

CRESST REPORT 776

Patti West
Daisy Wise Rutstein
Robert J. Mislevy
Junhui Liu
Younyoung Choi
Roy Levy
Aaron Crawford
Kristen E. DiCerbo
Kristina Chappel
John T. Behrens

**A BAYESIAN NETWORK
APPROACH TO MODELING
LEARNING PROGRESSIONS
AND TASK PERFORMANCE**

AUGUST, 2010



National Center for Research on Evaluation, Standards, and Student Testing

Graduate School of Education & Information Studies
UCLA | University of California, Los Angeles

A Bayesian Network Approach to Modeling Learning Progressions and Task Performance

CRESST Report 776

Patti West
Cisco Systems, Inc.

Daisy Wise Rutstein, Robert J. Mislevy, Junhui Liu, Younyoung Choi
University of Maryland / CRESST

Roy Levy, Aaron Crawford
Arizona State University

Kristen E. DiCerbo
Independent Researcher

Kristina Chappel, John T. Behrens
Cisco Systems, Inc.

August, 2010

National Center for Research on Evaluation,
Standards, and Student Testing (CRESST)
Center for the Study of Evaluation (CSE)
Graduate School of Education & Information Studies
University of California, Los Angeles
300 Charles E. Young Drive North
GSE&IS Bldg., Box 951522
Los Angeles, CA 90095-1522
(310) 206-1532

Copyright © 2010 The Regents of the University of California

The work reported herein was supported in part by the Center for Advance Technology in Schools (CATS), PR/Award Number R305C080015, as administered by the Institute of Education Sciences, U.S. Department of Education.

The findings and opinions expressed in this report are those of the authors and do not necessarily reflect the positions or policies of the Center for Advance Technology in Schools (CATS), the National Center for Education Research (NCER), the Institute of Education Sciences (IES), or the U.S. Department of Education.

To cite from this report, please use the following as your APA reference:

West, P., Rutstein, D. W., Mislevy, R. J., Liu, J., Choi, Y., Levy, R., ... Behrens, J. T. (2010). *A Bayesian network approach to modeling learning progressions and task performance*. (CRESST Report 776). Los Angeles, CA: University of California, National Center for Research on Evaluation, Standards, and Student Testing (CRESST).

A BAYESIAN NETWORK APPROACH TO MODELING LEARNING PROGRESSIONS AND TASK PERFORMANCES¹

Patti West
Cisco Systems, Inc.

Daisy Wise Rutstein, Robert J. Mislevy, Junhui Liu, Younyoung Choi
University of Maryland/CRESST

Roy Levy, Aaron Crawford
Arizona State University

Kristen E. DiCerbo
Independent Researcher

Kristina Chappel, John T. Behrens
Cisco Systems, Inc.

Abstract

A major issue in the study of learning progressions (LPs) is linking student performance on assessment tasks to the progressions. This report describes the challenges faced in making this linkage using Bayesian networks to model LPs in the field of computer networking. The ideas are illustrated with exemplar Bayesian networks built on Cisco Networking Academy LPs and tasks designed to obtain evidence in their terms. We briefly discuss challenges in the development of LPs, and then move to challenges with the implementation of Bayesian networks, including selection of the method, issues of model fit and confirmation, and grainsize. We conclude with a discussion of the challenges we face in ongoing work.

Introduction

The overarching challenge of learning progressions (LPs) is to determine whether, then how, applying them can provide a unifying cognitive/substantive foundation for practical work in curriculum, assessment, and instruction. We believe that LPs have the potential to address this foundational challenge, and to help with specific challenges of task design, test data analysis, simulation design, reporting to students and instructors, improving curriculum, and modeling complex performances. However, in order to realize this potential, data-based models of LPs are required. It is necessary to develop a suitable framework of statistical

¹ Presented at the Learning Progressions in Science conference, sponsored by the National Science Foundation, organized by Alicia Alonzo and Amelia Gotwals, June 24–26, 2009, University of Iowa, Iowa City, IA.

modeling tools to link student performance on assessment tasks to learning progressions, in order to first validate the learning progressions and to subsequently inform decisions of students, instructors, and curriculum designers.

In order to facilitate inferences from an assessment system, the statistical or psychometric model of the system should be aligned with a substantive theory regarding cognition and the development of expertise in the learning progression (Borsboom, 2006; Mislevy, Steinberg, & Almond, 2003). Bayesian networks (Jensen, 1996; Pearl, 1988) represent a flexible approach to latent variable modeling of familiar and complex assessments (Almond, DiBello, Moulder, & Zapata-Rivera, 2007; Levy & Mislevy, 2004). As such they can be applied to the problem of modeling performance aligned with learning progressions in a given content area.

This report details the challenges we have experienced in constructing, calibrating, and applying Bayesian network models of assessments cast in terms of learning progressions. We will provide a brief background on the development of the learning progressions in the curriculum and the task design framework that enables this analysis, including a discussion of the challenges of combining expert analysis, curriculum, and assessment information to create preliminary LPs. We will then discuss the challenges faced in modeling LPs with Bayesian networks, including issues of grainsize and fit. We conclude with a brief comment regarding challenges in our ongoing work, including modeling performance on simulation-based assessments. A worked example of our Bayesian Network modeling of LPs is threaded throughout the report to illustrate some of the issues we faced in conceptualizing and implementing the approach.

This work takes place in the context of the Cisco Networking Academy, and addresses components of the 4-semester Cisco Certified Network Associate (CCNA) course sequence. The Cisco Networking Academy is a global program in which information technology is taught through a blended program of face-to-face classroom instruction, an online curriculum, and online assessments. Courses delivered at high schools, 2- and 3-year community college and technical schools, and 4-year colleges and universities. Since its inception in 1997, the Networking Academy has grown to reach a diverse population of approximately 600,000 students each year in more than 160 countries. Murnane, Sharkey, and Levy (2002) discuss the motivations and origins of the program, while Levy and Murnane (2004) describe issues related to technological application of the curriculum and assessment. Behrens, Collison, and DeMark (2005) discuss the assessment framework that drives the ongoing assessment activity in the program and provides the data for this work.

Challenges with Development of Learning Progressions

Relating curricular structure to LPs. Before modeling an LP, a preliminary structure of the LP must be developed. One of our first challenges involved understanding whether, and to what extent, the substantive structures of the current curricula and tasks could inform us about LPs. Given that the Cisco Networking Academy has been evolving for more than a decade, a wealth of research, subject-matter expertise and instructor expertise, and data from formative, chapter, and final assessments were available to us. We pursued an iterative strategy of identifying evidence of LPs that might underlie the practices as they have evolved, sharpening their focus, modeling them explicitly, and feeding the insights into improved curriculum and assessment design through the lens of the emerging LPs.

In 2007, the Cisco Networking Academy updated and redesigned the curriculum for their primary networking course offerings. Previously, the Academy offered a single four course series that focused on specific individual networking technologies, each course focusing on a specific technology: Physical networking and protocols, Routing, LAN switching, and Wide-area networking (WAN). Taken as a whole, the curriculum prepared students for entry level networking jobs and CCNA certification. As part of the redesign, two separate curriculum strategies were adopted. One strategy (used to create the Discovery course sequence) evolved from a whole task practice (van Merriënboer, 1997) design in which students were presented with the opportunity to build functional networks of increasingly larger and more complex designs as they progressed through each of the four courses. The other strategy (used to create the Exploration course sequence) updated the previous course offerings while maintaining the focus within each technology silo. In their own ways, both curricula were built on beliefs about learning in the domain as reflected by design choices about instructional sequences, learning activities, and assessment practices.

Informing the design of both curricula were the results of statistical analyses of millions of student exams taken over the life of the previous 4-course curriculum. From this analysis (employing classical and item response theoretic methods at the chapter exam level and at the summative final exam level), patterns emerged that indicated that the placement of certain assessment tasks targeted specific Knowledge, Skills, and Abilities (KSAs) at different points within the curriculum affected the performance (difficulty) of the task. These patterns, in combination with subject matter expert input, helped create the initial learning progressions framework. The Cisco Networking Academy participated in ongoing research into methods to identify the features that differentiated between novice and expert performance within the curriculum domain (DeMark & Behrens, 2004; Behrens, Frezzo, Mislevy, Kroopnick, & Wise, 2007). In addition to the Cisco Networking Academy input,

external research highlighting the real-world skill and knowledge necessary for various job levels was used to validate subject matter, expert opinion, and analysis results. Curriculum maps for the new courses used this initial learning progressions framework as a basis for developing chapter and course objectives. Within each chapter, the learning material, practice activity and formative assessment opportunities were designed to build on each other to present and reinforce KSAs within a section of the overall LP framework. As a result, it was determined that the curricula can in fact be viewed as being built around implicit notions of learning progressions.

