Mobile Application and Data Visualization for Sensei: Sensing Educational Interaction in Montessori Classrooms

by

Ayesha Bose

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

Author . . . .
Department of Electrical Engineering and Computer Science
May 20, 2016

Certified by . . . .
Sep Kamvar
Associate Professor
Thesis Supervisor

Accepted by . . . .
Dr. Christopher Terman
Chairman, Masters of Engineering Thesis Committee

Mobile Application and Data Visualization for Sensei: Sensing Educational Interaction in Montessori Classrooms

by

Ayesha Bose

Submitted to the Department of Electrical Engineering and Computer Science
on May 20, 2016, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Understanding early childhood development can help teachers individualize their methods to facilitate growth in children. The Montessori educational approach emphasizes independence and allows children to guide their own learning. Given this style of learning, a crucial part of this kind of education is observation of the students, so teachers can assist individuals and help address specific developmental needs.

Sensei aims to assist this observation through a dynamic range-based sensor network that detects proximity. These sensors are placed on the children's shoes, on lessons, and around the classroom. I developed the mobile application and the visualization dashboard for the Sensei system.

Teachers can maintain the sensor deployment in their classroom through a mobile application. This allows teachers to start the sensors in their classroom at the beginning of each day. They can also collect data from the sensors at the end of the day and view an initial graph of the collected data, showing the time they spent with each child.

With this unique data set, we also provide detailed visualizations to teachers so they can determine who children are spending their time with, what lessons they are spending time with, and what areas of their classroom are most active. With this data, teachers can, for example, determine the right time to introduce a child to a new lesson or re-arrange their classroom to facilitate learning. These visualizations are easily accessible for teachers in a web application.

Sensei helps discover insights for teachers that would have otherwise been lost. This system can help provide a deeper understanding of early childhood development for teachers, educators, and researchers.

Thesis Supervisor: Sep Kamvar
Title: Associate Professor

Acknowledgments

I would like to thank my advisor, Sep Kamvar, for his invaluable guidance throughout this entire year. His advice about both my thesis project and my extended career has helped me immensely. His words in his books, in his presentations, and in his meetings with me have left an immeasurable mark on my life.

I would like to also thank Dwyane George and Nazmus Saquib for working with me throughout this year. This thesis has truly been a group effort and would simply not have been possible without the both of them.

The entire Social Computing group has helped us throughout the course of this project. I would like to thank Yonatan Cohen, Kim Holleman, Gal Koren, Gloria Mark, RaMell Ross, Kim Smith and Jia Zhang who have helped us design and debug Sensei.

This work would not have been possible without support from the Wildflower foundation and the teachers at the Wildflower Schools. I would like to thank Mary Rockett, Erin McKay, Mardie Hinkley, Astrid Steverlynck, and Lindy McGrail for their time and patience.

Lastly, I would like to thank my parents, Shantanu and Krishna Bose, and my younger sister, Anoushka Bose, for their support through this year and the last 23 years of my life. This is because and for all of you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Understanding early childhood development has been a topic of subject of increasing interest over recent years. Early childhood has been shown to be the most important time in development [12, 4]. This development occurs at a rate that exceeds any other stage of life. A study sponsored by the National Academy of Sciences found that early enviroments are very influential, especially nuturing relationships formed during this stage [17]. A child's ability to learn and absorb is unparalleled, but still remains relatively complex to study.

Our work aims to provide a preliminary understanding of early childhood development by studying Montessori enviroments. This chapter describes the background and related work.

1.1 Background

Sensei, Sensing Educational Interaction, is a system to help teachers, educators, and researchers develop a greater understanding of early childhood learning in the Montessori classroom. The Montessori classroom emphasizes independence and teachers act as facilitators to guide learning. As part of this process, teachers spend a part of their day taking observations of the children. Their current methods of observation are difficult and error-prone, especially in a busy classroom. Sensei helps alleviate some of the need for constant observation by introducing an unobtrusive sensor network in the classroom to measure interaction.

Sensei has been developed by a team of graduate students: Nazmus Saquib, Dwyane George, and myself. The overall description of the system in this thesis has been joint work amongst our team. Some of the text from this thesis comes from a paper we wrote for ACM UIST 2016 [45]. My specific contributions to the system are the cellphone application and visualization dashboard, described in detail in chapters 3 and 4.

1.1.1 Montessori Education

A variety of different educational philosophies have emerged in recent years [22, 23]. The Montessori educational method is over 100 years old and is used in over 5000 schools in the United States [26]. The method is unique in its multi-age classrooms, specific educational materials, and student-chosed work. These materials are tactile, multi-sensory, self-teaching tools that helps students explore concepts in the areas of mathematics, language, sensorial work, and practical life. Studies have shown that Montessori children performed better on standardized tests, engaged in more positive interaction, and show more advanced social cognition. Montessori students have also indicated having a greater sense of community, as the method allows for more peer-guided learning than traditional educational environments [11].

The Social Computing group at the MIT Media Lab has started a network of schools around the world, known as Wildflower Schools. These schools are an open-source approach to Montessori learning, blending both traditional Montessori methods with new approaches. These schools have nine design principles: an authentic Montessori environment, a neighborhood-nested design, a seamless learning community, an artist-in-residence, an attention to nature, a role in shaping the city, a spirit of generosity, an open-source design, and a lab school. As lab schools, the Wildflower schools serve as a research setting dedicated to advancing the Montessori method. Our work with Sensei focuses on one aspect of the Montessori method: observation.

1.1.2 Observation

Dr. Maria Montessori, the creator of the Montessori Method, described observation as critical to the Montessori method [14]. Her own observations are what she used to develop the method and she described observation as something that should be practiced continually. Through observation, a teacher can better understand each student's interests, learning style, and needs. A teacher can use her observations to understand which lessons to introduce and when to introduce them. There are three main types of observation:

- Individuals: Teachers can observe an individual's progress with all the lessons in a classroom. For each lesson, they can track difficulty, concentration of the student (engagement with the lesson), and level of completeness. They can also assess an individual's levels of stress or mood when working with specific materials or with other students.
- Social: Teachers can determine patterns of social behavior, like learning a new lesson together or assisting others. They can track clusters of students and study their evolution over time.
- Environmental: Teachers can track which areas of the classroom or materials are used most often, adjusting the design of the room as needed to encourage students to explore new or important concepts.

Given all these different metrics, it can be very difficult for two teachers to accurately assess a 12-15 person classroom. Outside of observation, teachers also introduce new lessons to students, help students who are having trouble, and often intervene with students who are new to the Montessori environment (Figure 1-1). One teacher we spoke to said she would spend only about 20 minutes on observation out of a typical three-hour period. She would refer back to these observations to make weekly lesson plans for each student as well as reflect on her own methods. Every teacher we spoke to expressed a desire to spend more time on observation. Teachers currently use handwritten notes to record their notes, which can also be very difficult to synthesize and identify patterns.

Figure 1-1: This is a top view of one part of a Wildflower classroom. Here, a teacher is demonstrating a new material to three seated students, while other students are busy with other activities. The independence of the children in the Montessori environment makes observation very difficult using traditional methods.

Sensei helps aid in this observation by measuring proximity in the classroom.

1.1.3 Contributions

Sensei is a distributed sensor network system designed to study social interaction and learning outcomes in an early childhood environment.

My unique contributions to Sensei include the cellphone application and the visualization dashboard.

The contributions of the Sensei system are:

- A hardware and firmware design and architecture that places sensors in a minimally invasive way in the classroom that facilitates data collection with high frequency throughout a 8-hour school day

- A cellphone application designed for teachers to manage and collect data from the sensors

- A visualization dashboard for teachers to view interaction data about their classroom

Figure 1-2 shows the Sensei system and how these contributions fit together. The sensors in the classroom transfer their data through a mother node to the cellphone application. The application then posts this data to a centralized database. The API from this database can serve data to both the web application and a statistical modeling engine.

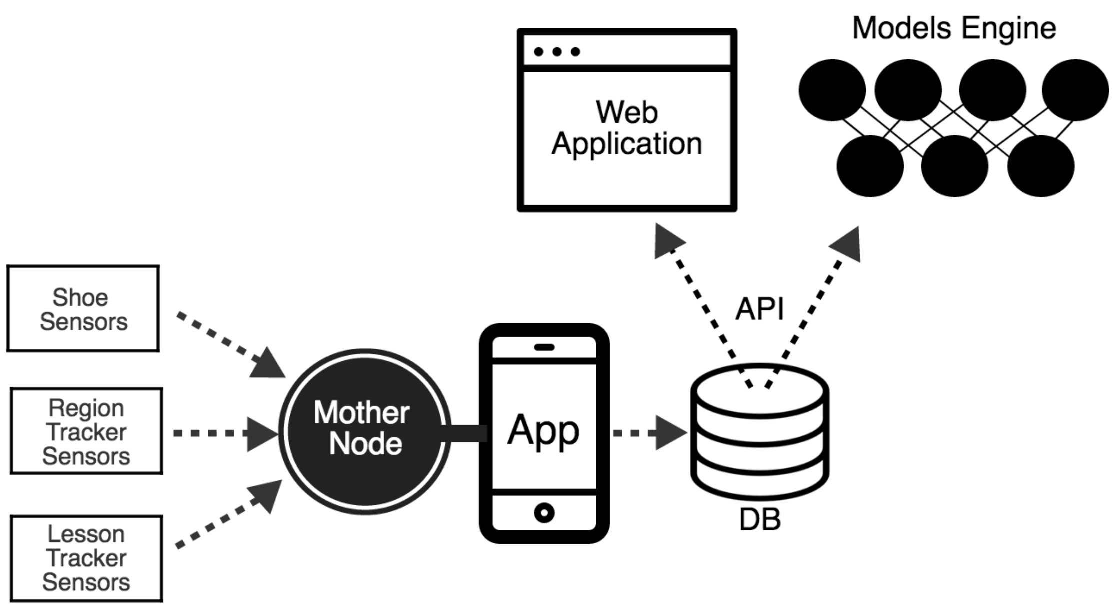

Figure 1-2: The Sensei system includes custom sensors, a cellphone application, and a web application.

We evaluated our system by conducting a system robustness study in a constrained setting, a proximity analysis using manual video annotation, and a qualtitative usability study with teachers to determine the usefulness of our system.

1.2 Related Work

To date, no other system exists using proximity sensors in a Montessori environment to study interaction.

Other observation techniques have focused on camera data or other data collection methods, in a different context. Other work in proximity data analysis have created unique visualizations, mostly intended for a technical user.

1.2.1 Camera Observation

Although we first attempted to augment observations with cameras, the Montessori classroom is too unpredictable of an environment for modern computer vision techniques. There have been several studies looking at tracking people in a calmer scene [3, 27]. Recent research has focused on using multilayer neural networks [1] or depth sensing decives like Kinect [28]. There is also some studies using thermal imaging [9] and infrared motion tracking [24].

However, the Montessori classroom is a busy environment where lighting can be uneven and furniture can occlude parts of the scene. Cameras are also more expensive and infeasible for a school to install. Our sensor network does not require any susbstantial changes in the classroom.

1.2.2 Other Data Collection Methods

With regards to data collection, some studies have shown the importance of studying social interactions in the classroom. Martinez et. al. have developed a mixed evaluation method to gather data through automatic sources, observations, focus groups, and questionnaires and draw conclusions to characterize classroom social interactions [29]. Unlike their classroom study, we cannot use either focus groups or questionnaires, given that the students in our Montessorri schools are ages 3-6.

Some work has been done to add sensors in classes to monitor different lesson interactions. Srivastava et. al. have developed an advanced network to cover an entire public school using Bluetooth modules [30]. Our system will be similar, designed for only one classroom. Much of their work is centered around communicating with the sensors through a High-speed Wireless LAN.

Sociometric badges have been used in contexts like office workspaces to understand social interaction over time [15]. These badges are much larger and worn around the neck, which is infeasible for a Montessori classroom.

1.2.3 Visualization Techniques

With regards to data visualuzation, there has been some work in understanding classroom dynamics in online classrooms. We use similar methods to evaluate our visualizations with teachers, coupled with qualitiative feedback. In constructing our visualizations, we use visualization methods that view a large amount of data over time, like systems that use multivariate numerical time series data or systems that can help visualize significant events from time data. Stein et. al. have developed a grid-like visualization to track changes in social network structure [31]. We combine many of these visualization approaches into a unique dashboard, built for teachers.

