# Thin Data, Thick Description: Modeling Socio-Environmental Problem-Solving Trajectories in Localized Land-Use Simulations

A. R. Ruis(✉), Yuanru Tan, Jais Brohinsky, Binrui Yang, Yeyu Wang, Zhiqiang Cai, and David Williamson Shaffer

Wisconsin Center for Education Research, University of WI–Madison, Madison, WI 53706, USA
arruis@wisc.edu

**Abstract.** Many learning technologies are now able to support both user-customization of the content and automated personalization of the experience based on user activities. However, there is a tradeoff between customization and personalization: the more control an educator or learner has over the parameters that define the experience, the more difficult it is to develop learning analytic models that can reliably assess learning and adapt the system accordingly. In this paper, we present a novel QE method for automatically generating a learning analytic model for the land-use planning simulation *iPlan*, which enables users to construct custom local simulations of socio-environmental issues. Specifically, this method employs data simulation and network analysis to construct a measurement space using nothing but log data. This space can be used to analyze users' problem-solving processes in a context where normative measurement criteria cannot be specified in advance. In doing so, we argue that QE methods can be developed and employed even in the absence of rich qualitative data, facilitating thick(er) descriptions of complex processes based on relatively thin records of users' activities in digital systems.

**Keywords:** QE methods · data simulation · network analysis · learning analytics · problem solving · trajectory analysis · environmental education

## 1 Introduction

Learning technologies are increasingly designed to be *adaptable* to the needs of learners. This has generally taken two forms, often hybridized: (1) *customization*, in which educators can modify digital environments for their local contexts, learning objectives, and learner populations, or learners themselves can make modifications based on their interests or personal preferences; and (2) *personalization*, in which the technology adapts itself automatically, determining content and progression or providing formative feedback in real time [1–3].

Although adaptable educational technologies can have positive impacts on learning (e.g., [4, 5]), there is a tradeoff between customization and personalization: the more


G. Arastoopour Irgens and S. Knight (Eds.): ICQE 2023, CCIS 1895, pp. 349–364, 2023.
https://doi.org/10.1007/978-3-031-47014-1_24

customizable the educational experience is, the more difficult it is to design learning analytic models or other automated assessment systems that can be reliably used to measure learning. In other words, constructing normative assessments of learning in complex learning environments when neither the content nor the context is standardized presents significant challenges [6].

While this is generally true, it is especially so in environmental education, which often deals with *socio-environmental systems*: complex interactions among human (social, political, economic) and natural (biophysical, ecological, environmental) processes [7]. To make complex socio-environmental systems, such as land use, water quality management, or climate change, more accessible to learners, one effective approach is to *localize* them [8–10]. This situates authentic, real-world problems in a real place, one that students know and care about. For example, the online platform *iPlan* [11, 12] enables educators (or learners) to construct simulated land-use planning problems in which the location and the social and environmental issues are selected by the user. Some features are also non-deterministic, such that each simulated problem is different, even when the input choices are the same. While this enables educators to create localized simulations of realistic land-use problems that are well adapted to their curricula and contexts, it is particularly challenging to construct learning analytic models of student problem solving because each problem is unique and sufficiently complex to support manifold appropriate solution pathways.

Like most digital learning environments, *iPlan* records user activities in log files. This provides a large amount of data but a relatively low amount of information about learning and problem solving; that is, it is big data but not necessarily *thick* data. This presents further challenges to constructing learning analytic models, one that is too often solved by throwing all the data at a carousel of models in the pursuit of statistical significance.

In this paper, we present a pilot approach to modeling problem-solving processes in *iPlan* using only clickstream data. This approach leverages data simulation [13] and a novel network modeling method to create a normative measurement space that enables analysis of individual problem-solving trajectories. This method, which could inform the construction of assessment models for both *iPlan* and other highly customizable problem-solving spaces, makes it possible to produce interpretable representations of solution trajectories from click data. We argue that in the absence of richer ethnographic data, such as recordings of think alouds or interviews that can be used to assess how individuals develop solutions to complex problems, this approach enables construction of thick(er) descriptions from the relatively thin records of users' activities in digital systems.

## 2 Background

### 2.1 iPlan

*iPlan* [11] enables users to construct a realistic, localized land-use planning simulation for any location in the contiguous United States. The process of simulation construction is explained in detail elsewhere [14], but in brief, users (a) select a location using a

Google Maps interface and (b) choose five ecological and socio-economic indicators—measures of air and water pollution, greenhouse gas emissions, wildlife population levels, agricultural production, commercial activity, and housing—to include in the simulation. Based on the location and indicators selected, *iPlan* generates a land-use map of the selected region with at most 200 parcels and nine virtual stakeholders—business owners, activists, and concerned citizens—who advocate for different issues that the indicators reflect. *iPlan* uses a set of optimization routines to divide the selected region into parcels, assign an appropriate land-use class to each parcel, and set stakeholder thresholds—minimum or maximum satisfactory values—for the selected indicators. Collectively, this process results in localized, reduced-form simulations that are realistic and appropriately complex for non-specialists, who can use *iPlan* to explore some of the scientific and social challenges involved in land-use planning and management [12] (see Fig. 1 for an example of the map interface).

In *iPlan*, the goal is to produce a new land-use plan for the modeled region that satisfies as many stakeholders as possible. To do this, learners use a map interface to model the effects of specific land-use changes on the selected indicators. They then create *land-use scenarios* [15] and submit them to the virtual stakeholders for feedback. Learners have a limited number of feedback requests, so they are challenged to conduct experiments, or *stated preference surveys* [16], that help them determine with more precision the changes each stakeholder will accept.

