Sensei: Sensing Educational Interaction

by

Nazmus Saquib

BA, Bard College, 2011

MS, University of Utah, 2013

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning,

in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2016

Signature of author: Signature redacted

Nazmus Saquib
Program in Media Arts and Sciences

Certified by: Signature redacted

Sepandar Kamvar
Associate Professor of Media Arts and Sciences
Thesis Supervisor

Signature redacted

Accepted by:

Pattie Maes, Ph.D.
Academic Head
Program in Media Arts and Sciences

Abstract

We present Sensei, the first system designed to understand social interaction and learning in an early-childhood classroom using a distributed sensor network. Our unobtrusive sensors measure proximity between each node in a dynamic range-based mesh network. The sensors can be worn in the shoes, attached to selected landmarks in the classroom, and placed on lessons. This data, accessible to teachers in a web dashboard, enables teachers to derive deeper insights from their classrooms. Further, the anonymized data can be used in large-scale research in early childhood. Sensei is currently deployed in three Montessori schools and we have evaluated the effectiveness of the system with teachers. Our user evaluations have shown that Sensei helps discover insights that would have otherwise been lost.

Sensei: Sensing Educational Interaction

by

Nazmus Saquib

The following served as a reader for this thesis.

Signature redacted

Alex (Sandy) Pentland
Professor, Media Arts and Sciences
Massachusetts Institute of Technology

Sensei: Sensing Educational Interaction

by

Nazmus Saquib

The following served as a reader for this thesis.

Signature redacted

Mitchel Resnick
LEGO Papert Professor of Learning Research
Program in Media Arts and Sciences
Massachusetts Institute of Technology

Acknowledgements

Sep Kamvar has become an influential figure in my life and the future course it is going to take. Not many people have significantly influenced my decisions and philosophy in life, but I am proud to say that he has. It's been a privilege having him as my mentor for the last few years. Sandy Pentland and Mitch Resnick provided valuable feedback about my work and served as readers for the thesis.

Sensei would not have been possible to develop at scale and deploy without the help of two fellow graduate students: Ayesha Bose and Dwyane George. Their dedication to the project was unparalleled and inspiring.

Yonatan Cohen, an artist and designer in the Social Computing group, has been a good friend, mentor, and collaborator for the past two years. He was an early enthusiast of Sensei, and assisted us in many aspects of the project, starting from helping us design and build lesson trays, to producing various ideas of customizing shoes for deploying the sensors.

Fellow Social Computers who assisted in developing Sensei deserve a big thank you. Kim Smith, Gal Koren, Salman Ahmad, Jia Zhang, Noah Wall, Caroline Jaffe, and Elizabeth Christoforetti helped me with valuable feedback and inspiration during the project. RaMell Ross helped us produce a great video that describes Sensei with playful details.

Deeni Fatiha has always been there during the ups and downs in my life at MIT. Thank you for tolerating and encouraging me. My parents and sisters have always done the same, I am indebted to them forever.

The MIT Bangladeshi Students' Association community has been an integral part of my MIT life. Special thanks to Khalid Jawed and Syed Arefinul Haque who patiently listened about my projects during buffet lunches. Mehrab Bakhtiar, Tamzid Islam, and Salman Saqib peppered my life with humor every now and then - thanks for being good friends to me.

Last, but not the least, Linda Peterson and Keira Horowitz have saved my life multiple times! They always find a way to accommodate the needs of MIT Media Lab students. Thank you for being there.

Dedicated to everyone who filled my life with love.

I am grateful for all the love I have received during the critical times of my life. The least I can do is to dedicate my work to people who were there during my difficult times. It requires only a few words to dedicate my work, but the magnitude of love and support I have received surpassed all expectations by many folds. I always thought that three things essential to life are dreams, hope, and imagination. As I go through different life experiences over the years, I have decided to add one more ingredient to the list of essentials: people who shower you with unconditional love. This may sound like a cliché, but it is only by going through difficult experiences that we internalize such ingredients for living.

Here is to loving and living.

Table of Contents

1 Introduction

Understanding early childhood development has received increasing attention in the recent years. Early childhood has been shown to be the most important time in development [1, 2]. Early environments are very influential, especially nurturing relationships formed during this stage [3]. A child's ability to learn and actively construct knowledge is unparalleled, but it is still a field that remains relatively complicated to analyze and study.

Our work aims to create a set of tools for a preliminary understanding of early childhood development. 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. It is more effective in alternative forms of education, such as Montessori [4], Reggio [5] etc. where children have freedom of movement and exploration in the classroom.

The Montessori classroom emphasizes independence and teachers act as facilitators to guide learning. As facilitators, teachers are to spend a good part of their day observing the children [6]. Current methods of such 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 educational interactions.

Sensei has been developed by a team of graduate students at MIT: myself, Ayesha Bose, and Dwyane George. I conceived the notion of Sensei after attempting several technical solutions to automate and augment observation in Montessori classrooms, creating the early RFduino prototypes of Sensei in the process (Chapter 4). Over the past year, this system has grown in its breadth and maturity as my teammates joined to enhance and scale it. Specifically, I designed the hardware and developed the firmware for the current version of Sensei, along with assisting the development of the smartphone app and web dashboard. This thesis briefly chronicles the journey and design decisions behind Sensei, and describes the current version in detail.

1.1 Montessori Education

A variety of different educational methods have emerged that challenge the traditional notion of education, teachers, and students [7, 8]. The Montessori educational method is over 100 years old and is used in over 5000 schools in the United States [6]. The method is unique in its multi-age classrooms, educational materials, and the students' freedom to choose their lessons.

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 educational method allows for more peer-guided learning than traditional classroom environments [7, 9].

The Social Computing group at the MIT Media Lab has recently started a network of schools, known as Wildflower Schools. These schools are an open source approach to Montessori learning, blending both traditional Montessori methods with new enhancements. As lab schools, the Wildflower schools serve as a research setting dedicated to advancing the Montessori method. Our work with Sensei focuses on enhancing an important component of the Montessori method: observation.

1.2 Manual Observation

Dr. Maria Montessori, the creator of the Montessori Method, described observation as a critical component in a Montessori classroom [4]. Through observation, a teacher can better understand each student's interests, learning style, and individual needs.

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 concentration of the student (engagement with the lesson), level of completeness, and the mood associated with working on the lesson. They can track the mood associated while working with other students too.

Social: Teachers can determine patterns of social behavior, like learning a new lesson together or assisting others. They can track recurring social groups of students and study their learning progress over time.

Environmental: Teachers can track which areas of the classroom or materials are used more, adjusting the design of the room as needed to encourage students to explore new or important concepts.

Wildflower classrooms usually have two teachers and 12 – 15 students. Given all the distinct events for observation, it can be very difficult for two teachers to accurately assess such a classroom. Using these observation, teachers also help students who are having trouble, and often intervene with students who are new to the Montessori environment. A common theme that was apparent in our initial interviews of teachers (Chapter 2) was how they want to spend more time on observation. Teachers currently record observations with handwritten notes, which are harder and more time consuming to analyze when there are many of these notes.

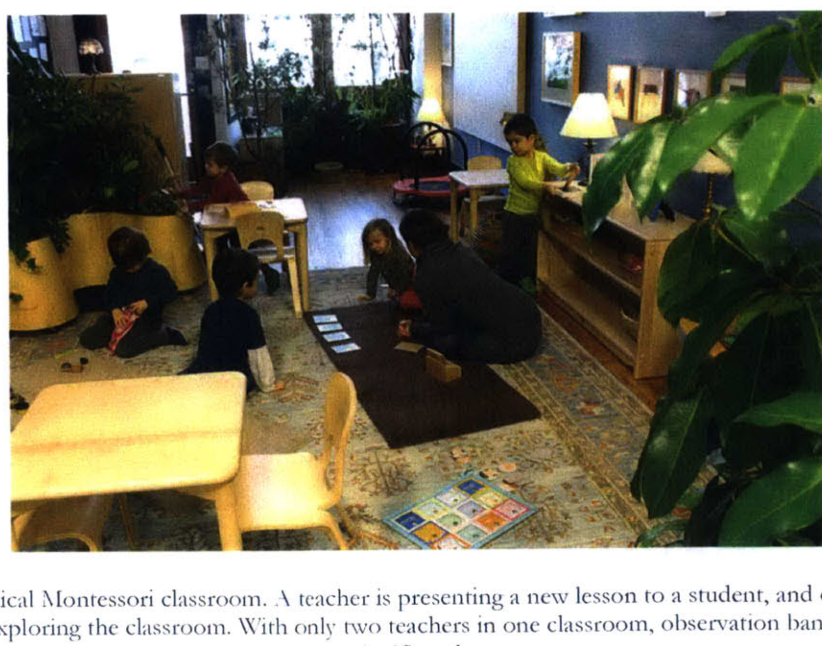

Figure 1.1. A typical Montessori classroom. A teacher is presenting a new lesson to a student, and other students are independently exploring the classroom. With only two teachers in one classroom, observation bandwidth decreases significantly.

1.3 Contribution

Sensei enhances the fields of learning analytics and Montessori observation methods by the following technical contributions:

- Sensors designed to instrument a classroom to capture proximity and motion data in an accurate, low-cost, and minimally-invasive manner
- A network event scheduling scheme that enables data collection at a reasonably high sampling rate for social interaction, in a battery-preserving manner
- A smartphone application to facilitate data collection by teachers
- An interactive web application designed to visualize social, material, and classroom interactions

Together, the system allows the study of both social interaction and learning at scale.

