Reconciling the Promise and Pragmatics of Enhancing Computing Pedagogy with Data Science

Austin Cory Bart, Dennis Kafura, Clifford A. Shaffer, and Eli Tilevich

Virginia Tech

Blacksburg, Virginia

{acbart,kafura,shaffer,tilevich}@vt.edu

ABSTRACT

Data science keeps growing in popularity as an introductory computing experience, in which students answer real-world questions by processing data. Armed with carefully prepared pedagogical datasets, computing educators can contextualize assignments and projects in societally meaningful ways, thereby benefiting students' long-term professional careers. However, integrating data science into introductory computing courses requires that the datasets be sufficiently complex, follow appropriate organizational structure, and possess ample documentation. Moreover, the impact of a data science context on students' motivation remains poorly understood.

To address these issues, we have created an open-sourced manual for developing pedagogical datasets (freely available at <https://think.cs.vt.edu/pragmatics>). Structured as a collection of patterns, this manual shares the expertise that we have gained over the last several years, collecting and curating a large collection of real-world datasets, used in a dozen of universities worldwide. We also present new evidence confirming the efficacy of integrating data science in an introductory computing course. As a significant extension of our ongoing work, this study not only validates existing positive assessment, but also provides fine-grained nuance to the potential of data science as a motivational educational element.

CCS CONCEPTS

• **Social and professional topics** → **Computing education**; • **Information systems** → *Information integration*;

KEYWORDS

Datasets, Data Science, Computational Thinking

ACM Reference format:

Austin Cory Bart, Dennis Kafura, Clifford A. Shaffer, and Eli Tilevich. 2018. Reconciling the Promise and Pragmatics of Enhancing Computing Pedagogy with Data Science. In *Proceedings of the 49th ACM Technical Symposium on Computing Science Education, Baltimore, MD, USA, February 21-24, 2018 (SIGCSE '18)*, 6 pages. <https://doi.org/10.1145/3159450.3159465>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCSE '18, February 21-24, 2018, Baltimore, MD, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5103-4/18/02...\$15.00

<https://doi.org/10.1145/3159450.3159465>

1 INTRODUCTION

Data science—using datasets to explore and answer real-world questions—can compellingly contextualize introductory computing experiences. Applicable across a wide array of problems, data science techniques and methods present exciting opportunities for advancing a variety of disciplines. In introductory courses, pairing non-computing majors with data relevant to their long-term career interests improves the resulting learning experience's usefulness and relevance [4]. Further, many of the techniques used in processing datasets align closely with introductory computing course content. Although we have developed datasets for novice non-major students, pedagogical datasets can be used at any level of learning.

However, data science can be extremely challenging to integrate into courses. First, one must locate appropriate datasets of appropriate size and quality. These datasets then need to be cleaned and organized into a form suitable for novice learners. Finally, beginners should be able to discover and understand these datasets in a timely fashion. All of these activities have constraints and complications associated with them.

To ameliorate this problem, we introduce a new guide, “The Pragmatics of Pedagogical Dataset Development”. The *Pragmatics* is an open-sourced, web-based document (freely available at <https://think.cs.vt.edu/pragmatics>). The goal of the document is to help developers create “pedagogical datasets”—datasets specifically targeted at novices for learning purposes. The *Pragmatics* is an organized collection of design issues, affordances, considerations, suggestions, and the authors' experiences in creating a large repository of pedagogical datasets.

In addition to discussing key results related to our experience with preparing pedagogical datasets, we present the results of a study that explored introductory learners' experience with pedagogical datasets. This study significantly extends our prior work [4], with a larger body of participants. Although confirming our prior results, here we highlight the newly discovered insights.

This paper makes the following three major contributions:

- It concisely reviews the use of data science in computing education,
- It details the open-source *Pragmatics* manual, and
- It further validates the efficacy of data science as an introductory context.

The audience that can benefit from reading this paper includes instructors and curriculum developers who are interested in preparing pedagogical datasets or using them in their courses. These developers might be planning to submit to open repositories, or be developing datasets specifically for their own students.

2 DATA SCIENCE

In the past two decades, the field of Data Science has emerged as a popular area at the intersection of computer science, statistics, mathematics, and a number of other fields [8]. Data Science answers questions by building, exploring, and processing datasets. There are many models that define the term more strictly, but in general it can be described as an iterative model of collecting, sanitizing, processing, rendering, and interpreting information.