To make these implicit notions explicit, a learning progression for the skill of Internet Protocol (IP) addressing was developed by an expert in the field. This progression, presented in Appendix A, contains five levels, from 1 (*Novice*) to 5 (*Expert*). Within each level is a set of KSAs that individuals at that level would be expected to possess as their understanding develops. Tasks reflect these beliefs, and, as we will see next, we continue to get test data and student instructor feedback about what works within these progressions to foster learning.

Relating Assessment Data Back to Learning Progressions

Our next challenge was to understand how existing data from multiple-choice items (tasks) on assessments could inform us about LPs. Understanding this relationship between tasks and LPs required examination of response data, discussions with subject matter experts, and understanding of curriculum maps. Tasks in chapter test focus on the KSAs within the chapter. Each task on a chapter test is generally aimed at one level of one LP. These are generally built to evoke evidence about one targeted level of one LP by means of task features that are keyed to the targeted level in the LP (although knowledge and skills presumed to be mastered earlier in instruction may be required as well). Analyses of data from exams that accompany the new curricula enabled us to refine both the learning progressions and the assessment design. We found that some unexpectedly difficult tasks incorporated ideas from KSAs outside the targeted LP and required skills either not yet studied, currently being studied, or previously studied but interacting with the targeted skills in novel ways. These situations increased the difficulty relative to tasks that assessed only the intended LP. Subject matter expert review of similarly designed tasks that performed differently isolated the features of tasks that affected the difficulty in the originally unanticipated ways. The process helped us to define the LPs and to create the assessment task design patterns to target the KSAs at each LP level. It alerted us as well to the eventual necessity of modeling performance on tasks that require skills from multiple LPs.

Evidence centered design (ECD; Mislevy et al., 2003) is a framework for designing assessments to support desired inferences about students similar to other approaches that explicitly incorporate theories of cognition into the design process (e.g., Embretson, 1998). ECD guides the assessment design process via addressing a series of questions:

- “What claims or inferences do we want to make about students?”
- “What evidence is necessary to support such inferences?” and
- “What features of observable behavior facilitate the collection of that evidence?”

ECD is applied to develop tasks and scoring rules for measuring a student’s proficiencies through the perspective of learning progressions.

Of particular assistance in this process is an ECD tool called design patterns (Mislevy et al., 2003), which were used to develop tasks and scoring rules for various types of assessments in the Cisco Networking Academy (Mislevy et al., 2003; Wise-Rutstein, 2005). Design patterns provide a structured model of the knowledge and skills required as needed in a particular task. A design pattern outlines the knowledge, skills, and abilities to be measured, the type of evidence needed to measure these skills, and the methods for determining how this evidence reflects on the skills. While a design pattern may specify the requirements of a particular assessment, it provides support for developing multiple tasks in the skill area in question. While these tasks may be similar, they can be varied in difficulty and other aspects in order to reflect the purpose of the assessment.

One feature of tasks that could be varied is the amount of previous knowledge required, as seen through the lens of LPs. Careful attention to task features shows how two seemingly similar items actually assess different levels of a learning progression. Below is an example of two such items:

Variant A	Variant B
<p data-bbox="212 264 763 411">It is necessary to block all traffic from an entire subnet with a standard access control list. What IP address and wildcard mask should be used in the access control list to block only hosts from the subnet on which the host 192.168.16.43/24 resides?</p> <p data-bbox="354 438 621 464">A. 192.168.16.0 0.0.0.15</p> <p data-bbox="354 493 621 518">B. 192.168.16.0 0.0.0.31</p> <p data-bbox="354 548 621 573">C. 192.168.16.16 0.0.0.31</p> <p data-bbox="354 602 621 627">D. 192.168.16.32 0.0.0.15</p> <p data-bbox="354 657 621 682">E. 192.168.16.32 0.0.0.16</p> <p data-bbox="337 711 638 737">**F. 192.168.16.0 0.0.0.255</p>	<p data-bbox="813 264 1364 411">It is necessary to block all traffic from an entire subnet with a standard access control list. What IP address and wildcard mask should be used in the access control list to block only hosts from the subnet on which the host 192.168.16.43/28 resides?</p> <p data-bbox="954 438 1222 464">A. 192.168.16.0 0.0.0.15</p> <p data-bbox="954 493 1222 518">B. 192.168.16.0 0.0.0.31</p> <p data-bbox="954 548 1222 573">C. 192.168.16.16 0.0.0.31</p> <p data-bbox="938 602 1239 627">**D. 192.168.16.32 0.0.0.15</p> <p data-bbox="954 657 1222 682">E. 192.168.16.32 0.0.0.16</p> <p data-bbox="954 711 1222 737">F. 192.168.16.0 0.0.0.255</p>

The change in the stem from /24 to /28 requires students to perform a more advanced IP addressing skill, namely, subdividing one of the octets. This moves the question from one that distinguishes novices at a lower level who know nothing to one that distinguishes individuals who are at a higher level (Level 3 in terms of the learning progression described on Appendix A). Even changes such as this that seem minor on the surface must be accounted for in task design when they affect demands related to the learning progression.

Overall, design patterns and other tools in ECD can aid in developing an assessment that will support inference about where a student is located on scales defined in terms of learning progressions. This information can be used in turn to draw inferences about the skills a student has, and by implication what learning activities may be appropriate for further learning.

To return to our example, the content expert who specified the IP addressing LP examined the end of our chapter exams for the first course in the Discovery course sequence in order to identify tasks that map to the levels in the IP Addressing progression. These end-of-chapter exams are traditional multiple-choice exams that average around 20 questions per exam. The first Discovery course contains nine chapter exams. The analysis led to the identification of 4 items at the novice level, 9 items at the basic level, 12 items at the intermediate level, and 11 items at the advanced level. The items at each level are those that should differentiate best between that level and the one below it. This intention is affected by the choice of features in the task and the expectations for performance, both suggested by the description of the targeted level in the LP (although, as noted above, this intention can be thwarted by demands for additional knowledge at higher levels in the targeted LP or

requirements from other LPs). An individual at the basic level should have a lower probability of mastering the intermediate items than an individual at the intermediate level. It is not surprising that no items were identified at the expert level, given that the course being examined is the first in a series of four. The items came from five different chapter exams, and in most cases a given chapter exam yielded items at multiple levels.

We next sought to use the end-of-chapter exam data to validate the number of expert-identified skill levels and identify the exam items that best discriminate between these levels. As such, a cross-sectional sample of data was taken, as opposed to a longitudinal sample. In the future, in which a goal might be to model individual students' progressions through the LP, a longitudinal sample might be taken. However, in this case, data from all of the end-of-chapter exams taken in November 2007 were included in the analysis. This month was selected due to the high volume of exams taken. In cases where a student took exams on multiple days in the month, only the exams(s) taken on the first day were included in the data. This resulted in a sample of 3827 student records.

The number of data points for each chapter is shown in Table 1. In any given chapter, data from at least 198 students were taken. In addition, 86 students took all of the chapter exams on the same day. Since it is assumed that no learning occurred during that day, all items taken in one day should reflect the students' appropriate level of the learning progression, so all data for these students were used.

Table 1
Number of Examinees for Each Chapter and Each Chapter Grouping

Chapter	Chapter				
	3	4	5	6	9
3	1992				
4	374	1621			
5	217	331	745		
6	140	154	247	336	
9	86	89	99	113	198

Note. The number in any cell corresponds to the number of people who took the column chapter through the row chapter (e.g., 140 people took Chapters 3 through 6).

Initial analysis was performed to provide insight into the nature of the items and their relationships. Classical difficulty values, or percents-correct, were calculated in order to identify items that might not be appropriate for further analysis (see Table 2). One level 1

item (from Chapter 6) was found to have a difficulty value of 1 (everyone obtained a correct answer for that item) and was therefore not used in the analysis. On average the items were seen to increase in difficulty as they increased in levels. While this is to be expected, caution should be taken in the interpretation of this finding, because the samples on which the item difficulties are based may differ and therefore comparisons across items may reflect differences in the population as well as differences in the items.

Polychoric correlations were calculated to study patterns of relationships among items. It was expected that items that measured the same levels would have higher correlations than items that measured different levels. While a few items followed this expected pattern, there was a high level of correlation among across items from all levels.

Table 2.

Item Difficulty Value

Chapter	Level	Item	Difficulty
4	1	1	0.924
6	1	2	1.000
6	1	3	0.833
6	1	4	0.818
6	2	5	0.732
5	2	6	0.651
3	2	7	0.619
3	2	8	0.630
9	2	9	0.692
5	2	10	0.894
3	2	13	0.699
3	3	14	0.738
3	3	15	0.354
3	3	16	0.611
3	3	17	0.467
3	3	18	0.841
5	3	19	0.741
5	3	20	0.734
5	3	21	0.850
5	3	22	0.643
5	3	23	0.710
5	3	24	0.711
9	3	25	0.833
9	4	27	0.778
3	4	28	0.654
5	4	29	0.631
5	4	30	0.773
5	4	31	0.647
5	4	32	0.387
5	4	33	0.532
5	4	34	0.790
5	4	35	0.556
5	4	36	0.514

On average the correlations between any two groups of items were between .35 and .44, and it was not always the case that items of the same level had the highest average correlation with other items of the same level. In general there were, however, relatively higher correlations between adjacent levels than remote levels. For example, Level 3 items on average have higher correlation with Level 2 and Level 4 items than with Level 1 items (see Figure 1). The similarities across levels of correlations may in part be due to the fact that all of the items should be measuring the same underlying skill. Items were also examined to determine if items from the same chapter had higher correlations than items from differing chapters. Again no strong patterns of correlations were found (see Figure 2).

