

Introducing Data Science to School Kids

Shashank Srikant
Aspiring Minds
shashank.srikant@aspiringminds.com

Varun Aggarwal
Aspiring Minds
varun@aspiringminds.com

ABSTRACT

Data-driven decision making is fast becoming a necessary skill in jobs across the board. The industry today uses analytics and machine learning to get useful insights from a wealth of digital information in order to make decisions. With data science becoming an important skill needed in varying degrees of complexity by the workforce of the near future, we felt the need to expose school-goers to its power through a hands-on exercise. We organized a half-day long data science tutorial for kids in grades 5 through 9 (10-15 years old). Our aim was to expose them to the full cycle of a typical supervised learning approach - data collection, data entry, data visualization, feature engineering, model building, model testing and data permissions. We discuss herein the design choices made while developing the dataset, the method and the pedagogy for the tutorial. These choices aimed to maximize student engagement while ensuring minimal pre-requisite knowledge. This was a challenging task given that we limited the pre-requisites for the kids to the knowledge of counting, addition, percentages, comparisons and a basic exposure to operating computers. By designing an exercise with the stated principles, we were able to provide to kids an exciting, hands-on introduction to data science, as confirmed by their experiences. To the best of the authors' knowledge, the tutorial was the first of its kind. Considering the positive reception of such a tutorial, we hope that educators across the world are encouraged to introduce data science in their respective curricula for high-schoolers and are able to use the principles laid out in this work to build full-fledged courses.

1 Introduction

Data-driven decision making has become ubiquitous in businesses. One of the reasons for this is that businesses have become 'digital' - customer acquisition, product/service development and delivery happens through the internet. A majority of our communication and social engagement also happens on the web. Unlike before, lots of data is now

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 '17, March 08 - 11, 2017, Seattle, WA, USA

Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4698-6/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/3017680.3017717>

created, automatically recorded and is amenable to experiments. This allows for data driven techniques to be more naturally applicable and efficient [9, 10].

In the last decade, machine learning performing feature engineering to build predictive models has become a major addition to the methods used in the industry to analyze information [2]. Earlier methods of data analysis comprised studying group differences by ANOVA or doing a linear regression. Feature engineering has become naturally important to derive insights from unstructured data like text, voice, videos, etc. Modern analysis may take the form of both - supervised and unsupervised learning, the former being tractable for use by industries and often more useful for quick solutions. This new field of data analysis is loosely called data science.

In the near future, data science is likely to have a varying degree of influence on almost every available job. While some jobs will require the ability to record data in amenable formats and infer using predictive models, others will involve creating models and tools for analysis and insights [4]. For instance, a personal assistant will use an online tool to understand how his/her manager has spent her time in previous weeks, offer her a visualization and could also predict meeting trends ahead in time and plan her schedule better. Such a reporting task will probably require recording data properly, cleaning and bucketing it, being able to visualize it and deriving features for deeper insights. Likewise, a sales manager will be interested in answering a number of questions - what sales pitches work better, when to call a customer and what habits are common to productive sales executives. To learn all of this, the sales manager should trust the fact that data-based methods can answer such questions and be able to collaborate with a data scientist. She needs to record data in a form convenient for analysis and know how to interpret results from such an analysis in the context of sales and marketing process knowledge, which only s/he has. With such a growing demand for data science in various professions, we think that it will become a general employability skill. It needs to be developed early on, much like basic computer skills.

We decided to make a first attempt at teaching data science to school kids from the 5th to 9th grades (10-15 years old). Our goal was to give kids a hands-on experience of the full cycle of a supervised learning task - data collection, data entry, feature engineering, data visualization, model building, model testing and handling data permissions. We didn't want them to spend time staring at video tutorials or looking at an instructor working her way through data. We believed

a student would learn best when she would solve a problem herself using data science; she would be able to think of other problems she could similarly solve and question the ifs and whys of the different steps involved in the process. Such an approach draws from different pedagogic theories like experiential learning [8], problem-based learning [12], cooperative learning [7], cognitive apprenticeship [3] and blended learning [5], the combination of which has shown to improve student learning rates and decrease teacher stress [1].