There are many overlaps between Data Science and topics typically covered in computing curricula. Data Science can be used as *content* or *context* [7]. The size and organization of the data can be used to motivate and exercise a variety of algorithmic and data structuring techniques as well as software engineering practices that are ideal for Computer Science classes. As a context, data science becomes a means to an end – the goal is to motivate students to learn by anchoring the learning in a compelling experience. Data science as a context does not mean that students should learn to become data scientists, anymore than it is the goal of a framework like Media Computation to teach students how to be professional computational artists [13]. Incorporating a context into a course incurs pedagogical penalties, requiring time on material that potentially distracts from core learning objectives. An instructor can downplay the focus on these topics, or emphasize subject matter's strengths (e.g., a statistics major might find it interesting to use their mathematical background to strengthen their investigation).

Data Science has been used as both an introductory context and content in a wide range of courses and research studies. Undergraduate curriculum specifically targeted at teaching Data Science have become increasingly popular [1, 10, 17]. Thus, a number of instructors have worked to make datasets available for individual projects [2, 9, 14], and some researchers have even explored using datasets in non-formal learning experiences such as “Datathons” [3]. Research projects have provided collections of datasets [4], visualization tools for those datasets [19], and analysis of the efficacy of data science in classrooms [20]. Readers interested in a comprehensive overview of the field of Data Science, including its roles in both industry and classrooms, can refer to Cao’s survey paper [6].

3 PEDAGOGICAL DATASETS

When teaching data science, projects and assignments will typically be structured around the use of datasets. These “pedagogical datasets” are differentiated from regular datasets in that they are specifically targeted at learners. Such datasets can be scaffolded or organized for a wide range of tasks: to make it easier to process, to highlight a particular class of problems, or to provide an authentic learning experience for the student, for example. Meanwhile, conventional dataset preparation is typically only concerned with processing datasets strictly to obtain meaning and significance for a stakeholder. The difference in design goals necessitates different approaches in building pedagogical datasets.

3.1 Prior Work

There has been little prior work in the literature on practical design and development of pedagogical datasets. The UCI Machine Learning Repository aimed to collect and organize many pedagogical datasets suitable for introductory Machine Learning topics [16].

Although an impressive collection, the materials are targeted for tasks in advanced computing topics. The STARS project collected real-world datasets available for introductory statistics courses [5]. Both projects seem to have stopped production some time ago.

More crucially, prior projects did not attempt to articulate the formal process of developing pedagogical datasets, to create enduring lessons for future developers who wish to pick up where these projects left off. Radinsky et al. [18] pose design principles for integrating datasets into inquiry-based learning curriculum, including a call for cultivating data, but stop short of best practices for organizing and disseminating said data. Verbert describes organizational principles for pedagogical datasets for educational data, specifically for Learning and Knowledge Analytics [21], but not data meant directly for student use. None of the formerly described projects codify best practices of preparing pedagogical datasets.

Although pedagogical datasets involve many distinct problems and require specialized expertise, we acknowledge that there is significant overlap with traditional dataset development. The design and development of datasets is a fundamental topic in Data Science. Entire textbooks have been dedicated to the topic, and introductory curriculum spend significant amounts of time covering subproblems in this field. Readers interested in learning more general-purpose techniques are recommended to consult a text such as the one by Fry [11]. Although this particular text was discovered by the authors after the creation of the *Pragmatics*, Fry incorporates many of the same principles in his descriptions of the data science process, such as the need for managing data over time or the importance of understanding your audience. Our goal in writing the *Pragmatics* is to convey enough fundamentals that a novice dataset developer would not need to refer to external texts, while focusing on topics that require effort unique to pedagogical datasets.

3.2 Design Considerations

The design criteria for pedagogical datasets diverges from the criteria for conventional datasets, although there are significant overlaps. Fundamentally, the goal of datasets is to prepare data for 1) consumption by a general audience, 2) transmission to specific stakeholders, or 3) immediate processing into a visualization or summarization. Therefore, the proper design of the datasets is usually aimed towards organizing the form of the dataset for the audience without damaging its contents. However, pedagogical datasets can be distinctive in fulfilling the objectives.

