

Module 1. Introduction, Challenges, and Framework

- What is your past experience with statistics, data mining, and machine learning? (If any)
- What is your past history of work, practice, or research in digital learning platforms?
- Fast-forward 5 years into the future. What kind of work are you doing to make education better, using learning analytics in digital learning platforms?

Module 2. Data and Measurement Validity

Each of the following prompts focus on the concepts of (1) construct validity, (2) statistical conclusion and other types of validity, (3) internal validity and causal inference, (4) external validity and generalization. For each concept, respond to one prompt of your choice and do so in a well-developed, two-paragraph format. Also, choose one other post and respond to your peer's comments.

(1) Construct validity

1. Why are concepts related to construct validity important (to researchers or to the general public of consumers of research)?
2. What are ways to strengthen the construct validity of a tool or predictive model? Are there limitations in terms of time, resources, knowledge to doing so?

(2) Statistical conclusion and other types of validity

1. What kinds of practices or considerations researchers should engage in and make when considering ways to strengthen the validity of statistical conclusions?
2. What are some ways to discourage practices that jeopardize statistical conclusion and other forms of validity?
3. What role do you believe discussions of validity have in your area of research and experience?

(3) Internal validity and causal inference

1. What are the most common threats to internal validity in your experience? How do you think about causal relationships in your everyday experiences?
2. Do you think causality and the way that we assess it will change in the future? How so?
3. What are the ways that causality differs from correlational relationships? Why is this important?

(4) External validity and generalization

1. How important are elements of external validity in your everyday life and experiences?
2. What role do you think ideas of external validity will play in your future work in education?
3. One area of research that has been underdeveloped is how results potentially generalize over time. How do you think research or findings can generalize in time? Is it possible to do so?

Module 3. Feature Extraction and Feature Engineering

1. How do you personally approach data preprocessing? Share the specific techniques you apply to effectively handle and manage the data in your own work.
2. What specific challenges do you encounter in the process of feature extraction and feature engineering? Elaborate on how you overcome these challenges and share insights into the creative or innovative solutions you've implemented.

3. Are there any topics of feature extraction and feature engineering not covered in the module that you find useful? If so, which specific elements would you like to explore further for a more comprehensive understanding?

Module 4. Prediction Modeling and Metrics

- Are you aware of any digital learning platforms or technologies that are currently implementing prediction models? What is typically the focus of these prediction models?
- How much would you trust the application of a prediction model in practice? What is the threshold by which you would consider a model “good enough” for deployment? Is this true for all applications?
- What do you believe are the most important things (e.g. outcomes, behaviors, constructs) that we, as researchers/developers/practitioners, should be focusing on to predict/detect to have the largest positive impact on education?

Module 5. Causal Reasoning

- What are the most important uses of propensity score analysis in data science?
- What are some challenges for using propensity score analysis to evaluate digital learning platforms?
- What are some practical considerations for using regression discontinuity designs to evaluate digital learning platforms?
- What are some challenges of using machine learning for detecting treatment effect heterogeneity?

Module 6. Ethics, Equity, and Algorithmic Bias

1. **Pick a model:** Name one predictive model that is being used in your educational context. What does it predict, what data does it use, and who acts on the output? Keep this short, this is just set up for everything below. If nothing in your context fits, take a case from the readings or the program up to this point: an at-risk or dropout prediction model, an automated essay scorer, or an affect or engagement detector.
2. **Sources of bias:** Bias can enter at any stage of the system. Consider the pipeline inspired by Suresh and Guttag (2021):
World → Data collection → Data Preparation → Model development → Evaluation → Deployment → [World]
For your model, identify at least one stage where bias could occur. Try to be as specific as you can.
3. **Potential Harms - who could this model fail, and what does that cost them?** Based on the source of bias you identified above, which population(s) could this model produce negative outcomes for, due to bias? What is the nature of the negative outcomes? Does it impact students immediately? Can the negative impacts be longer term?
4. **Challenges in Auditing for Bias - would you ever find out?** What factors would make it difficult to determine whether your model is biased? Try to be as specific as you can. You can consider:
 - **Demographic data:** do you have the group labels you'd need to test for disparities, and are they accurate and complete?
 - **Unknown groups:** Which possible populations in your context are unlikely to appear in any audit at all, due to either being rare or unevenly distributed across contexts?
 - **Group size:** Are some populations present but too small in your data to produce a reliable estimate?

- **Ground truth:** Are the labels or outcomes you'd audit against themselves products of biased human judgment?
5. **Progression:** The questions so far have been about recognizing bias. This one is about acting on it. Consider the progression introduced in Baker and Hawn (2021). Where does your example model currently sit on that progression? What is the single most immediate, realistic step you could take to move toward the next stage?
 6. **Privacy vs bias trade-off:** Auditing for bias often requires collecting demographic data about students, but this data is sometimes unavailable due to steps taken to protect privacy. For your setting, what is the right balance between protecting privacy and maintaining fairness? How would you make this case to parents, students, or other stakeholders?

