of the transition probability matrix for state space {1, 2, 3, 4, 5, 6 (year at university), D (Drop out), G (Graduate)} to represent the conditional probability of the system transitioning to a final state (column), given that the system starts in a particular initial state (row).

<https://doi.org/10.1371/journal.pone.0287775.t001>

in state j by the end of that year. Then the combined estimate of the transition probability from state i to state j is

$$P_{AB,ij} = \frac{N_{A,j} + N_{B,j}}{N_{A,i} + N_{B,i}}. \quad (2)$$

Using Eq (2) enables the Markov model to include the most up-to-date data. For example, as of this writing, students who matriculated in Fall 2021 have only been at university for one full year. Using traditional methods, a SYGR estimate that includes these students will not be available for another five years. Their data, however, can be incorporated into the Markov model for SYGR by influencing the estimates of P_{12} , P_{1D} , and P_{1G} . For such a “partial data” cohort, they do not influence any of the subsequent transition probability estimates.

Validation. We first perform a “positive control” experiment to demonstrate that a reduced form of the Markov model produces results that match the traditional SYGR calculation. For each of the cohorts matriculating in Fall 2013, Fall 2014, and Fall 2015, we construct a reduced form of the Markov model by only using the Fall 2013, 2014, or 2015 cohorts’ data (respectively) to construct each Markov transition probability matrix (Eq 1). In this way, these models employ the same information that traditional estimates of SYGR use for those same cohorts. We compare the graduation rate estimates from these reduced form Markov models to traditional graduation rate estimates as a way to validate that the Markov model faithfully reproduces the traditional baseline.

In contrast to this positive control, the full form of the Markov model involves a single full data cohort, and five partial data cohorts. For example, a program manager in Fall 2021 could use the full six years of data for the Fall 2015 cohort in order to estimate the SYGR using the traditional approach. They could also employ the Fall 2016, 2017, 2018, 2019, and 2020 cohorts’ information in an estimate of SYGR using the Markov chain approach. Further, the Markov chain model provides a natural structure (Eq 2) to combine all of the available information into a single estimate of SYGR by also including the Fall 2013 and 2014 cohorts’ data. This constitutes the “full” Markov model for SYGR.

Bootstrapping uncertainty characterization

To test our hypothesis that a Markov model will lead to more confident estimates of graduation rates than a traditional model, we quantify uncertainty in our estimates of SYGR by constructing confidence intervals using bootstrap resampling. We do this for both the traditional

and the Markov chain models. We resample with replacement rows from the original overall student data set, resampling a number of rows equal to the original size of the data set. For each resample, we compute the SYGR using the traditional (or Markov) approach. We use ensembles of 1,000 resamples and estimate a 95% confidence interval as the 2.5–97.5% range in the ensemble of computed SYGR. Experiments using different sizes of resample ensembles indicate this ensemble size is sufficiently large that our results are not sensitive to this choice.

Subsets of students

We demonstrate the strength of the Markov model for SYGR for studying relatively small cohorts of students by examining how SYGR differs for students in courses supported by undergraduate LAs and those without such LA support. Further, we examine how this impact of LA support may have differential effects for underrepresented minority students and first generation college students. All of these groups (LA-supported, underrepresented minority, and first generation) are relatively small on their own, so their intersections provide an ample proving ground for the benefits of the Markov model for SYGR (e.g., LA-supported first generation college students). For these experiments using smaller subgroups of students, we do not present any results using the traditional approach to compute SYGR because the benchmarking experiment (Sec. 2.3.2) will demonstrate the fidelity with which the Markov model estimates SYGR. We use Cohen's h to quantify the effect size differences in graduation rates between the student subsets who have participated in LA-supported courses (p_{LA}) and those who haven't (p_{noLA} ; Eq 3) [16].

$$h = 2[\arcsin(\sqrt{p_{LA}}) - \arcsin(\sqrt{p_{noLA}})] \quad (3)$$

To control for the variety of students' major disciplines, and the variety of courses that LA support is incorporated into, we include the LA-supported group only students who encounter LA support in high-attrition courses within the College of Science at UNI. These include introductory physics and mathematics courses, such as precalculus and calculus (see Supporting Information). We note that other high-attrition courses are present, but few or no sections that incorporate LAs. For these experiments, we restrict our data set to only students who take at least one of these courses between Fall 2013 and Summer 2021. We note that i) the LA and no-LA groups of students are demographically similar, ii) students have no advance knowledge of which sections of a course will be LA-supported when they enroll, and iii) aside from LA support, there are no further systematic differences between LA-supported sections of a course and non-LA-supported sections (e.g., sections use the same syllabus).