Pedagogical datasets support assignments and activities, which in turn support learning objectives. The audience for any pedagogical dataset will be learners, ideally at a known level, which means that the designer needs to design for their prior and desired skills. These skills should be used to determine the appropriate complexity of the *navigability* and *understandability* of the datasets: Is the dataset formatted in such a way that students will know how to navigate it? Do students understand the data without needing to rely heavily on the documentation? In particular, navigation of the dataset should require the desired skills, but should otherwise not involve skills not identified as prior skills. For example, consider a lesson to teach students how to use homogeneous collections (e.g., lists), where learners had previously only learned about a single primitive type (e.g., numbers). A dataset that supports this lesson

Figure 1: A high-level process view of preparing pedagogical datasets

could be composed of that type, but should not involve heterogeneous collections or other data types, lest they provide distractions.

A distinctive feature of pedagogical datasets is the level of control that the developer has over the narrative of the dataset. A dataset's narrative can be understood as the potential knowledge that can be learned from its contents about the dataset's subject; for example, what can be learned about crime from the contents of a crime statistics dataset. This narrative is framed by the context surrounding the provenance of the dataset—how it was collected, who collected it, their purpose in collecting it, and many other important details. A dataset can have multiple narratives, and is partially dependent on the interpretation of its reader.

The ultimate goal for a conventional dataset is to be transformed into more consumable knowledge. Typically, this means creating a visualization, a statistical summarization, or some other simplification of the data. However, in a pedagogical dataset, the developer is intentionally not the agent who finally transforms the data. Although the developer should be informed by the intended approaches that the learner will use, in order to maximize its *manipulability*, they also need to focus on the narratives in the dataset.

3.3 Dataset Process

Figure 1 gives a high-level overview of a general process for preparing pedagogical datasets. This is similar to processes for conventional dataset preparation (such as [12]), in particular the middle phases related to collecting and editing the data. However, the process for pedagogical datasets has distinctive elements. First, this process emphasizes the role of the learner and the learning objectives in the design of the dataset; the first phase is to analyze these elements to establish the target audience and complexity of the dataset. Second, the penultimate phase of data dissemination replaces typical final phases that discuss presentation and visualization of data. Finally, the last phase ("Evaluate") is used to show how the generated dataset is meant to be used, observed, and modified to better suit the needs of the learners, as opposed to a frozen artifact.

4 PRAGMATICS

The *Pragmatics* text is divided into six major chapters, each of which is subdivided into 5-9 sections as shown in Table 1. These chapters correspond only approximately to the process outlined before, since some of the advice is not exclusive to a particular phase. In this section, we summarize each chapter, and highlight elements

that are particularly unique to pedagogical datasets compared to conventional datasets.

4.1 General Advice

The *Pragmatics* begins with a chapter describing advice that either begins the preparation process or cuts across the entire process. Two pieces of advice in this section encourages the reader to consider the learner (as described earlier) and the context. Perhaps more so than in regular dataset preparation, the developer may not be a subject-matter expert with regards to the narratives embedded in the data. If authenticity is an important design criteria, then learning more about the nature of the data becomes crucial.

"Find and reach out to subject-matter experts in order to resolve questions. Ensure that the dataset you end up preparing aligns with the kind of data that professionals might expect to see, at least in nature if not shape or format."

4.2 Collecting Data

The second chapter gives tips and strategies for finding datasets. Typically, data scientists already have a dataset in mind for a particular project. However, instructional developers may be interested in a wide range of pedagogical datasets; rather than finding one for a specific context, they may be searching for one with a particular set of properties. For example, to satisfy an assignment with learning objectives associated with string processing tools, a text-heavy dataset could be desired. Unfortunately, search tools can still be limited in searching beyond context keywords. This chapter proposes a number of methods for finding data sources. Besides advice on using search engines and reviewing aggregate lists, there are also recommendations on customizing existing data sources. This can be done by scraping web sites, synthesizing existing data sources, or even through mining a real-time data source:

"Mining a real-time data source is an easy way to translate a high-velocity, low-volume dataset into a low-velocity, high-volume dataset. The idea is to take a data source that updates regularly and to retrieve data from it on a consistent schedule, aggregating the data over time."