Other than describing the contributions in detail, a quantitative evaluation of our custom wireless protocol's robustness is done, along with a qualitative evaluation and interviews with teachers about the use cases of the system. Chapter 2 gives an overview of our initial interviews with teachers and manual data collection tests we have done in the schools. Chapter 3 gives an overview of the system design. Chapter 4 shows hardware design iterations of Sensei, and the reasoning behind each iteration. My custom sensor network protocol and other firmware related developments are described in Chapter 5. Chapter 6 describes the smartphone application and the web visualization dashboard. Chapter 7 demonstrates the use of some statistical and network analysis methods as a way to research and model the Sensei data further. Chapter 8 describes the use cases interviews we have done with Montessori teachers after our final pilot study was conducted in three different schools. Related works are described in Chapter 9, and the exciting future work section concludes the thesis in Chapter 10.

2 User Interviews and Class Observation

The idea of Sensei was conceived and motivated by interviewing teachers in Wildflower Montessori schools. Interviews were designed to inquire into the nature of classroom observations, and the current limitations teachers face in their everyday classroom activities. By aggregating and analyzing answers from three different schools, we came up with several concrete use cases where technology could have a potentially impactful contribution. The design decisions for the hardware, firmware, and visualization components were also motivated by the teacher interviews.

Between the time of teacher interviews and actual deployment, we also deployed some manual observation cards in the classroom that provided us some data and insights about these concrete use cases. In this section, we detail our interview process and questionnaire, along with the deployment of manual observation cards that informed the intended use cases of Sensei.

2.1 Initial Teacher Interviews

Six teachers were recruited from three different Montessori Schools in Cambridge. They were all females, certified by the American Montessori Association to teach at the pre-K level. The schools were: Wildflower Montessori School (children aged 3 – 5), Dandelion Montessori School (children aged 2 – 4), and Violeta Montessori School (children aged 2 – 4). Each interview typically lasted for 20 minutes. Handwritten notes on paper were taken for each question (appendix A), and the interviews were recorded on audio recorders. The key answer parts were transcribed later into text. Teachers were usually given an initial prompt, and we let them speak about their individual classroom scenario without leading the discussion.

2.1.1 Questions

The goal of the prompts was to find out what social interactions and learning interactions are most important in the Montessori classrooms, and determine the key information teachers look for day to day to gauge learning.

The following questions were the general prompts for the user interviews. As teachers answered our questions, sometimes we asked additional clarifying questions.

1. What events lead/motivate you to an observation at different times of the day?
2. What motivates you to write down these observations? How do you use these notes later?
3. I want you to walk me through a day in the school. What do you observe at different times in relation to the activities in the school?
4. What are some key lessons/categories that are housed in lesson trays?
5. What leads you to interact with a child?
6. How do you define focus or engagement for a child? For example, is this the length of time spent on the lesson, or the intensity?
7. What are the different kind of social interactions that happen in the classroom?
8. Tell us about specific patterns that seem to be recurring activities in the classroom. Is it important for you to observe this rhythm/pattern?
9. How can we help you observe? Or enable you to observe certain things?
10. If you see a child is not interacting much with a category of lessons, do you intervene or try to encourage them to use it?

2.1.2 Summary of Key Insights

We received varied responses from different schools. It was interesting to see that depending on different age groups of the students, the problems faced by teachers were somewhat different. There were also some common limitations of current observation methods that they pointed out.

2.1.2.1 Observation Criteria

There are many factors that prompt a teacher to make observations and write them down. Some key observations made by teachers are:

- Individual student work cycles
    - Level of focus and pace
    - Some lessons deal with physical engagement: e.g., table washing
    - Some are mental engagement: mentally engaged and usually involves fine motor movement
    - Purposeful work is most important: they should not be distracted
- Social connections and group work activities.
- Activity level of teachers: how much time they spend observing and working with individual child as opposed to moving around in classroom.
- Most active areas of classroom. As the classroom space is organized according to different categories of lessons, what shelves and activities are specific children choosing?
- Is there a group flow that happens throughout the day?
- Difference between movement of younger and older children.
- Gauge whether students are interrupting each other.
- How many pieces of work they took off the shelf every day.

2.1.2.2 Frequency and Importance of Observation

All teachers emphasized on the importance of observing students and their activities as recommended by Maria Montessori in her book [4]. Lesson planning, how much time teachers spend per student, individual child's progress tracking, what category of lessons are children interested in, these are a few common themes that were of concern to most teachers. Given the unobtrusive nature of the Montessori Method, teachers generally like to allow the children to work on their own and act as a passive influence to stir them towards new lessons and class interactions. Keeping notes about student activities allow them to design informed teaching methodology for each child. Record keeping lesson progress for each child is a common activity for teachers. Sometimes the frequency of repeating the same lesson is also noted by teachers, as they deem repetition is important.

A teacher spends most of her day giving lessons to individual students. When not giving lessons, they observe students and write down these observations in notebooks. The aim is to aggregate these observations for each student to plan individual curriculum for them. Thus many times teachers would like to observe how a child is performing a lesson right after they explain the lesson to him/her. Typically, teachers usually shift their roles several times a day from actively giving lessons and helping students to passively observing them. Sometimes, between two teachers in a classroom, one dedicates her time on presenting lessons, and the other observes and takes notes all day.

2.1.2.3 *Challenges*

When asked "what is the most difficult thing to keep track of," all teachers said that it is hard to glean the full picture for the observation data that recorded different moments in the classroom. The only way to track a child right now is to observe a single child all day. Hence, they are usually forced to decide on one particular child who seems to have trouble picking up lessons or exhibits a specific behavior.

Additionally, it is also quite difficult to find time to go back to the notes and read them, or to even find notes about a specific child or certain event. Teachers need to report progress to parents periodically, and the observation notes serve an important role in such reports. Compiling such a report is also troublesome from handwritten notes. It seemed that teachers are currently using their notes within a time window of a few days. The notes and associated observations remain fresh in their minds only within that period, hence their planning activities remain bounded to a week or two at maximum. They also did not seem too interested to digitize their notes as it would be a painstaking task that would not solve all of the existing challenges.

2.2 Manual Data Collection Experiment

Before we started designing a technological solution for the existing observation challenges in Montessori schools, we decided to gauge teachers' interest in collecting lesson progress and social interaction data by giving them some manual observation cards. The goal was to collect their feedback after a few days to see if such data would be useful for them in their daily planning activities.

2.2.1 Observation Cards

The following observation cards were given to teachers. They were asked to fill them up every 20 minutes during the active class period.

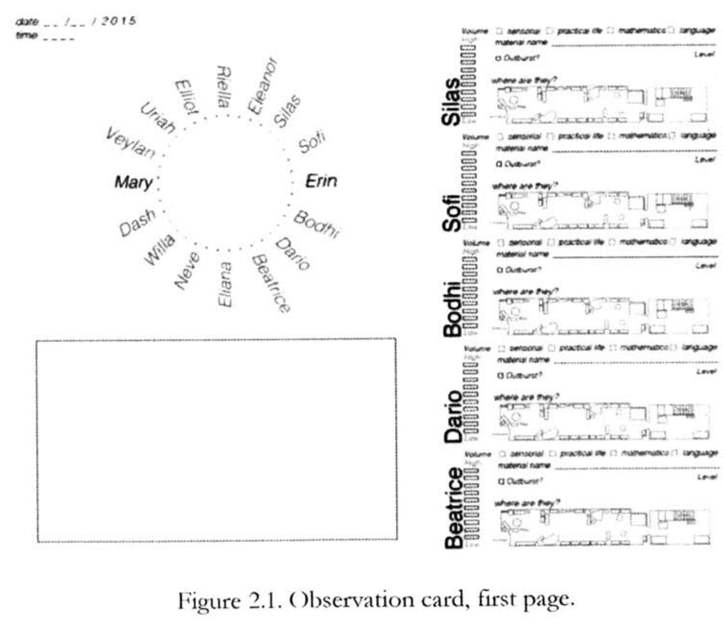

Figure 2.1. Observation card, first page.

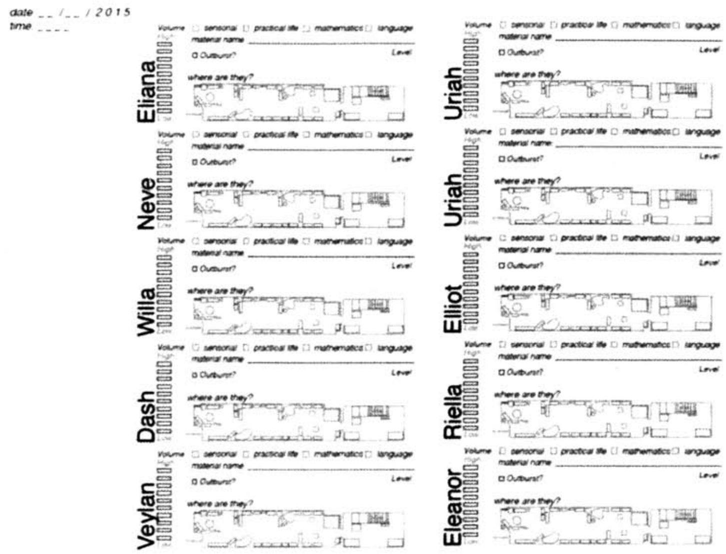

Figure 2.2. Observation card, second page.

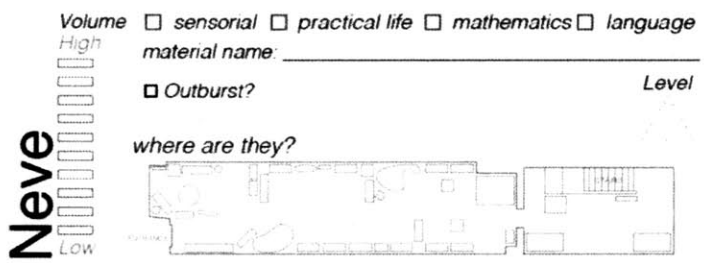

Figure 2.3. A closer look at the observation card for a particular child.

Appendix B shows some filled up cards we collected from the schools.