Students who encounter LA support in at least one of the aforementioned courses contribute to the Markov-based calculation starting in the first year in which they have LA support. In this way, the transition probabilities in the Markov transition matrix for the LA-supported students should only represent transitions which would have been affected by the LA support "treatment". By contrast, the "control" group of students includes all students who have no exposure to LA support in the high-attrition science courses (see above). The numbers of students and class sections each semester, and the breakdown between LA-supported and non-LA-supported, is provided in Supporting Information (S1 Table in [S1 File](#)).

We note that separating students based on LA status (LA-supported or not) and examining the groups' differences in graduation rates establishes a correlative or associative relationship, but cannot prove a causal one. This is, of course, the nature of all such treatment-versus-control type experiments (e.g., vaccine trials). The deep existing literature documenting the benefits of LA support for students across different institutional profiles and disciplines—including transferable and durable gains in higher-order cognitive skills, conceptual learning, and

Table 2. Validating graduation rates from the reduced-form Markov model.

Cohort	Method	2.5th percentile	Median	97.5th percentile	95% CI width
Fall 2013	Traditional	66	71	75	8
	Markov	66	71	75	9
Fall 2014	Traditional	66	71	75	9
	Markov	67	71	75	9
Fall 2015	Traditional	69	73	77	9
	Markov	69	73	77	8
All combined	Traditional	69	72	74	5
	Markov	69	72	74	5

Estimates of six-year graduation rate (%) from the traditional calculation and the reduced form (full-data cohorts only) Markov model. Numbers are rounded to the nearest percentage point, so the 95% CI width reported may not appear to equal the difference between the 97.5th and 2.5th percentiles.

<https://doi.org/10.1371/journal.pone.0287775.t002>

problem-solving skills (see Sec. 1)—provides strong reason to hypothesize that students' exposure to LA support should be related causally to their likelihood of graduation within six years.

Results

Benchmarking against traditional graduation rate calculation

When only cohorts with a full six years of data are used to compute SYGR using the Markov method, estimates of GR_{ma} from the reduced form Markov model faithfully match the traditional graduation rate estimates, GR_{tr} (Table 2). To within a percentage point, the medians and 2.5–97.5% bootstrap confidence interval bounds (and widths) are all consistent for the Fall 2013, Fall 2014, and Fall 2015 cohorts, as well as when all three cohorts' data are combined (using Eq 2). These results provide strong evidence that the Markov model is a valid approach to estimate SYGR. For the remainder of this work, the Markov model uses all student data that would be hypothetically available at the time when a traditional SYGR estimate could be computed. For example, the Markov-based SYGR for the Fall 2013 cohort will include five years of data for the Fall 2014 cohort, four years of data for the Fall 2015 cohort, and so forth. The “all combined” model includes all available student data in our data set from Fall 2013 to Summer 2021 (the last academic term in our data set). We use “percentage point” to refer to ranges and absolute differences in the SYGR to avoid confusion with relative changes or relative differences in SYGR.

Quantification of uncertainty

The Markov model produces estimates of SYGR that are substantially more confident than the estimates of SYGR using the traditional approach (Fig 1). For the Fall 2013–2015 cohorts, using the Markov method results in 95% CI widths that are tightened by between 33% and 44% relative to the CI widths following the traditional calculation. Importantly, all of the median SYGR estimates are consistent across the two methods to within 1 percentage point: 71% for Fall 2013 for both methods; 71% for Fall 2014 for both methods; and 73% for Fall 2015 using the traditional approach, as compared to 72% when using the Markov model (Table 3). When using all available data, the reduction in 95% CI width is somewhat lower (34% relative to the traditional calculation) but still appreciable. This is attributable to the fact that using all available data means that the amount of information from partial-data cohorts (Fall 2016 to present) is relatively smaller compared to the full-cohort information (Fall 2013–2015).