Chapter 2

Hardware and Firmware

The Montessori environment puts unique constraints on the design of the sensors, requiring a custom sensor board. These sensors capture proximity and motion data in an accurate, low-cost, and minimally invasive manner. These sensors also run custom firmware that allows them to collect data frequently in the classroom and communicate with the cellphone application. This enables data collection at a high sampling rate that preserves battery life. This chapter discusses the hardware and firmware aspects of Sensei.

2.1 Hardware

The Montessori environment requires the sensors to be:

- low-cost to remain feasible for a small school to afford
- have a long battery life so they run throughout the entire school day, without requiring a teacher to frequently recharge the sensors
- small enough to be minimally invasive in the classroom, so children are not interrupted

These constraints required a custom board. Each sensor measures the RSSI signal strength of incoming data packets to approximate proximity.

We experimented with two main types of modules that measure proximity using RSSI: RFDuino [32] and Simblee [33]. The RFDuino module has significantly more documentation and development, but we found that individual modules produced slightly different RSSI values, making data collection dependent on individual sensors. We moved to the Simblee transceiver radio, released in December 2015, and found much more consistent proximity values. With the Simblee radio, we designed three types of boards: shoe sensors, lesson/region sensors, and a mother node. We manufactured the boards using a factory overseas and assembled the boards ourselves, using a reflow oven.

2.1.1 Shoe Sensors

The shoe sensors are 2 by 2.5 cm and contain the Simblee radio, a real-time clock (RTC), a microphone, and an accelerometer (Figure 2-1). The sensors can be used with either a 3.3 V coin-cell battery or a rechargable lithium-polymer battery. The sensors are enclosed in a 3-D printed custom made casing for safety.

The sensors are placed in the velcro strap of the shoe (Figure 2-2). We have found that most students are unaware of the sensor, from our deployments in Wildflower schools. The shoes were a pre-existing part of the Montessori classroom, and we made a slight modification to the shoe to insert the sensor in the strap. These sensors can be plugged into a custom USB recharging strip at the end of the day.

2.1.2 Lesson and Region Sensors

We designed different boards for the lesson and region sensors, so they could remain in the classroom for a longer period of time without recharging. These boards could also be larger, as long as they were placed in an unobtrusive manner. These boards are very similar to the shoe sensor boards but contain an external ROM, have charging LED indicators, and have a larger lithium-polymer battery (Figure 2-3). This battery allows them to remain in the classroom for up to 20 days without needing to recharge.

The sensors were placed around the room in round wooden casings that could

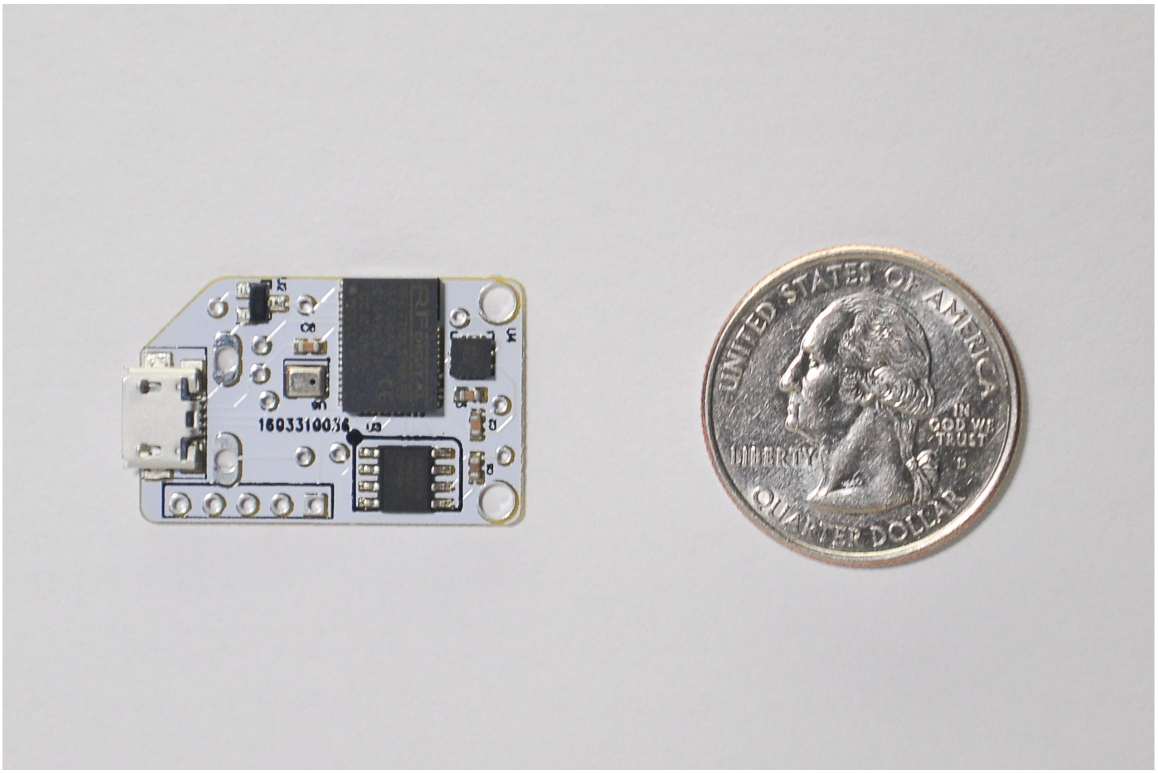

Figure 2-1: This shoe sensor contains a Simblee radio, an accelermeter, a RTC, a microphone and a USB port that can be used for both data collection and recharging.

be attached underneath shelves (like a smoke detector). The Montessori classroom is arranged by curriculum area, so the data gathered from these sensors can identify how much time a student is spending on a given subject. They could also be placed in the feet of a lesson tray or within a custom-built tray (Figure 2-4). Most Montessori materials are contained within trays, so this could help us gauge what lessons a particular child is spending time with.

With the shoe sensors and the lesson and region sensors, we can instrument a Montessori classroom to gather proximity data over a long period of time.

2.1.3 Mother Node

This sensor was identical to the the shoe sensor but also included LEDs and a separate FTDI microcontroller (Figure 2-5). This sensor is plugged into the cellphone and can communicate with all the other sensors on the network. The mother node starts all the sensors and collects data from all the sensors.

Figure 2-2: Sensors are placed in a minimally invasive way in the strap of the shoe.

2.2 Firmware

The firmware running on the sensors was written in Arduino. We wrote libraries in C++ to handle the ROM and time. We used an exisiting library for the Simblee communication.

The sensors communicate through the SimbleeCOM protocol in a mesh network. The network event scheduling scheme controls when sensors collect data, write data, and sleep. One sensor would be used as a mother node sensor to communicate with the cellphone application to start all the devices and transfer all the data.

2.2.1 Network Event Scheduling Scheme

Although the Simblee radio comes with functions to delay and get the current millisecond counter, we found that the radios often drifted. To combat this, we included an external RTC on the sensor boards. We needed this precision to ensure that sensors could sleep for an extended period of time. This sleep time is neccesary to preserve

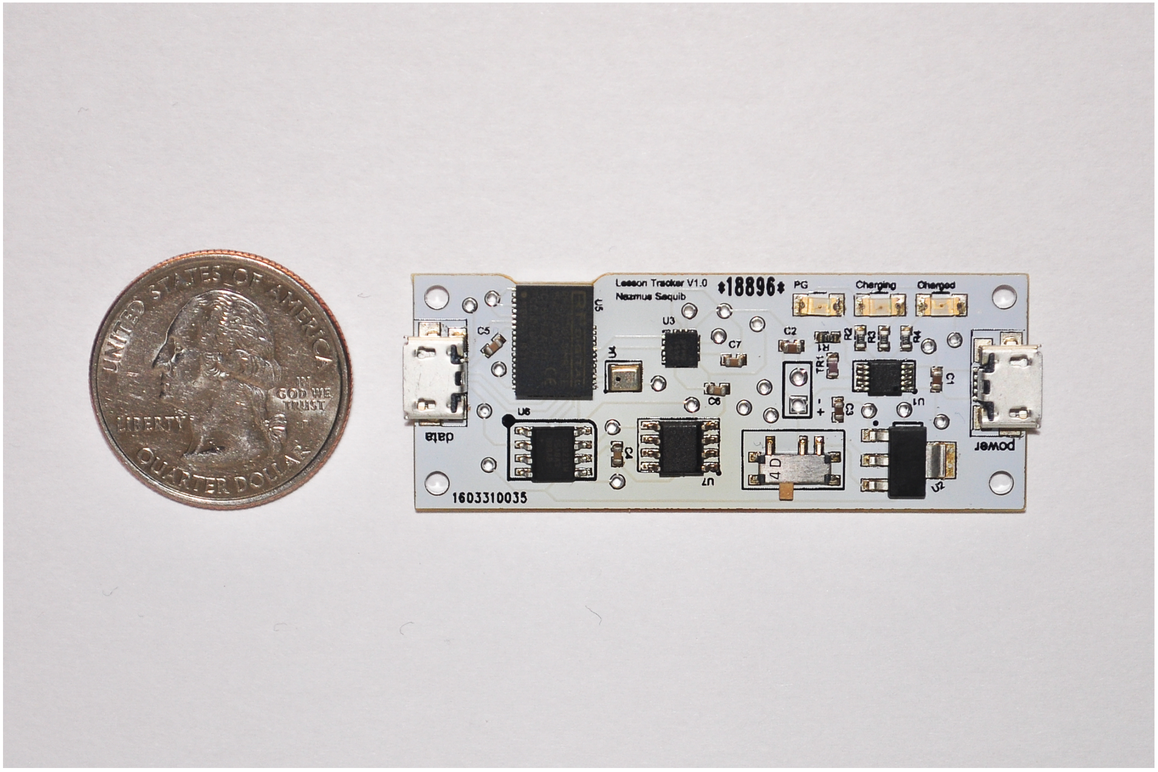

Figure 2-3: This region and lesson sensor is similar to the shoe sensor, but contains a larger battery and external ROM to run for more days in the classroom.

battery life. The Simblee radios have a data transmission time of 3 milliseconds, which allowed us to precisely synchronize the time between the devices. Figure 2-6 shows how the sensors send and receive data for 500 milliseconds. The sensors also read data from the accelerometer and microphone at that time to measure motion and ambient sound volume. These values are then written to memory and then the sensors sleep for the remainder of the cycle.

The transfer power level of the radios is set to -12. We conducted several experiments in both a lab setting and in the classroom to identify the best level to approximiate interaction. With a higher power level, the sensors would transmit packets at a very large range. We found the best power level so the sensors only transmit data when the sensors are within three to four feet of each other, which we identified as a reasonable range for interaction by observing the classrooms. The shoe sensors both send and receive data, while the lesson/region sensors only receive data (to preserve battery life).

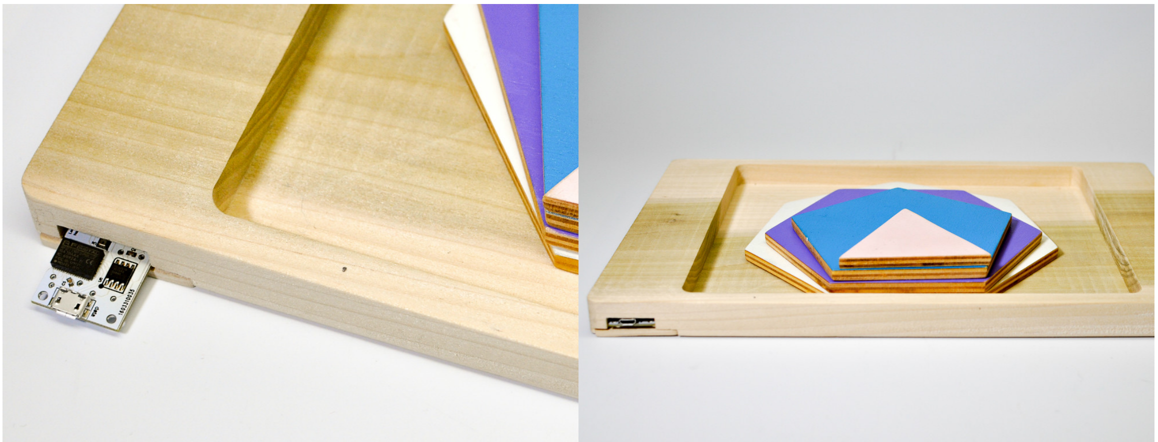

Figure 2-4: Sensors are placed in a minimally invasive way in a custom-built lesson tray.