We have conducted such a half-a-day long data science tutorial in four cities so far - New Delhi, Bangalore and Pune in India and Urbana Champaign in Illinois, USA in which a total of 71 students have attended. In our experience, we felt that the students were able to perform all the steps in the flow and understand the various insights the exercise had to offer. We have made available the design of the experiment, dataset, mentor and student experiences on <http://datasciencekids.org> for the community to use and develop further.

With the aim to keep the material accessible to students given their training in math and computers and their cognitive development, we made a considerable number of decisions when designing this tutorial. This included choosing an apt dataset, its construction, choosing a modeling technique and the software platform used to implement this exercise. We feel that one of the reasons for the successful execution of this tutorial was our ability to control the complexity of the exercise while giving the kids a rich experience of doing something new and exciting.

Specifically, the paper makes the following contributions

- We lay out a set of design principles to choose a problem statement and create a dataset to provide a hands-on experience to students in the 10-15 years' age group of the whole flow of a supervised machine learning task.
- We also propose a simple interpretable supervised learning model, analogous to a game, which the kids can easily understand, build themselves and see in action.
- We demonstrate the design principles by actually constructing an exercise which requires the kids to know only counting, addition, percentages and comparisons. Having ensured restricted pre-requisites for the exercise, the kids get an exciting, rich experience of being able to visualize effects, identify features and predict an unknown.
- We believe that this is a first successful demonstration of teaching young kids data science and should be an encouragement to the community to investigate and build full-fledged courses containing modules which follow the design considerations discussed here.

This paper is organized as follows - §2 discusses the design decisions which were considered to ensure that the exercise was accessible to kids. §3 details the predictor which the kids had to design. §4 describes the various steps involved in the entire exercise. §5 discusses observations from the tutorial and describes student experiences. Concluding remarks and future work is discussed in §6.

2 Design Considerations

The following constraints were considered while designing the tutorial. The aim was to provide a gentle introduction to the core concepts behind data science and supervised machine learning while ensuring that the material was readily

comprehensible, intuitive and the pre-requisites for participation were minimal. In addition to discussing the design considerations for a comprehensive introductory tutorial on data science for kids, this section also sheds light on design choices which can be introduced as modules in a longer, full-fledged data science course. In the remainder of this paper, we refer to the intended audience for such a data science tutorial, students in the 5th to 9th grades (10-15 years old), as our *target group (TG)*. We also interchangeably refer to the dependent variables as output variables and independent variables as input variables respectively.

2.1 Problem Statement

Full Data Cycle The exercise should provide the *TG* with a hands-on exposure to the full cycle of a typical supervised learning approach - data collection, data entry, data visualization, feature engineering, model building, model testing and data permissions. Introducing unsupervised learning would be harder to relate to and understand and is hence avoided.

Relatable Dataset The *TG* should be able to relate to and be interested in the data set used. They should also find what the model may predict to be exciting. For instance, commodity prices and stock market information would be rich in data but wouldn't be appropriate. On the other hand, predicting the weather based on the clothes people wore could seem obvious, thereby underplaying the role of data science. In summary, use a dataset which is relatable, interesting and can lead to some exciting prediction.

Pre-built Datasets There are several datasets available on the web, including those specifically built for educational purposes. Avoid using these curated datasets. The *TG* should be exposed to the process of data collection and entry, which is an important component requiring time and attention when solving real world problems. In the larger scheme of applying data science to their surroundings, such an exercise would be an essential first step. Moreover, being involved in collecting and entering data would also give them greater ownership around the exercise and enhance the activity element.