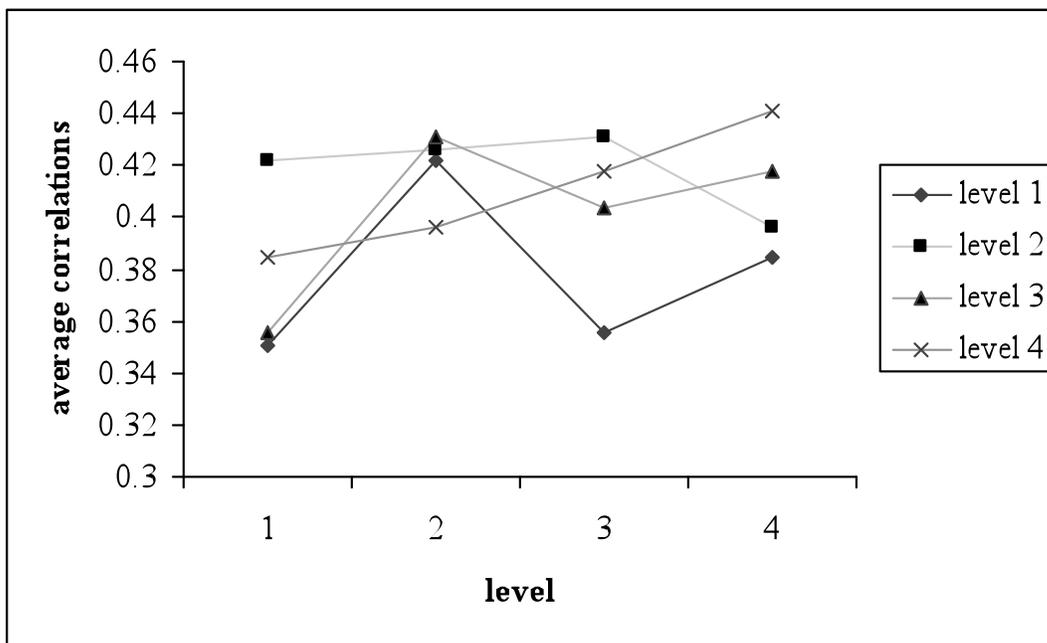


Figure 1. Average correlations of items across levels. Each point is the average correlation of items from the level specified on the x axis with the items from the level of the line the point is on.

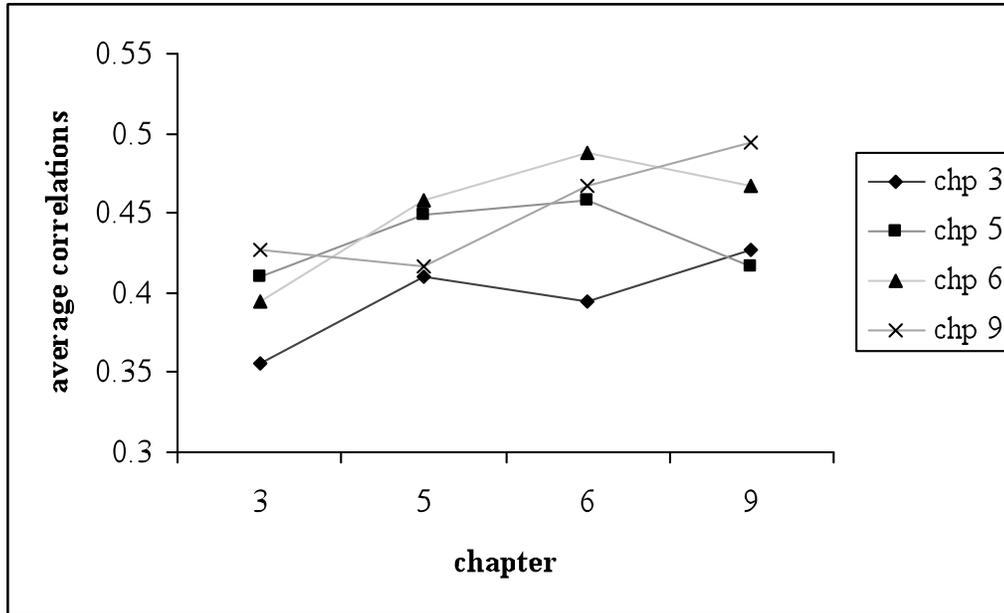


Figure 2. Average correlations across chapters. Each point is the average correlation of items from the chapter specified on the x axis with the items from the chapter of the line the point is on. Chapter 4 does not appear to as there was only one item from that chapter. Other chapters did not have items for that skill.

A factor analysis of the polychoric correlations yielded strong evidence for one dominant factor (see Figure 3). This finding was also not unexpected, as all of the items should be measuring the same underlying skill.

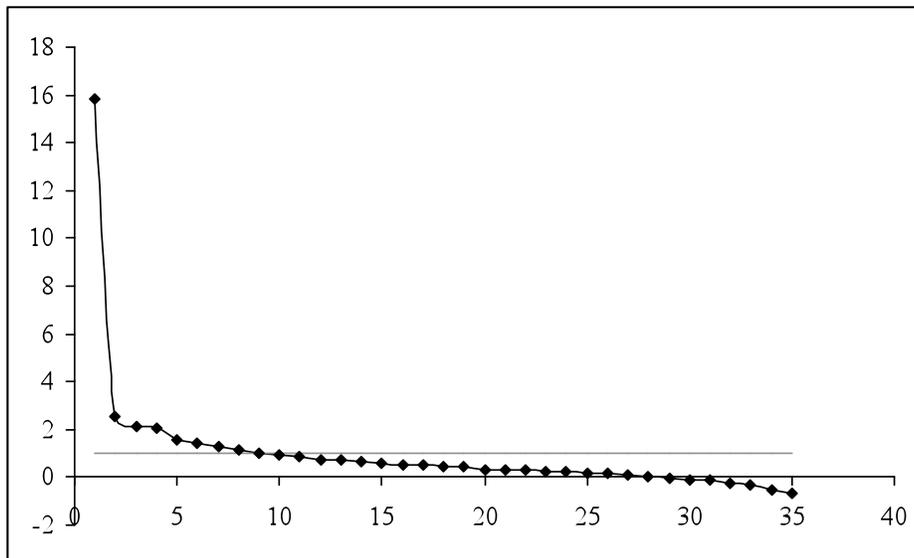


Figure 3. A Scree plot for factor analysis. There is a line at 1 on the y axis.

In general, while the exploratory data analysis provided evidence that the items were related to each other and that they were all measuring the same general skill set, there still did

not seem to be evidence one way or another regarding whether the items themselves were labeled at the appropriate level. For this, further analysis using Bayes Networks (BNs) was conducted.

Challenges in Implementing Bayes Networks

Selection of BN method. The proposed learning progressions and the ECD-based assessments lead naturally to the question of drawing inferences from assessment performances about the students' status in learning progressions. The challenge was to select a modeling tool that would allow for these probabilistic inferences in the LP framework. BNs had been used in the past to model assessment data in this domain (Levy & Mislevy, 2004), and it was surmised that they might also be useful in this context. However, folding together the curriculum and assessment information required in this LP modeling task was a new challenge in our work with BNs.

BNs leverage connections between probability theory and graphical models to represent the probabilistic relationships among a large number of variables. As flexible modeling tools, they have been employed across a wide variety of applications, including education and related settings such as diagnostic and expert systems (Spiegelhalter, Dawid, Lauritzen, & Cowell, 1993). In education, BNs have been used in complex assessment systems (Almond et al., 2007; Levy & Mislevy, 2004; Reye, 2004) and frequently have been used in the context of intelligent tutors to create models of an individual student's knowledge and provide information based on that model (Conati, Gertner, & VanLehn, 2002; Murray & VanLehn, 2000). Important for this application is that they allow for a representation of the theory of the relationships in a domain and use probability theory to characterize and examine the strength of those relationships. As shown in the following examples, BNs used in educational assessment typically include unobservable or latent variables that characterize aspects of students' knowledge and skill, and observable variables that characterize features of students' task performances.