This network event scheduling scheme ensures that the shoe sensors can run for up to two school days on a 90 mAh battery. The lesson and region tracker sensors can run for up to twenty days on a 1000 mAh battery.

2.2.2 Communication with Cellphone Application

One sensor node is not used for data collection and is not placed in the shoes, lessons, or around the room. This node is used for communication between the cellphone application and the network of sensors in the classroom. This sensor sets the initial clock time for all the sensors in the network and collects all the data from the sensors. The mother node receives the current time from the cellphone and shares that time with the other sensors. The mother node receives confirmation from the sensors as they come online and relays that to the cellphone user, as seen in the next section. During data collection, the mother node reads the data from each sensor's memory, waiting for acknowledgement packets before requesting a new packet. The data collection process takes about one minute for each sensor in the network. However, as seen in the next section, these internals are abstracted away from the users and a teacher can simply collect the data while she continues with her other classroom activities.

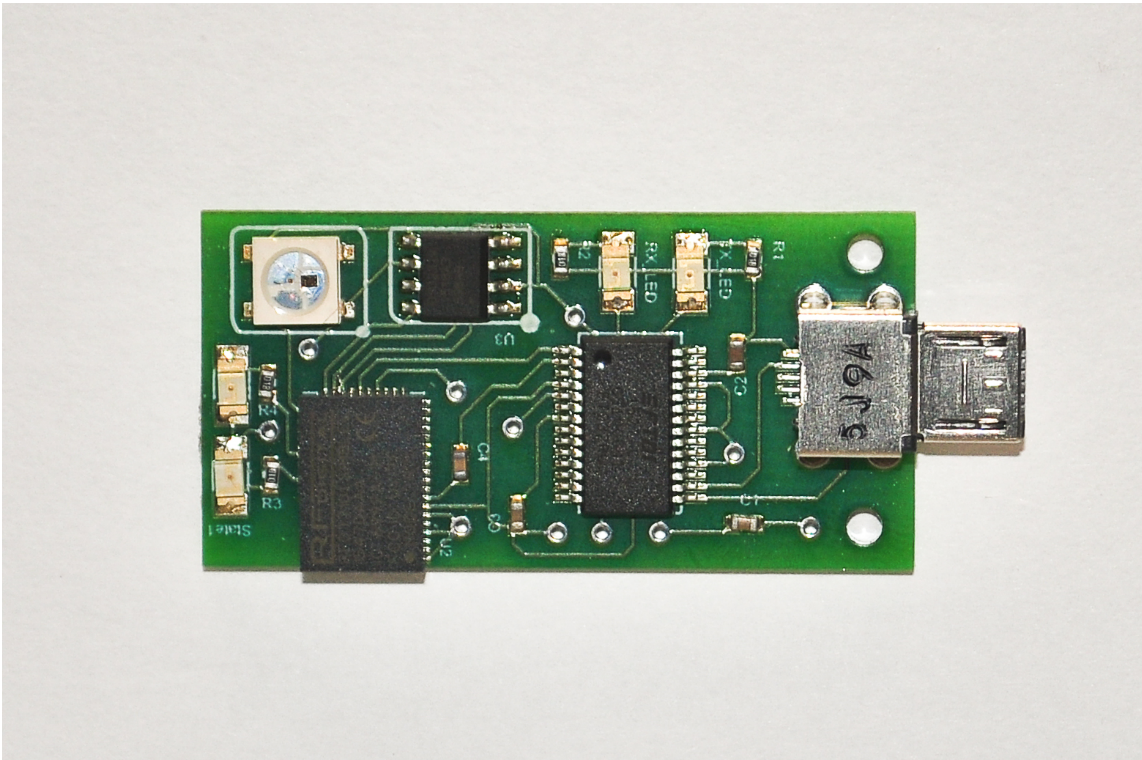

Figure 2-5: The mother node sensor can be plugged into a cellphone and can interface with all the sensors in the network.

2.3 Proximity Analysis

One limitation of the Sensei system is our approximation for interaction using proximity. Our sensors can determine distance, based on the received RSSI values, which does not always translate to interaction. To determine how often students were actually interacting when our sensors detected proximity, we conducted a manual annotation of video data of three hours of a typical classroom day. Three observers recorded times when they thought interaction was occuring between students. They also recorded times when students were just in close proximity but not interacting. We found that 84.9% of the time, students were interacting when they were in close proximity. As we gather more data over time, we can help reduce this error margin by observing higher-level patterns in the classroom.

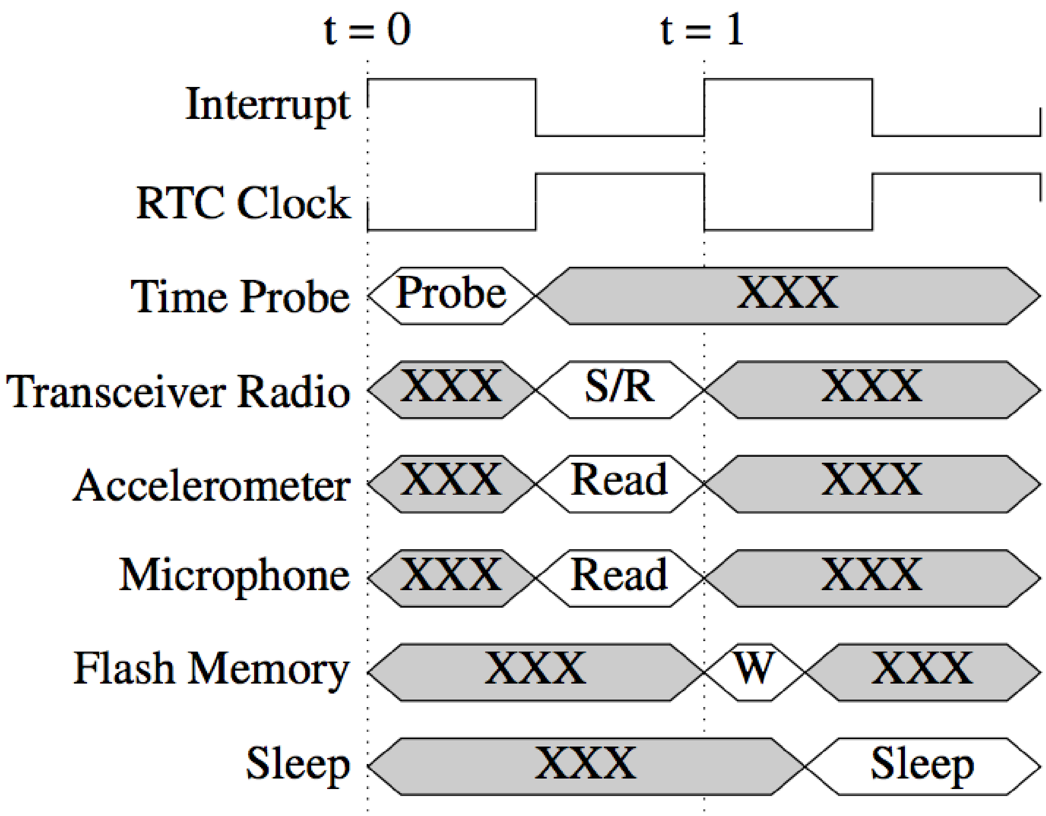

Figure 2-6: Sensors send and receive (S/R) data, read data from the accelerometer and microphone, write data (W), and then sleep.

2.4 Evaluation

To evaluate the firmware, we conducted a system robustness study in a lab setting. The system robustness study quantifies how reliably sensors receive packets. To measure how well the sensors capture proximity, we placed five shoe sensors nodes and five lesson/region tracker nodes within three to four feet (the range required for transmission). Sensors were modified to record the number of packets received from all other nodes. For this experiment, the shoe sensors ran for 12 hours and the lesson/region tracker nodes ran for 25 hours. Figure 2-7 shows the distribution of packets received from both types of sensors. The lesson/region trackers collected more packets as they do not send any data and hence had less frequent data collisions. The shoe sensors received packets 95.98% of the time and the lesson/region sensors received packets 99.75% of the time. Given the high sampling frequency, this does not affect the quality of the data.

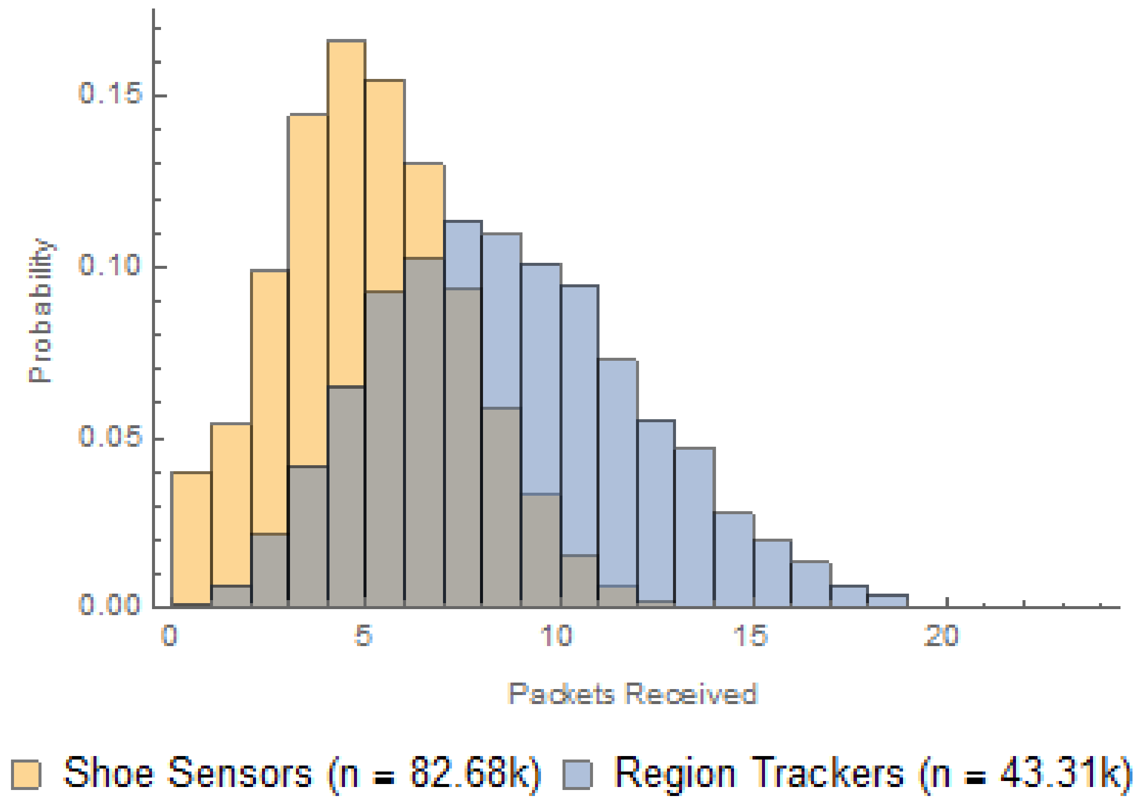

Figure 2-7: This graph shows the normalized distributions of packets received by the sensors over a 12+ hour time period.

Chapter 3

Cellphone Application

The cellphone application can be used by teachers to manage the sensors in the classroom by:

- Starting the sensors
- Collecting the data from the sensors
- Viewing an initial graph of the data most recently collected

This functionality abstracts most of the technical details but still allows a teacher to manage the network. This chapter first dicusses the communication with the mother node through the USB port with an FTDI chip. This allows a teacher to directly plug in the mother node into their phone and interface with the network. We then discuss the interface of the application. The application is designed for Android, given its increased compability with physical computing boards.

3.1 FTDI Communication

Our Android application leverages the Physicaloid library [34], an open-source Android library for communicating with physical-computing boards, including Arduino. This library does not require devices to be rooted, an important criteria as we did not want teachers to have to modify their own phones substantially. We used this library for USB-Serial communication and we focused on the FTDI protocol (though the library can support CDC-ACM and Silicon Labs CP210x). The application opens a connection to the device connected and listens for new data in the serial data buffer. The mother node is programmed to print its data to the serial monitor and the Android application parses this into a user-friendly interface, as seen in the next section. When a user closes the application, the connection to the physical device is also closed.