Binary Variables To get a sense of the input variables, the *TG* should be able to visualize them and infer whether they discriminate the output. If the inputs are continuous-valued, they would need to plot probability distributions to see, say, group differences, which would be hard. Plotting a scatter between the dependent and an independent variable to intuit a pattern, an increasing or decreasing relationship, in a typically noisy graph would also be hard. Discretizing the continuous inputs, a process which would require understanding thresholds and their effect, would make the overall exercise complex. Hence, to keep the exercise simple, use discrete binary values for the dependent and independent variables. This reduces to a two-class classification problem with binary features. Here, visualizations could be made by merely counting, say, how many times a feature was represented (was valued 1) in the two classes of the output. It would reduce to a simple bar graph where one compared the heights of the different bars. Additionally, have not more than 3-4 independent variables in the dataset. The *TG* should be able to clearly understand the relationship between the dependent and independent variables.

Problem statement	Relatable dataset	The dataset used in the exercise must be relatable to a high-schooler. Some bad examples - oil prices, stocks etc.
	Data collection	Participants need to collect and enter data. This would give them a sense of how data is retrieved and collected in real applications.
	Prediction	The final prediction ought to have an aha-moment. It should not be something obvious, such as predicting the weather by looking at ones' clothes.
Dataset	Data-type	It is best if the dataset has discrete-valued variables. It makes analyzing the data much easier.
	Independent variables	The dataset should contain at most 3-4 independent variables.
	Balanced dataset	To make feature engineering intuitive, ensure that each feature is represented equally, creating a fully balanced dataset.
	Unseen data	Let the unseen dataset again be balanced, with each feature being represented equally in it.
Model	Model building	Participants should be able to design a simple, interpretable model from the dataset.
	Arithmetic involved	The math involved to design such a model should be constrained to operations in counting, addition, percentages and simple if-then logic.
	Model design	The design of the model needs to be amenable to high schoolers. It needs to have intuitive properties like a point-based system which adds up when the most discriminating feature is present in a sample.
Platform	Easy tech	A spreadsheet software should be the most the participants use. A full-fledged programming language like R or Python would be high on pre-requisites.
	Manual override	The exercises should be designed such that it does not rely on formulas alone. Filtering, counting, pivoting information in any spreadsheet software should be demonstrable by manually performing these actions as well.

Table 1: Design principles for a hands-on exercise in data science for kids

Balanced Dataset Ensure that the two categories in each independent variable are represented equally in the training set and also within each category of every other independent variable in the dataset. This would mean having 2^n unique data-points for n independent variables. This has a couple of advantages. First, when the effect of each independent variable on the output is visualized, it would likely represent the actual trend and not a result of an over/under representation of categories of another variable.¹ This increases the chance of succeeding in being able to demonstrate an intuitively correct result. Second, and more importantly, the *TG* has to just see how the ratio of the categories of an independent variable moved from the expected 1:1 (See §4 for more details). This reduces the overall complexity without compromising the accuracy of the experiment design. Additionally, the training and validation set sizes should be at least in the ratio of 2:1. This would mean having at least $(2^n \times 3)$ data points in total.

2.2 Predictive Model

The *TG* should be able to design a predictive model themselves and also understand the intuition behind its working. Learning a linear logistic model manually was ruled out and so were any methods involving complicated means and standard deviation calculations. We were inspired by the work of Hunt et al. [6] which suggests that concept creation based on existing knowledge in human subjects can be modeled as decision trees. The realization of this concept can be thought about as an ensemble of single-node decision trees which vote

¹This is not entirely true since we balance only known variables and categories. Unknown variables or categories with varying distributions can also create an imbalance

together. A score of +2 is given if the feature exists and -2 if it doesn't. This is done for each feature and the final score is the sum of individual scores. The final classification is done on whether the sum is greater than 0 or not. With a plus/minus point system, we reduce the complexity of taking out averages and doing these calculations for each class and then comparing. The *TG* has to only count, add and compare. This model is fairly easy to understand - games have a point-based system where an accumulated score decides a win or loss. Here, points given to each class decide whether it finally wins.

2.3 Technology

Any spreadsheet software (like Microsoft's Excel) could be used as the technology platform for the exercise. The *TG* would have to familiarize themselves to be able to enter data into the spreadsheet, count manually, draw graphs and be able to write an IF-condition formula and copy it. The *TG* could additionally use filters to make the counting easy. Those not comfortable with using spreadsheet commands can be guided to perform such tasks manually. In our experience, kids picked up the software fairly quickly and were able to learn by example. We do think that building a better tool specifically for data science, something akin to Scratch for programming [11], will be very useful and a promising area of research.