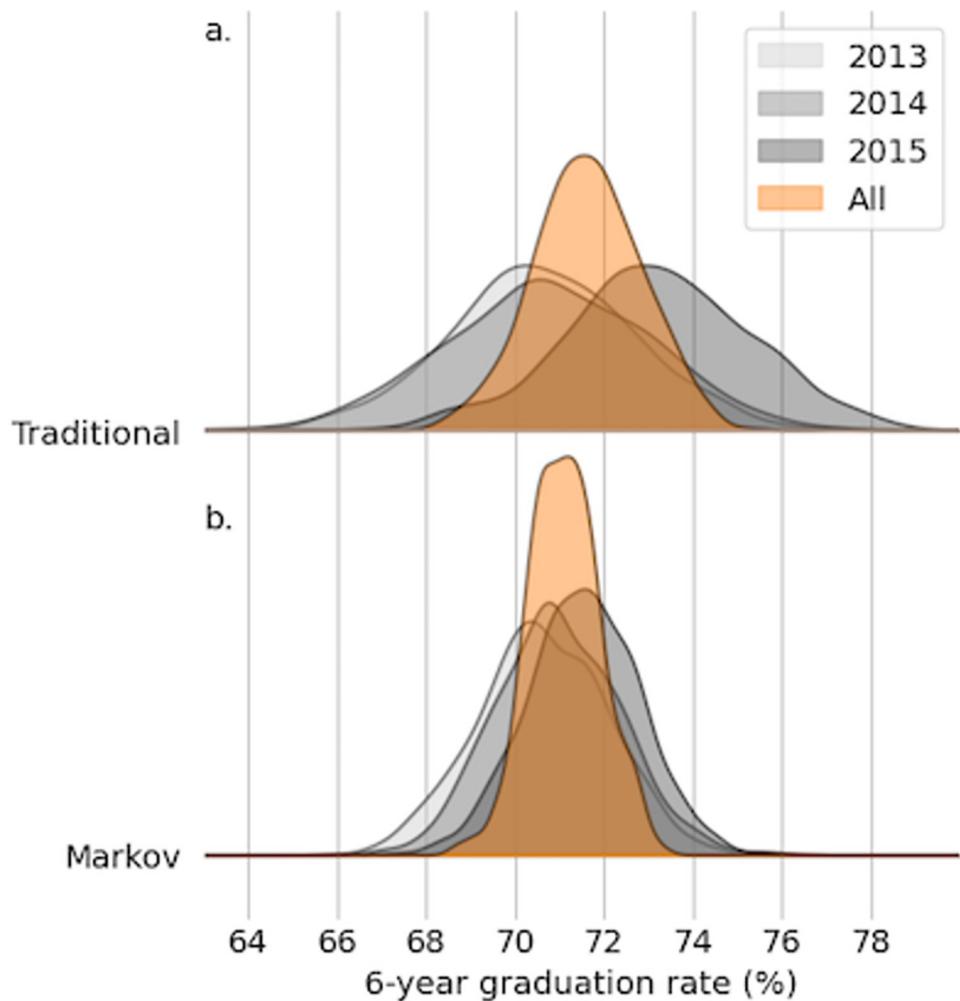


Fig 1. Six-year graduation rates. Using a Markov model to estimate SYGR (bottom panel) leads to tighter 95% confidence intervals than when a traditional approach is used (top panel). Shown are kernel density estimates for six-year graduation rate for the Fall 2013 (light gray), Fall 2014 (gray), and Fall 2015 (dark gray) cohorts and all three of these cohorts combined (orange).

<https://doi.org/10.1371/journal.pone.0287775.g001>

Table 3. Graduation rate estimates and confidence intervals.

Cohort	Method	2.5th percentile	Median	97.5th percentile	95% CI width
Fall 2013	Traditional	66	71	75	8
	Markov	68	71	73	6
Fall 2014	Traditional	66	71	75	9
	Markov	68	71	74	6
Fall 2015	Traditional	69	73	77	9
	Markov	69	72	74	5
All combined	Traditional	69	72	74	5
	Markov	70	71	73	3

Estimates of six-year graduation rate (%) from the traditional calculation and the full-form Markov model. Numbers are rounded to the nearest percentage point, so the 95% CI width reported may not appear to equal the difference between the 97.5th and 2.5th percentiles.

<https://doi.org/10.1371/journal.pone.0287775.t003>

Table 4. Graduation rates and confidence intervals for small subgroups.