3.2 User Interface

The Android application has three main screens that a user can select. Each screen extends the *Fragment* class in Android and the Main activity creates a new instantiation of each *Fragment* when a user selects from the menu. Each *Fragment* calls a new instantiation of the Physicaloid object to communicate with the mother node (with the exception of the graph screen, which does not require the presence of the mother node).

We selected a color scheme and design to reflect the nature of the materials in the Montessori classroom. Although a teacher would never be using the cellphone application during the school day, we wanted to make sure that it felt natural to the environment and we used shades to reflect the wooden materials and shelves. A teacher would use the application at the beginning of the day to start the sensors and at the end of the day to collect and view data.

3.2.1 Starting the Sensors

StartSensors begins by inflating the xml file and cycling through the network size and names to populate a checkbox interface (Figure 3-1). After a teacher clicks the 'Start Sensors' button, these checkboxes will tick off as the mother node confirms they are online.

The general loop for StartSensors is:

if *startButtonClicked* **then**

$time \leftarrow getCurrentTime()$

$toSend \leftarrow startCode + time$

$writeToSerial(toSend)$

$start \leftarrow true$

end if

while *start* **do**

if $serialDataReceived.matches(pattern)$ **then**

$device \leftarrow parseDataForDeviceID$

$checkbox[device].check()$

end if

end while

The user interface abstracts away these details and a teacher can see all of the students' sensors come online. This makes it very easy for a teacher to troubleshoot a problematic sensor: they simply find that student's shoes and replace the sensors inside the strap, in the case of failure.

3.2.2 Collecting Data

CollectData is very similar to StartSensors: it begins by inflating the xml file and cycling through the network size and names to populate a checkbox interface (Figure 3-2). After a teacher clicks the 'Collect Data' button, these checkboxes will tick off as the mother node confirms that it has received all the packets from a given sensor. The application writes each sensor to an individual text file, which is later parsed into the centralized database that powers the web application.

The general loop for CollectData is:

while *notAllDataCollected* **do**

$toSend \leftarrow devicesLeft()$

$writeToSerial(toSend)$

if $serialDataReceived.matches(pattern)$ **then**

$device \leftarrow parseDataForDeviceID$

Figure 3-1: This screen of the cellphone application allows a teacher to start the sensors in her classroom.

$writeFile(device)$

$updateDevicesLeft()$

end if

if $devicesLeft() == 0$ **then**

$notAllDataCollected \leftarrow false$

end if

end while

The user interface abstracts away these details and a teacher can see all of the students' sensors that have transferred their data. Again, this makes it very easy for a teacher to troubleshoot a problematic sensor.

3.2.3 Viewing Data

Our first prototype of the Android application just allowed teachers start the sensors and collect the data. Our next version aimed to give teachers more immediate feedback, so they could use results directly from their day. At the end of the day, after collecting the data, a teacher can immediately view a graph of her time spent with each child (Figure 3-3). We chose to display this result because it was easier to view on a smaller screen and teachers repeatedly highlighted this use case as one that was the most important for them. Using this data, a teacher can immediately see which children need more attention and time at a glance. This allows teachers to quickly iterate on the information: if they spent a lot of their time with Alice on a given day, the next day they can focus more on Bob, for example.

The graph is constructed using the GraphView library for Android [35]. This library allowed for customization, while still providing functionality like scrolling and scaling.

Figure 3-2: This screen of the cellphone application allows a teacher to collect the data from the sensors in her classroom.

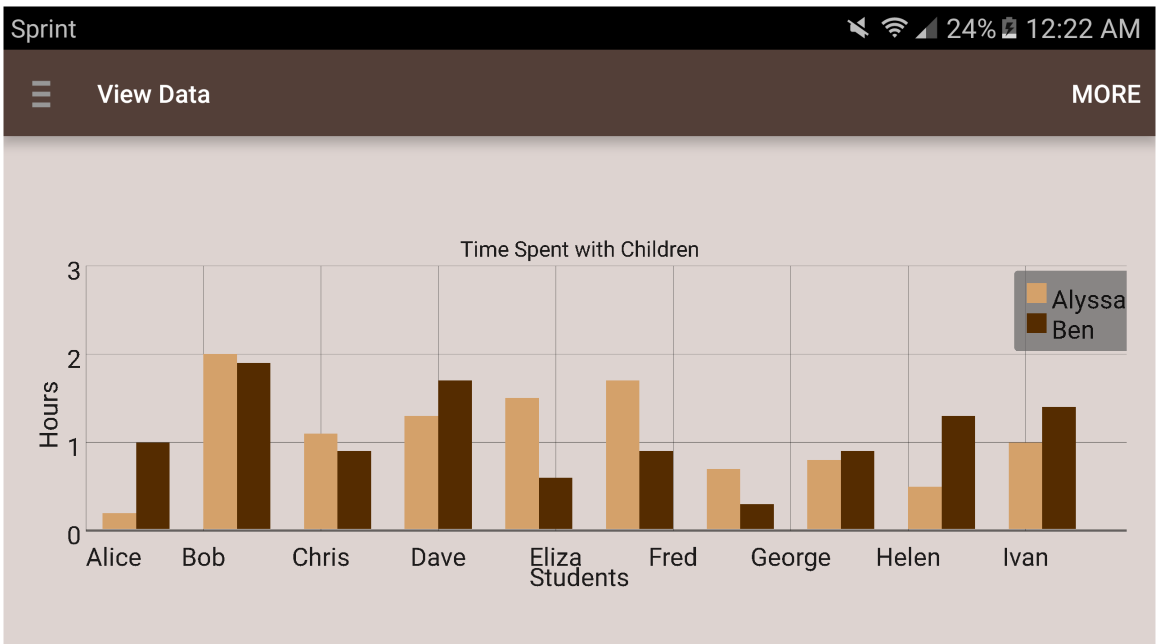

Figure 3-3: This screen of the cellphone application allows a teacher to see how much time she is spending with each child in the classroom.

Chapter 4

Dashboard and Visualizations

Teachers can access a classroom summary and the visualizations from an online web application. Their unique account credentials show them only data about their own classroom. The application is built using the Flask microframework [36] and is hosted on Heroku [37]. The centralized database uses SQLAlchemy [38]. The visualizations are constructed using visualization libraries like d3.js [39] and Google Charts [40].

4.1 Dashboard

We first constructed low-fidelity wireframes using Balsamiq [41] to do initial user tests with teachers. Figure 4-1 shows an initial mock of the homepage, which would present teachers with high-level updates about their classroom. We interviewed two teachers and presented them with these interactive wireframes. We observed their interaction with the website and also conducted an exit interview. Most of the interface was very intuitive to them, but they wanted to see more data about their own interaction. After these interviews, we added a 'teachers' tab to the main menu, to address these concerns.

Figure 4-2 shows the live homepage. The design of the page reflects wildflowerschools.org, an interface the teachers are already familiar with as it is the general information website about the Wildflower schools. This design includes a custom font file as well as unique color scheme.

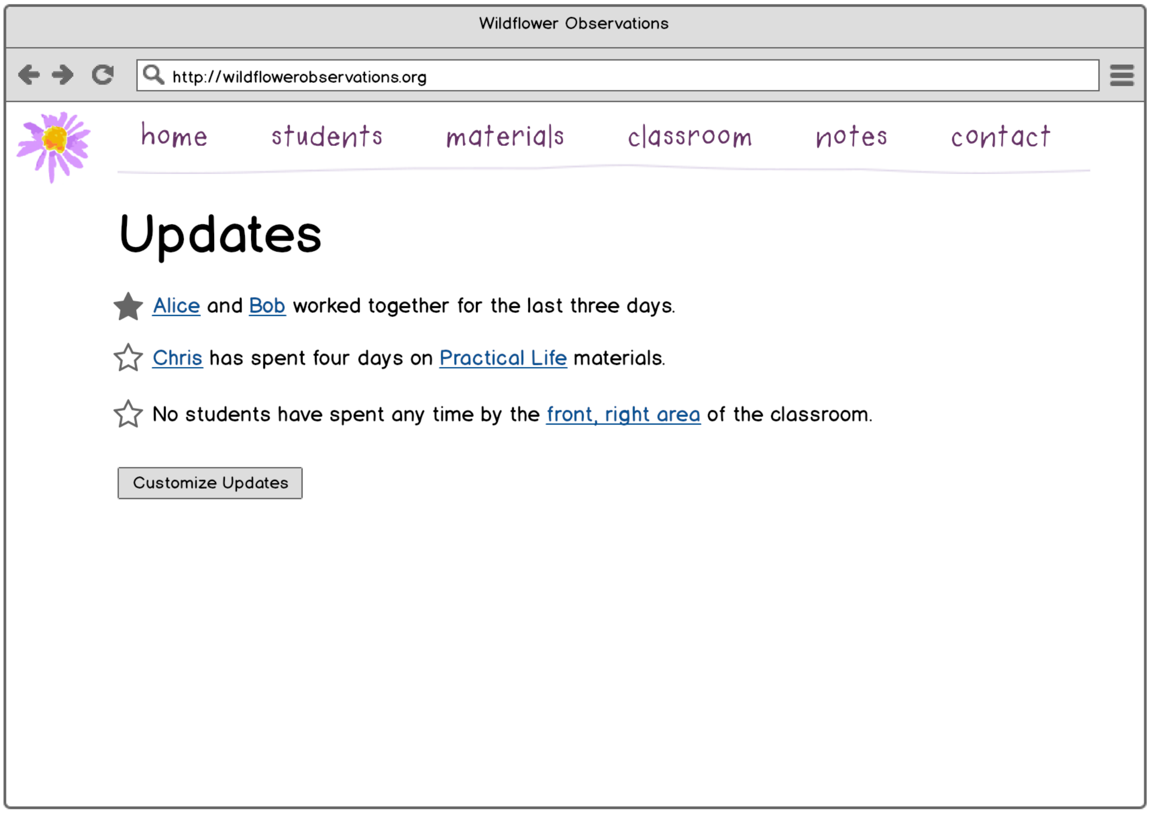

Figure 4-1: This is a screenshot of the interactive mocks that we used to get initial feedback.

Figure 4-2: This is a screenshot of the homepage of the web application.

The website currently uses the Bootstrap framework [42] and JQuery [43]. The visualizations use more complex JavaScript libraries.

4.2 Visualizations

As teachers click on 'students' or a specific student's name, they are presented with a tabbed interface that allows them to explore social, material, and classroom interaction. For each of these types of interaction, we explored several different types of visualizations, finally settling on the ones that were most intuitive for teachers, as seen in our usability study.

4.2.1 Social Interaction

We constructed three main types of visualizations to study social interaction in the classroom:

- A stacked-bar graph that allows teachers to see fine-grained time resolution
- A cluster diagram that allows teachers to see aggregate interactions more intuitively
- A radar diagram that allows teachers to see their interaction with children and easily determine lows, highs, and overlaps

These three types of visualizations vary different visualization techniques as well as time resolution to give teachers a more complete view of their classroom.

Figure 4-3 shows a stacked bar chart where teachers select a student from a drop-down menu and view that student's interaction with other students in the class. The stacked bar chart highlights periods when students are clustered together. This visualization clearly illuminates the periods when the shoes are returned to a shoe rack at the start and end of the school day. Teachers can identify patterns during other scheduled group activities. This particular student, Dave, spends most of his day with Chris, occasionally spending time with Helen.