2.4 Risks in the Design

- One risk with using an on-the-fly dataset is to not find any interesting trends or good predictions. This can be countered in a couple of ways. One way could be to design a dataset where there is already evidence of a relationship, an extreme example being collecting data

for a known physical law. The other way could be to pilot the dataset beforehand with a few kids to see if something interesting comes out. We followed a mix of both these approaches.

- One objection could be that we have conveniently simplified our experimental setup to ensure that it is always successful. These techniques are not entirely correct and would fail for real world datasets. We are aware of this. Our aim was to show one successful application of data science to the kids through which we could instill in them that they could solve problems themselves and encourage them to explore more sophisticated techniques. We might be at the risk of being inaccurate in our process of simplifying things, but we think this is in line with *arundhati nyaya*, a useful eastern pedagogical tradition. It describes the act of teaching an approximately correct but palatable idea first before teaching a fully developed, correct and non-trivial idea.

3 Problem Statement

Considering the design choices mentioned, we decided that our *TG* should design a *friend predictor*. We arrived at this decision after eliminating a few other choices. For instance, we considered a movie preference predictor, but found that it did not work well in cases where kids had not seen a movie, creating buggy or missing data in the process. Such advanced use-cases can be considered when designing material for a full-fledged course.

The Friend Predictor Each *TG* member got a set of images containing kids’ faces along with their names and hobbies. By looking at just an image and the description provided therein, they had to decide whether they would befriend the kid shown in the image. This data would then be used to design a predictive model to predict if a new kid was friend-worthy.

We avoided showing the *TG* faces of real kids as that would have added variance which our small sample size wouldn’t have done justice to. There were only four dimensions of variance in our sample set which our *TG* could implicitly consider in deciding whether they would befriend the kid in the image or not. These were - *gender*, *hobby*, *name* and *facial expression*. Hobbies could broadly be categorized into two - indoor and outdoor activities. Names in some geographies have evolved. There’s a noticeable difference between old-sounding names from the new. We hoped the *TG* would take to these names differently. We made some of the boys and girls in the images look gloomy and some cheerful. We wanted to see if the *TG* was affected by such moods. This entire exercise was posed as a problem in supervised learning, which had four independent variables and one dependent variable - the rating provided by the *TG* participants.

4 Exercise Workflow

In this section, we describe the sequence of operations which our *TG* went through as part of the data-exercise. We list out the nuances involved in each step.

Training and Validation Sets The data set was presented to the *TG* as flash-cards, each of which had an image, a description of names and hobbies and some space in its corner

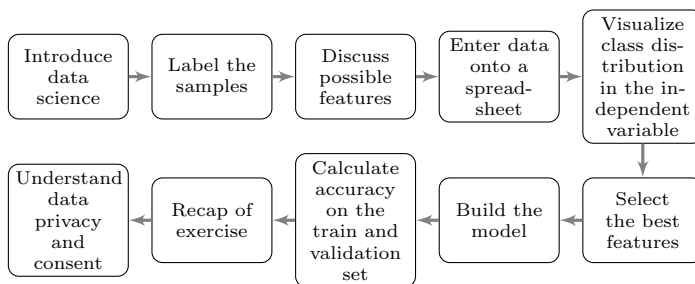


Figure 1: Workflow of our data science tutorial

where it could be rated (see Figure 2). Considering the design principles discussed in §2, there were a total of $(2^4 \times 3)$ cards to be labeled. We also had placed 8 additional cards for the *TG* to practice. Each *TG* member saw the flash-cards in the same order. The cards were ordered as follows - the first eight cards were for practice. The next 32 was the training-set, which contained two sets of the 16 unique permutations of the 4 features. The last 16 was the hold-out set kept for validating the model.