At the core, BNs are a set of conditional probabilities in which the probability of one event, for example success on a given assessment task, is conditional on the probability of a previous event, for example success on a previous task. However, instead of focusing only on the relationship between two variables, BNs and related graphical models structure relationships across multivariate systems. A BN (Jensen, 1996; Pearl, 1988) models the relationships among a set of variables by specifying recursive conditional distributions in order to structure the joint distribution. The networks are so named because they support the

application of the Bayes' theorem across complex networks by structuring the appropriate computations (Lauritzen & Spiegelhalter, 1988; Pearl, 1988).

A BN may also be represented as a graphical model (see Figure 4), consisting of the following elements (Jensen, 1996):

- A set of discrete variables represented by ellipses or boxes and referred to as nodes. Each variable has a set of exhaustive and mutually exclusive states.
- A set of directed edges (represented by arrows) between nodes indicating the probabilistic dependence between variables. Nodes at the source of a directed edge are referred to as parents of nodes at the destination of the directed edge, their children.
- For each exogenous variable (i.e., a variable without parents such as the student proficiency variables Connectivity and IP Addressing in Figure 4), there is an associated unconditional probability distribution where the probabilities over the states sum to one.
- For each endogenous variable (i.e., a variable with parents such as ConTask 1 in Figure 4, an observable task response posited to depend on students' proficiency with regard to network connectivity), there is an associated set of conditional probability distributions corresponding to each possible combination of the values of the parent variables, where the probabilities of the states in each conditional distribution sum to one (see Figures 5 and 6).

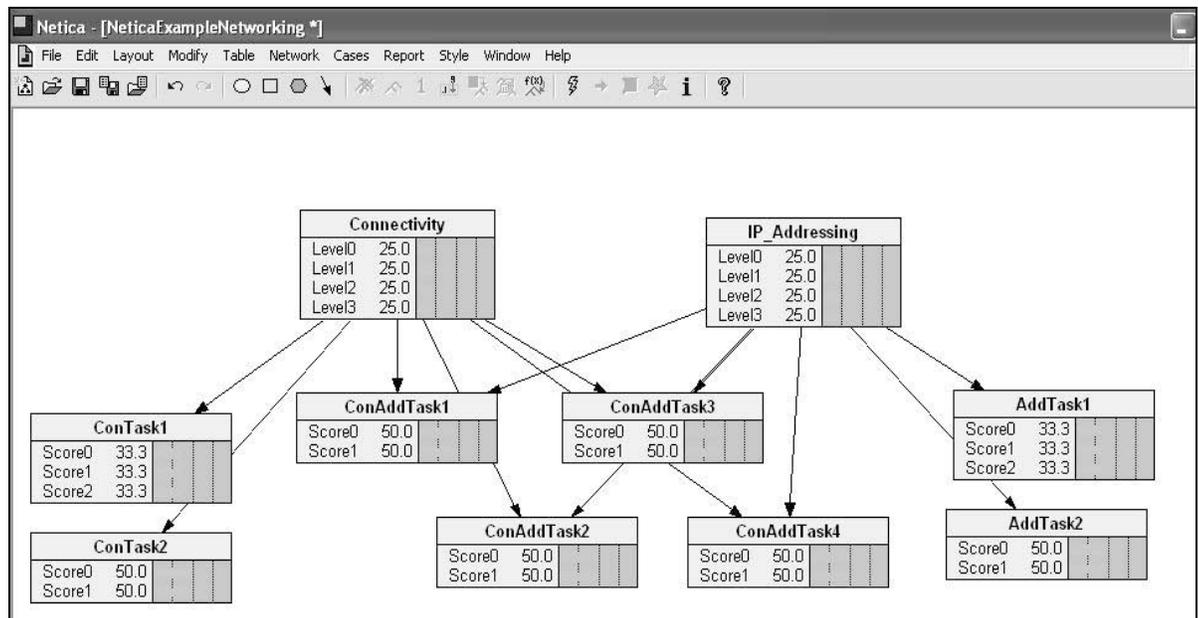


Figure 4. A sample Bayes net with two student model variables (SMVs: Connectivity and IP Addressing), each embodying a 4-level learning progression, and eight observable variables (OVs). By construction around salient task features and requirements, the OV's depend on one or both SMVs and are targeted to discriminate at specified values. Figure obtained using the Netica Bayes net program.

Connectivity	Score0	Score1	Score2
Level0	85.000	10.000	5.000
Level1	60.000	30.000	10.000
Level2	30.000	50.000	20.000
Level3	10.000	60.000	30.000

Figure 5. Sample Netica Output—conditional probabilities of the observable variable ConTask1. Its possible values are 0, 1, and 2 from a partial credit scoring scheme. Connectivity is the SMV parent of this OV. Each row is the conditional probability distribution for the OV given a value of the SMV Connectivity. This is a task meant to discriminate best at Level 2, the level at which there is a 70% probability of scoring at least a 1.

Connectivity	IP_Addressing	Score0	Score1
Level0	Level0	90.000	10.000
Level0	Level1	90.000	10.000
Level0	Level2	90.000	10.000
Level0	Level3	90.000	10.000
Level1	Level0	90.000	10.000
Level1	Level1	90.000	10.000
Level1	Level2	20.000	80.000
Level1	Level3	20.000	80.000
Level2	Level0	90.000	10.000
Level2	Level1	90.000	10.000
Level2	Level2	20.000	80.000
Level2	Level3	20.000	80.000
Level3	Level0	90.000	10.000
Level3	Level1	90.000	10.000
Level3	Level2	20.000	80.000
Level3	Level3	20.000	80.000

Figure 6. Conditional probabilities of the observable variable ConAddTask1. Its possible values are 0 and 1 (unsuccessful and successful solution). Both Connectivity and IP Addressing are the SMV parents of this OV. Each row is the conditional probability distribution for the OV given a combination of value of the two parents. By construction, this task has features for which understanding and carrying out a solution uses concepts at Level 1 of the Connectivity learning progression and Level 2 of IP Addressing. The conditional probability distributions thus show only 10% probability of a successful solution for all SMV combinations in which these levels are not reached, and 80% probability at all combinations with at least these levels.

The variables and the directed edges together form an acyclic directed graph (frequently referred to as a directed acyclic graph; Brooks, 1998; Jensen, 1996; Pearl, 1988). These graphs are directed in the sense that the edges follow a “flow” of dependence in a single direction; in contrast to other graphical modeling traditions (e.g., Bollen, 1989) the arrows are always unidirectional rather than bi-directional. The graphs are acyclic in that following the directional flow of directed edges from any node it is impossible to return to the node of origin. The structure of the graph conveys the patterns of dependence and (conditional) independence among the variables in the joint distribution and corresponds to the computations involved in constructing the joint distribution of the variables in the system and subsequently conducting Bayesian inference to yield posterior distributions for the unknown variables once data have been observed (Lauritzen & Spiegelhalter, 1988; Pearl, 1988). In connection with this last point, BNs properly and efficiently quantify and propagate the evidentiary import of observed data on unknown entities, thereby facilitating evidentiary reasoning under uncertainty as is warranted in psychometric and related applications (Almond & Mislevy, 1999; Almond et al., 2007; Levy & Mislevy, 2004; Mislevy, 1994; Mislevy & Levy, 2007; Spiegelhalter et al., 1993).

BNs may be employed to model the hypothesized structure if multiple learning progressions, where discrete latent variables correspond to the skills and the categories of latent variables correspond to the different levels of the skills. The pattern of dependence of the observables on the latent variables reflects the hypothesized structure of the manner in which performance depends on the students’ status with respect to the progression. Possible sources by which to model the relationships among the latent variables include exploratory path analyses of scores on the exams and subject matter experts’ beliefs about the domain and students’ learning progressions. Such models also support modeling of observable variables (OVs) as dependent on multiple latent variables. Figure 6 displays conditional probabilities for an OV in a multidimensional BN. This item has two student model variable (SMV) parents (IP Addressing and Connectivity), which combine to form the probability distribution for the OV. Building out networks of combinations of OVs and SMVs allows us to map out complex and interrelated learning progressions.

One challenge in working with Bayesian networks is that, although they are very flexible in terms of inputs and modeling, they rely on an already coordinated system of Learning Progressions and assessments. BNs by themselves, for example, cannot make up for an assessment system that does not match to a learning progression. Aside from content mismatch, assessments, and LPs, it can also fail to match on the level of grain size.