Teachers filled up these cards over a period of a week. We transcribed and aggregated this data, and provided them a basic mockup of some visualizations that we could potentially build using such data. The responses were quite positive, and formed the inspiration and intended use cases of Sensei. We will present a full use case interview of the Sensei system in Chapter 8.

3 System Overview

We designed Sensei to augment teachers' observations and solve some of the challenges faced by teachers currently. In this section, we present the description of the pipeline and some of the intended use cases of this system.

3.1 Core Concept

The core concept of our system relies on a wearable sensor board equipped with a Bluetooth radio that is capable of creating a mesh network with other sensors. The perceived signal strength value (RSSI) of each data packet received from other nodes in this network can be used to create an approximation of social proximity. These nodes will be placed in children's shoes, and the microcontroller in the sensor board will log the proximity data of other shoes present within a certain radius of the individual child (Figure 3.1). This will provide us data about group work and social interactions. If teachers carry these wearable boards, then we can also get data about how much time a teacher spends with each child over a day.

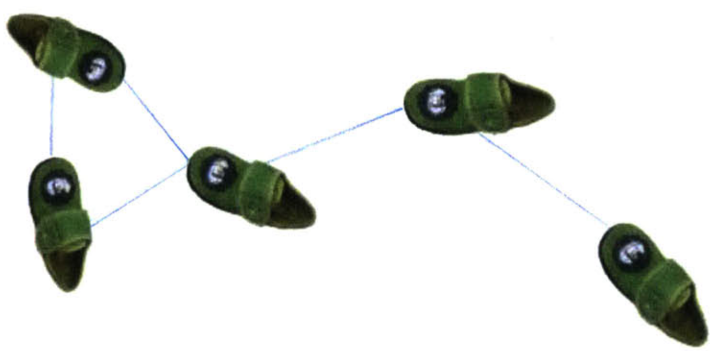

Figure 3.1. Shoes create a dynamic and physical range based mesh network between themselves.

The same idea can be used to track lesson progress for the lessons that are housed in lesson trays. A similar sensor board attached to a lesson tray can be used to receive pings from nearby shoe sensors, measuring the approximate time a student spends near a lesson tray. To avoid scanning for pings when the trays are idle on the shelves, we modify the design to include motion based activation, so the trays only start scanning when they are taken out from the shelves.

The lesson tracker idea can also be extended so that we place them on shelves, and continuously monitor the shelves to know which students spend how long in the vicinity of a shelf.

3.2 System Overview

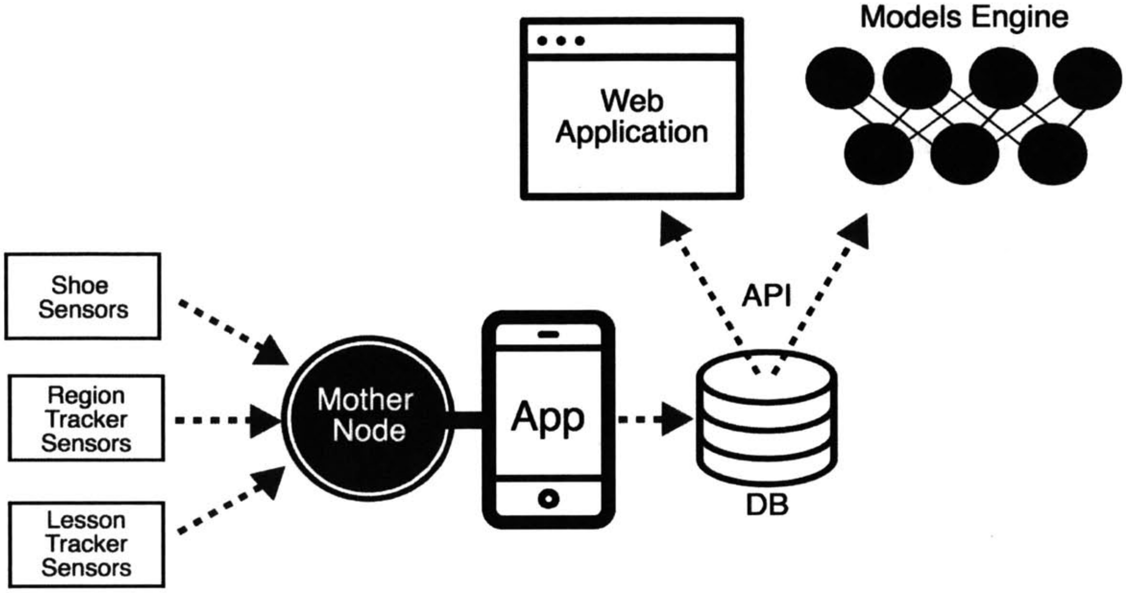

Figure 3.2. Sensei system overview

Figure 3.2 shows the components of Sensei. Shoe sensors log social proximity, region trackers are used to track which shelves are visited more, and lesson trackers are attached to lesson trays and scan for nearby shoe sensors to log how much time children spend working on particular lessons. The logged data from each of these sensor boards can be collected through a mother node attached to an app. The app visualizes some preliminary data and also uploads all the data to a server. This forms the backbone of a web visualization dashboard available to teachers. The data can also be used to study longer term patterns in education using statistical models.

3.3 Intended Use Cases

Teachers may expect to use Sensei in several ways to augment their classroom observations. Some of the use cases are:

- Measuring a child's progress through the curriculum.
- Sharing such progress evaluation with parents.
- Identifying students in need of more teacher time.
- Identifying lessons that all students have mastered and can be removed.
- Understanding group work and also the evolution of social groups among students.

3.4 Potential Pitfalls

Proximity does not always mean interaction. Proximity among children may not mean they are working together or chatting. Proximity of a student to a lesson tray that is taken out from the shelf does not guarantee that she/he is working on the lesson actively. However, our aim in designing this system is to approximate educational interactions. Tracking focus, concentration, and real interactions without interrupting natural interactions in the classroom is left as future work.

4 Hardware Design

Sensei has several sensor boards that were designed over a few iterations. Some of these iterations were deployed in the classrooms, and we received intermediate feedback from teachers to improve, and at times entirely change our designs. In this section, we describe the design decisions and the constraints imposed by the classroom environment, and an overview of the hardware developed during major iterations.

4.1 Design Decisions

Initially, instead of sensors that can be embedded in shoes, we planned to design a badge or a wristband that would house the same proximity tracking circuitry. Some teachers during a follow up meeting suggested that badges or wristbands may distract students. After some brainstorming, we decided that the school shoe could be a possible candidate for placing a sensor board in a minimally invasive fashion. Children come in the morning and change their shoes to wear a particular kind of shoes for the rest of the day. At the end of the day, they leave the shoes to a shoe rack before leaving school.

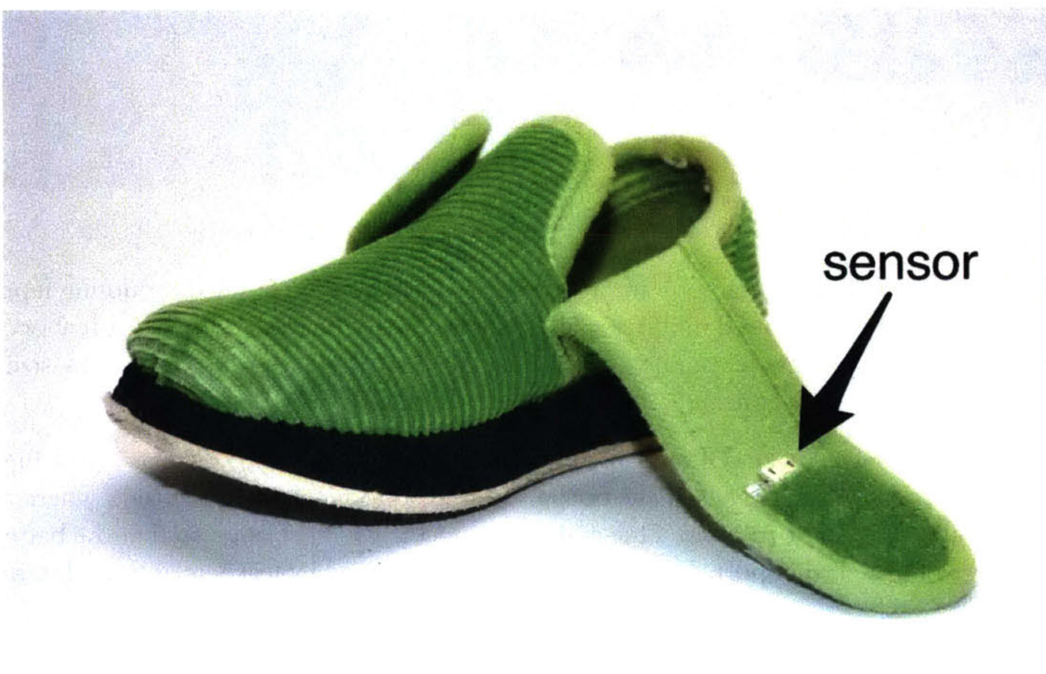

Figure 4.1 Modified shoe that can house a 2cm x 2.5cm PCB in a minimally invasive fashion.

In our final pilot deployment, each shoe was modified by cutting the velcro seam (Figure 4.1), creating a pocket to insert the sensor PCB in a minimally invasive fashion. This worked well during the deployment, with no reported incidents of student distractions or diversions from the teachers (Chapter 8).

Not all Montessori lesson materials are housed in trays. However, a good portion of the lessons are. We decided to focus on tracking those particular lessons. Initially, we created lesson feet that can house a 5cm x 1.5cm PCB, which is the size of our lesson tracker, and a 3.7V 1000mAh capacity

Lithium-Polymer (Li-Po) battery. Figure 4.2 shows four lesson feet that can be attached to a lesson tray on four corners, one of the feet carrying the PCB and battery.