Because the simulated stakeholders have different and often conflicting demands, learners must identify and negotiate trade-offs. For example, one stakeholder may advocate for an increase in jobs, which is easiest to accomplish by rezoning parcels for commercial or industrial use, but another stakeholder may want a decrease in greenhouse gas emissions, which will increase with commercial or industrial expansion. Thus, *iPlan* models not only the *effects* of land-use change on socio-economic and environmental indicators but also the *acceptability* of land-use change to various civic interest groups, and it is generally impossible to satisfy everyone simultaneously. That is, *iPlan* constructs simulations that help people learn about the scientific *and* civic practices through which land-use planning is managed and contested, helping them understand land-use management as a complex socio-environmental system.

### 2.2 Assessing Problem Solving in *iPlan* Simulations

Educators have found *iPlan* to be a useful pedagogical tool across a number of learning contexts [14], and localization of socio-environmental learning using simulation is a powerful technique for improving learning and civic engagement [9]. However, the complexity of the solution space in *iPlan* makes developing learning analytic models—and thus providing personalized scaffolding based on formative assessment—particularly challenging.

Each land-use simulation created by the *iPlan* system is a unique result of user selection (region, indicators) and non-deterministic optimization (parcelization, stakeholders' preferences), and there are no strong constraints on user actions. Moreover, there are many interaction effects: the impact of land use on indicators is determined by both the type of land use and the area (i.e., parcel size), and each land-use class influences multiple indicators.

Because of this, the simulations produced by *iPlan* present challenges to modeling problem-solving processes in two primary ways. First, while the division of the user-selected region into 200 parcels is not arbitrary—the boundaries are determined by Census boundaries—the parcels are formed based on an optimization algorithm that minimizes land-use assignment error and avoids significant asymmetries in parcel size. Thus, even if a user selected approximately the same region a second time, the resulting parcelized map may not be exactly the same. Moreover, learners can change any of 200 parcels to one of 10 other land-use classifications, resulting in $11^{200}$ possible land-use scenarios that can be constructed for a given land-use simulation. In other words, the problem space is large and relatively unbounded, and the exact features of any land-use map will not be available until it is created.

Second, the preferences of the virtual stakeholders are set such that the resulting land-use problem space is neither too simple nor too complex for teenagers and non-specialist adults. The system selects nine out of 57 possible stakeholders for each land-use simulation based on the indicators chosen and a prioritization algorithm, then runs an optimization routine to set the indicator threshold for each stakeholder. Because learners are trying to satisfy as many of the nine stakeholders as possible by making strategic land-use changes, and the thresholds determine *how much change* in the indicators is needed to satisfy each stakeholder, there are many possible solution pathways that cannot be specified in advance or even optimized mathematically.

As a result, it is difficult to assess user actions in the simulation. There are some universally useful strategies; for example, it is always helpful to begin by submitting the initial map, with no land-use changes, to all of the stakeholders in order to determine what the stakeholders want and to identify whether any of them are already satisfied. However, successful solution strategies will generally differ depending on the particular features of the simulation and the ways different learners negotiate the tradeoffs and challenges intrinsic to the problem.

In what follows, we present a method for automatically constructing a measurement space that accounts for the unique features of a given *iPlan* simulation and enables meaningful interpretation of the land-use scenarios that users construct and submit to stakeholders in that simulation. This method extracts information from log files that is otherwise inscrutable—that is, summary information about the type and amount of land-use changes made—and enhances understanding of both solution processes and outcomes. We then present two constructed cases to illustrate some affordances of the method and discuss future directions for this pilot work.

## 3 Methods

To model learners' problem-solving processes in *iPlan*, we developed an analytic procedure with three main components: (a) a *data simulation algorithm* that uses the features of a given *iPlan* simulation to construct a large and diverse set of land-use scenarios representative of the kinds of proposals that learners might submit to the virtual stakeholders; (b) a *measurement model* that uses the simulated data to construct a metric space into which learners' land-use scenarios can be projected, producing a summary measure of their decisions over the course of the simulation; and (c) a *coordinated visualization* that facilitates interpretation of learners' problem-solving trajectories. This process transforms unreadable click data into a meaningful, unified representation of learners' problem-solving approaches.

### 3.1 Data Simulation

Because the primary goal in *iPlan* is to construct a land-use scenario that pleases as many stakeholders as possible, it is important that the simulated data contain scenarios that cover all the stakeholders' preferences (and dispreferences). To do this, the data simulation algorithm produces 100 scenarios for each of the nine stakeholder, 50 that satisfy the stakeholder and 50 that do not. This set of 900 land-use scenarios is used as the input data for the measurement model.

Each stakeholder in *iPlan* is associated with one indicator, and their satisfaction or dissatisfaction with a given scenario is determined by whether that indicator is above or below their threshold. The effects of the 11 land-use classes on indicators are computed using a set of equations that take into account the area of each land-use class and the magnitude of the effect a given land-use class has on a given indicator per unit of area. In other words, for any given scenario submitted to the stakeholders, *iPlan* computes the indicator values and compares them to the stakeholders' thresholds to determine whether the stakeholders are satisfied or dissatisfied with that scenario.