The issue of level of specificity and detail at which to model learning progressions is an important one, and one that is forced in our project by the existence of two distinct curricula for the same target set of knowledge and skills. As discussed previously, some items differed in difficulty because they tapped increasingly complex knowledge and skill along a progression of concepts. The learning progression shown in Appendix A illustrates this idea. But we also saw that the relative difficulties of two items could be reversed in the two curricula at a given point in time simply because of the order in which they were introduced. We could either try to define two different coarsely-grained LPs for the two curricula, or to define fine grained LPs that would maintain their integrities in both curricula. We opted for the latter, reasoning that learning progressions defined by increasing conceptual complexity were preferable to ones defined by both complexity and the coincidence of topic order. In other words, we decided that LPs should not be so specific that they apply to only one curriculum; the notion of an LP is to specify the progression of big ideas to be mastered in a content area, not merely in a given curriculum. Fortunately, the ECD approach forces assessment design to specify the ways in which evidence from OVs depend on higher level variables like SMVs. Importantly, the choice of grain size is more an issue for the coordination between (a) desired inferences and (b) evidence that will be available. BNs can handle a variety of different forms of evidentiary inputs and structures for facilitating inference, but, like any other statistical or psychometric modeling tool, unlikely to be useful if (a) and (b) are not coordinated. LPs such as the example in this report can be combined with other LPs to create larger LPs, as we will discuss in the Conclusion and Future Challenges section. This may allow for variation in the grain size modeled.

Implementation of Bayes Nets

Once grain size is determined, a preliminary LP structure has been developed based on expert input and assessment data, and initial statistical analyses are complete, the BN modeling process can begin. The BN modeling approach is a relatively new one, and all the details of implementation have not yet been worked out. Therefore, in implementation, we were on occasion faced with challenges of determining the best way to proceed with model-building decisions.

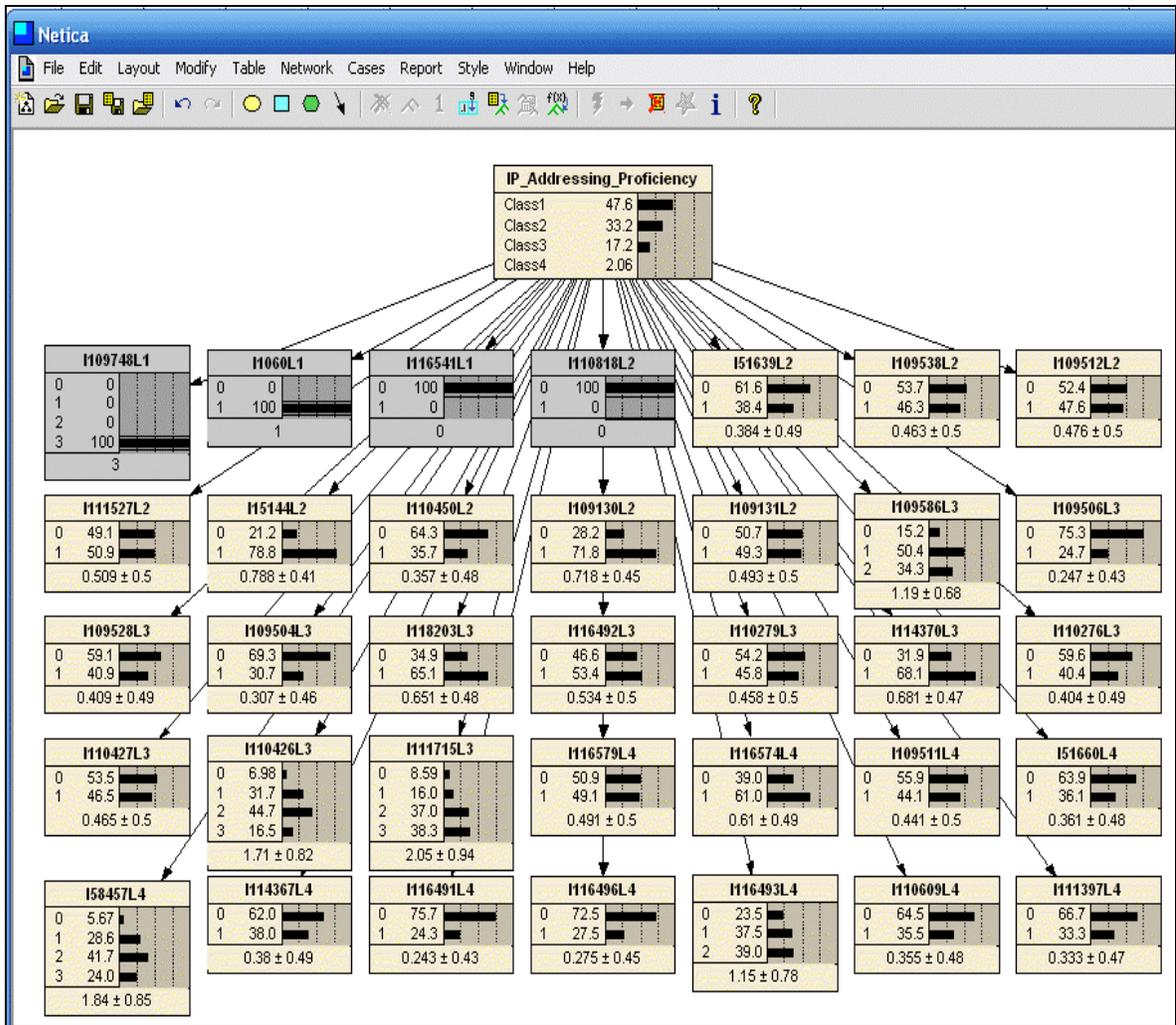


Figure 7. Little Johnny Bayes Net. Bayes net for Student 320 with data on the first four items only. One SMV: IP Addressing, which embodies a 4-level learning progression, and 35 OVs. The posterior distribution indicates that this student is more likely to be a member of Class 1 than any one of the other classes. Figure obtained using *Netica* Bayes net program.

In our example, the BN contains a single discrete latent variable modeled as the parent for the discrete OVs (i.e., scored item responses), the children, as graphically depicted in Figure 7. The BN model described here is equivalent to a latent class model (Dayton & Macready, 2007; Lazarsfeld & Henry, 1968). A latent class analysis was conducted using the polka package (Linzer & Lewis, 2007) in R (R Development Core Team, 2008) using multiple start values to determine the optimal solution. Given that the item pool contained items corresponding to four levels of the expert-based learning progression pool, it was anticipated that ideally a 5-class model would be supported, where, in addition to the four levels defined by the expert represented in the items, a fifth class would emerge representing a knowledge state below the first (novice) level. The implication for the data analysis is that, by defining the levels in terms of what students know or can do (Appendix A), there is an

additional, implicit baseline level (Level 0) representing essentially below novice knowledge. That is, the novice level items would discriminate between students at the novice level (or beyond) and students who had not yet learned the novice material (i.e., level 0). Similarly, items at the advanced level would discriminate between the students at the intermediate level and students who were at the advanced level (or beyond). The lack of items at the expert level in this sample precluded us from expecting the 6-class model that might otherwise have been suggested theoretically.

To allow for the possibility that the data provided better support for a model with a different number of levels, we compared 2-class, 3-class, 4-class, 5-class, and 6-class models. The 4-class model demonstrated the best fit to the data, based on statistical fit in terms of the BIC (Schwarz, 1978) and the bootstrapped likelihood ratio test (McLachlan & Peel, 2000; Nylund, Asparouhov, & Muthén, 2007) conducted in Mplus (Muthén & Muthén, 1998–2006). In addition, this model offered the best interpretability of the classes in terms of class membership proportions and consistently ordered patterns of class performance across items. That is, the four classes identified in the analysis correspond to increasing levels of performance on the items and are interpreted as increasing levels of knowledge, skills, and proficiency. A BN representation of the 4-class model was then constructed in Netica (Norsys Software Corp., 2007).

The lack of support for a 5-class model was apparently due to the small number of items at the novice level (Level 1), as well as an absence of student below the novice level. This is unsurprising, given that the items used in this analysis are drawn from assessments administered after instruction has occurred. In other words, to discriminate well between students who are essentially ignorant of the material and those who have achieved novice understanding of the material, students would need to be measured earlier, perhaps with a pre-test that included more novice items. As discussed below, many of the items functioned in ways consistent with the expert-based expectations of the learning progression. Thus, the four classes are interpreted as the first four levels of the learning progressions, where the first class is perhaps a mixture of students at and below the first level of the progression.

Inferences regarding assessment items. One goal of the analysis included the modeling of the items' locations along the learning progression. Specifically, an item was classified as being "at the level" of a certain class if it supported an interpretation that students reaching that level would be able to solve or complete the task whereas students at lower levels would be unlikely to be successful. To classify items, the conditional probability tables were examined. For each item, the odds of answering the item completely correctly were calculated in each class and odds ratios were calculated to compare adjacent classes.

These odds ratios capture the power of the items to discriminate between classes. To construct an odds ratio for the first class, the probability that a complete novice would get the item right was defined as the probability of getting the item right by guessing. Each item was assigned to a level based on considerations of (a) the size of these odds ratio between successive classes, (b) the criterion that the probability the responding correctly at the assigned level exceeded .50 for dichotomously scored items, and (c) the distribution of probability across the response categories for polytomously scored items.