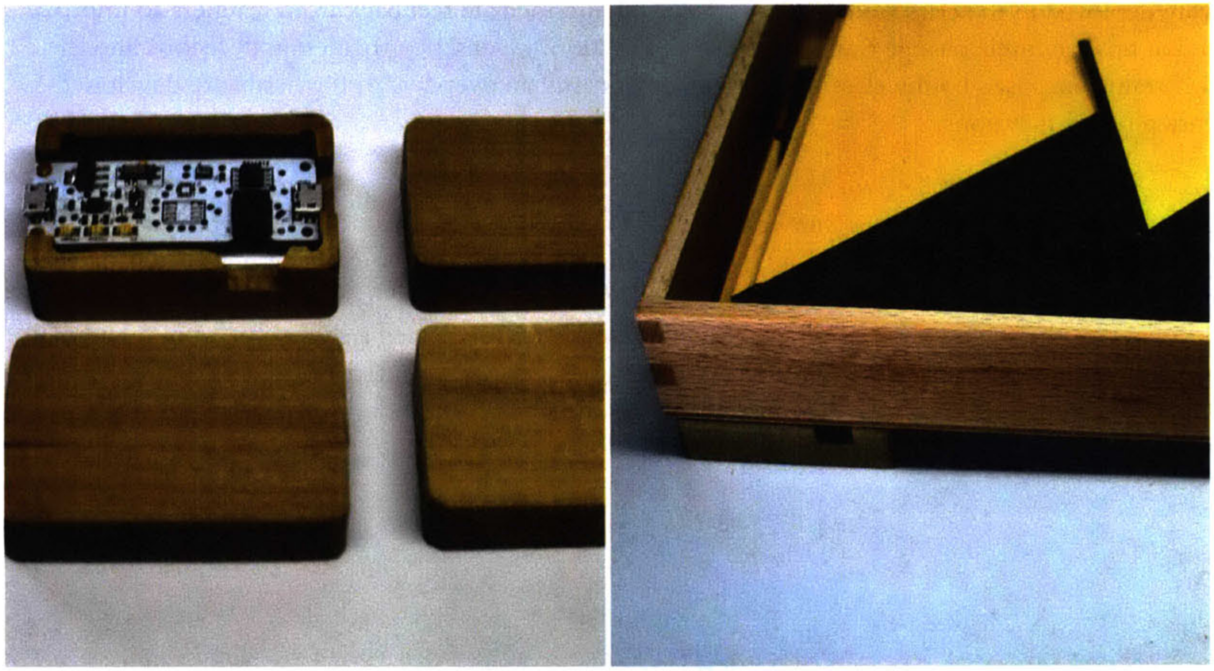

Figure 4.2. Lesson tray feet that were deployed during initial pilot studies.

This design was not particularly famous among teachers. We used double sided mounting tapes to mount these feet, and this required modifying all existing trays in the classroom. On a few occasions (during our initial pilots), a few feet came off some lesson trays. Also, lesson trays vary by size. Some small sized trays did not look appealing to teachers with big lesson feet.

In this case, we decided to make our own lesson trays after consulting with some teachers. Figure 4.3 shows the new lesson trays. These have pockets on the bottom corner of the trays where the PCB and the battery could be inserted. One end of the pocket was kept open so that the battery could be charged through that port. This design was better received by teachers and used during our final pilot study.

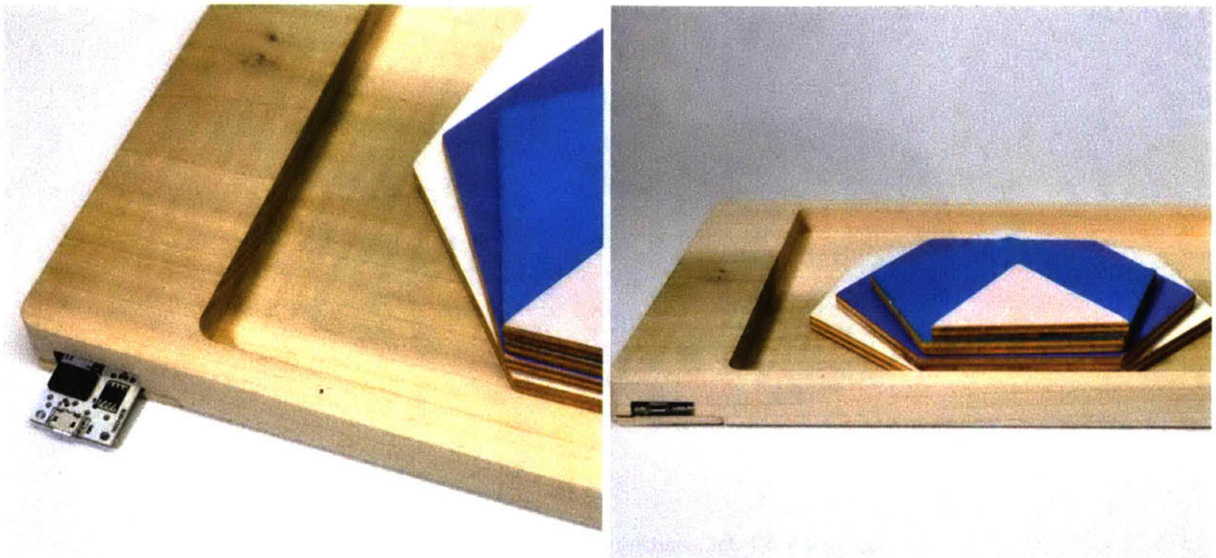

Figure 4.3. Newly designed lesson trays that house the lesson tracker PCB with a battery inside the pocket.

These were the major design decisions that we had to make during our hardware development process. Some of the other decisions and improvements were solely based on low power optimization. This will be appropriately described in the subsequent sections.

4.2 Design Iterations

4.2.1 Shoe Sensor

RFduino Prototype: The first prototype was designed with the RFduino module [10]. The RFduino is a 1.5cm x 1.5cm Bluetooth module with an embedded antenna. Our prototype had a SPDT switch (to turn scanning on and off) and a coin cell battery holder at the back.

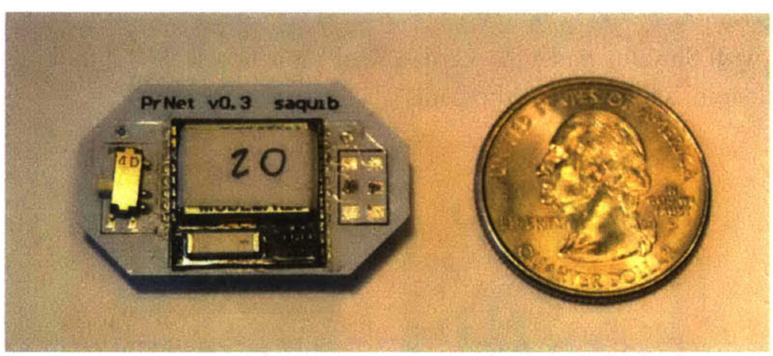

Figure 4.4. First PCB prototype with RFduino module.

The PCB was 3cm x 2cm in size. They were initially deployed in the classrooms by stitching a badge-like pocket above the shoe.

Figure 4.5. RFduino prototype on shoe (without cover).

Even though this worked well initially, we started noticing drops in our Bluetooth ping and RSSI data during initial pilots. We also eventually required mesh networking in order to scale up from the star network topology of Bluetooth to a larger classroom. Hence we decided to change our platform from RFduino and started researching mesh networking SoC (System-on-Chip) modules. Our research led us to an nRF51 based module called Simblee. This led to our next series of iterations.

Simblee Prototypes: Simblee [11] is a 1cm x 0.7cm QFN module that has a smaller form factor compared to RFduino, 26 GPIO pins (so we can accommodate more sensing modules), and an embedded antenna. Our experiments suggested that the RSSI fall off with respect to distance is smoother for Simblee compared to RFduino. Moreover, Simblee offers a mesh networking protocol (SimbleeCOM) to the developers, which would be useful for scaling to a large number of units in the classroom.

The first prototype with Simblee had a SD card, a Real Time Clock (RTC) module, a kxtc209 accelerometer, and a coin cell battery at the back.

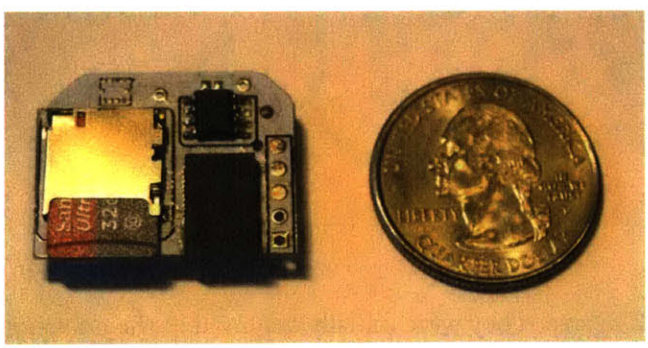

Figure 4.6. Simblee shoe sensor with SD card and RTC.

Our initial Simblee prototype looked a little different. We did not have the RTC in the beginning. However, time syncing issues arose with the internal 32.768 kHz clock when we tested a mesh network with 20 units. The sleeping cycles to conserve battery would add a variable offset value to the timer when using the internal clock, hence all units in the mesh network would be off sync within a few minutes of powering up. To tackle the issue, we added a DS3231M RTC module to our design. It fires an accurate interrupt every second, which allows the Simblee to sync its time even while sleeping.

Initially we planned that the proximity data will be logged in the SD card. The SD card required SPI communication from the Simblee module at a high rate of current (90 mA). The 20mm CR2032 coin cell battery we used had a capacity of 225 mAh, and using the SD card was not a good idea in terms of saving battery life.