4.3 Restructuring Data

When a learner encounters a pedagogical dataset, their first observations will be colored by the structure of the dataset. This

1. General Advice	1.1. Have a plan 1.2. Build for your audience 1.3. Iterate 1.4. Standardize your process 1.5. Keep a clean workspace 1.6. Manage dataset health 1.7. Beware breaking convention 1.8. Work in phases 1.9. Understand the context
2. Collecting Data	2.1. Hunting sources 2.2. Working with file formats 2.3. Scraping web data 2.4. Mining real-time data 2.5. Legality of your data 2.6. Synthesizing datasets
3. Restructuring Data	3.1. Choose your target structure 3.2. Layering columnar data 3.3. Converting XML to JSON 3.4. Working with indexes 3.5. Collapsing fields 3.6. Stacking data 3.7. Redundant total field
4. Manipulating Data	4.1. Standardize fields 4.2. Names are important 4.3. Working with bad data 4.4. Cleaning up by hand 4.5. Reshaping data 4.6. Extending a dataset
5. Working with Data Types	5.1. Numbers 5.2. Textual 5.3. Dates and times 5.4. Measurements 5.5. Locations 5.6. URLs 5.7. Enumerated data
6. Knowing Data	6.1. No one reads documentation 6.2. Learning the structure 6.3. Learning the distribution 6.4. Disseminating materials 6.5. Monitor usage

Table 1: Outline of the *Pragmatics*

chapter covers several strategies for organizing data to minimize students' cognitive load and make it easier for them to comprehend the structure of the data, and pitfalls to avoid when making design trade-offs. Although much of the advice would benefit the creation of any dataset, the advice becomes more urgent for pedagogical dataset design because of the limited abilities of the learners.

“When working with spreadsheets that are particularly wide (i.e., have a lot of columns), chunking columns can be an excellent way to help users navigate the data’s structure. These chunks of columns can represent sub-abstractions that group related fields. For example, an address can be grouped with latitude and longitude fields under a “location” field.”

4.4 Manipulating Data

This section contains a number of recommendations to encourage uniformity, simplicity, and readability in the data. Unless the learning objectives specifically require non-uniform data, every diverging element should be cleaned, lest they distract the learner from the intended lesson.

“Be as consistent as possible across fields names, types, and their values. Ensure that every field name has the same style of capitalization, spelling, use of symbols, and punctuation, and make sure there are no errors in any of the above. Make sure that every instance of a field has the same type and uses the same kinds of units.”

This section also covers techniques to change the shape of the data, ideally without changing its nature. These techniques can include imputation (smoothing out missing values), mathematical transformations, and even cleaning bad data.

4.5 Working with Data Types

This chapter is devoted to the pedagogical idiosyncrasies of working with common data types and formats. There can be surprising nuances to different kinds of data. For instance, representing dates and times so that beginners can use them is not straightforward: is a string representation sufficient? What about as a Unix epoch timestamp? Although the data can themselves have a natural representation, this isn't necessarily ideal either:

“When designing for a specific audience, it can be helpful to use a measurement format they are already comfortable with. Inversely, putting data into an awkward format can provide an opportunity to have students practice manipulating data. In other situations, the context may demand or benefit from a certain format (e.g., using metric for scientific data, or sortable dates).”

4.6 Knowing Data

The final chapter includes a number of important points for considering students' encounters with pedagogical datasets. For example, how will students retrieve and store the necessary data? How will learners be expected to learn about the structure of the data? What documentation needs to be available to supplement the dataset? This last question is particularly problematic, since many novice learners seem to be predisposed to not read the documentation, no matter how conveniently available.

“Make sure that the documentation is always readily available, easy to access, and able to answer whatever questions they have. The fewer the barriers, and the better the documentation is at answering questions, the more students will learn to use it and build confidence in documentation as a concept. One method for improving documentation is to log what questions students typically ask about a dataset, and using that data to improve the existing documentation.”

5 DATASETS IN THE CLASSROOM

As part of our exploration of the efficacy of using pedagogical datasets in the classroom, we have replicated an experiment by Bart et al. [4]. The previous study was relatively small, with only 50 participants, but found a number of important results: particularly, that data science could be a source of motivation to participate in a course, but that it may not be an effective way to encourage students to continue learning computing.

5.1 Methodology

The replication study was conducted over two semesters of an “Introduction to Computational Thinking” course for non-computing majors, representing a population similar to that of the original study. A survey was administered at the end of each semester to all students. There were 191 students total, and 176 gave consent for their responses to be used for research purposes (a 92.1% response rate). Student demographic data was only partially retained, but the gender ratio was roughly 60% female. The students skewed towards freshmen and sophomores (roughly 30% each) and the remainder roughly split between juniors and seniors (roughly 20% each).