	Cohort	Method	2.5th percentile	Median	97.5th percentile	95% CI width
AALANA	Fall 2013	Traditional	40	55	68	28
		Markov	54	62	71	17
	Fall 2014	Traditional	48	59	70	22
		Markov	57	64	73	17
	Fall 2015	Traditional	59	70	83	24
		Markov	62	69	77	15
	All combined	Traditional	55	62	69	14
		Markov	61	65	70	9
First-generation	Fall 2013	Traditional	61	70	81	19
		Markov	64	70	76	11
	Fall 2014	Traditional	57	67	77	20
		Markov	61	68	76	15
	Fall 2015	Traditional	63	72	81	17
		Markov	65	70	75	10
	All combined	Traditional	64	70	75	11
		Markov	67	70	73	7

Estimates of six-year graduation rate (%) for AALANA (top set) and first-generation (bottom set) science majors, using both the traditional calculation and the full-form Markov model. Numbers are rounded to the nearest percentage point, so the 95% CI width reported may not appear to equal the difference between the 97.5th and 2.5th percentiles.

<https://doi.org/10.1371/journal.pone.0287775.t004>

Small subsets of students

As expected, the CI widths for smaller subsets of students (AALANA students and first-generation college students) are notably wider than for the general student population (Table 4). This is due to the fact that the sample sizes for these underrepresented groups are smaller, by their nature as underrepresented groups (see S2 Table in [S1 File](#)). Importantly, similar benefits in terms of reduced CI widths are seen in the Markov-based estimates of SYGR for AALANA students and first-generation students in College of Science majors (Table 4).

The Markov approach improves the confidence interval widths by between 25 and 42% (tightening) relative to the traditional model (Fig 2A and 2C, and Table 4). It is notable that the CIs for AALANA science majors using the traditional approach for computing SYGR are consistently more than 20 percentage points wide, with sizable variation from year-to-year in the median estimate (55% in Fall 2013, 59% in Fall 2014, and 70% in Fall 2015). This interannual variability when the traditional approach is used is problematic, as students and academic program managers may make decisions based on SYGR information that varies widely from year to year. For example, when using the Fall 2013 cohort, the 95% CI for AALANA science major SYGR spans from 40 to 68%, which does not even include the median SYGR when using the Fall 2015 cohort (70%). While this does not necessarily show statistically significant differences from year to year, these results highlight that for small subsets of students, there are substantial variations in best estimates of SYGR when using the traditional approach.

By contrast, using the Markov model incorporates more information as it becomes available. Thus, the estimates of SYGR are better centered on the long-term median: 62% in Fall 2013, 64% in Fall 2014, 69% in Fall 2015, and 65% overall, for AALANA students. Further, the CI widths are consistently below 20 percentage points. Similar gains in terms of tighter CIs and lower interannual variability in estimated SYGR are seen when the Markov model is applied to the subset of first-generation science majors (Fig 2B and 2D, and Table 4).