The results indicate that many of the items discriminate strongly between classes. For example, Figure 8 contains the conditional probability table for an item where it is clearly seen that only students in the fourth (highest) class are likely to successfully solve it. Statistically, this item aids in distinguishing students in the fourth latent class (level) from the remaining classes (levels). Substantively, the item captures one aspect of what it means to be at the fourth level of the learning progression. Students at the fourth level have learned the knowledge ad skills necessary to correctly answer this item; students at the lower levels have not.

IP_Addressing_Proficiency	Score0	Score1
Class1	84.200	15.800
Class2	74.420	25.580
Class3	62.440	37.560
Class4	12.480	87.520

Figure 8. Conditional probabilities of a clearly discriminating item (Item 31). Its possible values are 0 and 1 from a dichotomous scoring scheme. IP Addressing Skill is the SMV parent of this Observable Variable. Each row is the conditional probability distribution for the Observable Variable given a value of the SMV IP Addressing Skill. This is a task that discriminates best at Level 4, the level at which there is a 87.5% probability of scoring 1.

IP_Addressing_Proficiency	Score0	Score1	Score2
Class1	40.300	49.000	10.700
Class2	11.870	36.280	51.850
Class3	2.350	12.050	85.600
Class4	0.000	5.110	94.890

Figure 9. Conditional probabilities of a more ambiguous item (Item 33). Its possible values are 0, 1, and 2 from a partial credit scoring scheme. IP Addressing Skill is the SMV parent of this Observable Variable. Each row is the conditional probability distribution for the Observable Variable given a value of the SMV IP Addressing Skill.

Still other items were more ambiguous in terms of their levels. For example, Figure 9 contains the conditional probability table for an item where it is seen that students in the second class have a probability of .88 for earning partial or full credit, but only a .52 probability of earning full credit. A simple classification of this item in terms of one level is insufficient to fully capture its connection to the classes. A richer characterization of the item, recognizing that it discriminates well between multiple adjacent classes, states that once a student reaches Class 2, she is very likely to earn at least partial credit but needs to reach Class 3 (or 4) in order to be as likely to earn full credit.

The results were largely consistent with the expert-based expectations regarding the items. Ten items exhibited clear and distinct patterns in which they distinguished between classes exactly as predicted by experts. That is, these items were “located” at the level as expected. Figure 8 is an example of one such item; the expert prediction of this item as a Level 4 is strongly supported by the results. Five items distinguished roughly well at the level predicted by experts and at one other level; that is, they appeared to be located at the expected level and one other level. Eighteen of the items were located at a level adjacent to where they were predicted to be located (e.g., an item expected at Level 4 was located at Class 3). One item was located at one class adjacent to the predicted class and another class not adjacent. The results for this item are given in Figure 9. This item was expected to be a Level 4 item. As discussed above, the polytomous scoring of this item makes it possible to view it as being located at Class 2 or Class 3. Only one item was clearly located at a class that was not equal to or adjacent to the predicted level. As noted earlier, inspection of any items found to have empirical odds ratios that differed from their intended levels can provide insights about features that make them spuriously hard or easy for reasons unrelated to the LP they are meant to assess.

Inferences regarding students. The conditional probability tables also reveal how inferences regarding students are conducted in the BN. For example, observing a correct response for the item in Figure 8 is strong evidence that the student is in Class 4; observing an incorrect response for the item in Figure 8 is relatively strong evidence that the student is not in Class 4. The use of a BN approach supports inferences regarding students by collecting and synthesizing the evidence in the form of observed values of variables. That information is then propagated through the network via Bayes’ theorem to yield posterior distributions for the remaining unknown variables (Pearl, 1988), including the latent class variable corresponding to the skill level. For example, Figure 7 contains the BN for a student who has completed four of the items. The student correctly answered the first two items and incorrectly answered the latter two items. On the basis of this evidence, the posterior

distribution for their latent skill variable indicates that this student has a probability of being in Classes 1–4 of .476, .332, .172, and .021, respectively. From this, we may infer that the student is most likely in one of the first two classes (i.e., is at one of the first two levels of the skill progression) but that there still remains considerable uncertainty. The collection and inclusion of more data would lead to a more refined inference, as illustrated in Figure 10, which contains the BN for another student who has completed all of the items. The posterior distribution for this student is quite clear in supporting an inference that the student is in Class 3 (posterior probability equals .997); that is, the student is in the third level of the learning progression.

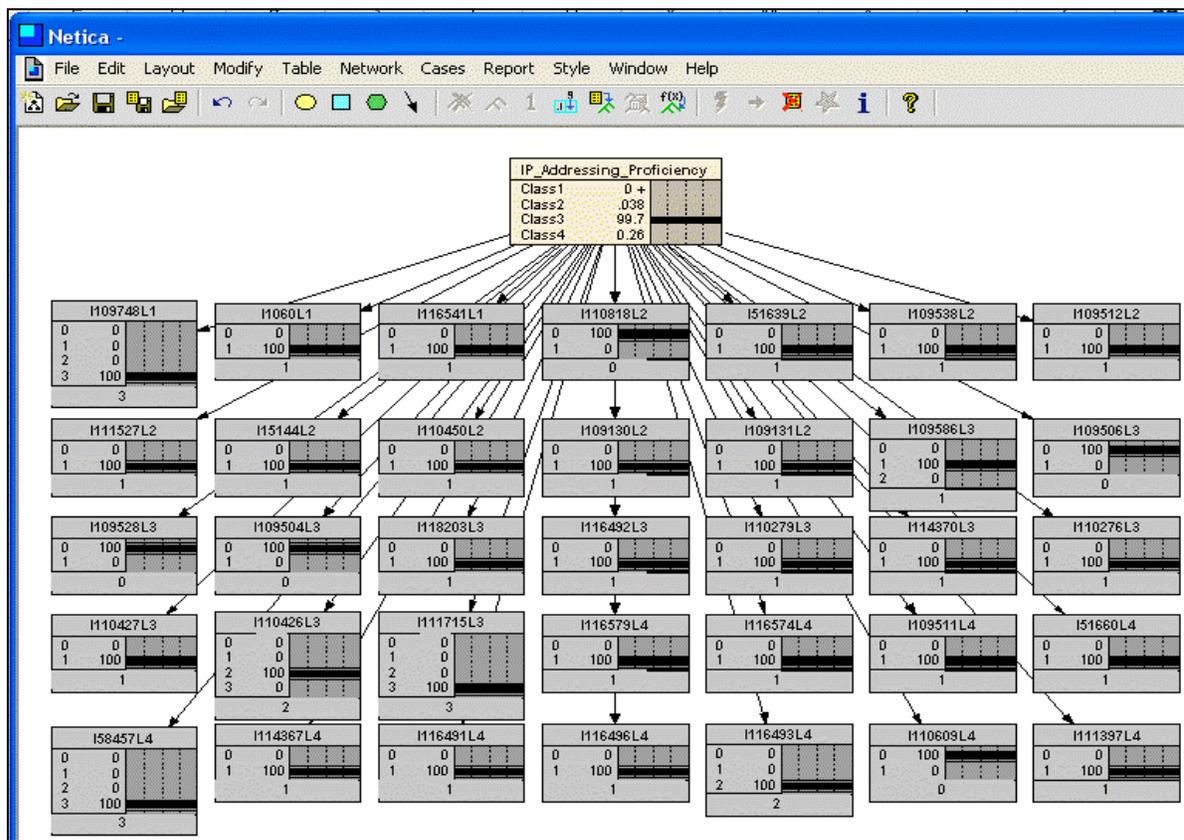


Figure 10. Little Sally Bayes Net. Bayes net for Student 67 with data on all 35 of the items. One SMV: IP Addressing, which embodies a 4-level learning progression, and 35 observable variables (OVs). The posterior distribution indicates that this student is almost certainly a member of Class 3. Figure obtained using *Netica* Bayes net program.

Conclusions and Future Challenges

This study describes the applications of a variety of techniques centered around Bayesian network modeling of a real-world example of learning progressions. LPs defined by experts matched with ECD-based assessment tasks completed by thousands of students

provide the basis for this analysis. The example demonstrates how assessment data can be used to validate a learning progression using statistical modeling in the form of a BN. Assessment items that discriminated between various levels in the progression were identified. In addition, it was demonstrated how students could be classified into levels based on their assessment results.

The results of the modeling offer a data-based interpretation of the development of skills that constitute the learning progression. In most cases, the results for items serve to confirm the expert-based expectation. For other items, the results were more ambiguous or offer an alternative explanation to that from the experts. Taking a comprehensive perspective on assessment of learning progressions, the results of the statistical analyses will be taken back to the subject matter experts for consultation and possible refinements in terms of the definition of the learning progression (Appendix A), the items that assess the aspects of the learning progression, and the utility of the additional items for modeling students' progression.

The BN modeling approach facilitates probability-based reasoning about students in terms of their learning progression. Assessment data (e.g., scored item responses) enter the network in the form of OVs. Synthesizing the evidentiary import of the data, the posterior distribution of class membership, interpreted as the level of the learning progression, governs the inferences regarding the student.