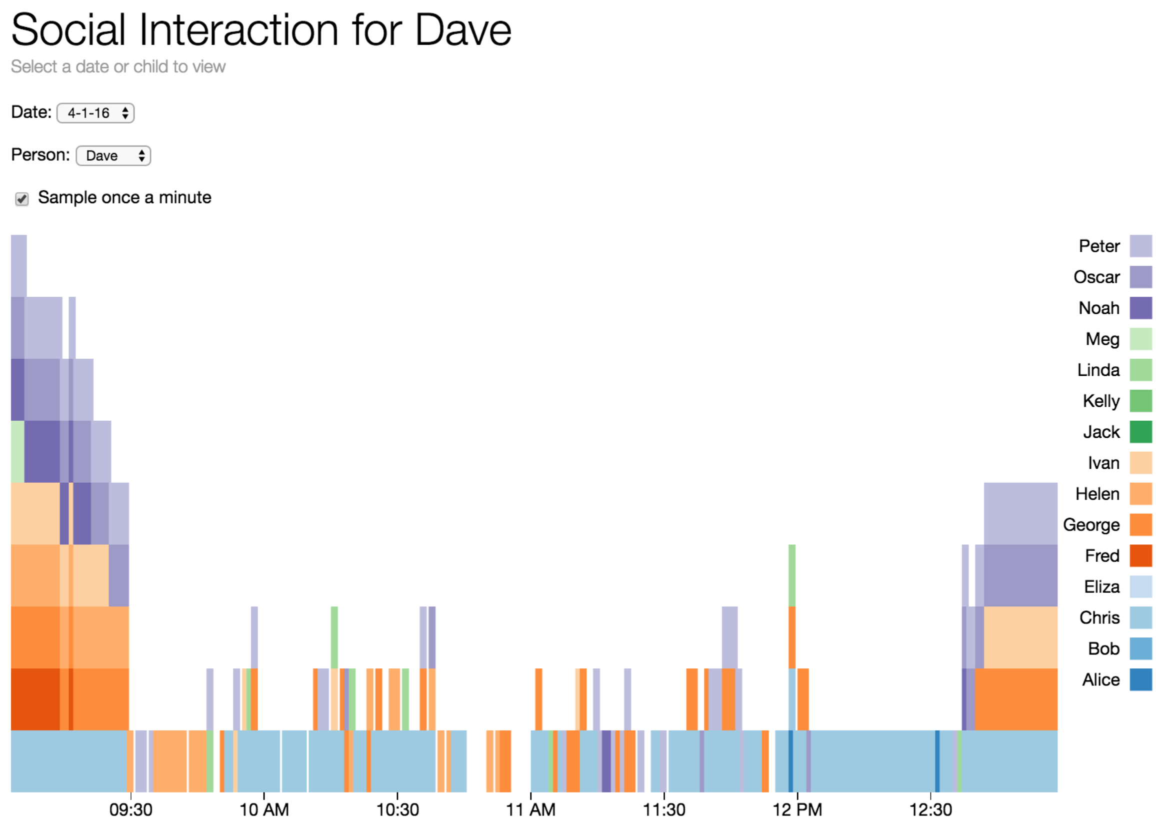

Figure 4-3: This visualization shows who children are spending time with in the classroom. In this chart, Dave spends most of his day with Chris.

In order to present a smoother chart, we first sample the raw sensor data, selecting one value for each minute. We then smooth these by calculating a moving average with window size of 3. If students are together for this time period, we consider that to be interaction. For each time stamp, we generate a bar and color the bars based on the device ids seen by the sensor. Although we initially set the bar height to be the raw RSSI value, we decided to set the height to a standard as there was very little variation in bar height that was distracting and not informative. Therefore, the presence of the colored bar indicates interaction.

Figure 4-4 shows connections between students over a day of activity. This visualization shows how much time students spend together through the thickness of each chord, allowing teachers to get a sense for aggregate interaction. As a teacher hovers over a name, they can view that student's interaction.

To construct this visualization, we took a traditional chord diagram and modified it[44]. The arcs are restricted to only flow to the other side; chords from the left side only connect to the right side. To do this, we passed in the data as a matrix where the upper left and lower right quadrants of the matrix were zero matrices. The sensor data was summed over the course of the entire day and the data was copied in both the upper right and lower left quadrants of the matrix. The resulting matrix ensured that the chords only connected opposite sides of the diagram.

In order to add curvature to the chords and create a clear distinction between the two halves, we added an empty, dummy region in between the two halves of the diagram. We created two arcs of arbitrary thickness that were only connected to each other. To reflect this in the data, we added a row and column of zeroes to the center of matrix, offsetting the four quadrants. We also added a row and column of zeroes to the end of the matrix.

We then caluclated the start angle and end angle of each arc by calculating the offset based on the sum of the matrix. From this, we can get a clustered chord diagram.

To reveal which students teachers spend the least time with, we created a visualization, shown in Figure 4-5. This shows the percentage of time two teachers spent with their students, so they can adjust accordingly. As a teacher hovers over a student, they can see the exact percentage of their time they are spending with this student. The overlapping area shows how a teacher distributes their time with respect to the other teacher. In this chart, both teachers are spending a lot of time with Bob, Chris, George, and Noah.

We use a variation of the D3 radar chart that uses rounded strokes and applies a gradient to the different data sets corresponding to the two teachers. These are mostly stylistic changes, intended to improve the appearance of the visualization.

4.2.2 Material Interaction

Teachers can also see a child's interaction with specific lessons. We can visualize how long one student spends time with a lesson material, as shown in Figure 4-6. Teachers

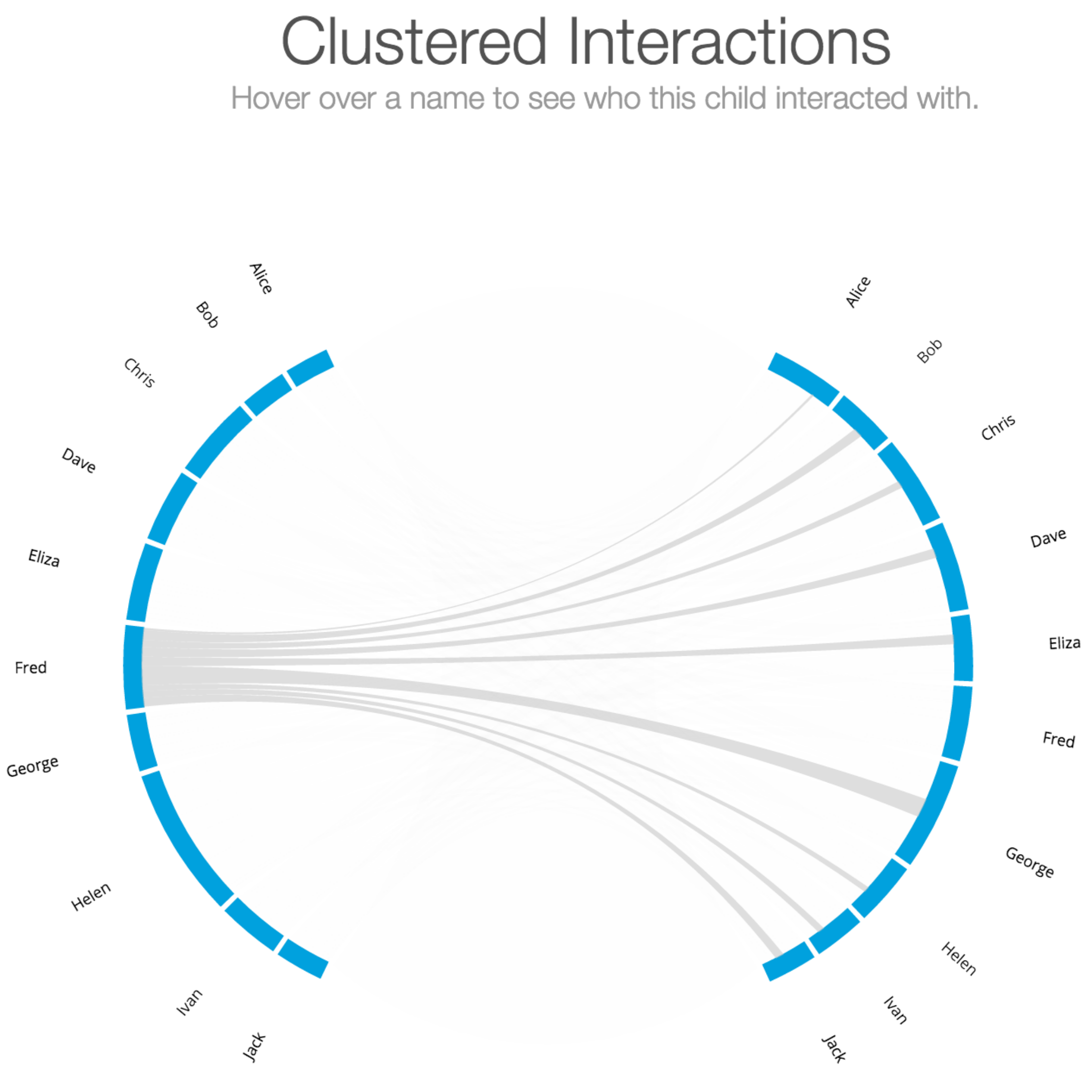

Figure 4-4: This visualization shows who children are spending time with over a day. Here, Fred is spending the most time with George.

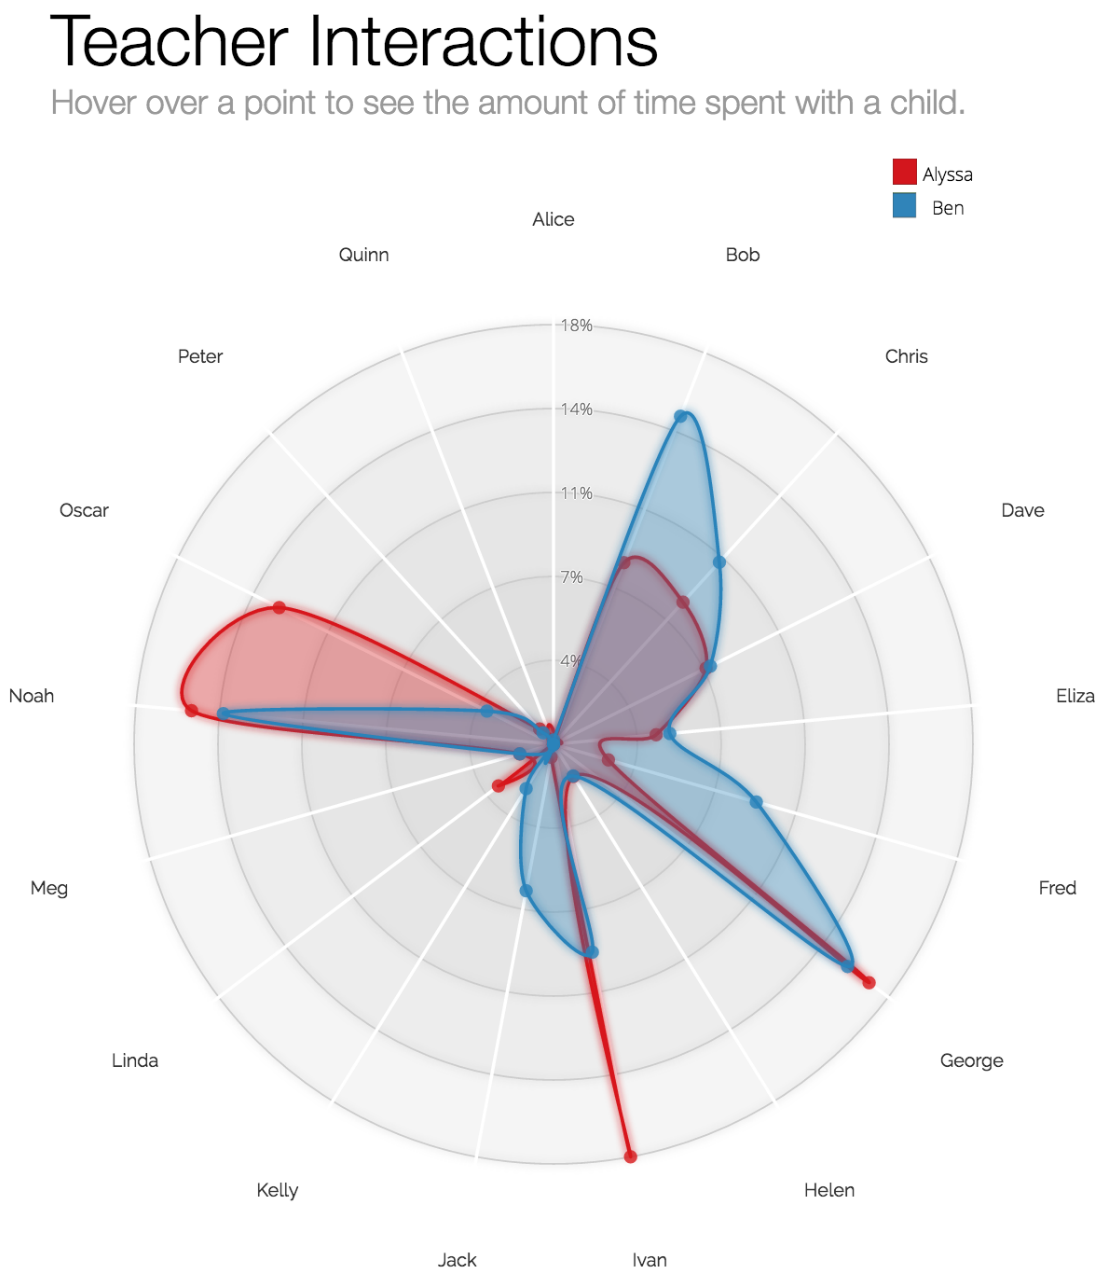

Figure 4-5: This visualization helps teachers distribute their time more evenly. Here, both teachers are spending very little time with Alice, Helen, Meg, Peter, and Quinn.