To generate simulated scenarios that are likely to satisfy (or dissatisfy) a given stakeholder, the data simulation algorithm uses the features of the specific *iPlan* simulation—the area and initial land-use class of each parcel; the indicators selected and the models that relate land-use classes to those indicators in that location; and the stakeholders' preferences (thresholds and directionality) as inputs. Because the effects of each land-use on each indicator are known, for each indicator, we constructed two lists of land-use classes: the first list (List A) contains land-use classes with a large effect on the indicator; the second list (List B) contains land-use classes with a small or no effect on the indicator. In cases where a given land-use class has a moderate effect on the indicator, that land-use class is included in both lists. (For the purposes of data simulation and model construction, two of the land-use classes—*limited use* and *conservation*—are combined because they have identical effects on all indicators. Thus there are 10 land-use classes included in the lists.) These lists are constructed such that changing a parcel with a land-use class in List A to a (different) land-use class in List B will generally increase the value of the indicator, and making changes in the opposite direction (from land-use classes in List B to those in List A) will generally decrease the value of the indicator. Using this information as inputs, the algorithm used to generate 100 scenarios for each stakeholder is as follows:

1. A number, $k$, between 1 and 17 is randomly selected to determine the number of parcels whose land-use class will be changed. This range was selected because approximately two-thirds (64%) of the more than 1,300 scenarios submitted by users in a one-year period contained fewer than 18 land-use changes, and also because a secondary goal of the simulation is to maximize the number of stakeholders satisfied *while minimizing the amount of land-use change.*
2. The algorithm randomly chooses a land-use class from List A, randomly chooses a parcel with that land-use class, and changes it to a (different) randomly selected land-use class from List B. This process is repeated $k$ times.

3. The value of the indicator that the stakeholder cares about is computed and compared against the stakeholder's threshold. If it is above the threshold, it is saved. If it is below, it is discarded.
4. Steps 1–3 are repeated until 50 scenarios are generated with indicator values above the stakeholder's threshold.
5. Steps 1–4 are then repeated in the opposite direction, that is, making land-use changes from List B to List A (Step 2) and saving scenarios with indicator values below the stakeholder's threshold (Steps 3–4).
6. Depending on whether the stakeholder wants the indicator to be below or above the threshold, one set of 50 scenarios represents satisfaction and the other represents dissatisfaction.

This process is conducted for each of the nine stakeholders, resulting in 900 scenarios.

Each land-use scenario is represented by a vector with 100 terms, where each term represents a unique ordered pair of 10 different land-use classes, and the value of each term is the total amount of area changed relative to the starting land-use map. Because there are a maximum of 17 land-use changes in any given scenario, most of the terms in the vectors are zeroes.

### 3.2 Constructing a Measurement Model

To construct a measurement model, we use the simulated data to parameterize a metric space into which learner-generated land-use scenarios can be projected. To do this, we (a) sphere normalize the 900 vectors generated by the data simulation algorithm; (b) construct nine dimensions, where each dimension maximizes the difference between the 50 scenarios that satisfy a given stakeholder and the 50 scenarios that dissatisfy that stakeholder; (c) perform a dimensional reduction using *singular value decomposition* (SVD); and (d) project the 900 scenarios into the reduced space formed by the first two SVD dimensions, which account for the most and second-most variance in the data, respectively. The details of this process are as follows.

Let $M$ be a $900 \times 100$ matrix, where each row corresponds to a simulated land-use scenario, and each column corresponds to a dimension in the feature space, that is, a unique ordered pair of the 10 land-use classes. Each row is a vector, $S$, of length 100 that either satisfies ($S_i = 1$) or does not satisfy ($S_i = -1$) the corresponding stakeholder. The vectors are sphere normalized by dividing each term by the total map area, converting the raw areas into proportions of total area. This accounts for differences in length between vectors. The normalized vectors are then represented as points in a 100-dimensional space.

For each stakeholder $j$ ($j = 1, 2, \ldots, 9$), the subset of points in the high-dimensional space where $S_j = 1$ ($n = 50$) and the subset where $S_j = -1$ ($n = 50$) are used to define a dimension. Specifically, the space is rotated (rigid-body rotation) so as to maximize the difference between the 50 points representing $S_j = 1$ and the 50 points representing $S_j = -1$. This results in nine dimensions, one for each stakeholder, each of which maximizes the difference between land-use scenarios that stakeholder likes and those they dislike.

An SVD is performed to construct a reduced set of dimensions that relate the type and magnitude of land-use change to stakeholder satisfaction across all nine stakeholders.

Each scenario is thus represented by a set of SVD scores, and the first two dimensions can be used to define a normative metric space into which other scenarios can be projected[1].

### 3.3 Visualizing and Interpreting Problem-Solving Trajectories

All 900 simulated land-use scenarios are represented as points in the two-dimensional space formed by the first and second SVD dimensions. The ten land-use classes are placed as nodes in the space using the *ordered semantic co-registration layout* (OSCL), the same layout used in ordered network analysis (ONA) [18]. (For more on the mathematics and affordances of co-registration, see [19].) Then, each land-use scenario can be visualized as an ordered network graph using the same visualization as ONA, and the nodes (i.e., the land-use classes) can be positioned in the space such that the centroid of each network corresponds with the location of the corresponding scenario in the reduced space. This results in two coordinated representations: (1) one in which each land-use scenario is summarized by a single point, and (2) one in which each land-use scenario is represented as a directed network graph that indicates the type and proportional magnitude of the land-use changes made.

This space can be interpreted not only based on the node positions, as in ENA or ONA, but also by where in the space different stakeholders are satisfied. The mean, 95% confidence interval, and range of the points representing land-use scenarios that satisfy each stakeholder can be computed, providing a mapping of the space based on stakeholder preferences. Stakeholder satisfaction is, in effect, sets of land-use changes that produce desired results, and so clustering points based on the stakeholders' preferences provides an additional means of interpreting the space.