In our next iteration (Figure 4.7), we discarded the SD card and relied on the internally available 128KB flash ROM memory in the Simblee. We compressed the data by using some bit shifting techniques (more on this in the Firmware section), and the memory proved to be adequate for a day worth data. We also changed our accelerometer to ADXL337 analog accelerometer module since it had better energy saving features. A few 0603 0.01uF SMD capacitors were added to create low pass filters on the analog three axes outputs of the accelerometer. A mini USB port was added so we could upload our firmware using a custom USB programmer, and also transfer data from the internal ROM through UART communication (serial port).

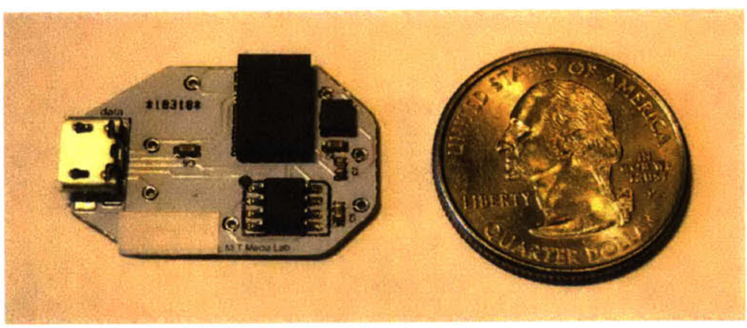

Figure 4.7. Second design iteration with Simblee.

Our final design (Figure 4.8) involved adding some nice-to-have features. We added options for using either a coin cell or a small Lithium-Polymer battery. A small 3.3V voltage regulator was added so we can use the 3.7V Li-Po battery. A MEMS microphone was added to our design so we could collect sound level data if needed.

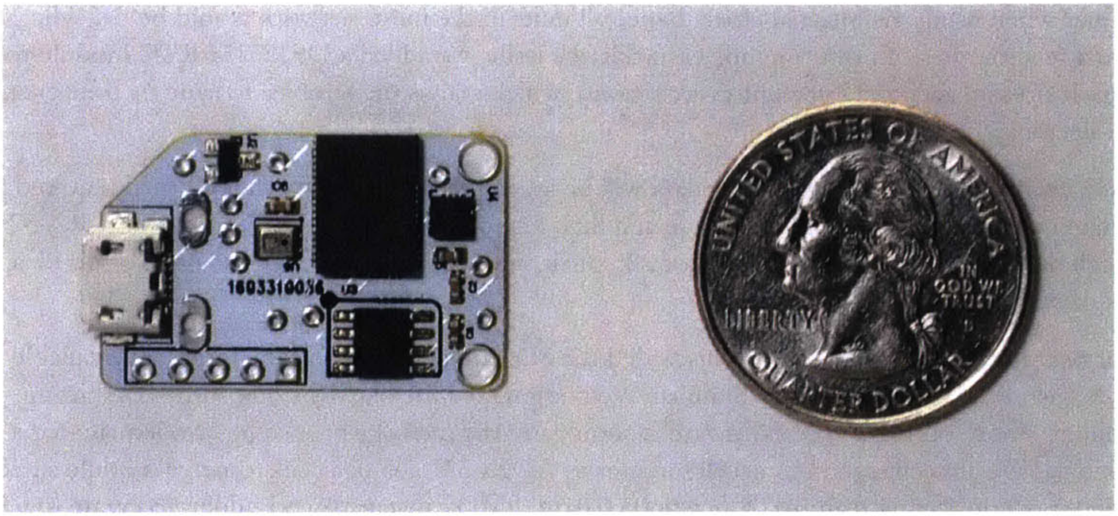

Figure 4.8. Final shoe sensor design.

4.2.2 Lesson/Region Tracker

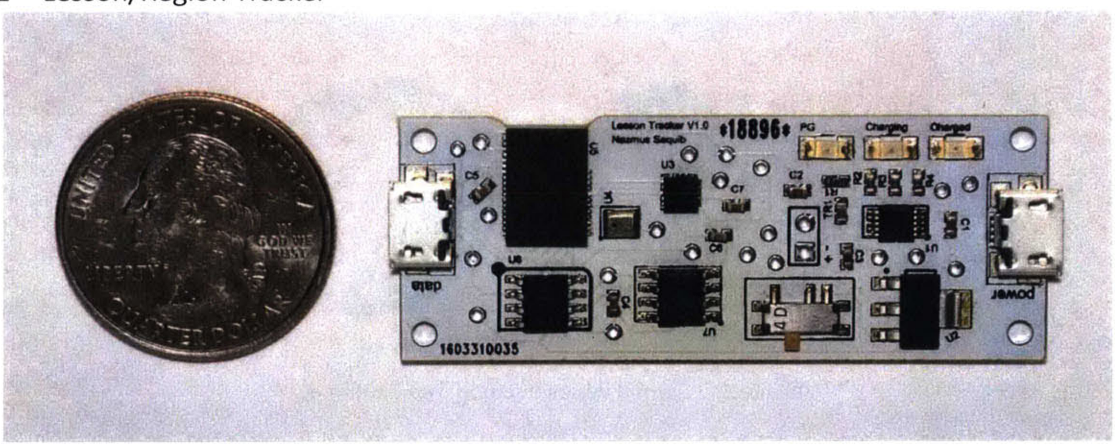

Figure 4.9. Lesson/Region tracker PCB.

After iterating on our designs for the shoe sensor, prototyping the lesson/region tracker board was easier. These could be made considerably bigger (5cm x 1.5cm) compared to the shoe, as there was minimal size constraints with our custom made lesson trays and region tracker enclosures. These boards were similar in size with the 1000 mAh rechargeable Li-Po battery we used, so they fit on the same stack.

Other than the Simblee module, RTC, and MEMS microphone, we added an external flash ROM, a power switch, and a whole new section for recharging circuitry. We used Texas Instrument's MCP73833 Li-Po charging IC for charge and power management. A nice feature of this chip helps us control or stop charging instantaneously as the circuit heats up. A 10k thermistor helps us control the charging rate in this way. The power management section allows us to connect the circuit to a 5V USB wall power supply for charging, and also simultaneously provide power to the sensing section.

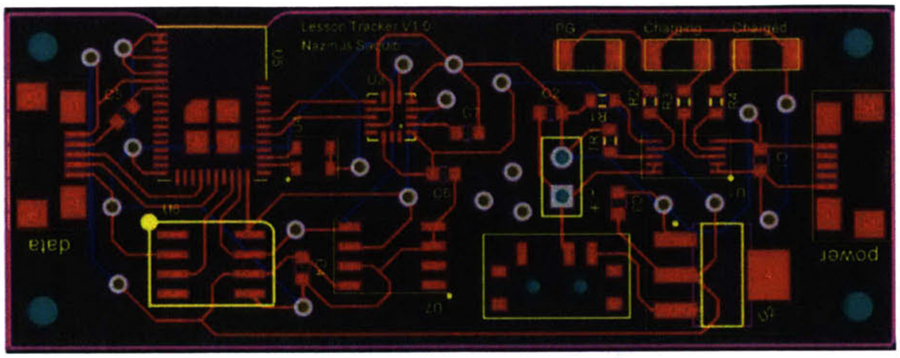

Figure 4.10. PCB trace routing for the lesson tracker board.

We also added a new accelerometer MMA8451 (from NXP Semiconductors) that has motion activation capabilities. This is utilized in our firmware code to carry on minimal operations under a deep sleep state until an interrupt is fired by the accelerometer in the event of a big shake or movement. This feature is used in lesson trays to conserve battery and scan for other units only when a tray is taken out from a shelf.

The region tracker (that sits on shelves to track shoes that are within 3 – 4 feet) has a custom enclosure that we designed and built. It houses the same circuit and the battery, and can be attached under shelf racks using mounting tapes. It is made by cutting and carving wood in MIT Media Lab's machine shop. The choice of material was inspired by teachers' recommendation to camouflage our equipment in the classroom environment.

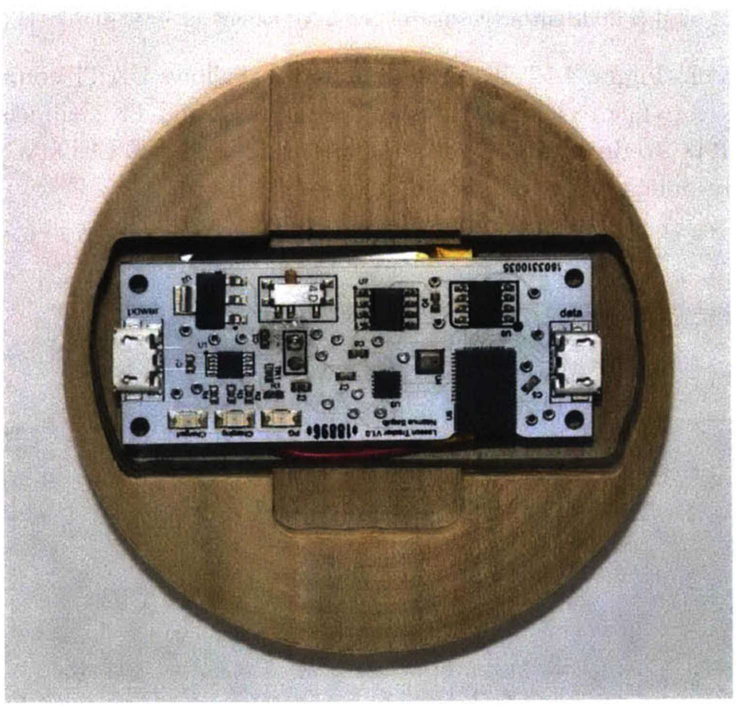

Figure 4.11. Region Tracker enclosure made from wood.