Module 7. Cluster Analysis

- Please replicate the code and markdown provided in the 2022 Bowers et al. reading for the mtcars dataset to visualize it using hierarchical cluster analysis heatmaps following the recommendations from the readings, including using uncentered correlation. <https://doi.org/10.7916/r1mg-yn37>
- To use the hopach library, if it is not available to install directly from CRAN, the latest version can be found here and installed with this code:

```
https://bioconductor.org/packages/hopach/
if (!requireNamespace("BiocManager", quietly = TRUE))
  install.packages("BiocManager")
BiocManager::install("hopach")
```

Then, select at least three different options/alternatives from ComplexHeatmap's first three chapters <https://jokergoo.github.io/ComplexHeatmap-reference/book/>, and apply them to the mtcars dataset. For example, such as using different clustering algorithms (k-means, etc.), separation of heatmaps, stacking, different ways to annotate (columns and rows and different graphing options beyond black/white blocks), etc.

1. For this discussion post then please post at least one of the cluster analysis heatmaps that you've generated. Please provide both the cluster analysis heatmap image and the code block used to generate the image. Then please discuss options/alternatives you selected, why this was interesting, and any challenges you experienced with running the visualizations. Provide a few sentences to interpret each new heatmap visualization.
2. For your current or future work, how do you envision applying cluster analysis or cluster analysis heatmaps to help you pattern and describe the data available to you? Elaborate on the benefits but also the challenges you would need to solve to use this technique in your work.
3. What further topics in cluster analysis would you find useful to explore that were either not discussed in this module or that were just touched on? Please provide specific issues and how exploring those issues could help you in your current or future work.

Module 8. Network Analysis

- How may network analysis be applied in your own work, especially in projects related to digital learning platforms? What specific network analysis techniques would be useful for your work?

- Drawing on the readings, what analytical decisions you may need to make if you adopt network analysis in one of your projects? What challenges do you anticipate?

Module 9. Sequential Pattern Mining and Temporal Analysis

- How may sequential or temporal analysis be applied in your own work, especial in projects related to digital learning platforms?
- Choose one project you are interested in. Explain how you could apply temporal or sequential analysis to investigate a learning or education phenomenon. What is the key theoretical construct? How is the construct captured in data? How can the construct be analyzed using temporal analysis methods?
- Are there other temporal or sequential analysis methods we should learn about? Please explain why and tell us how we could learn more about them.

Module 10. Neural Networks and Deep Learning

- The ability of deep learning models to automate parts of the feature engineering process has led many to rely on this approach over the traditional method of hand-crafting features (despite some research suggesting there may be performance benefits). Discuss your perspective on the benefits and drawbacks of these feature engineering approaches. Are there scenarios where one approach may be more appropriate or advantageous than the other?
- While often powerful, the complexity of most deep learning models often makes it difficult or even infeasible to understand how they are producing their predictions. Many argue that there is greater need for “explainable” AI. In practice, explainable methods often underperform in comparison to their deep learning counterparts. What are your thoughts on the trade-off between explainable methods and prediction accuracy in the context of applying these models in practice?
- Although it depends on the set of features available and how a task is structured, deep learning models typically require larger datasets to train than more traditional approaches. Despite this, there are often conflicting goals in regard to the collection of large educational datasets due to concerns over privacy, consent, and the ethical use of student information. Please share your thoughts on this issue and consider how developers, educators, and others may navigate these conflicting goals (i.e. what might be the risk of allowing or restricting access to education data)?

Module 11. Natural Language Processing

- What applications do you see for natural language processing in your own work?
- Natural language processing is more successful when more training data is available. What training data would you like to collect to address educational inequities using natural language processing?
- What elements of natural language processing that were not discussed in the module would you like to learn more about?

Module 12. Transformer and Foundation Models

- Can you identify and discuss real-world applications where models like BERT have significantly improved performance or enabled new capabilities? How do these applications benefit from the specific features of transformer models?

- What are some ethical considerations and potential limitations when using large language models like BERT in real-world applications? How might these models inadvertently perpetuate biases, and what steps can be taken to mitigate such issues?

Module 13. Knowledge Tracing

- What are the most important uses of knowledge tracing in digital learning platforms? (Particularly digital learning platforms that you are working on or see yourself working on)?
- What are some practical considerations for using knowledge tracing algorithms in real-world systems?
- Contemporary algorithms are usually more accurate at predicting future student correctness, but most digital learning platforms still use Bayesian Knowledge Tracing. Why? What steps do we need to take to see wider use of contemporary algorithms? (Or is BKT good enough?)

Module 14. Knowledge Graphs

1. How familiar are you with the Knowledge Graphs? Could you share an example of a knowledge graph you have observed? Why do you think this qualifies as a knowledge graph?
2. We have reviewed POKS, PFA and Q-Matrices as a few of the several techniques for creating and refining knowledge/skill prerequisite graphs. Choose one of these techniques and discuss briefly one of the limitations of your chosen method. Suggest an approach for resolving the limitation.

Module 15. Data Visualization

- What are the most important things we can do with data visualization?
- Think about a concrete example of a learning experience (i.e., a specific course on a specific subject for a specific group of learners). What kind of information do we need to change how a student is performing?
- Beyond checking for colorblindness, what other ways can we improve the accessibility of our visualizations?

Module 16. Data Management and Database Access

NA