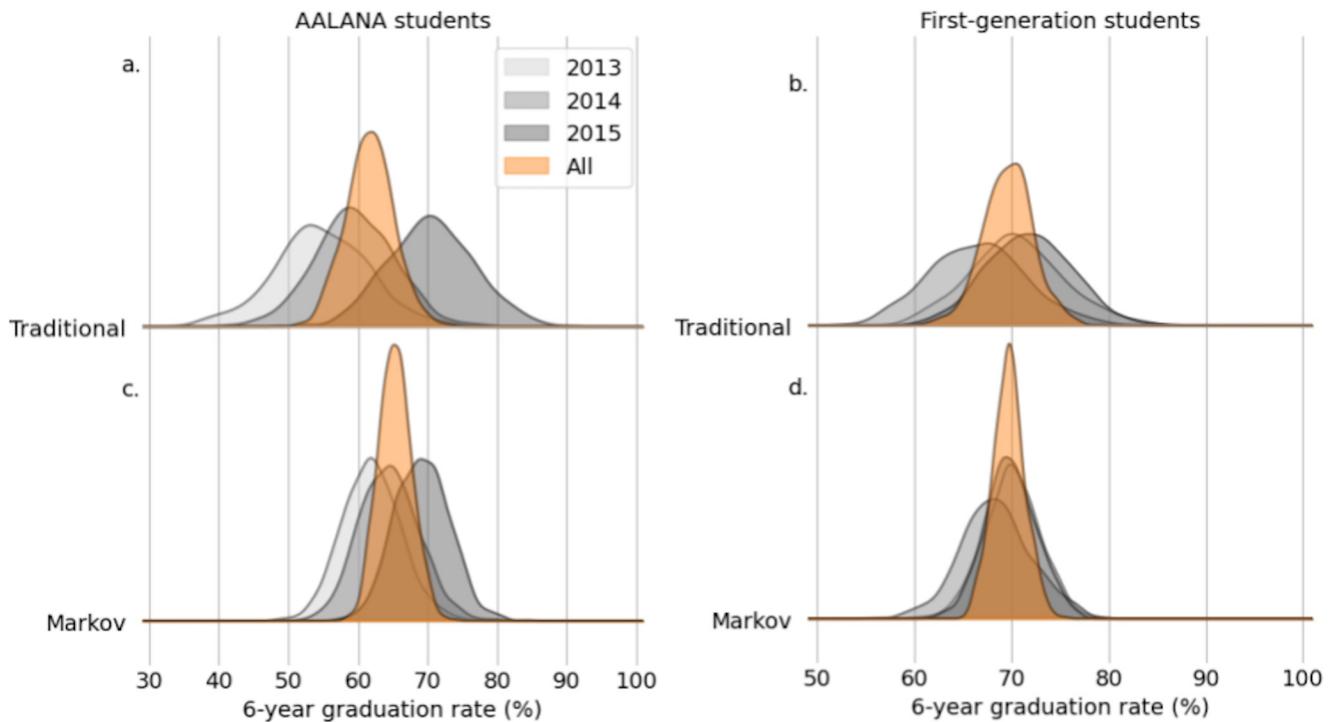


Fig 2. Six-year graduation rate estimates for small subgroups of students. Using a Markov model (bottom panels) to estimate SYGR for AALANA (left panels) and first-generation (right panels) science majors leads to tighter 95% confidence intervals and decreases year-to-year variability as compared to a traditional approach (top panels). Shown are kernel density estimates for six-year graduation rate for the Fall 2013 (light gray), Fall 2014 (gray), and Fall 2015 (dark gray) cohorts and all three of these cohorts combined (orange).

<https://doi.org/10.1371/journal.pone.0287775.g002>

Proof of concept: Evaluating the impacts of LA support

We focus now on students who, at some point during their time at UNI, were enrolled in at least one of the high-attrition science courses and had declared a science major (see Sec. 2.5). We separate this group of students into two groups: the “LA group” of students had an undergraduate Learning Assistant support at least one of their high-attrition science courses; the “no-LA group” were enrolled in sections of these high-attrition courses that were not supported by an LA. All results presented in this section include all available student data from Fall 2013 to Summer 2021 (the “all combined” case from the previous Results sections), and employ the Markov model.

We find that exposure to LAs in high-attrition science courses is associated with a 9 percentage point improvement in SYGR, relative to students who were enrolled in sections of these courses that did not have LA support (Fig 3A). LA-supported students have a median SYGR of 77% (95% CI of 73–81%). This is statistically significantly higher than the non-LA group, which has a median SYGR of 68% (95% CI of 66–70%). While the difference in graduation rates between LA-supported students (77%) and non-LA-supported students (68%) corresponds to a “small” effect size $h = 0.20$, a 9 percentage point difference in graduation rates is noteworthy in practice.

These gains are even more substantial for underrepresented groups of students. For AALANA science majors enrolled in high-attrition courses, LA support is associated with an improvement of 21% in SYGR (Cohen’s $h = 0.49$; medium effect size). LA-supported AALANA students had a median SYGR of 81% (71–93%), as compared to the non-LA group