Labeling the Data Each *TG* member was shown the set of 56 images to rate. They were asked to assign their ratings on a scale of 1-5, where 1 meant they would certainly not befriend the person in the image and 5 meant they were very sure of befriending. They were given 10-12 minutes to complete the annotation exercise. Once rated, the practice images were removed, the 32 images meant for training and the last 16 meant for hold-out were placed in separate envelopes. In the remainder of this paper, we refer to a set of images belonging to a participant as *sample*. The envelope containing training images will be referred to as the *training sample* and the hold-out set as the *validation sample*.

Introducing Features Through a lively discussion, the *TG* was encouraged to think out aloud possible features which could signal whether a person would befriend another. Amidst suggestions of befriending those who played the same video games and liking the same flavor of chocolates, the *TG* consistently concluded and saw reason in the four features we had set up and were ready to investigate the effect of each of these further.

Data Entry Each *TG* member was presented with the training sample of another *TG* member to analyze. Each of them had been provided with a computer containing a template spreadsheet workbook which had some pre-filled information of the samples. The *TG* had to identify the features in each data-point of the training sample and enter them into the workbook. The order of entries presented in the workbook matched the order of the flash cards the *TG* had labeled - this made data entry convenient and confusion-free. A snapshot of the sheet presented to the *TG* is illustrated in Table 2.

Flashcard details			Features (To be filled by participants)			
Card #	Name	Hobby	Gender	Name type	Hobby type	Rating
1	Tanya	Badminton	female	new name	outdoor	4
2	Natasha	Painting

Table 2: Illustration of a spreadsheet used

Visualization Once the information was entered in the sheets, we wanted the following questions to be answered

- Was a given sample “friendly”? Did the sample show an inclination towards befriending people or not.
- What features contributed in the friend-making process?

We considered visualizing the histogram of the ratings in the sample to help answer the first question (Figure 2). The *TG* manually counted the occurrences of each of the 1-5 ratings. They were then also shown how to apply filters to count the same. In order to readily see whether a sample had endorsed more friends than non-friends, we had to have the output in a boolean form. The *TG* took to it intuitively and could reason why any rating greater than 3 (or 4) could be binned as “will befriend” and the rest as “would not befriend”. Once the output was binned into these two classes, the histogram was re-done.

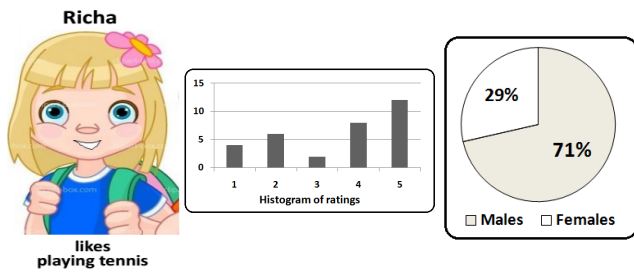


Figure 2: .a. A sample image which kids labeled .b. A histogram of the ratings .c. Pie-chart visualizing the number of males to females in those labeled as ‘friends’ in a sample

The second question was answered by visualizing each feature’s representation in the ‘friend’ and ‘non-friend’ classes of the output. This also suggested which features were to be used in modeling the data.

Feature Selection By looking at pie-charts (Figure 2), the *TG* had to decide which features helped in discriminating between the ‘friend’ and ‘non-friend’ classes. For instance, if the ratio between the two categories of an independent variable in any one class of the output (we consistently considered the ‘friend’ class since the goal was to design a friend predictor) was, say, between 50:50 to 60:40, then the variable was considered to be unable to discriminate between the two classes. If the skew in the ratio was higher, it was considered as a discriminating feature and was used in the subsequent modeling process.

As a note, we would like to highlight here the advantage of having the same number of categories for each feature in the training sample. For instance, if we considered analyzing *gender*, the instructional statement to the *TG* would be - “Let us see what percent males are in the ‘friend’ class. 80% are males in this class as compared to 50% in the whole set; it can be thus inferred that this person prefers making male friends over females.” In the absence of a balanced number of categories, this would have been - “Let us see what percent males are in the ‘friend’ class and in the ‘non friend’ class. 80% are males in the former and 50% in the latter; it can thus be inferred that this person prefers making male friends over females”. The presence of equally represented

categories eases the instructional overhead and the cognitive load of this exercise.