It is argued that BNs are well positioned to support inferences at fine-grained levels aligned with rich substantive theories, and as such are powerful statistical tools for modeling and structuring substantive inferences and feedback to students, instructors, and curricular designers. However, we also expect challenges as we proceed with using BNs in increasingly complex ways.

Growing The Progression

The worked example shown in this report focused on OVs related to one SMV (IP Addressing). The results from this analysis, however, can be “plugged in” to a much larger model that displays the relationships between SMVs (see Figure 11). This allows for the modeling of the influence of mastery of one area on the mastery of another area. However, it presents the challenge of modeling variables with multiple parents and the determination of conditional dependence/independence of variables from a given task. Larger investigations of the progression(s) will also entail longitudinal modeling of student performance and learning over time using BNs (Reye, 2004).

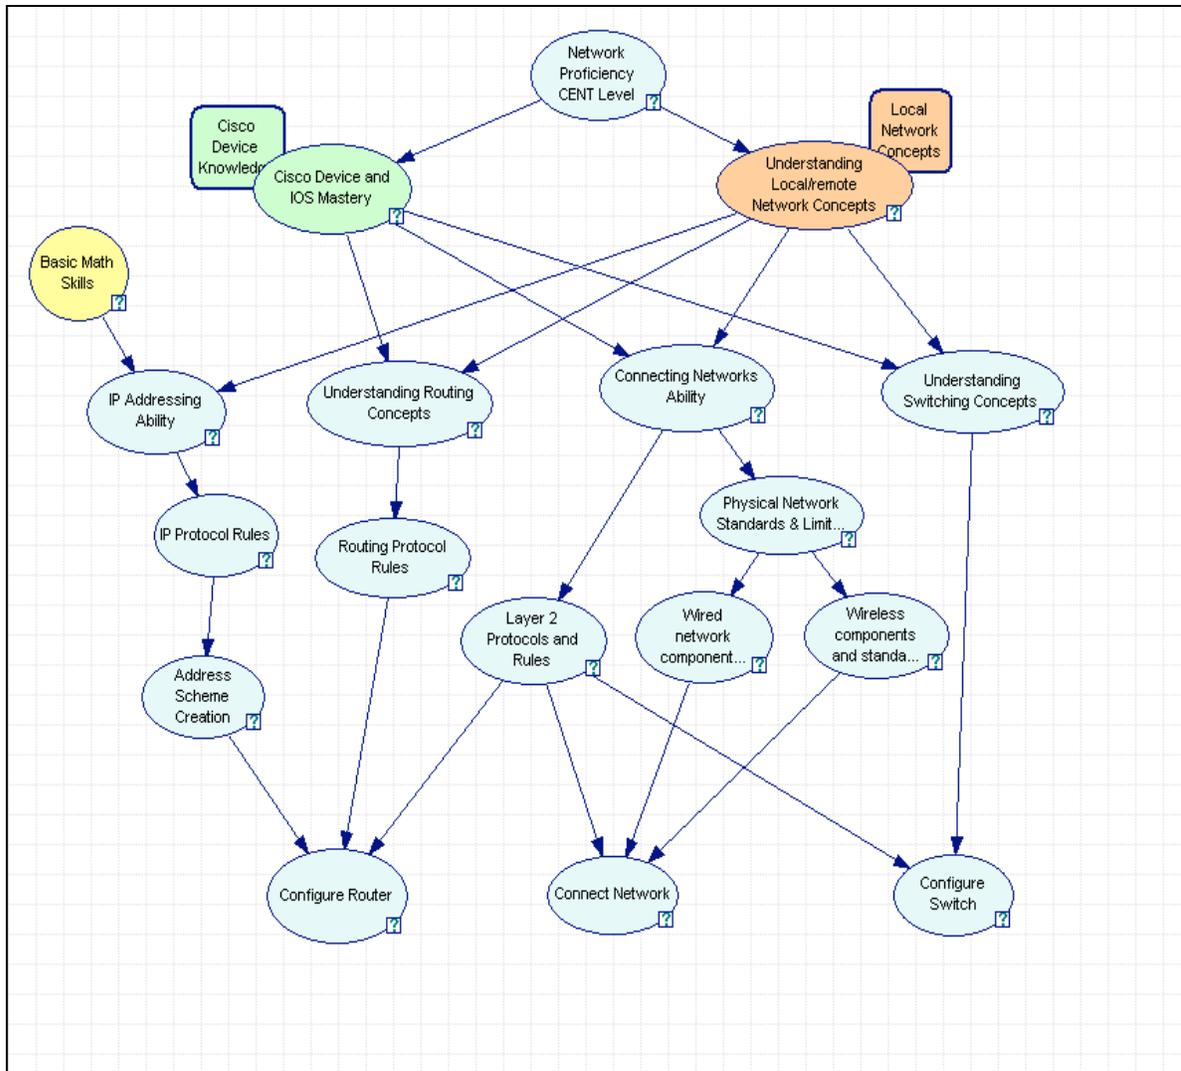


Figure 11. Large map of the Discovery curriculum. This model displays the relationship between different networking skills of which IP Addressing Skill is a small part.

Scaling from Small Success

There are a number of directions in which we could proceed following our small successes with BNs and LPs. One challenge will lie in knowing when to apply this method and when it may not be necessary. We need to determine whether to use these methods primarily in research to inform curriculum and assessment design or move them into operations, where they would provide feedback into the system of more than 17,000 instructors and 700,000 students a year. Such a move into an operations setting would require the automated construction of Bayesian networks in the four-process delivery system (Frezza, Behrens, Mislevy, West, & DiCerbo, 2009). It would also entail constructing Bayesian networks in near-real time. Although this has been done before in intelligent

tutoring applications, it is a challenging assignment. The benefits to instructors and students would need to be weighed against the resources needed to make this a reality. Overall, the challenge lies in knowing when to use these tools and when something simpler might be nearly as effective.

Additional Assessment Types

This report focused on analysis using traditional multiple-choice exams. In the future, we will be using the same process in the analysis of assessments from a Networking Academy tool called Packet Tracer (PT). PT is a comprehensive simulation, visualization, collaboration, and micro-world authoring tool for teaching networking concepts (Frezzo, Behrens, Mislevy, West & DiCerbo, 2009). PT assessments are being constructed using design patterns and task templates to create complex tasks at appropriate levels (Wise-Rutstein, 2005; Frezzo, et al., 2009). These design patterns and templates additionally provide structure for conditional probabilities in the Bayesian networks and thus cast the interpretation of performance in terms of the LPs through the SMVs. Data from a field trial of an earlier prototype of Packet Tracer called NetPASS were successfully modeled using Bayesian networks (Levy & Mislevy, 2004), although not in the framework of learning progressions.

In the next phases, we will be seeing the larger-scale deployment of PT as an assessment tool, providing automated scoring of performance assessments. We will be applying the methods described in this report, and others as needed, to the assessment information resulting from students completing those tasks. We anticipate that the inclusion of this rich information will provide new insight into students' learning progressions. However, it will also present challenges in trying to model this more complex data. For example, we will be attempting to use Natural Language Processing (NLP) extraction of features from student command logs to build observables. In addition, we will attempt to model the PT tasks using the same LP SMVs as their parents.

Closing the Feedback Loop

In this project, teacher and subject matter experts have served as inputs into the modeling process. In the future, we need to continue to close the loop to that information resulting from the modeling then feeds back to inform future instruction and curriculum design. As such, methods of communication of both student-level and aggregate results need to continue to be refined. With these and other improvements, the Bayesian network modeling of learning progressions will play an important role in understanding and improving student outcomes.