Using the rotation matrix produced by the measurement model, other land-use scenarios produced under the same simulation, such as ones constructed by learners, can be projected into this space and interpreted by how they locate relative to the stakeholders' areas of satisfaction. Series of land-use scenarios produced by learners thus form trajectories through the space, providing insight into the problem-solving approach that learners take and facilitating meaningful interpretation of decision-making beyond what can be determined based on the outcome of each submission (i.e., which stakeholders were satisfied).

[1] For readers who may be wondering why we don't simply apply the SVD to the set of normalized vectors directly and omit the step involving the construction of a dimension for each stakeholder, this is in part because SVDs do not perform well on relatively sparse matrices, i.e., matrices in which many or most of the coefficients are zeroes [17]. Attempts to do this produced dimensions with low variance explained (generally $< 3\%$) and poor co-registration (see §3.3). While there are many techniques specifically designed to decompose sparse matrices, we took an approach, inspired by means rotation in epistemic network analysis (ENA), that both addresses the sparse matrix problem and facilitates meaningful interpretation of the resulting space based on stakeholder preferences, which is useful given that the goal in *iPlan* is to maximize stakeholder satisfaction.

# 4 Proof of Concept

In what follows, we illustrate the method described above using one *iPlan* land-use simulation, and we show several constructed scenarios to demonstrate the affordances of the method.

## 4.1 Measurement Model of One *iPlan* Simulation

To construct a measurement model for evaluating problem-solving processes in *iPlan*, we developed a land-use simulation for Eugene, Oregon, in the northwestern United States. The simulation includes the indicators birds, runoff, greenhouse gas emissions, jobs, and population (see Fig. 1). This location was chosen because it exhibits a range of parcel sizes and land-use types, both developed and not. The central area contains mostly high-density housing, commercial, industrial, and recreation land, while the periphery contains mostly low-density housing, cropland, pasture, and land with limited human use.

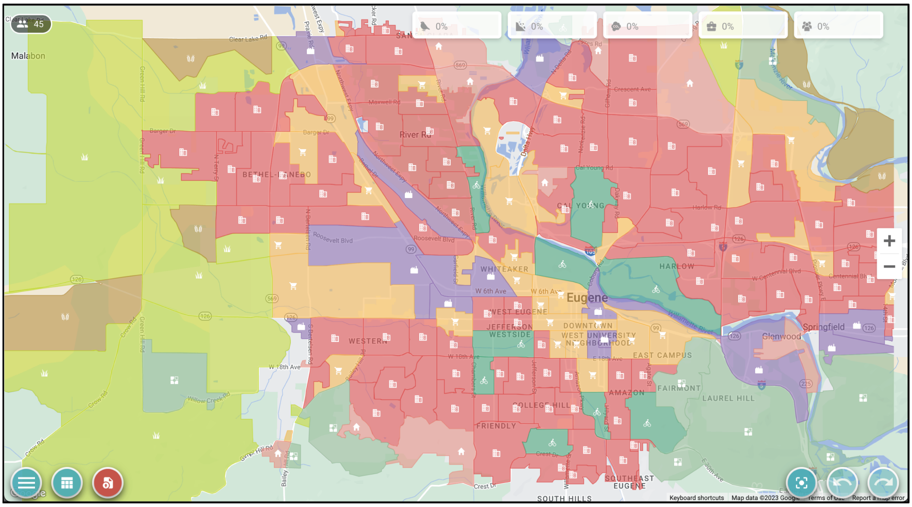


**Fig. 1.** *iPlan* simulation for Eugene, Oregon. Users can click any parcel(s) to change the land-use class, and the resulting effects on indicators (percentage change) are indicated at the top.

This simulation was used to construct the metric space shown in Fig. 2. The first ($x$) dimension accounts for 24% of the variance in the land-use scenarios, and the second ($y$) dimension accounts for 18% of the variance, indicating that the reduced space captures salient differences in the 900 land-use scenarios. The goodness of fit, or the correlation of the SVD scores with the corresponding network centroids, which is a measure of the extent to which the node positions can be used to interpret the space, is high: Pearson's and Spearman's $r > 0.95$ for both dimensions.

The first dimension generally distinguishes *low-intensity development* (low-density housing and recreation) from *high-intensity development* (commercial and industrial).

This is also the dimension that distinguishes the stakeholder who wants to increase bird populations (Grace) from the stakeholders who want to increase jobs (Ezra and Said). This makes sense, as commercial and industrial expansion significantly increases jobs, while low-intensity development, such as single-family homes, parks, and golf courses, provides ideal habitats for American robins (*Turdus migratorius*), the bird species that is modeled in *iPlan.*

The second dimension generally distinguishes *minimal development* (all of the land-uses in quadrant two) from *high-intensity development* (industrial, commercial, and high-density housing). This is also the dimension that distinguishes the stakeholders who want greenhouse gas emissions to decrease (Maya and Javier) from the stakeholders who want to increase jobs (Ezra and Said). This makes sense, as high-intensity development produces the most greenhouse gas emissions by area, while the lowest greenhouse gas emissions by area come from relatively undeveloped land.