In order to verify how accurate the insights drawn from these visualizations and ratios were, members from the *TG* who had rated a sample were quizzed on whether these indeed were his/her tastes. This helped them get a sense of how the independent variables were indeed able to predict the dependent variable.

Model Building Once the *TG* had visualized which features were able to differentiate the dependent variable, they went ahead with the simplified model building exercise, detailed in §2. For each discriminating feature, the category which was represented higher in the ratio in the ‘friend’ class (referred to as ‘dominated feature’ in the subsequent sections) was counted in every point in the training sample. The occurrence of a dominated feature was given a score +2, its absence was given a -2. A threshold of > 0 was applied to classify the data-point in the ‘friend’ class. This was implemented by using an IF-logic formula in the spreadsheet. The formula was demonstrated by the mentors for the first 2-3 data points of the training sample. The remainder was done by the *TG*.

Validation Once the *TG* had built their models, they tested how well it generalized on the hold-out set. The validation sample corresponding to each training sample was retrieved. An element of theatricality was added at this stage. The mentors built it up like an act of magic where they challenged the *TG* if the model could really predict. This created a sense of suspense and excitement in the group. The if-else deciding classifier was applied to the unseen points and the classification accuracy was noted down and tabulated.

Consent Once the *TG* noted the accuracy of their first classifiers, the entire exercise was summarized to ensure that they absorbed the key ideas. This was followed by an explanation of the importance of data privacy and being cognizant about the right to own the data they create. The importance of anonymizing information was explained. The *TG* then signed consent forms to allow their ratings and analysis be made public in an anonymized form so that the academic community could analyze it further.

5 Learning

We report here some experiences from our tutorials. We first see whether we succeeded in making moderately accurate and generalizable models with our dataset and models. This was a necessary condition for our success. We then discuss mentors’ assessment of how the kids learnt and the difficulties they faced, followed by the kids’ experiences.

Results The exercise was successful in that our simple features were able to differentiate the output classes. Of the 71 models built, 53 models had at least one discriminating feature and 12 models had two such features. The average train accuracy across the 71 models was 78.2% and the average validation accuracy was 62.1%.

Mentors’ observations Working professionals and graduate students with an engineering background volunteered as mentors. They were trained for close to three hours to get them familiar with the exercise. Each was assigned a cohort of 2-4 students depending on the turnout at the event. Each mentor provided feedback about his/her cohort in a focused

discussion at the end of each tutorial. These notes have been aggregated and summarized in Table 3. One common observation was that kids’ understanding greatly improved on demonstrating the tasks with an example. The mentors reported that kids were largely able to follow and perform the different steps in the exercise. They reported 85% students to have completed the whole exercise (till the last step of validating their models). Around 11% students found the material hard while less than 5% did not find it interesting. The primary difficulty for students in grade 5 and 6 was in understanding percentages. Two kids exercised their right to not consent.

Location	Event date	#Kids attended	#Not interested	#Not signed consent	#Found the material hard
New Delhi	June '15	14	1	1	2
Bangalore	Sep '15	19	0	0	3
Pune	Oct '15	18	1	1	2
Urbana Champaign	April '16	20	0	NA*	1
Total		71	2	2	8

*This exercise could not be completed due to a paucity of time

Table 3: Mentors’ notes from the tutorial

Kids’ experiences They found the problem statement to be fun and exciting. They were comfortable in using a spreadsheet and picked it up rather quickly. They enjoyed making graphs and using different styles even though it wasn’t required in the exercise. At the end of the tutorial, each of them was asked to blog what s/he learnt and also fill a short questionnaire containing open-ended and Likert-scale based questions on their experience.² Around 65% of them responded to the questionnaire. Table 4 summarises their responses.³ The kids mostly strongly agreed the tutorial was interactive, was not theoretical, was clear, paced right, was not difficult to comprehend and conveyed new concepts - all of which were part of our guiding design principles. They also had a choice to fill in subjective comments, in which most senior students asked for a follow up to the topics introduced. One suggested he’d use data science to predict his pocket-money expenditure and another her school grades.