Figure 4-6: Teachers can see how long a child has spent with lesson.

can see a spike when the lesson is introduced and can use this to determine when a child has mastered a lesson. They can also use this data to decide when someone is ready to be shown a new lesson. In this case, Alice has probably mastered the wood polishing lesson. Teachers can select a lesson and a child to view data.

This visualization is constructed using a simple D3 bar chart function. We annotate this bar chart when we find the first strong peak of time spent with a given lesson and approximate that as when the lesson was first introduced.

4.2.3 Classroom Interaction

Teachers can also view the interaction with the classroom. Figure 4-7 shows the interaction data gathered from the region tracker sensors. As teachers hover over an active area of the classroom, they can see which students spend the most time there in the pie chart to the right. Montessori classrooms are arranged by subject material, so a teacher can immediately see a student's particular attention to a given subject.

To construct this visualization, we overlayed a hexagonal grid, constructed using

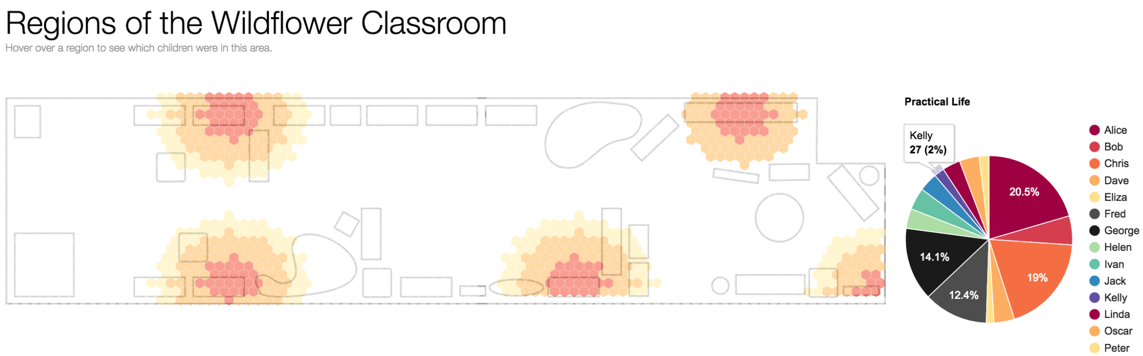

Figure 4-7: The map on the left side of this visualization highlights the areas where region trackers have been placed. The chart on the right shows which students have spent time in this area.

D3 hexbin, on a floorplan of the classroom. We calculated the color of each grid tile by looking at the aggregate RSSI values found around each region tracker location, adjusting the hue based of signal strength. We used these aggregate RSSI values to populate the pie chart for each region tracker's data. The pie chart is constructed using Google Charts.

4.3 Usability Study

We interviewed ten teachers from Wildflower schools to evaluate Sensei. The teachers were from three different types of schools: ages 1.5-3, ages 3-6, and bilingual schools. This variety in school type allowed us to get a larger diversity of perspectives. We conducted a qualitative study with these teachers to determine how useful Sensei would be for their classrooms. We found that the teachers were more inclined to discuss how they could leverage the data, rather than specific details about the dashboard and the visualizations. Teachers were shown all of the visualizations shown in this section and we recorded video of the 30-minute interviews, tracking their interaction with the dashboard and their comments.

The teachers were very enthusiastic about augmenting their observations with data from Sensei. One teacher mentioned that "specific quantifying elements" would help her improve her classroom immensely, because the network would capture what she was unable to observe in the classroom. One teacher commented on how her classroom was divided with walls, preventing her from observing all children at once. A teacher who taught younger students, ages 1.5-3, thought this data would be especially helpful for her classroom. Younger students are often more active in the classroom - moving between lessons, regions, and other students more frequently. These classrooms are even more challenging to observe and the teachers are often even more engaged with the students, as they have to assist them more often. Teachers echoed how they often resorted to their own memory of previous days to plan future lessons, because they don't have time to record all their notes during the school day. Another teacher mentioned that this data "would be especially meaningful around parent-teacher check-in time", as they can have more anchored and quantative conversations with parents.

When shown the stacked bar visualization, the teachers were very drawn to the relationship between Chris and Dave. Although we showed each set of teachers these anonymized visualizations, the two teachers from the classroom where we had collected this particular set of data were immediately able to identify both Chris and Dave in their classroom as these two students spend most of their time together.

When shown the cluster diagram, the teachers found this more intuitive to understand relationships between children. Although we had initially assumed that teachers would like to see the fine-grained time resolution present in the stacked bar graph, we learned that their weekly planning is more influenced by aggregate information, like that present in the cluster diagram. From this diagram, they can easily identify children who are isolated whom they might encourage to interact with their peers.

The teachers were very excited to see the radar chart, detailing their own interactions. One teacher reflected on the visualization, mentioning how she "need(s) to get these two kids more independent from us. One of the key Montessori principles is independence." In the bilingual classroom, the teachers were particularly interested in seeing who spent more time with the Spanish-speaking teacher vs. the English-speaking teacher. Nearly all the teachers drew comparisons between the anonymized student names and their own classrooms, often commenting on how the children who they spent most of their time with (as those children are often more disruptive and unfocused in the classroom, requiring a little more guidance).

Teachers from 3-6 year old classrooms were more interested in the lesson material graph. Teachers from younger classrooms did not find this as relevant, as those students did not spend as much time working on specific lessons. The other teachers were very interested in using this to inform their weekly lesson planning. They remarked on how the sensors could capture quick interactions, like a child removing a lesson from a shelf and then returning it when they realized that it is too difficult for them. From seeing data from this interaction, a teacher could identify this and introduce this lesson to the students, since they've indicated an interest. They also reflected on how this would help them adjust their classroom: "If these lessons aren't being used at all, then my class is done with them." Teachers did have some questions on how to determine if a child is actually interacting with a lesson. We plan to address this by adding more sensors to the lesson trays to identify touch.

When shown the visualization using classroom interaction data, teachers were very drawn to the pie chart. The teachers identified a use case where they might encourage a student to explore a new subject area if they see a student spending too much time in a given region of the classroom.

Montessori teachers are usually very hesistant about introducing technology in the classroom, and most Montessori schools do not have devices like tablets or laptops that one might see in a more traditional school. Despite this, Sensei was well received by the teachers and they felt that it did not change the classroom experience for the students.

Chapter 5

Conclusion

We have built Sensei: Sensing Educational Interaction to better understand learning and early childhood development in the Montessori classroom. We have designed and build custom sensor nodes that are placed in shoes, lessons, and around the room and can communicate with a custom mother node sensor. This mother node sensor can connect to our Android application and start all the sensors and collect all the data from the sensors. The firmware is optimized to preserve battery life with a scheduling scheme that cycles through data collection and sleep. Once the data is collected from the cellphone application, teachers can view this data in a web application. We have constructed five main visualizations to study social, lesson material, and classroom interaction. Teachers were very enthusiastic about using Sensei in their classroom and identified several key use cases for the data.

5.1 Future Work

In the future, we hope to improve upon every aspect of the Sensei system. This includes inductive charging, increasing sleep time to improve battery life, adding a screen to the cellphone application to view the battery life of the sensors, and using the data collected by the accelerometer and microphone to augment the visualizations. Outside of these improvements, we also hope to build a statisical modeling engine, leveraging the rich dataset we have gathered.

Using mathematical modeling, we hope to better study how clusters form in the classrom and what influences these have on learning outcomes, using pagerank-based algorithms [6, 16]. We can study periods where children are most receptive to learning, “sensitive periods”, using Hidden Markov Models [18]. We can also uncover latent themes amongst lessons by observing themes found using algorithms like probabilistic latent semantic indexing [7]. While our primary focus has been on teachers, we want to expand use of this dataset for educators and researchers who can look more deeply at early childhood development.

Appendix A

Tables

None

Appendix B

Figures

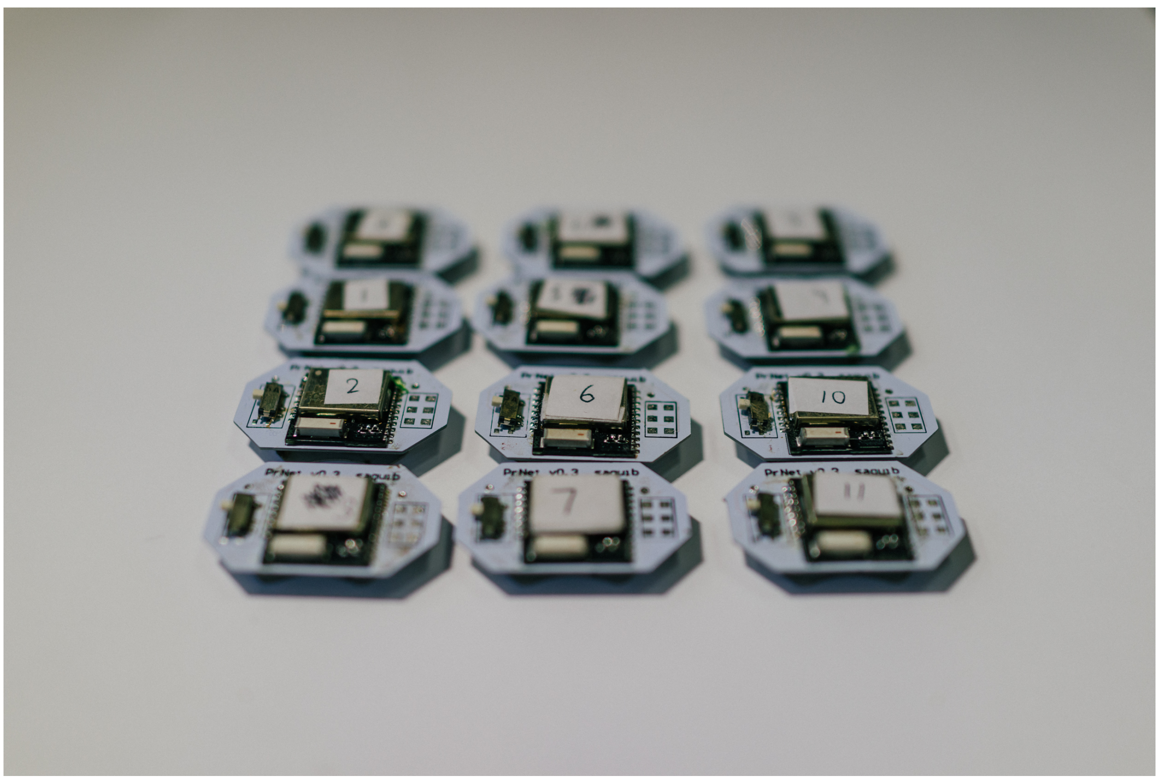

Figure B-1: These sensors were an earlier prototype version that used the RFduino module. We found these to be inconsistent in RSSI measurement.