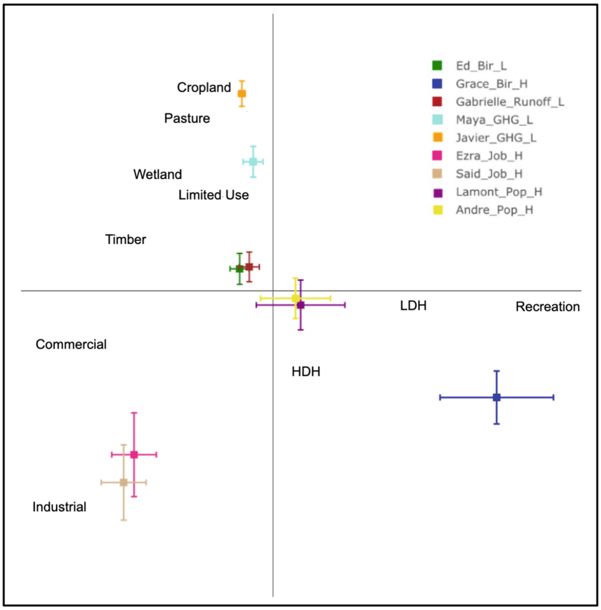


**Fig. 2.** Metric space for the Eugene, Oregon, *iPlan* simulation (Fig. 1). Colored squares with 95% confidence intervals are the mean locations of the scenarios that satisfied the associated stakeholders (key in upper right); the locations of the nodes (land-use classes) are labeled. HDH = High-Density Housing; LDH = Low-Density Housing; Limited Use includes Conservation

To evaluate the potential use of this modeling method as an assessment tool, we constructed two hypothetical cases using the Eugene simulation and projected the resulting land-use scenarios into the same metric space.

### 4.2 Case 1: Same Mountain, Different Paths

We constructed a set of submissions for each of two hypothetical *iPlan* users, User A and User B. Users A and B both adopt a similar, effective problem-solving strategy, but they carry out this strategy in distinct ways. Specifically, they both submit the initial map to gauge stakeholder preferences and then employ an *accretion strategy*, in which they attempt to make progressive land-use changes such that each submitted scenario contains all the changes from previous scenarios, the goal being to increase stakeholder support without losing any stakeholders who were previously satisfied. Both Users A and B have identical outcomes, in the sense that each scenario they submit satisfies exactly the same stakeholders (see Fig. 3). In other words, Users A and B employ the same basic strategy with the same outcome, but they make different decisions for each submission, which indicates different ways of thinking about the problem.

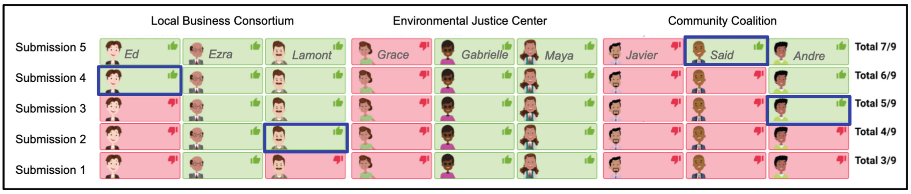


**Fig. 3.** Land-use scenario submission outcomes for User A and User B. Green indicates stakeholders satisfied with a submission, red indicates stakeholders dissatisfied with a submission, and blue boxes indicate stakeholders whose rating changed compared with the previous submission (Color figure online).

The problem-solving trajectories of Users A and B are shown in Fig. 4. (Note that the points are jittered to aid legibility.) Both users submit the initial map as their first scenario to gauge stakeholder preferences; because there are no land-use changes, the Submission 1 points of both users appear at the origin.

The biggest difference in the trajectories is in Submission 2 (see Fig. 5). User A made only one type of change, from land with limited human use to high-density housing, and as a result was able to satisfy Lamont, one of the stakeholders who is concerned with planning for an increasing population. User B also satisfied Lamont with Submission 2 by increasing high-density housing, but did so by proposing *infill*—replacing low-density housing with high-density housing rather than expanding to less developed land. In addition, User B also changed cropland to recreation. Given that this change has no impact on population, User B was most likely attempting to simultaneously satisfy Grace, the stakeholder who advocates for bird populations. While this attempt was not successful, it indicates that where User A may have preferred to focus on one indicator at a time, User B was likely attempting to satisfy multiple stakeholders across more than one indicator at once.

This case suggests that even when two users employ similar strategies and achieve identical outcomes, there are meaningful differences in their actions that an educator or the system itself could use as a basis for providing encouragement or additional scaffolding, or as an opportunity for broader discussions of land-use planning and civic

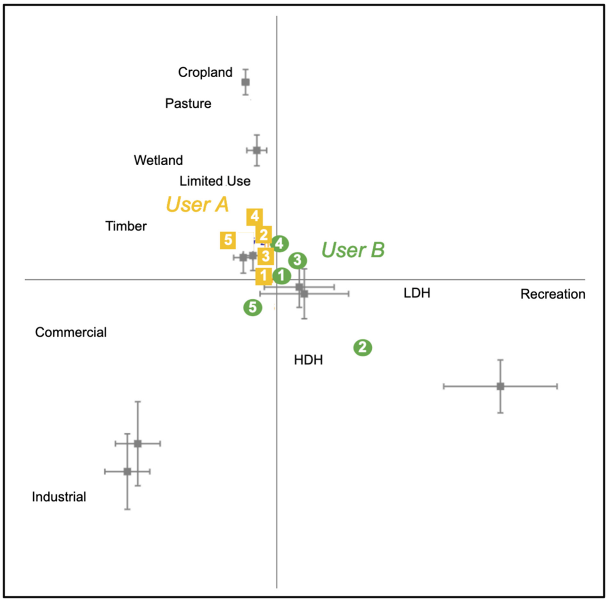


**Fig. 4.** Problem-solving trajectories of User A (yellow squares) and User B (green circles). The points are jittered to aid legibility (Color figure online).