Statement	Strongly disagree	Disagree	Agree	Strongly agree
I could understand what the tutor was explaining	0%	1%	32%	67%
I understood how data science is applied to problems	0%	0%	3%	97%
The tutorial was interactive	0%	0%	21%	79%
The tutorial was boring	79%	21%	0%	0%
The tutorial was theoretical	97%	3%	0%	0%
The tutorial was difficult	79%	19%	1%	1%
The tutorial was fast for me	65%	34%	1%	0%
I learned new concepts	0%	0%	2%	98%

Table 4: Kids’ feedback on the tutorial

² A shorter questionnaire was sent out at UC due to a paucity of time.

³ Given this is a self-report, it is a good measure of engagement. A better measure of learning comes from mentor experiences.

6 Conclusion and Future Work

The paper describes a half-a-day long data science tutorial that was organized for kids from grades 5 through 9. The camp was designed to expose them to the full cycle of a typical supervised learning approach and act as a primer to a more involved treatment of the subject. During the course of this tutorial, kids experienced how they could use data science to successfully solve a relevant problem, and in doing so, also appreciate the power and the applicability of such a technique. We share a set of design principles to help choose a problem statement, create a dataset and create a modeling technique such that it maximizes participation while ensuring minimal pre-requisites from the kids. The exercise was fairly successful, as signaled by both, the mentors’ and kids’ feedback, and also led to some learning for further design. In future, we plan to use more methods to gauge student learning.

We believe that our work presents a good starting point for educators to explore this further. We strongly believe that data science needs to be inculcated in school curricula and we see this as a resource which educators and curricula designers could take inspiration from. From a research position, a Scratch equivalent for data science poses to be a promising area of research at the intersection of HCI, programming languages and machine learning. Finally, we would also like to explore how we could scale the organization of these camps so that we could maximize its outreach to school students.

7 Acknowledgments

The authors thank Prof. Una-May O’Reilly from CSAIL, MIT for very helpful discussions. Thanks to members of Aspiring Minds’ research group for their extensive help in executing these tutorials. Thanks to all volunteers (listed on <http://datasciencekids.org>) who mentored the kids. Special thanks to Aspiring Minds’ Bangalore team, Dr. Neeran Karnik from Pune and Narendra Gupta from UIUC for making it happen in their respective cities.

8 References

- [1] R. Ben-Ari, R. Krole, and D. Har-Even. Differential effects of simple frontal versus complex teaching strategy on teachers’ stress, burnout, and satisfaction. *IJSM*, 10(2):173, 2003.
- [2] L. Breiman. Statistical modeling: The two cultures. 2003.
- [3] A. Collins and et al. Cognitive apprenticeship: Making thinking visible. *American educator*, 15(3):6–11, 1991.
- [4] M. Diehl. Why marketers really need to know about data science, August 2015.
- [5] D. R. Garrison and H. Kanuka. Blended learning: Uncovering its transformative potential in higher education. *The internet and higher education*, 7(2):95–105, 2004.
- [6] E. B. Hunt, J. Marin, and P. J. Stone. Experiments in induction. 1966.
- [7] D. W. Johnson et al. *Cooperative learning in the classroom*. ERIC, 1994.
- [8] D. A. Kolb, et al. Experiential learning theory: Previous research and new directions. *Perspectives on thinking, learning, and cognitive styles*, 1:227–247, 2001.
- [9] S. Lohr. For today’s graduate, just one word: Statistics.
- [10] J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. H. Byers. Big data: The next frontier for innovation, competition, and productivity. 2011.
- [11] M. Resnick and et al. Scratch: programming for all. *Communications of the ACM*, 52(11):60–67, 2009.
- [12] J. R. Savery and T. M. Duffy. Problem based learning: An instructional model and its constructivist framework. *Educational technology*, 35(5):31–38, 1995.