4.2.3 Mother Node

To transfer data to a smartphone application, we designed and manufactured a mother node that can be attached to a smartphone with OTG (On The Go) host capabilities. This allows all the sensor boards (shoes, lesson tracker, and region tracker) to send and receive commands from a smartphone when within Simblee radio's maximum range, and transfer data from their flash ROM to the smartphone.

Figure 4.12. The mother node enables a smartphone to communicate to all sensors in the classroom.

The mother node circuit (Figure 4.13) has an FTDI chip that allows UART communication to a host machine. It is used to talk to its host through the serial RX and TX channels of a Simblee module. There is an RTC to drive time synced operations and an RGB LED (WS2812) to give the user a sense of current operations performed by the node.

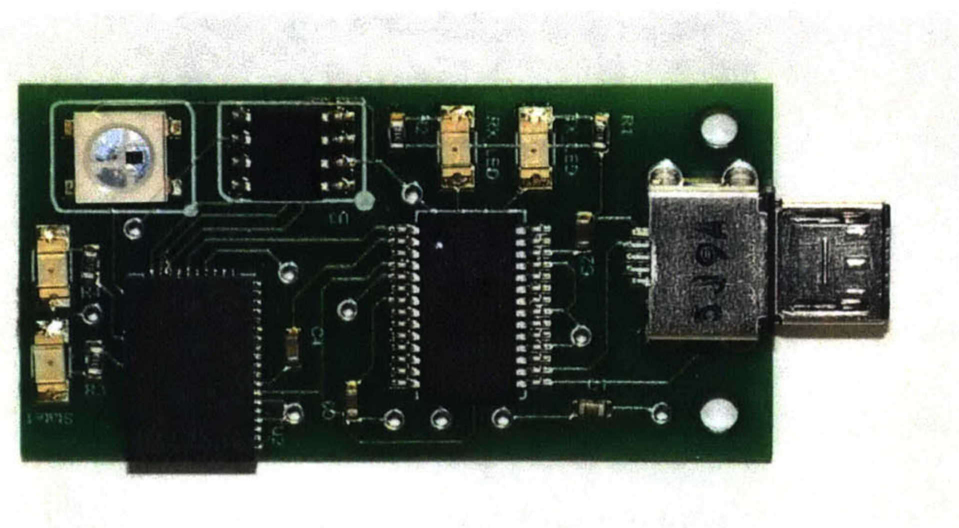

Figure 4.13. Mother node attachable to a smartphone through USB.

4.2.4 Programmers

Programmers to upload firmware code from the computer to the sensor boards also evolved as the project evolved over different design iterations. Initially, we used an FTDI based commercial programmer for RFduino.

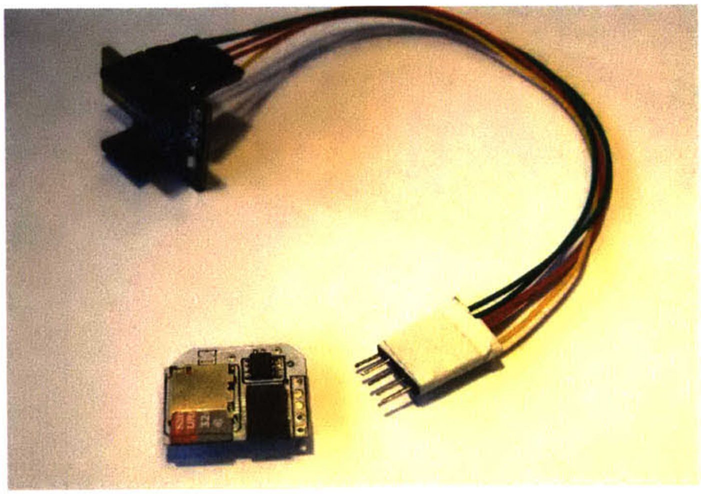

Figure 4.14. Initial programmer solution.

However, we soon moved onto adding USB ports in our sensor boards, and designed and manufactured a custom programmer for this purpose with an FTDI chip. We can upload firmware both from a computer or a smartphone in this case.

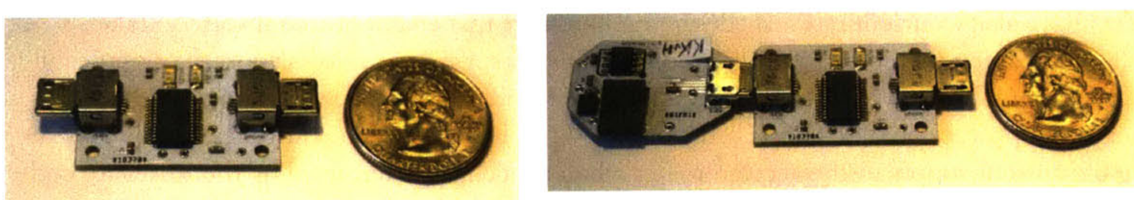

Figure 4.15. FTDI based custom programmer.

This proved to be particularly useful when we needed to rapidly test and upload programs to sensor boards during field trips. The loose pin contacts in our previous versions would make it hard to program sensors otherwise.

4.2.5 Recharging Strip

At the end of the day, the teachers can recharge the sensors that have run low in battery. We have designed a 5V USB wall port based recharging strip where shoe sensors can be plugged in. This is an extension of the MCP73833 based recharging section we designed for the lesson/region tracker circuits, but it has six ports for simultaneous recharging.

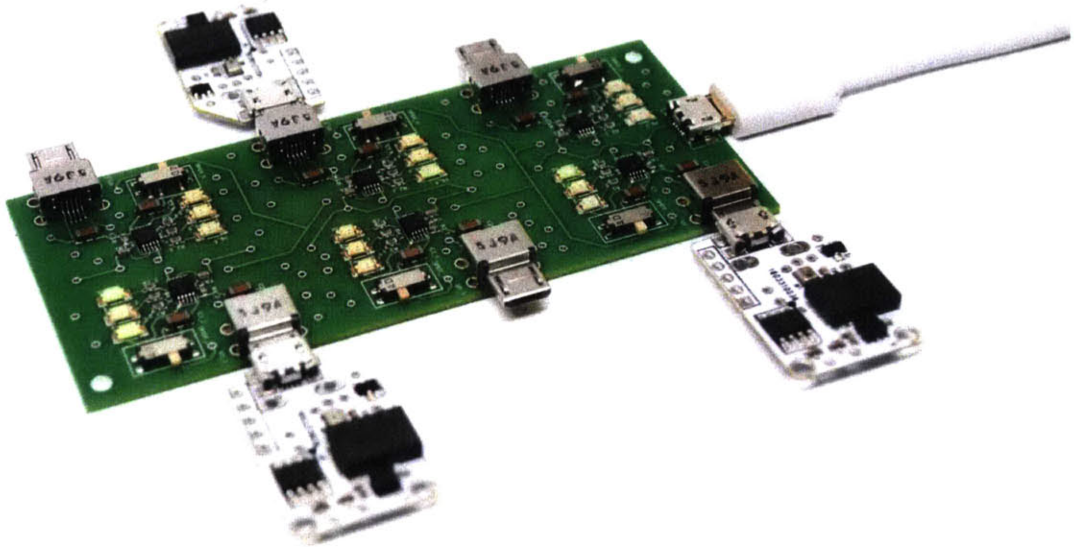

Figure 4.16. USB recharging strip that can recharge six shoe sensors at a time.

4.3 Low Power Considerations

Throughout the design iterations, one of our main focus was optimizing for battery conservation. The modules used in our sensor boards were carefully chosen after researching competing modules' datasheets for leakage current rates. By designing hardware that enables the firmware to take advantage of sleeping the microcontroller most of the times, we were able to conserve significant amount of battery. For example, the ADXL337 accelerometer was chosen based on its price and also because its power could be controlled from a PWM pin of the microcontroller, completely shutting off current supply to the accelerometer when needed, as determined by the firmware code. We moved away from a SD card solution to ROM memory for logging data over time, to conserve battery. There are several tricks employed in the firmware design that also helps ensure longer battery life.

5 Firmware Development

Sensei has two main firmware components. In this section, we will describe the design behind these firmware algorithms.

- A distributed sensor network protocol to sense proximity and establish social interaction context by logging motion and ambient sound level data.
- A time synchronized network event scheduling scheme that allows efficient battery consumption, while maintaining a proximity sampling rate of 10 seconds to distinguish between ephemeral contacts and longer social interactions.

5.1 Proximity Sensing Protocol

The Simblee transceiver radio provides a suit of firmware for different networking protocols. We utilize the SimbleeCOM mesh networking protocol for communicating between the sensors. Under this protocol, each sensor can engage with another sensor for a duration of three milliseconds to exchange a packet of data and an acknowledgement.

In our proximity sensing protocol, every 10th second each sensor in the network repeatedly broadcasts a data packet containing a unique identifier, for 500 milliseconds. The Simblee module has only one radio chip for broadcasting and receiving. It cannot do both simultaneously. Thus, between each broadcast, we insert a random delay so we enable the radio to listen to others who are broadcasting. Other than the 10th second, we program the microcontroller to enter a deep sleep mode for 9 seconds, conserving battery in the process.

After collecting data packets from all sensors who were within range (range is defined by the transmission power of the radio), each node records an entry for each device that sent a packet. As described in the hardware design section, we have to rely on a constrained memory space to save our data with low power consumption. Thus, each entry for a node only takes 4 bytes of space, and we compress our data to fit in that space.

Each proximity entry is a 4 byte integer of the form DDRRHHMMSS, where DD is the device identifier, R is the mean RSSI of all packets received from the particular sender DD during that one second period, and H, M, and S are the hour, minute, and second components of the timestamp respectively. We use bit shifting techniques to compress all of this data to this 4 byte memory. The resulting data set is a time series identifying when a student is in proximity to another student or to another region/lesson tracker sensor.