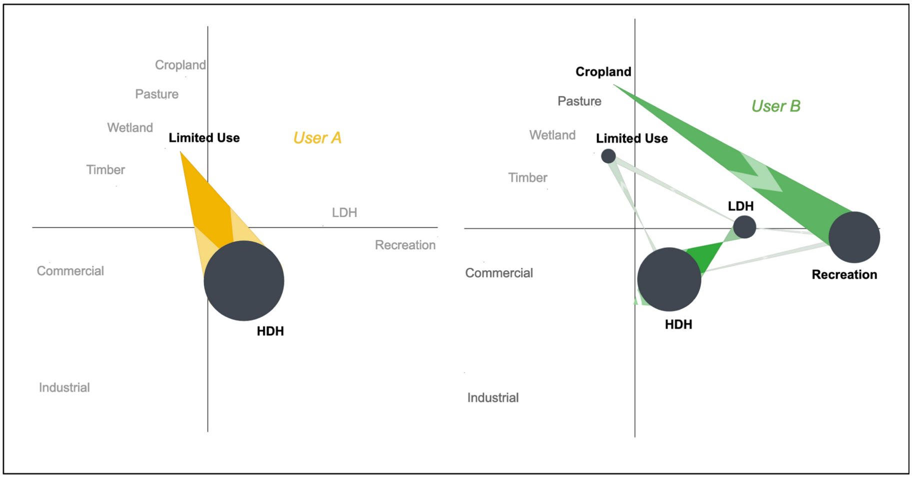


**Fig. 5.** Network graphs showing the land-use changes made by User A (left, yellow) and User B (right, green) for Submission 2 (Color figure online).

practices. For example, this one pair of submissions could spark a discussion of different planning strategies for addressing growing populations (e.g., infill vs. expansion) as well as broader discussions about complex problem solving when there are tradeoffs among indicators (e.g., the advantages and disadvantages of isolating one indicator at a time).

### 4.3 Case 2: If By Chance You Do Succeed…

We constructed a set of submissions for each of two hypothetical *iPlan* users, one who makes *random* land-use changes to submit to stakeholders and one who makes *systematic*, strategic land-use changes. The user making random changes finishes with a better result (6/9 stakeholders satisfied) than the user making systematic changes (5/9), and both sequences of submission outcomes appear to indicate progressive improvement (see Fig. 6, left). However, their trajectories reveal key differences (see Fig. 6, right).

The trajectory of the user making random submissions (purple squares) is erratic, moving to a different part of the space with each submission. This is indicative of a user simply clicking bunches of parcels and changing them without any particular goal in mind. The trajectory of the user making systematic submissions (blue circles) looks quite similar to that of User A in Case 1 (yellow squares in Fig. 4), with generally small-to-moderate changes and a clearer solution pathway. Note that both trajectories end in a similar place (Submission 5), but with slightly different results.

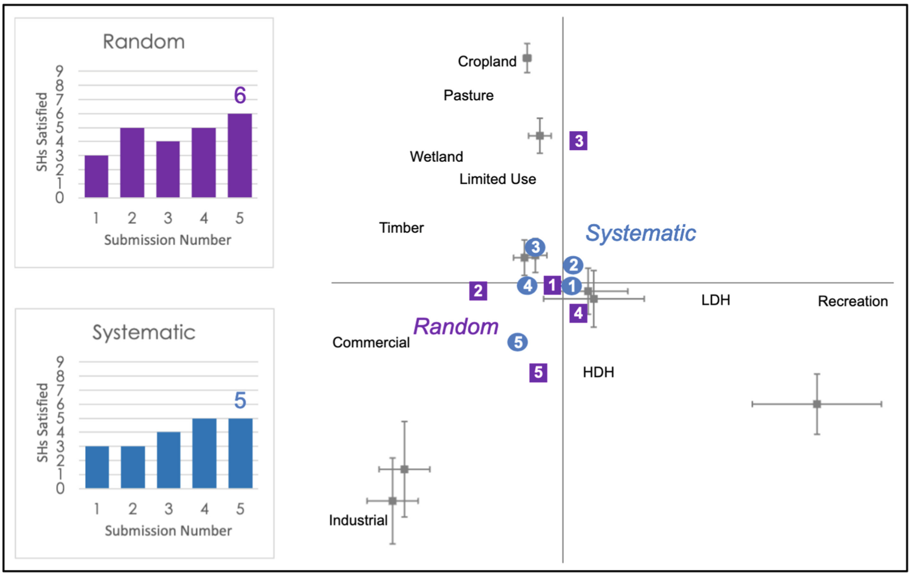


**Fig. 6.** Bar graphs (left) indicating the number of stakeholders satisfied by each scenario submitted and problem-solving trajectories (right) of two users: a user making *random* changes to construct land-use scenarios (purple squares) and a user making *systematic* changes (blue circles) (Color figure online).

This case suggests that the proposed method for modeling problem-solving trajectories could help educators and the system itself distinguish the activities of learners who are employing a good strategy but may require some additional scaffolding from the activities of learners who may be simply goofing around, even in cases when the outcomes do not disambiguate the two. That is, it can help distinguish those users who are trying unsuccessfully from those who are not trying but succeed anyway.

## 5 Discussion

At this point, you might be asking yourself if this study is actually quantitative *ethnography*: "Where is the *E* in QE?" While the proposed method for assessing problem-solving processes in a complex land-use simulation utilizes some of the mathematical and visual techniques commonly used in QE research, the only data come from log files, which in this case document the parameters and features of a given land-use simulation and record the land-use changes made and submitted to stakeholders by users. There is nothing that looks like a typical qualitative analysis, or even anything that looks like traditional qualitative data.

But while the data are, admittedly, thin, the clicks they represent are nonetheless a record of key decisions made and a latent reflection of the processes by which an individual attempts to solve a complex socio-environmental problem. The method we describe here enables a richer, *thicker* description of problem-solving activities to be constructed from data that are otherwise inaccessible to traditional qualitative analysis. That is, we use data simulation and modeling techniques to measure and visualize problem-solving trajectories by representing a key subset of the information captured in the log files in such a way that a meaningful story can be told about the behaviors they document.