The survey instrument used was almost identical to the previous study’s instrument. The instrument for our current work included a number of questions not relevant to this paper, but were included in order to match the previous research protocol. The relevant questions were divided into three sections. The first section (novel to our current survey) described seven potential introductory computing contexts and asked students how much they would prefer or avoid:

- “Working with data sets related to your major”
- “Working with pictures, sounds, movies”
- “Making games and animations”
- “Making websites”
- “Making scientific models of real-world phenomenon”
- “Controlling robots or drones”
- “Making phone apps”

The second section asked students how likely they were to continue to learn about computing, apply what they have learned, and to recommend the course to others; in the original study, these three subquestions were asked as a single question. The third section was a series of 25 questions organized into 5 groups. Each group corresponded to a component of the MUSIC Model of Academic Motivation, a 5-factor model that hypothesizes that students become motivated when one or more of the following are present [15]:

- (1) feel eMpowered to direct their own learning,
- (2) find the material Useful,
- (3) believe they will be Successful,
- (4) think that the material is Interesting, and
- (5) believe that the course staff Cares for them.

Each factor was connected to one of the 5 core course components: learning about abstraction, learning to write programs, learning about ethics in computation, learning to work with real-world data, and working within a cohort of students (4-person groups). The result of combining these 5 factors and 5 components were 25 statements such as: “I believe it was interesting to learn about abstraction”, and “I believe it was useful to my long-term career

goals to learn to write computer programs”. Students stated, on a 7-point Likert scale, how much they agreed or disagreed.

5.2 Results

Figure 2 shows the results of the first question. A one-way ANOVA test reveals that student preference for the Data Science context is significantly higher than any other context. Of course, there is little actual difference in student preference for some of the contexts: particularly, Media Computation was only slightly less preferred. Also notable is students’ increased neutrality on the other contexts—perhaps symptomatic of the ambiguity of asking students to consider contexts that they have not personally encountered.

Figure 2: Student Preference for Introductory Contexts

Most of our results for the second and third survey question sections, asking students to agree or disagree with the statements regarding motivation, corroborated the results found in the original study. To demonstrate this, Figure 3 shows students’ perception of the usefulness for each course component. These data show that a significantly higher number of students felt that working with data was more useful to their long-term career goals than learning to program or work with abstractions. This lends further evidence to a major, original hypothesis from the original study: students find a data science context more motivating than the core course content.

Another major result from the original study was the identification of a strong correlation between students’ intent to continue learning computing and their attitudes towards learning to program and other core course components (but, notably, not with their attitudes towards the context). Table 2 presents data that diverges somewhat from the original study. This table shows the Pearson correlation between students’ stated intent to continue learning computing, and their self-reported motivation towards each component of the course. The original high correlations are now considerably lower, although many of the pairwise statements still reach significance. Previously, the strongest correlations were found in students’ sense of usefulness of learning to program and working with abstractions. In this study, however, we find that usefulness is on par with students’ self-efficacy and interest – at least in terms of learning to program. However, we reiterate the interpretation of this result from the original study: this does not demonstrate

Figure 3: Student Perceptions of Usefulness of Course Components

that Data Science is not a motivating context for learners in a course, but instead that, when the goal is to convince students to learn more about computing, a data science context might be less motivating than a fostering an inherent desire to learn to program.

Component	M	U	S	I	C
Abstraction	0.198*	0.230*	0.174**	0.131	0.274*
Cohort	0.003	0.114	-0.032	0.024	0.115
Data	-0.005	0.057	-0.014	0.110	0.147
Ethics	0.052	0.146	0.029	0.050	0.207*
Programs	0.281*	0.397*	0.363*	0.434*	0.252*

Table 2: Correlation between Students' Intent to Continue Learning Computing Vs. Components of the Course with Respect to Motivational Components (at End of Semester)

5.3 Threats to Validity

Although we have improved on the original study, there are still some threats to validity present in our work. The first threat is confirmation bias. The students surveyed had already completed an introductory computing course with Data Science; their success might suggest a predisposition to the context. It is possible that surveying a more general population might reveal that a different course context is more widely preferred. A second threat is that using parametric tests such as the Pearson Correlation on non-continuous data, such as that collected from a Likert scale, carries risk. We believe that our large survey population and the granularity of the survey instrument make the statistical tests suitable.

6 CONCLUSION

In this paper, we have argued for the continued importance of Data Science as an introductory context for novice learners. To support instruction using this context, we have authored, and presented here, a guide for educational developers to create new pedagogical datasets. This guide covers strategies, tips, pitfalls, and examples from the authors' own experiences. In addition to presenting the guide, this paper also presents the results of a much larger replication study to evaluate the efficacy of Data Science as a motivating

context. With four times the number of participants, this replication study largely confirms the original study, albeit with nuances.