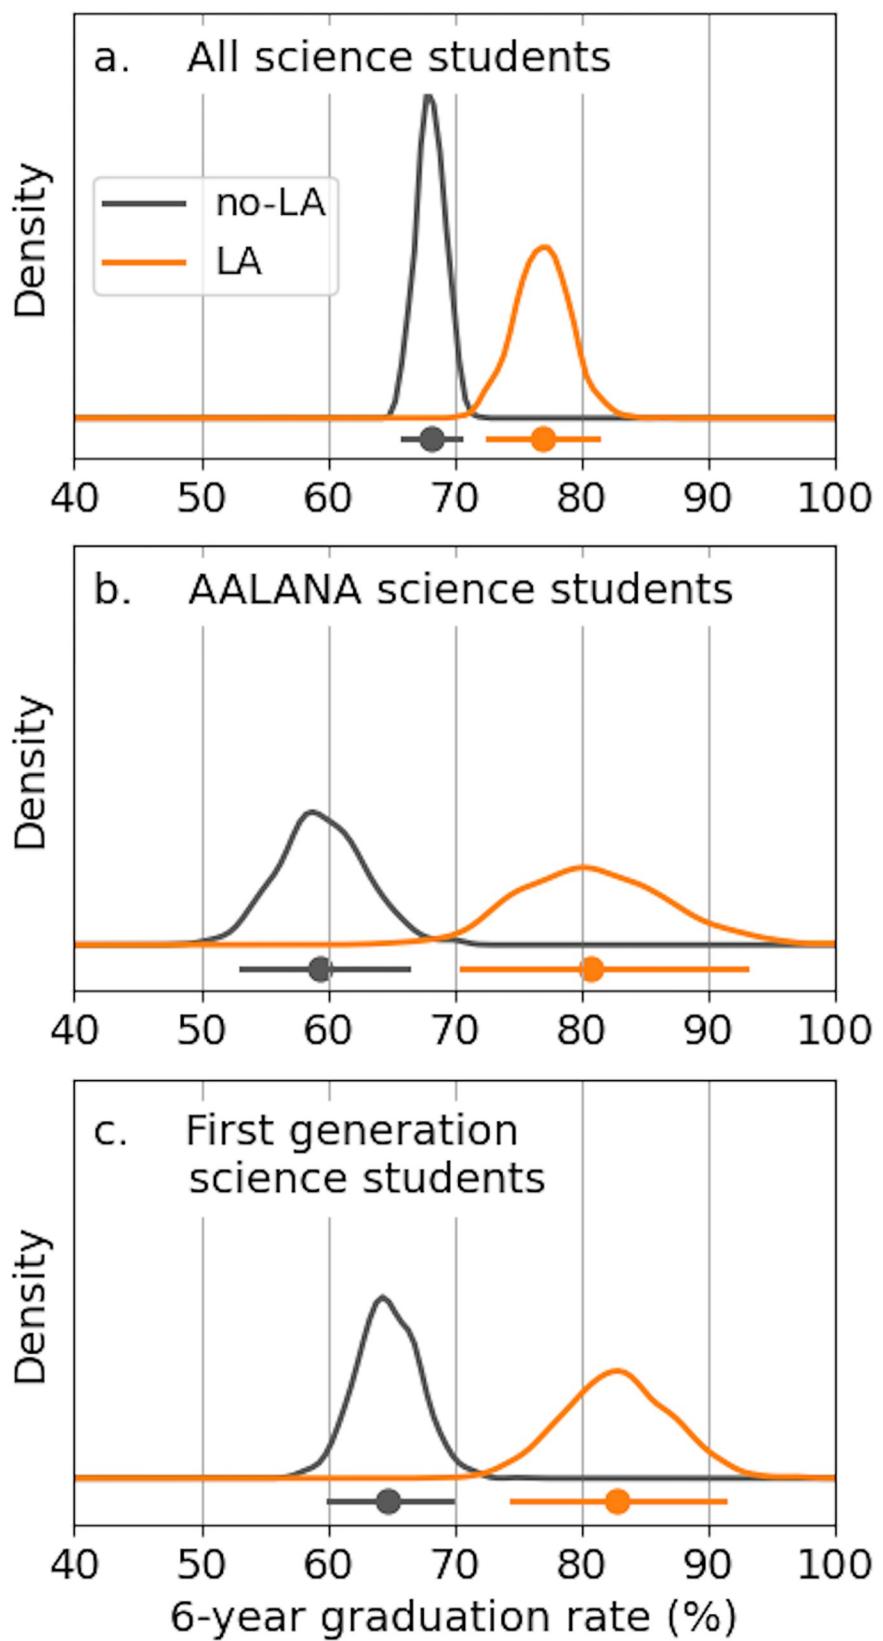


Fig 3. Improvement in graduation rate from learning assistant support. Learning Assistant support is associated with a 9 percentage point improvement in six-year graduation rates for science majors in general (a). These gains are even larger for AALANA science majors (b; 21 percentage point increase in SYGR) and for first-generation college students with science majors (c; 18 percentage point increase in SYGR).

<https://doi.org/10.1371/journal.pone.0287775.g003>

of AALANA students, who had a median SYGR of 59% (53–66%). Similarly, for first-generation college students with science majors, the no-LA group had a median SYGR of 65% (60–70%), while the LA group had a median SYGR of 83% (75–91%), for an improvement of 18 percentage points (Cohen's $h = 0.42$; medium effect size).