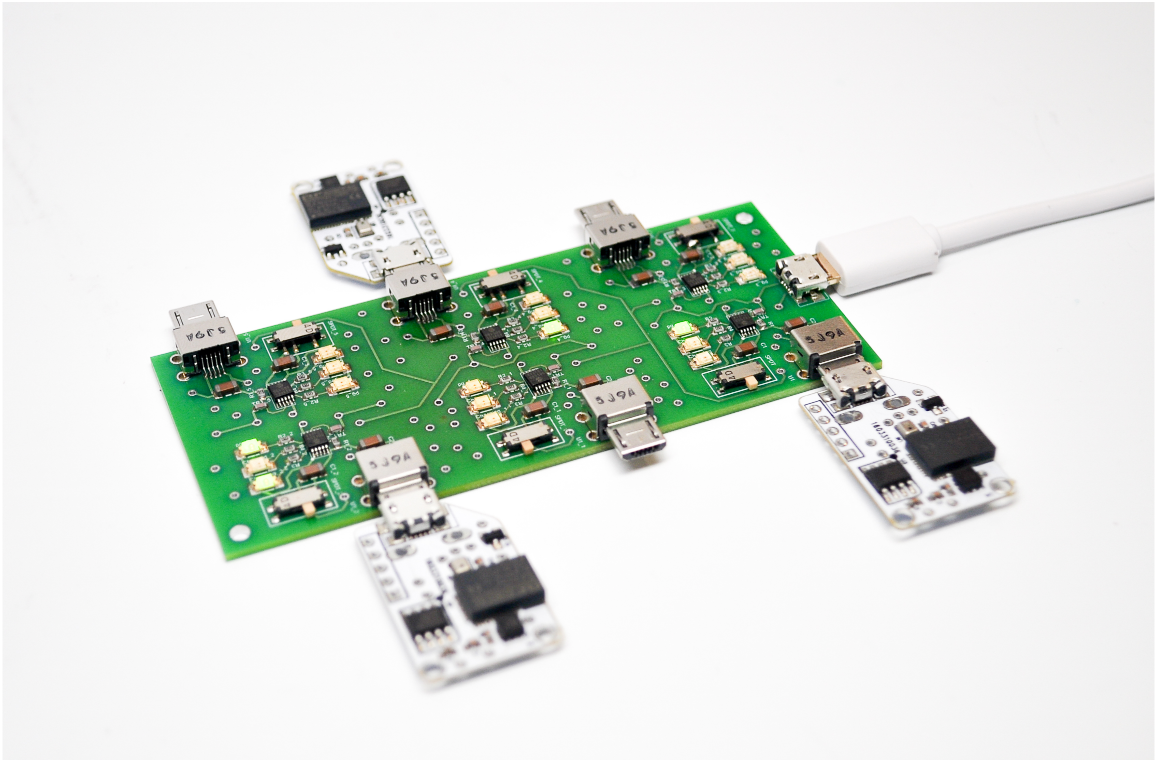

Figure B-2: Shoe sensors could be plugged into a USB recharging strip so they can collect data in the classroom for the next day.

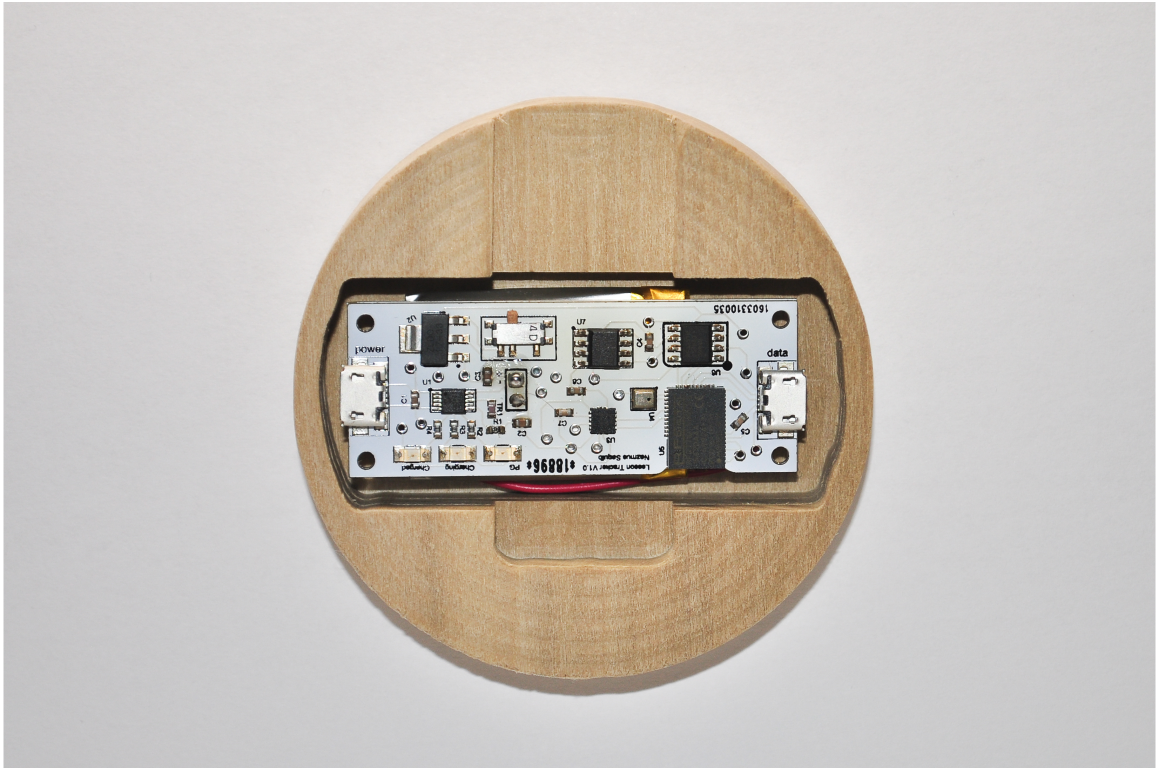

Figure B-3: Sensors can be placed around the room to track specific regions of the classroom. This is a wooden casing that can be attached underneath shelves to measure proximity.

Figure B-4: Sensors can be placed in the feet of a lesson tray. These feet can be attached to any pre-existing lesson in the classroom.

Figure B-5: The mother node sensor can be connected to the cellphone to start the sensors and collect data from the sensors.

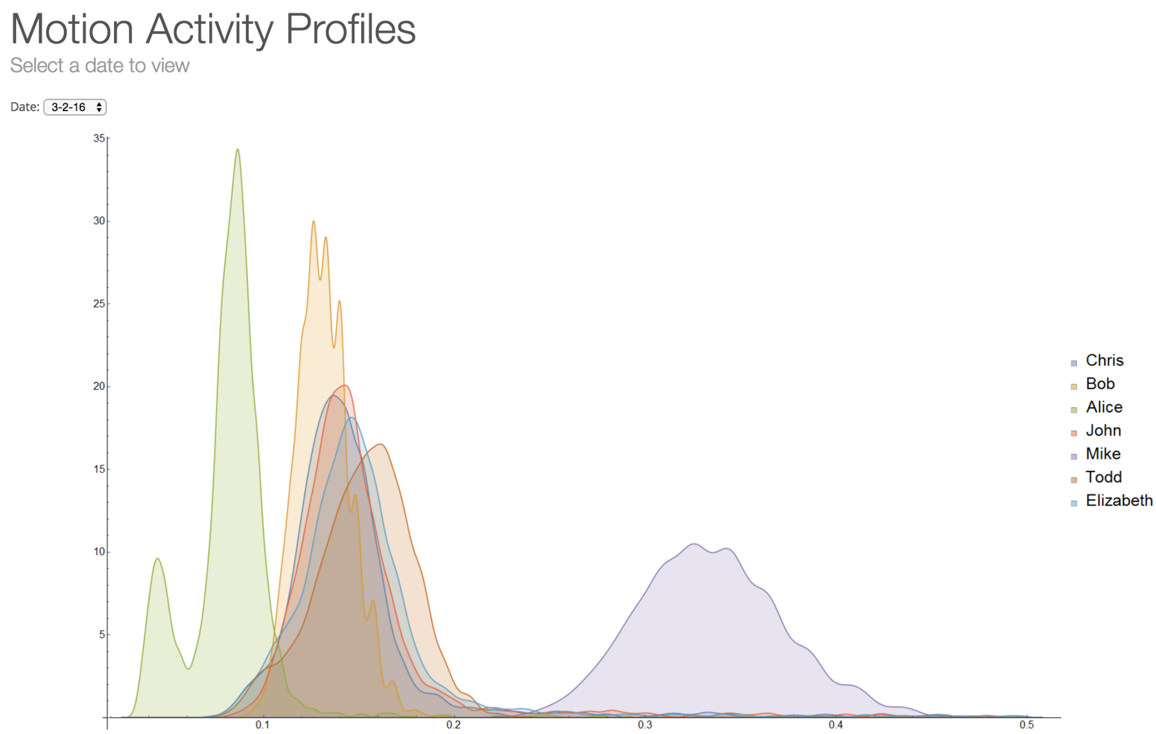

Figure B-6: We can use the data from the accelerometers to augment the proximity data. Here, we show a histogram visualization of just the motion data and can identify particularly active children, like Mike.

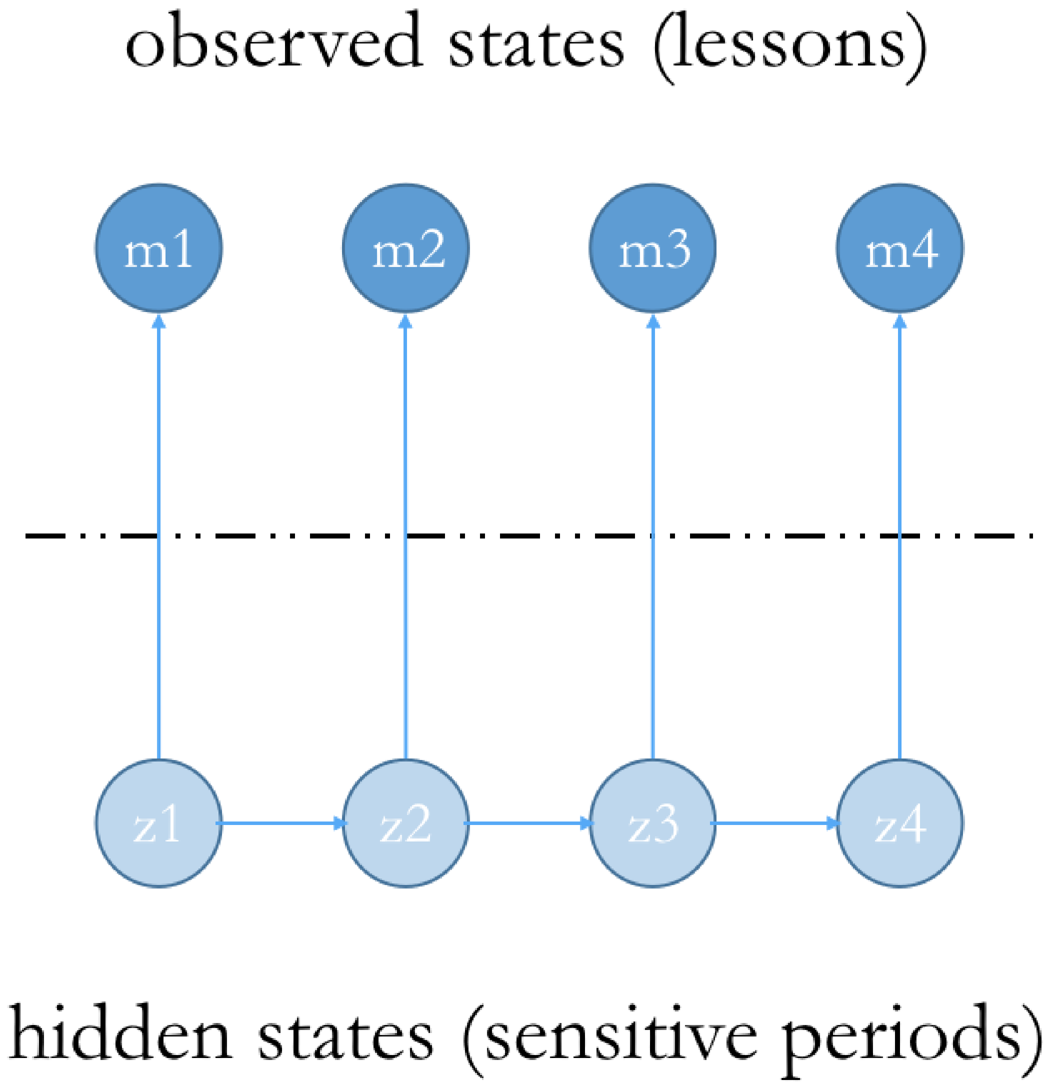

Figure B-7: For future work, we can leverage Hidden Markov Models to understand early childhood development by identifying sensitive periods when a child is particularly receptive to learning.