5.1.1 Packet Reception Evaluation

In our experiments, we found that a random sized delay between 5ms and 50ms works well for maximizing the number of packets received during the proximity sensing period (every 10 seconds). However, there are times, even with the random delay, when two or more radios within range would be simultaneously broadcasting and not listening. In those cases, the proximity data would be lost between these devices due to collisions.

To evaluate how well our protocol captures proximity with other nodes in range, we conducted some laboratory experiments to check packet reception success. Five shoe sensors were placed within a range of 3 – 4 feet (our estimated social proximity range in these schools), all of them placed inside shoes. Region trackers were also placed to measure how well they capture proximity of these shoe sensors. The shoe sensors ran for 12 hours and sent 82, 680 packets altogether. The region trackers ran for a total of 25 hours, and collected 43, 310 packets.

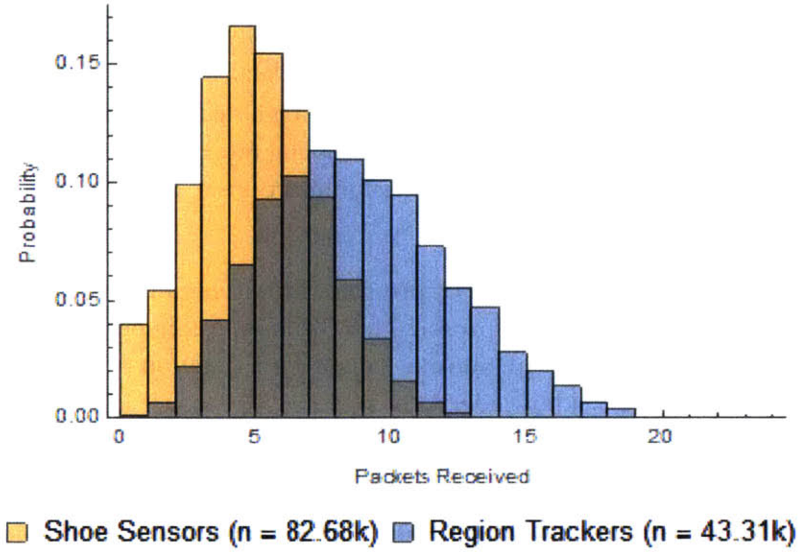

Figure 5.1. Packet reception statistics during our protocol robustness study.

The mean number of packets received by the shoe sensors was 4.6, during each 500ms window every 10th second. The region trackers collected 7.6 packets on average. The probability of detecting other shoes by a shoe sensor can be defined as the probability that at least one packet was received from each shoe in that 500ms time window. Using this metric, shoe sensors have a probability of 95.98% to detect each other when in range. For region trackers, the probability is 99.75%. Since region trackers do not broadcast but just listen for broadcasts from shoes, the probability of success is higher.

5.2 Synchronized Network Event Scheduling Scheme

As mentioned before, all nodes must have their radios active for 500ms period of time in every 10th second in order to measure proximity with each other. Sleeping for 9 seconds and waking up at the same instant is a challenging task if we consider many nodes in the network. The clock capabilities provided by the Simblee module is decent. However, for a millisecond level precision task, this does not suit our needs. An easier solution is to keep the transceiver radio active for a longer amount of time, but this consumes battery at a rate that does not allow us to collect proximity data for the duration of a full day.

Microcontroller and Transceiver State	Current draw (mA)
Active microcontroller	4.1
Active transceiver	12 – 19.8
RTC interrupt callback (background process)	0.2
Sleep (with background processes on)	0.45

Sleep (clock and most processes off, except interrupt checking)	0.005

When the microcontroller is active, it spends 4.1 mA. When the transceiver radio is active, the current draw is a few times higher. An interrupt trigger wakes the microcontroller up for a few milliseconds, consuming 0.2 mA. The only effective sleep state that ensures a longer battery life (enough to go on for a day or two) requires most background processes and the clock to be turned off.

The DS3231M RTC module provides accurate timekeeping capabilities with ±5 ppm accuracy. An external 1 Hz interrupt allows us to accomplish an event scheduling scheme for each shoe sensor, so they wake up at the same time (with some inherent offset of a few milliseconds, which does not have any significant effect for the 500 ms time window), communicate with each other for proximity mapping and read values from the accelerometer and the microphone. Figure 5.2 shows the event schedule in details.

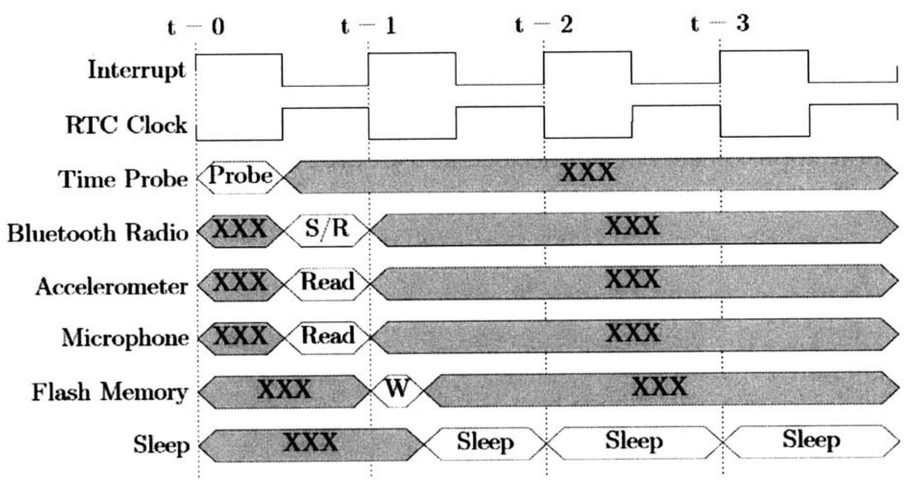

Figure 5.2. Synchronized event scheduling scheme that occurs in each shoe sensor.

At the rising edge of the interrupt, the MCU wakes up and the software clock updates itself from the external clock by probing the accurate time from the RTC. The interrupt is a 500 ms long square wave, during which the probing happens. In the remaining 500 ms window, the radio sends and receives pings using random delays to accommodate these operations. The random delay is between 5 ms to 50 ms, and the random number seed is chosen by probing for an unattended analog input pin for hardware generated random noise. During the proximity mapping operations, the accelerometer and microphones are sampled for analog values too. At the end of the 500 ms window, the proximity and sensor data are written to the flash memory (W), and the MCU goes back to the deep sleep mode. When new interrupts are fired every second, the MCU checks the software clock (0.2 mA) and go back to sleep if it is not the next 10th second yet (0.005 mA). This scheme allows maximum battery conservation (around 2 days of battery life with 80 mAh rechargeable Li-Po battery) while maintaining a decent proximity mapping rate of six probes per minute.

5.2.1 Transmission Power Level Tuning

The transceiver radio can transmit data packets at higher ranges at the cost of more current expenditure. For our purpose, I tuned the radio such that the effective range is only 3 – 4 feet, the social proximity range we have agreed on after manual observation of video feeds of interactions. Any RSSI trace means that the shoe is within the social proximity range.

Transmission Power Level (dBm)	Approximate Range (m)	Current Expenditure (mA)
4	40	19.8
0	25	14.6
-4	12	12.5
-8	4	11.5
-12	1	10.8
-16	0.2	10.3
-20	0.07	10

The approximate range is measured manually in laboratory setting by keeping one shoe sensor fixed in place and taking another one away from it until the signal is lost. Based on the current expenditure and the approximate range, I decided to use -12 as the appropriate power level for the transceiver radio.

6 Smartphone Application and Web Dashboard

A smartphone application facilitates data collection by teachers. A web dashboard enables them to look at visualized data and understand classroom dynamics in details. In this section, I describe the main components of the smartphone application and web dashboard, and how they serve the intended use cases of Sensei. These components in the Sensei pipeline were designed by me and developed in collaboration with my research group members. The student and teacher names are anonymized in these visualization demos for the purpose of presenting the thesis, but teachers see real names in their dashboard.

6.1 Smartphone Application

The app allows teachers to collect data at any time of the day, start the shoe sensors by sending a start command (otherwise the sensors sleep without receiving that command every day), visualize how much time they have spent with children for a particular day, and also view an estimated battery life statistics for the sensors present in the classroom to make plans for recharging the units.

The smartphone app utilizes the Physicaloid library [12] for Android to communicate to the mother node's FTDI chip using the UART protocol. By sending and receiving messages with appropriate header elements, the phone and the Simblee module in the mother node communicate through the FTDI chip.

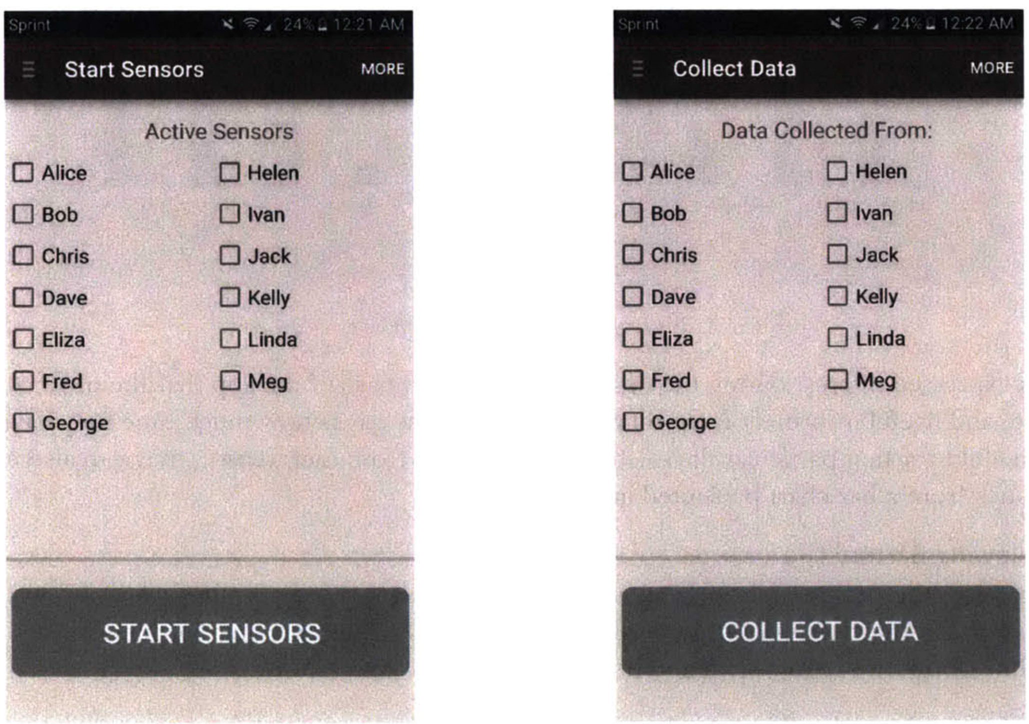

Figure 6.1. Starting sensors and collecting data screens from the app.