Critically, we modeled only a very thin slice of the clickstream data, namely the land-use scenarios that users submitted to stakeholders. We omitted, among other things, the land-use changes users explored but didn't include in a submitted scenario, the use of the graphing tool to explore stakeholder preferences in more detail, and the resources that provide basic information about the land-use classes, indicators, and virtual stakeholders. While each of these data types could enhance understanding of learners' problem-solving strategies, they are more likely to occlude than clarify because it is harder to reliably link the digital record to a specific aspect of the problem-solving process. For example, a user could access a resource but not read it, read it but misunderstand it, or not access a resource yet still possess the knowledge it contains. All we know is whether they accessed a given resource, when, and for how long, and thus it is difficult to interpret the access records reliably enough for the purposes of understanding problem-solving strategies. In other words, without strong theoretical grounds or prior empirical work that would guide interpretation of this information, it is likely to add more noise than signal. Put another way, we gain more by pruning information than we do by adding it because this boosts the signal relative to the noise. This deep, theory-based engagement with the data, we argue, is a key feature of good quantitative ethnography and of good learning analytics [20]. And indeed, this entire modeling approach could only be developed because we have a deep understanding of *iPlan* and how it is used.

But this still leaves the question of how to close the interpretive loop; that is, how to warrant that the model is, indeed, well aligned with the original data and that re-interpretation of the original data in light of the model does not change the story. In this case, we didn't start with a qualitative interpretation of the data—that isn't often possible with clickstream data—but with a theory about different ways that a learner could solve the land-use problems simulated in *iPlan* and what that might look like as a series of submitted land-use scenarios. (And, because we constructed the case studies, we knew what users were trying to do.) Because our modeling process employs simulated data to construct the measurement space and projects real data into that space, closing the interpretive loop involves demonstrating both that the measurement space itself is aligned with the features of the specific *iPlan* simulation *and* that there is alignment between the model's representation of users' submission trajectories and the problem-solving strategies they represent. In this pilot study, we demonstrated that the data simulation and modeling process produces a sensible representation of the problem space based both on the node positions (land-use classes) and the locations of stakeholders in the space. We then showed that users' problem-solving trajectories were well aligned with the strategies they employed based on two constructed case studies: one comparing two users with similar strategies and identical outcomes, and one comparing two users whose outcomes may be misleading as indicators of effective problem solving. In both cases, the method was highly sensitive to key differences in problem-solving processes.

Because this is a pilot study, further work involving data from learners using *iPlan* in typical educational contexts and carefully designed studies that can link reported intentions and observed actions to model outputs will be needed to fully validate this approach. Such studies could also explore the inclusion of the other clickstream data types described above, or additional process data, by generating the theoretical basis for thoughtfully integrating that data with the submitted scenarios. Nonetheless, this initial study makes several contributions to QE methodology.

First, as Shaffer and Ruis have argued, QE is *not* just ENA [21]. While the method described here utilizes the ordered semantic co-registration layout from ONA [18] to generate directional network graphs that are co-registered with the metric space, the techniques used to accumulate connections and construct the metric space itself are novel and emerged from the specific challenges of analyzing decision-making and problem-solving in *iPlan*. This study thus adds to the growing methodological toolkit of QE research.

Second, this study provides an example of the use of data simulation in QE research, and indeed a use for data simulation beyond the four cases described by Swiecki and Eagan [13]. Learning technologies with highly customizable inputs and non-deterministic outputs can construct problems for which optimal solution pathways do not exist (or can't be reasonably derived). This means that norms cannot be established *a priori* for formative assessment of learner activities. Data simulation provides a mechanism for constructing a normative space into which learner activities can be projected and measured, providing the system with a basis to better scaffold learning, providing teachers with the information needed for just-in-time intervention and encouragement, and providing researchers with a powerful model for studying learning in a complex problem-solving context. While this is not the first example of such projection—Siebert-Evenstone, for example, projected planned and enacted curricula into a space constructed using the Next Generation Science Standards [22]—it demonstrates the utility of data simulation for constructing such spaces in contexts where prior data cannot be obtained.

Lastly, this study challenges, albeit indirectly, the assumption that more or richer data is necessarily better for modeling complex processes. The data used to model problem-solving in this study are quite thin, representing the types of land-use changes made and the amount of those changes at key points in the problem-solving process. That is, the data document only the decisions that users make when they submit land-use scenarios to stakeholders. But those data document arguably the most important decisions, and thus provide a good proxy for a learner's strategy. In other words, the model presented here was constructed based on a theory that links a key problem-solving behavior—choosing land-use scenarios to submit to stakeholders—to a broader problem-solving strategy. The choice of data—its thinness—was a strategic decision to minimize noise; that is, we discarded much of the information in the log files but also discarded most of the noise, leaving a high signal-to-noise ratio. In other words, we had less information, but it was more useful, and we argue that another key element of good QE research (and good learning analytics) is that every decision along the *primary modeling pathway* [20] is guided by theory about what to attend to and what to ignore.

This paper thus addresses a critical challenge in learning analytics and process modeling more broadly, namely the challenge of analyzing complex thinking or decision making in contexts where normative measurement criteria cannot be specified in advance. In doing so, we present a novel theoretical and methodological approach to generating thick descriptions from thin data.

**Acknowledgements.** This work was funded in part by the National Science Foundation (DRL-1661036, DRL-2100320, DRL-2201723, DRL-2225240), the Wisconsin Alumni Research Foundation, and the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin–Madison. The opinions, findings, and conclusions do not reflect the views of the funding agencies, cooperating institutions, or other individuals.