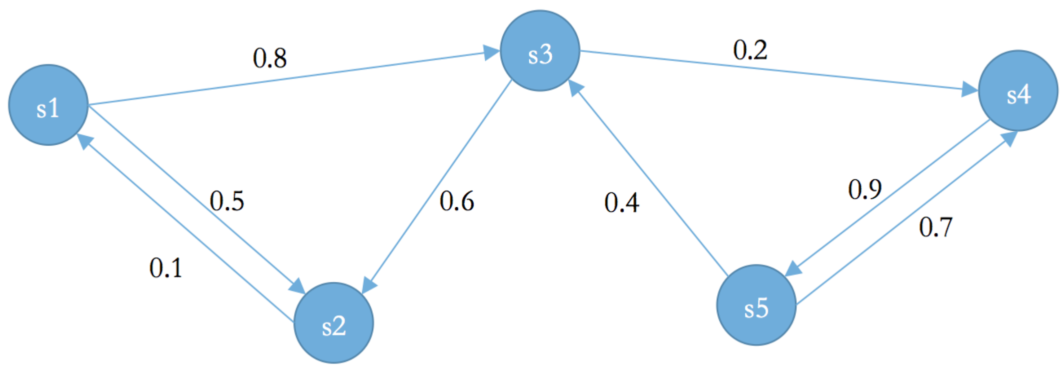

Figure B-8: For future work, we can leverage PageRank algorithms to understand the influence students have on each other and their learning outcomes.

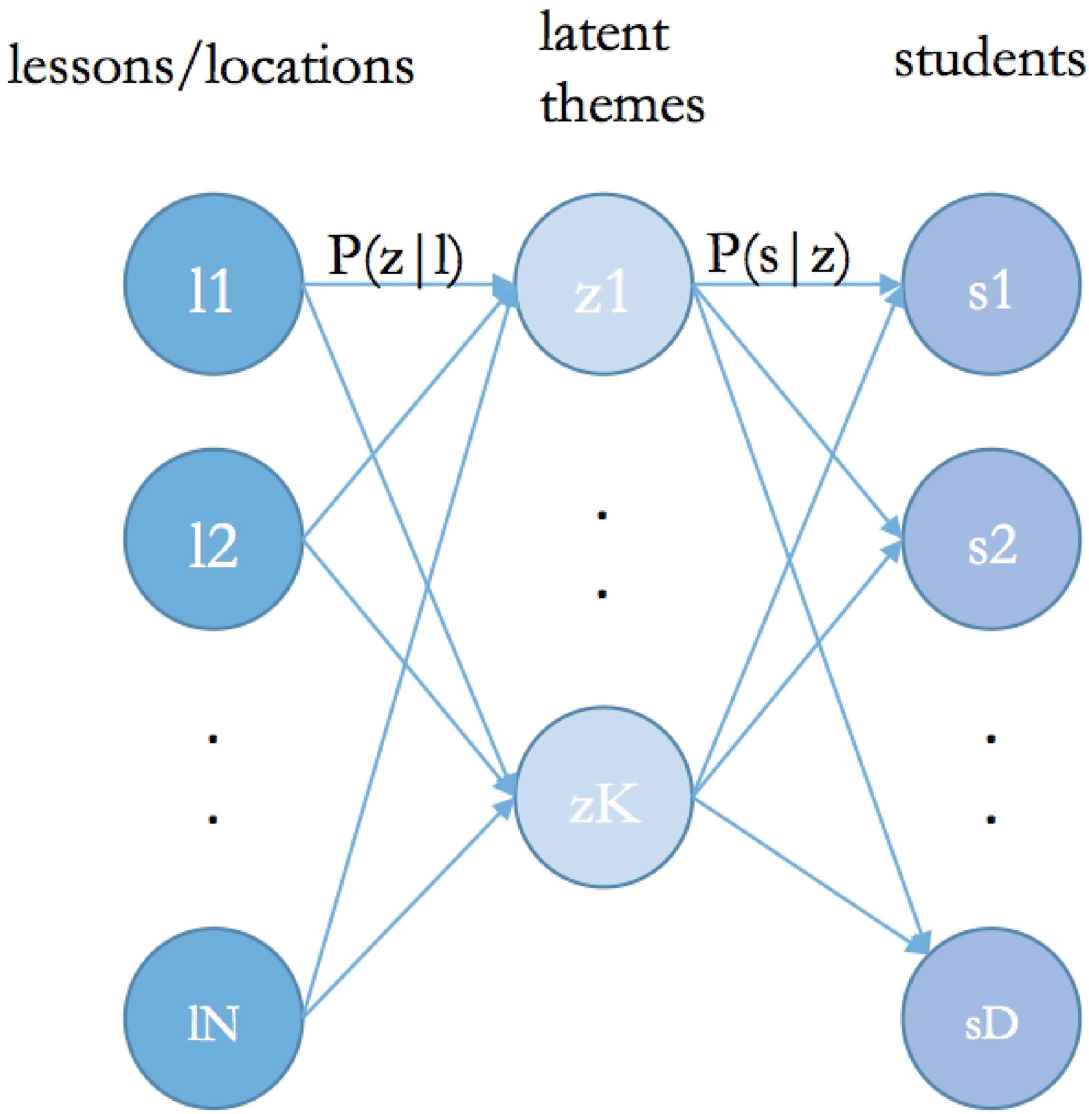

Figure B-9: For future work, we can identify connections betweens lessons and students to identify latent themes in materials, using probablisitic latent semantic indexing.

Bibliography

[1] Agrawal, P., Girshick, R., and Malik, J. Analyzing the performance of multilayer neural networks for object recognition. In Computer Vision–ECCV 2014. Springer, 2014, 329–344.

[2] Barrat, A., Cattuto, C., Colizza, V., Pinton, J., Van den Broeck, W., and Vespignani, A. High resolution dynamical mapping of social interactions with active rfid. ArXiv e-prints (Nov. 2008).

[3] Ben Shitrit, H., Berclaz, J., Fleuret, F., and Fua, P. Multi-commodity network flow for tracking multiple people. Pattern Analysis and Machine Intelligence, IEEE Transactions on 36, 8 (2014), 1614–1627.

[4] Dunn, J. The emanuel miller memorial lecture 1995 children's relationships: Bridging the divide between cognitive and social development. Journal of Child Psychology and Psychiatry 37, 5 (1996), 507–518.

[5] Epstein, P. An Observer's Notebook: Learning from Children with the Observation CORE. Montessori Foundation, 2012.

[6] Haveliwala, T., Kamvar, S., and Jeh, G. An analytical comparison of approaches to personalizing pagerank.

[7] Hofmann, T. Probabilistic latent semantic indexing. In Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval, ACM (1999), 50–57.

[8] Khoshelham, K., and Elberink, S. O. Accuracy and resolution of kinect depth data for indoor mapping applications. Sensors 12, 2 (2012), 1437–1454.

[9] Kumar, S., Marks, T., and Jones, M. Improving person tracking using an inexpensive thermal infrared sensor. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (2014), 217–224.

[10] Lillard, A. Montessori: The Science Behind the Genius. Oxford University Press, 2005.

[11] Lillard, A., and Else-Quest, N. Evaluating montessori education. SCIENCE-NEW YORK THEN WASHINGTON- 311, 5795 (2006), 1893.

[12] Lillard, A. S., Saracho, O., and Spodek, B. Playing with a theory of mind. Multiple perspectives on play in early childhood education (1998), 11–33.

[13] Mastrandrea, R., Fournet, J., and Barrat, A. Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys. PLoS ONE 10 (09 2015).

[14] Montessori, M. The montessori method. Transaction Publishers, 2013.

[15] Olguín-Olguín, D., and Pentland, A. Sensor-based organisational design and engineering. International Journal of Organisational Design and Engineering 1, 1-2 (2010), 69–97.

[16] Page, L., Brin, S., Motwani, R., and Winograd, T. The pagerank citation ranking: bringing order to the web.

[17] Phillips, D. A., Shonkoff, J. P., et al. From Neurons to Neighborhoods:: The Science of Early Childhood Development. National Academies Press, 2000.

[18] Rabiner, L. R. A tutorial on hidden markov models and selected applications in speech recognition. Proceedings of the IEEE 77, 2 (Feb 1989), 257–286.

[19] Stehlé, J., Voirin, N., Barrat, A., Cattuto, C., Isella, L., Pinton, J.-F., Quaggiotto, M., Van den Broeck, W., Régis, C., Lina, B., et al. High-resolution measurements of face-to-face contact patterns in a primary school. PloS one 6, 8 (2011), e23176.

[20] Stopczynski, A., Sekara, V., Sapiezynski, P., Cuttone, A., Madsen, M. M., Larsen, J. E., and Lehmann, S. Measuring large-scale social networks with high resolution. PloS one 9, 4 (2014), e95978.

[21] Stopczynski A, Sekara V, S. P. C. A. M. M. L. J. e. a. Measuring Large-Scale Social Networks with High Resolution.

[22] Uhrmacher, P. B. Uncommon schooling: A historical look at rudolf steiner, anthroposophy, and waldorf education. Curriculum Inquiry 25, 4 (1995), 381–406.

[23] Wurm, J. Working in the Reggio way: A beginner's guide for American teachers. Redleaf Press, 2005.

[24] Zappi, P., Farella, E., and Benini, L. Tracking motion direction and distance with pyroelectric ir sensors. Sensors Journal, IEEE 10, 9 (2010), 1486–1494.

[25] Hubbell, Elizabeth Ross. "Authenticity and Technology in Montessori Education." Montessori Life: A Publication of the American Montessori Society18.2 (2006): 16-20.

[26] "Montessori Education Provides Better Outcomes Than Traditional Methods, Study Indicates." UVA Today. N.p., 29 Sept. 2006. Web. 10 May 2016.

[27] Moeslund, Thomas B., and Erik Granum. "A survey of computer vision-based human motion capture." Computer vision and image understanding 81.3 (2001): 231-268.

[28] Han, Jungong, et al. "Enhanced computer vision with microsoft kinect sensor: A review." Cybernetics, IEEE Transactions on 43.5 (2013): 1318-1334.

[29] MartÄśnez, A., et al. "Combining qualitative evaluation and social network analysis for the study of classroom social interactions." Computers and Education 41.4 (2003): 353-368.

[30] Srivastava, Mani, Richard Muntz, and Miodrag Potkonjak. "Smart kindergarten: sensor-based wireless networks for smart developmental problem-solving environments." Proceedings of the 7th annual international conference on Mobile computing and networking. ACM, 2001.

[31] Stein, K.; Wegener, R. and Schlieder, C.: Pixel-Oriented Visualization of Change in Social Networks.Proceedings of the International Conference on Advances in Social Networks Analysis and Mining (ASONAM), IEEE Computer Society, 2010.

[32] RF Digital. N.p., n.d. Web. 10 May 2016.

[33] "SimbleeTM." Simblee. N.p., n.d. Web. 10 May 2016.

[34] "Ksksue/PhysicaloidLibrary." GitHub. N.p., n.d. Web. 10 May 2016.

[35] "Graph View." Graph View. N.p., n.d. Web. 10 May 2016.

[36] "Flask." Flask. Pocoo, n.d. Web. 10 May 2016.

[37] "Cloud Application Platform | Heroku." Cloud Application Platform | Heroku. N.p., n.d. Web. 10 May 2016.

[38] "The Python SQL Toolkit and Object Relational Mapper." SQLAlchemy. N.p., n.d. Web. 10 May 2016.

[39] "D3.js - Data-Driven Documents." D3.js - Data-Driven Documents. N.p., n.d. Web. 10 May 2016.

[40] "Charts | Google Developers." Google Developers. N.p., n.d. Web. 10 May 2016.

[41] "Balsamiq." . Rapid, Effective and Fun Wireframing Software. N.p., n.d. Web. 10 May 2016.

[42] "Bootstrap Âů The World's Most Popular Mobile-first and Responsive Front-end Framework." Bootstrap Âů The World's Most Popular Mobile-first and Responsive Front-end Framework. N.p., n.d. Web. 10 May 2016.

[43] "JQuery." JQuery. N.p., n.d. Web. 10 May 2016.

[44] "Hacking a Chord Diagram to Visualize Flows - VisualCinnamon." Visual Cinnamon. N.p., 28 Aug. 2015. Web. 10 May 2016.

[45] Nazmus Saquib, Ayesha Bose, Dwyane George, Connie Liu, Sepandar Kamvar. “Sensei: Sensing Educational Interaction.”