7 ACKNOWLEDGEMENTS

This work is supported in part by National Science Foundation grants DUE 1624320, DUE 1444094, and DGE 0822220.

REFERENCES

- [1] Paul Anderson, James Bowring, Renée McCauley, George Pothering, and Christopher Starr. An Undergraduate Degree in Data Science: Curriculum and a Decade of Implementation Experience. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education*.
- [2] Ruth E. Anderson, Michael D. Ernst, Robert Ordóñez, Paul Pham, and Steven A. Wolfman. Introductory Programming Meets the Real World: Using Real Problems and Data in CS1. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14)*. 465–466.
- [3] Craig Anslow, John Brosz, Frank Maurer, and Mike Boyes. 2016. Datathons: An Experience Report of Data Hackathons for Data Science Education. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*. ACM, 615–620.
- [4] A. C. Bart, R. Whitcomb, E. Tilevich, C. A. Shaffer, and D. Kafura. 2017. Computing with CORGIS: Diverse, Real-world Datasets for Introductory Computing. In *Proceedings of the 48th ACM Technical Symposium on Computer Science Education*.
- [5] Penelope Bidgood. 2006. Creating statistical resources from real datasets—the STARS project. In *Seventh International Conference on Teaching Statistics (ICOTS7)*, Salvador, Bahia, Brazil.
- [6] Longbing Cao. 2017. Data science: a comprehensive overview. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 43.
- [7] Jeong-In Choi and Michael Hannafin. 1995. Situated cognition and learning environments: Roles, structures, and implications for design. *Educational Technology Research and Development* 43 (1995), 53–69.
- [8] Thomas H Davenport and DJ Patil. 2012. Data Scientist: The Sexiest Job of the 21st Century—A new breed of professional holds the key to capitalizing on big data opportunities. But these specialists aren't easy to find. *Harvard Business Review* (2012), 70.
- [9] Peter DePasquale. Exploiting On-line Data Sources As the Basis of Programming Projects. In *Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education (SIGCSE '06)*. 283–287.
- [10] Roland DePratti, Garrett M Dancik, Fred Lucci, and Russell D Sampson. 2017. Development of an introductory big data programming and concepts course. *Journal of Computing Sciences in Colleges* 32, 6 (2017), 175–182.
- [11] Ben Fry. 2007. *Visualizing data: Exploring and explaining data with the processing environment*. " O'Reilly Media, Inc".
- [12] P Guo. 2013. Data science workflow: Overview and challenges. *blog@CACM, Communications of the ACM* (2013).
- [13] Mark Guzdial and Allison Elliott Tew. 2006. Imaginering inauthentic legitimate peripheral participation: an instructional design approach for motivating computing education. In *Proceedings of the second international workshop on Computing education research*. 51–58.
- [14] Olaf A. Hall-Holt and Kevin R. Sanft. Statistics-infused Introduction to Computer Science (SIGCSE '15). 138–143.
- [15] B. D. Jones. 2009. Motivating students to engage in learning: The MUSIC model of academic motivation. *International Journal of Teaching and Learning in Higher Education* 21, 2 (2009), 272–285.
- [16] M. Lichman. 2013. UCI Machine Learning Repository. (2013). <http://archive.ics.uci.edu/ml>
- [17] Aparna Mahadev and Karl R Wurst. 2015. Developing concentrations in big data analytics and software development at a small liberal arts university. *Journal of Computing Sciences in Colleges* 30, 3 (2015), 92–98.
- [18] Josh Radinsky, Ben Loh, Jennifer Mundt, Sue Marshall, Louis M Gomez, Brian J Reiser, and Daniel C Edelson. 1999. Problematising Complex Datasets for Students: Design Principles for Inquiry Curriculum. (1999).
- [19] Kalpathi Subramanian, Jamie Payton, David Burlinson, and Mihai Mehedin. 2016. Bringing Real-World Data And Visualizations Into Data Structures Courses Using BRIDGES. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*. 590–590.
- [20] David G. Sullivan. A Data-centric Introduction to Computer Science for Non-majors. In *Proceeding of the 44th ACM Technical Symposium on Computer Science Education (SIGCSE '13)*. 71–76.
- [21] Katrien Verbert, Nikos Manouselis, Hendrik Drachler, and Erik Duval. 2012. Dataset-driven research to support learning and knowledge analytics. *Educational Technology & Society* 15, 3 (2012), 133–148.