## References

1. Dede, C.: Customization in immersive learning environments: implications for digital teaching environments. In: Dede, C., Richards, J., (eds.) Digital Teaching Platforms: Customizing Classroom Learning for Each Student, pp. 282–297. Teachers College Press (2012)
2. Matuk, C.F., Linn, M.C., Eylon, B.-S.: Technology to support teachers using evidence from student work to customize technology-enhanced inquiry units. Instr. Sci. **43**, 229–257 (2015). https://doi.org/10.1007/s11251-014-9338-1
3. Shemshack, A., Spector, J.M.: A systematic literature review of personalized learning terms. Smart Learn. Environ. **7**, 33 (2020). https://doi.org/10.1186/s40561-020-00140-9
4. Littenberg-Tobias, J., Beheshti, E., Staudt, C.: To customize or not to customize? exploring science teacher customization in an online lesson portal. J. Res. Sci. Teach. **53**, 349–367 (2016)
5. Zhang, L., Basham, J.D., Yang, S.: Understanding the implementation of personalized learning: a research synthesis. Educ. Res. Rev. **31**, 100339 (2020)
6. Barab, S.A., Luehmann, A.L.: Building sustainable science curriculum: acknowledging and accommodating local adaptation. Sci. Educ. **87**, 454–467 (2003)
7. Elsawah, S., et al.: Eight grand challenges in socio-environmental systems modeling. Socio-Environ. Syst. Model. **2** (2020)
8. Gruenewald, D.A.: Foundations of place: a multidisciplinary framework for place-conscious education. Am. Educ. Res. J. **40**, 619–654 (2003)
9. Siebert-Evenstone, A.L., Shaffer, D.W.: Location, location, location: the effects of place in place-based simulations. In: Lund, K., Niccolai, G., Lavoué, E., Hmelo-Silver, C., Gweon, G., Baker, M. (eds.) A Wide Lens: Combining Embodied, Enactive, Extended, and Embedded Learning in Collaborative Settings: 13th International Conference on Computer-Supported Collaborative Learning (CSCL) 2019, pp. 152–159 (2019)

10. Smith, G.A., Sobel, D.: Place- and Community-Based Education in Schools. Routledge (2014)
11. Ruis, A.R., et al.: iPlan (2020). https://app.i-plan.us/
12. Ruis, A.R., et al.: Localizing Socio-Environmental Problem Solving. In: Weinberger, A., Chen, W., Hernández-Leo, D., Chen, B., (eds.) International Collaboration toward Educational Innovation for All: Overarching Research, Development, and Practices: 15th International Conference on Computer-Supported Collaborative Learning (CSCL) 2022, pp. 459–462. International Society for the Learning Sciences (2022)
13. wiecki, Z., Eagan, B.R.: The Role of Data Simulation in Quantitative Ethnography. In: Damşa, C., Barany, A., (eds.) Advances in Quantitative Ethnography: Fourth International Conference, ICQE 2022, Copenhagen, Denmark, 15–19 October 2022, Proceedings, pp. 87–100. Springer (2023). https://doi.org/10.1007/978-3-031-31726-2_7
14. Ruis, A.R., et al.: Iplan: A Platform for Constructing Localized, Reduced-Form Models of Land-Use Impacts (2023)
15. Xiang, W.-N., Clarke, K.C.: The use of scenarios in land-use planning. Environ. Plan. B Plan. Des. **30**, 885–909 (2003)
16. Tagliafierro, C., Boeri, M., Longo, A., Hutchinson, W.G.: Stated preference methods and landscape ecology indicators: an example of transdisciplinarity in landscape economic valuation. Ecol. Econ. **127**, 11–22 (2016)
17. Duff, I.S., Erisman, A.M., Reid, J.K.: Direct Methods for Sparse Matrices. Oxford University Press (2017)
18. Tan, Y., Ruis, A.R., Marquart, C.L., Cai, Z., Knowles, M., Shaffer, D.W.: Ordered Network Analysis. In: Damşa, C., Barany, A., (eds.) Advances in Quantitative Ethnography: Fourth International Conference, ICQE 2022, Copenhagen, Denmark, 15–19 October 2022, Proceedings, pp. 101–116. Springer (2023). https://doi.org/10.1007/978-3-031-31726-2_8
19. Bowman, D., et al.: The mathematical foundations of epistemic network analysis. In: Ruis, A.R., Lee, S.B. (eds.) ICQE 2021. CCIS, vol. 1312, pp. 91–105. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67788-6_7
20. Shaffer, D.W., Ruis, A.R.: Theories all the way across: the role of theory in learning analytics and the case for unified methods. In: Bartimote, K., Howard, S., Gašević, D., (eds.) Theory Informing and Arising from Learning Analytics. In press. Springer (2023)
21. Shaffer, D.W., Ruis, A.R.: Is QE just ENA? In: Damşa, C., Barany, A., (eds.) Advances in Quantitative Ethnography: 4th International Conference, ICQE 2022, Copenhagen, Denmark, 15–19 October 2022, Proceedings, pp. 71–86. Springer (2023).https://doi.org/10.1007/978-3-031-31726-2_6
22. Siebert-Evenstone, A.: A qualitative analysis of connection-making in the NGSS. In: Wasson, B., Zörgő, S., (eds.) Advances in Quantitative Ethnography: Third International Conference, ICQE 2021, Virtual Event, 6–11 November Proceedings, pp. 105–123. Springer (2022). https://doi.org/10.1007/978-3-030-93859-8_7