Discussion and conclusions

We have presented an approach to use Markov chains to estimate SYGR for cohorts of university students. We have used real-world data from UNI to fit this model, and have demonstrated that it faithfully reproduces graduation rate estimates that would be obtained by following a traditional approach. We find that the Markov model yields SYGR estimates with much tighter confidence intervals and lower interannual variability than a traditional approach. We attribute this to the fact that the Markov model leverages all available data, as opposed to only data for cohorts for which at least six years of data are available. This is particularly important for small cohorts of students, which we demonstrated using AALANA students, first-generation college students, and students who experience a specific academic intervention (LA support in high-attrition courses). Confidence intervals for SYGR for underrepresented groups are generally wider, owing to the smaller sample sizes. Even using the Markov model, incorporating all available data, for underrepresented groups who have had LA support in a course, CIs for SYGR can be more than 20 percentage points wide. In these cases, the CIs using a traditional approach are so wide as to lose practical value. Thus, the Markov model for SYGR is particularly valuable in assessing educational program efficacy, especially for underrepresented groups of students.

Graduation rates are frequently used to assess the efficacy of educational programs, but the uncertainties in these graduation rates are not typically addressed or included in these assessments. In this work, we have demonstrated that a Markov model can provide tighter and less variable confidence intervals by incorporating all student data as it becomes available. This is important for academic institutions and programs to provide the most up-to-date information as possible about program efficacy and student outcomes. This work indicates a path forward that addresses the issue that when a university shares graduation rate information, this information is out-of-date as soon as it is posted. Practically none of the students whose data went into the calculation of that graduation rate are still at that university.

Through our demonstration of the utility of the Markov model for small subsets of students, we find that Learning Assistant support in high-attrition gateway science courses is associated with a 9 percentage point improvement in SYGR among LA-supported students, as compared to their non-LA-supported counterparts. We acknowledge that our approach cannot directly correct for potential “instructor effects” based on individual instructors or their pedagogical choices. However, our results for the gains in SYGR associated with LA support are consistent from year to year, across all instructors and courses for the high-attrition introductory science courses studied here. Further, the integration of LAs into a course supports a variety of evidence-based pedagogies [6, 7], including, for example, more active learning approaches. In light of the demonstrated benefits of active learning [17–19], we interpret our findings as stemming from LA support enabling instructors to implement more evidence-based pedagogies, which in turn leads to improvements in student performance.

These improvements that are associated with LA support are substantially larger for underrepresented groups of students. LA-support is associated with a 21 percentage point improvement in SYGR for AALANA students and an 18 percentage point improvement for first-generation college students. The SYGR estimates for these small subsets of students have wide CIs, making it difficult to disentangle the impacts from an academic intervention such as undergraduate LAs. However, we have demonstrated here that the Markov approach can successfully constrain SYGR estimates for subsets as small as a few dozen students. In addition to the notable improvements to SYGR, LA support is associated with improvements in year-to-year persistence throughout a student's time at UNI (see Supporting Information). Our analysis indicates that undergraduate LA support in high-attrition introductory science courses may be a potentially fruitful pathway to address persistence and graduation rate gaps between non-underrepresented students and underrepresented students.

Supporting information

S1 File. Supplementary figures and tables. This file contains figures and tables with supplemental analyses to support the main text results.

(PDF)

Acknowledgments

We thank the members of the Science and Math Education Research Collaborative at the Rochester Institute of Technology for fruitful conversations.

Author Contributions

Conceptualization: Scott Franklin, Tony E. Wong.

Data curation: Emily K. Mehlman, Scott Franklin, Tony E. Wong.

Formal analysis: Mason N. Tedeschi, Tiana M. Hose, Scott Franklin, Tony E. Wong.

Investigation: Mason N. Tedeschi, Tiana M. Hose, Scott Franklin, Tony E. Wong.

Methodology: Tony E. Wong.

Project administration: Scott Franklin, Tony E. Wong.

Software: Mason N. Tedeschi, Tiana M. Hose, Tony E. Wong.

Supervision: Scott Franklin, Tony E. Wong.

Validation: Emily K. Mehlman, Tony E. Wong.

Visualization: Tony E. Wong.

Writing – original draft: Mason N. Tedeschi, Tiana M. Hose, Scott Franklin, Tony E. Wong.

Writing – review & editing: Mason N. Tedeschi, Tiana M. Hose, Emily K. Mehlman, Scott Franklin, Tony E. Wong.