References

- Almond, R. G., DiBello, L. V., Moulder, B., & Zapata-Rivera, J. D. (2007). Modeling diagnostic assessments with Bayesian networks. *Journal of Educational Measurement*, *44*, 341–359.
- Almond, R. G., & Mislevy, R. J. (1999). Graphical models and computerized adaptive testing. *Applied Psychological Measurement*, *23*, 223–237.
- Behrens, J. T., Collison, T. A., & Demark, S. F. (2005). The seven Cs of comprehensive assessment: Lessons learned from 40 million classroom exams in the Cisco Networking Academy Program. In S. Howell & M. Hricko (Eds.), *Online assessment and measurement: Case studies in higher education, K-12 and corporate* (pp. 229–245). Hershey, PA: Information Science Publishing.
- Behrens, J. T., Frezzo, D. C., Mislevy, R. J., Kroopnick, M., & Wise, D. (2007). Structural, functional, and semiotic symmetries in simulation-based games and assessments. In E. L. Baker, J. Dickieson, W. Wulfeck, & H. F. O’Neil (Eds.), *Assessment of problem solving using simulations* (pp. 59–80). New York, NY: Erlbaum.
- Borsboom, D. (2006). The attack of the psychometricians. *Psychometrika*, *71*, 425–440.
- Brooks, S. P. (1998). Markov chain Monte Carlo method and its application. *The Statistician*, *47*, 69–100.
- Conati, C., Gertner, A., & VanLehn, K. (2002). Using Bayesian networks to manage uncertainty in student modeling. *User Modeling and User-Adapted Interaction*, *12*(4), 371–417.
- DeMark, S. F., & Behrens, J. T. (2004). Using statistical natural language processing for understanding complex responses to free-response tasks. *International Journal of Testing*, *4*, 371–390.
- Dayton, C. M. & Macready, G. B. (2007). Latent class analysis in psychometrics. In C. R. Rao & S. Sinharay (Eds.), *Handbook of statistics* (Vol. 26, pp. 421–446). North-Holland: Elsevier.
- Embretson, S. E. (1998). A cognitive design system approach to generating valid tests: Application to abstract reasoning. *Psychological Methods*, *3*, 380–396.
- Frezzo, D. C., Behrens, J. T., Mislevy, R. J. (in press). Design patterns for learning and assessment: Facilitating the introduction of a complex simulation-based learning environment into a community of instructors. *Journal of Science Education and Technology*.
- Frezzo, D. C., Behrens, J. T., Mislevy, R. J. West, P., & DiCerbo, K. E. (2009). *Psychometric and evidentiary approaches to simulation assessment in Packet Tracer software*. ICNS, pp.555–560, *2009 Fifth International Conference on Networking and Services*. <http://doi.ieeecomputersociety.org/10.1109/ICNS.2009.89>
- Jensen, F.V. (1996). *An introduction to Bayesian networks*. New York, NY: Springer-Verlag.

- Lauritzen, S. L., & Spiegelhalter, D. J. (1988). Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society, Series B*, 50, 157–224.
- Lazarsfeld, P. F., & Henry, N. W. (1968). *Latent structure analysis*. Boston, MA: Houghton Mifflin.
- Levy, F., & Murnane, R. J. (2004). *The new division of labor: How computers are creating the next job market*. Princeton, NJ: Princeton University Press.
- Levy, R., & Mislevy, R. J. (2004). Specifying and refining a measurement model for a computer-based interactive assessment. *International Journal of Testing*, 4, 333–369.
- Linzer, D. A., & Lewis, J. (2007). *poLCA: Polytomous variable latent class analysis*. (Version 1.1) [Computer software]. Emory University.
- McLachlan, G., & Peel, D. (2000). *Finite mixture models*. New York, NY: John Wiley.
- Mislevy, R. J., Hamel, L., Fried, R. G., Gaffney, T., Haertel, G., Hafter, A., ... Wenk, A. (2003). Design patterns for assessing inquiry in science. *PADI Research Report #1*. Menlo Park, CA: SRI International.
- Mislevy, R. J., & Levy, R. (2007). Bayesian psychometric modeling from an evidence-centered design perspective. In C. R. Rao & S. Sinharay (Eds.), *Handbook of statistics* (Vol. 17, pp. 839–865). North-Holland: Elsevier.
- Mislevy, R. J., Steinberg, L. S., & Almond, R. G. (2003). On the structure of educational assessments. *Measurement: Interdisciplinary Research and Perspectives*, 1, 3–62.
- Murnane, R. J., Sharkey, N. S., & Levy, F. (2004). A role for the internet in American education? Lessons from the Cisco Networking Academies. In P.A. Graham & N. G. Stacey (Eds.), *The knowledge economy and postsecondary education: Report of a workshop* (pp.127–158). Committee on the Impact of the Changing Economy on the Education System, National Research Council. Washington, DC: National Academies Press.
- Murray, R. C., & VanLehn, K. (2000). DT Tutor: A decision-theoretic, dynamic approach for optimal selection of tutoring actions. In G. Gauthier, C. Frasson, & K. VanLehn (Eds.), *Intelligent tutoring systems. ITS 2000* (pp. 153–162). *Fifth International Conference Intelligent Tutoring Systems*. Montreal, Canada. New York, NY: Springer.
- Muthén, L. K., & Muthén, B.O. (1998–2006). *MPlus user's guide*. Fourth edition. Los Angeles, CA: Muthén & Muthén.
- Norsys Software Corporation. (2007). *Netica manual*. Retrieved from: <http://www.norsys.com/netica.html>
- Nylund, K. L., Asparouhov, T., & Muthén, B. O. (2007). Deciding on the number of classes in latent class analysis and growth mixture modeling: A Monte Carlo simulation study. *Structural Equation Modeling: A Multidisciplinary Journal*, 14, 535–569.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Kaufmann.

- R Development Core Team (2008). *R: A language and environment for statistical computing*. Retrieved from: <http://www.r-project.org/>
- Reye, J. (2004). Student modeling based on belief networks. *International Journal of Artificial Intelligence in Education*, 14, 63–96.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461–464.
- Spiegelhalter, D. J., Dawid, A. P., Lauritzen, S. L. & Cowell, R. G. (1993). Bayesian analysis in expert systems. *Statistical Science*, 8, 219–247.
- Van Merriënboer, J. J. G. (1997). *Training complex cognitive skills: A four-component instructional design model for technical training*. Englewood Cliffs, NJ: Educational Technology Publications.
- Wise-Rutstein, D. (2005, April). *Design patterns for assessing troubleshooting in computer networks*. Paper presented at the annual meeting of the American Educational Research Association, San Francisco, CA.

Appendix A

IP Addressing Skills Progression

Level 1: Novice—Knowledge/Skill (possible pre-course knowledge and skills)

- Student can navigate the operating system to get to the appropriate screen to configure the address.
- Student knows that four things need to be configured: IP address, subnet mask, default gateway and DNS server.
- Student can enter and save information.
- Student can use a web browser to test whether or not network is working.
- Student can verify that the correct information was entered and correct any errors.
- Student knows that DNS translates names to IP addresses.
- Student understands why a DNS server IP address must be configured.

Level 2: Basic—Knows Fundamental Concepts

- Student understands that an IP address corresponds to a source or destination host on the network.
- Student understands that an IP address has two parts, one indicating the individual unique host and one indicating the network that the host resides on.
- Student understands how the subnet mask indicates the network and the host portions of the address.
- Student understands the concept of local-vs.-remote networks.
- Student understands the purpose of a default gateway and why it must be specified.
- Student knows that IP address information can be assigned dynamically.
- Student can explain the difference between a broadcast traffic pattern and a unicast traffic pattern.

Level 3: Intermediate—Knows More Advanced Concepts

- Student understands the difference between physical and logical connectivity.
- Student can explain the process of encapsulation.
- Student understands the difference between Layer 2 and Layer 3 networks and addressing.
- Student understands that a local IP network corresponds to a local IP broadcast domain (both the terms and the functionality).
- Student knows how a device uses the subnet mask to determine which addresses are on the local Layer 3 broadcast domain and which addresses are not.
- Student understands the concept of subnets and how the subnet mask determines the network address.
- Student understands why the default gateway IP address must be on the same local broadcast domain as the host.
- Student understands ARP and how Layer 3 to Layer 2 address translation is accomplished.

- Student knows how to interpret a network diagram in order to determine the local and remote networks.
- Student understands how DHCP dynamically assigns IP addresses.³

Level 4: Advanced—Can Apply Knowledge and Skills in Context

- Student can use the subnet mask to determine what other devices are on the same local network as the configured host.
- Student can use a network diagram to find the local network where the configured host is located.
- Student can use a network diagram to find the other networks attached to the local default gateway.
- Student can use the PING utility to test connectivity to the gateway and to remote devices.
- Student can recognize the symptoms that occur when the IP address or subnet mask is incorrect.
- Student can recognize the symptoms that a default gateway is configured incorrectly.
- Student can recognize the symptoms that occur if an incorrect DNS server (or no DNS server) is specified.
- Student knows why DNS affects the operation of other applications and protocols, like email or file sharing.
- Student can use NSlookup output to determine if DNS is functioning correctly.
- Student can configure a DHCP pool to give out a range of IP addresses.
- Student knows the purpose of private and public IP address spaces and when to use either one.
- Student understands what NAT is and why it is needed.

Level 5: Expert—Can Readily Apply Advanced Skills

- Student can recognize a non-functional configuration by just looking at the configuration information, to testing of functionality required.
- Student can interpret a network diagram to determine an appropriate IP address/subnet mask/default gateway for a host device.
- Student can recognize the symptoms that occur if an incorrect subnet mask is configured on the intermediate routers or destination host.
- Student can interpret a network diagram in order to determine the best router to use as a default gateway when more than one router is on the local network.
- Student can evaluate a connectivity problem to determine if it could possibly be caused by an incorrect setting configured on the host.
- Student can propose changes to a host configuration to solve a connectivity problem.
- Student can make and test proposed changes to a host configuration to solve an identified connectivity problem.
- Student can implement NAT to translate private to public addresses.