References

1. Boumi S, Vela AE. Improving Graduation Rate Estimates Using Regularly Updating Multi-Level Absorbing Markov Chains. *Educ Sci*. 2020; 10: 377. <https://doi.org/10.3390/educsci10120377>
2. Nicholls MG. Assessing the progress and the underlying nature of the flows of doctoral and master degree candidates using absorbing Markov chains. *High Educ*. 2007; 53: 769–790. <https://doi.org/10.1007/s10734-005-5275-x>

3. Witteveen D, Attewell P. The College Completion Puzzle: A Hidden Markov Model Approach. *Res High Educ.* 2017; 58: 449–467.
4. Etkina E. Pedagogical content knowledge and preparation of high school physics teachers. *Phys Rev Spec Top—Phys Educ Res.* 2010; 6: 020110. <https://doi.org/10.1103/PhysRevSTPER.6.020110>
5. van Driel JH, Verloop N, de Vos W. Developing science teachers' pedagogical content knowledge. *J Res Sci Teach.* 1998; 35: 673–695. [https://doi.org/10.1002/\(SICI\)1098-2736\(199808\)35:6<673::AID-TEA5>3.0.CO;2-J](https://doi.org/10.1002/(SICI)1098-2736(199808)35:6<673::AID-TEA5>3.0.CO;2-J)
6. Pollock SJ. Transferring transformations: learning gains, student attitudes, and the impacts of multiple instructors in large lecture courses. *American Institute of Physics;* 2006. pp. 141–144.
7. Van Dusen B, Nissen J. Associations between learning assistants, passing introductory physics, and equity: A quantitative critical race theory investigation. *Phys Rev Phys Educ Res.* 2020; 16: 010117. <https://doi.org/10.1103/PhysRevPhysEducRes.16.010117>
8. Alzen JL, Langdon L, Otero V. The Learning Assistant model and DFW rates in introductory physics courses. 2017. pp. 36–39.
9. Alzen JL, Langdon LS, Otero VK. A logistic regression investigation of the relationship between the Learning Assistant model and failure rates in introductory STEM courses. *Int J STEM Educ.* 2018; 5: 1–12.
10. Close EW, Mailloux-Huberdeau J-M, Close HG, Donnelly D. Characterization of time scale for detecting impacts of reforms in an undergraduate physics program. 2018. pp. 88–91. Available: <https://www.per-central.org/items/detail.cfm?ID=14575>
11. Barrasso AP, Spiliotis KE. A scoping review of literature assessing the impact of the learning assistant model. *Int J STEM Educ.* 2021; 8: 12. <https://doi.org/10.1186/s40594-020-00267-8>
12. Sellami N, Shaked S, Laski FA, Eagan KM, Sanders ER. Implementation of a learning assistant program improves student performance on higher-order assessments. *CBE—Life Sci Educ.* 2017; 16: ar62. <https://doi.org/10.1187/cbe.16-12-0341> PMID: 29167224
13. Goertzen RM, Brewe E, Kramer LH, Wells L, Jones D. Moving toward change: Institutionalizing reform through implementation of the Learning Assistant model and Open Source Tutorials. *Phys Rev ST Phys Educ Res.* 2011; 7: 020105. <https://doi.org/10.1103/PhysRevSTPER.7.020105>
14. Herrera X, Nissen JM, Van Dusen B. Student Outcomes Across Collaborative-Learning Environments. 2018. Available: <https://www.per-central.org/items/detail.cfm?ID=14796>
15. Van Dusen B, Nissen J. Equity in college physics student learning: A critical quantitative intersectionality investigation. *J Res Sci Teach.* 2020; 57: 33–57. <https://doi.org/10.1002/tea.21584>
16. Cohen J. Statistical power analysis for the behavioral sciences. 2nd ed. Hillsdale, N.J: L. Erlbaum Associates; 1988.
17. Deslauriers L, McCarty LS, Miller K, Callaghan K, Kestin G. Measuring actual learning versus feeling of learning in response to being actively engaged in the classroom. *Proc Natl Acad Sci.* 2019; 116: 19251–19257. <https://doi.org/10.1073/pnas.1821936116> PMID: 31484770
18. Freeman S, Eddy SL, McDonough M, Smith MK, Okoroafor N, Jordt H, et al. Active learning increases student performance in science, engineering, and mathematics. *Proc Natl Acad Sci.* 2014; 111: 8410–8415. <https://doi.org/10.1073/pnas.1319030111> PMID: 24821756
19. Theobald EJ, Hill MJ, Tran E, Agrawal S, Arroyo EN, Behling S, et al. Active learning narrows achievement gaps for underrepresented students in undergraduate science, technology, engineering, and math. *Proc Natl Acad Sci.* 2020; 117: 6476–6483. <https://doi.org/10.1073/pnas.1916903117> PMID: 32152114