6.1.1 Starting Sensor Network

At the beginning of each day, teachers activate the network in the classroom with a single button press. The mother node needs to be attached to the smartphone in order to send this command to the rest of the nodes in the network. A real time view helps the teacher understand which sensors have acknowledged the receipt of the command.

6.1.2 Collecting Data

Typically, at the end of the day the teachers would be expected to collect data from the shoe sensors. A similar view shows the teachers if each sensor finished reporting their data to the mother node. The wireless protocol to transfer data remotely was designed by myself and developed by Dwyane George in my research group.

6.1.3 Viewing Time Spent with Children

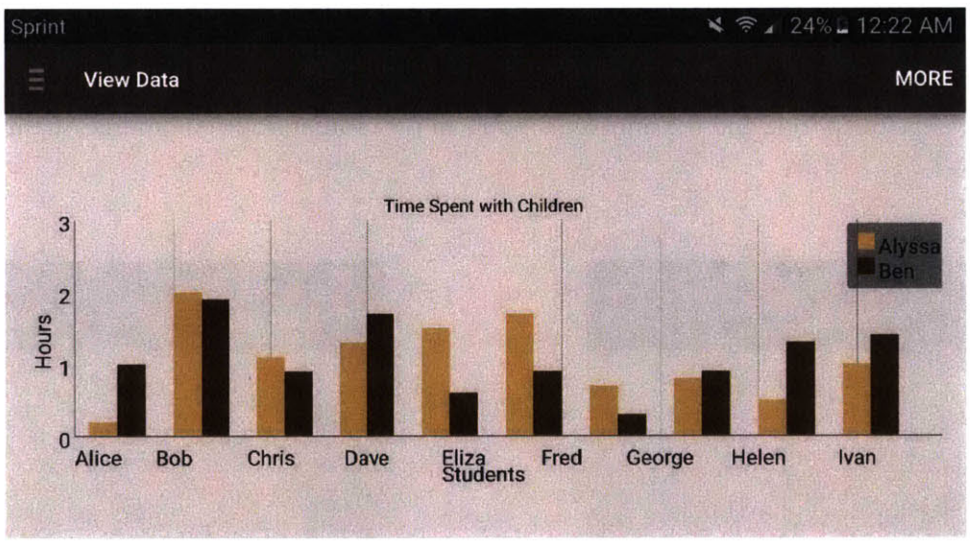

Figure 6.2. Viewing time spent with children over the day.

Teachers expressed interest during the design and mockup phase of the app that the most immediate and useful information for them at the end of the day is how much time they have spent with each child for that particular day. After collecting data from each sensor, they can also check this statistics from a bar chart presented in the app.

6.1.4 Viewing Battery Life

Currently a work in progress, the app allows teachers to view the sensors (shoes, lesson trays and region trackers in the classroom) that need recharging. This is approximated by checking the logged data for how long the sensors slept and how long they were active.

6.2 Web Dashboard

The web dashboard was developed (in collaboration with Ayesha Bose) with an aim to give insights to teachers that are usually lost due to the high frequency of social and lesson interaction going on in a dynamic classroom. When teachers enter the web portal, they are presented with a login screen, and the homepage after login allows them to view essential data visualizations categorized according to the use cases. Visualizations were created using the web engine D3 [13].

6.2.1 Homepage

Teachers get updates for the last few days from an automated summary generated from the weekly data. They can view data about students interaction, their own interaction duration with each student, progress of students on the materials that are on lesson trays, classroom dynamics (locations where students spent their time), and write notes and observations. Each teacher's dashboard are customized based on their classroom floorplan and student updates.

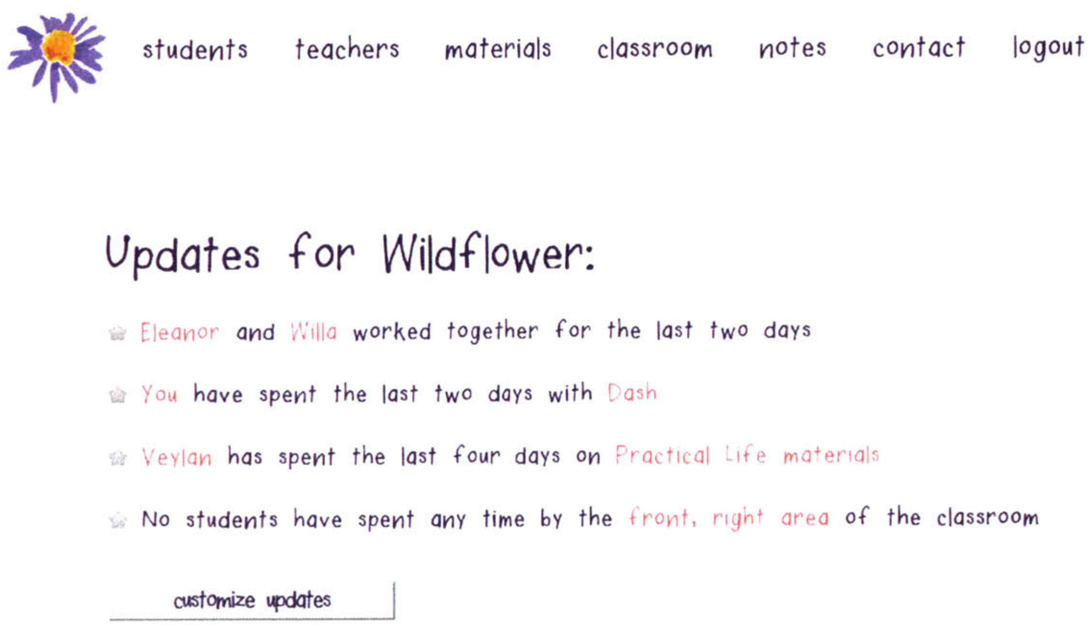

Figure 6.3. Customized homepage of Wildflower web portal.

6.2.2 Social Interaction Chart

Under the student interaction page (Figure 6.4), teachers are presented with a visual interface that allow them to see the social interactions for a full day between students. The stacked bar charts provide an understanding of group formation at different times of the day. They can hover their mouse over the charts to see individual interaction slices and the students associated in the interaction.

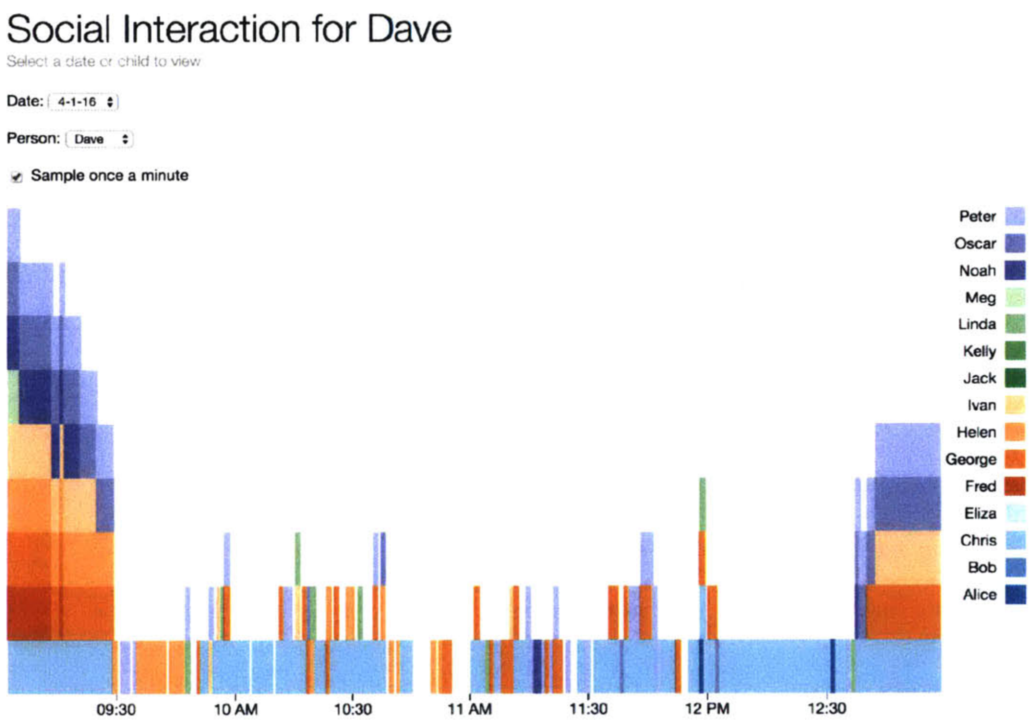

Figure 6.4. Social interaction visualization.

As seen from the particular dataset of Dave, he spent most of his time with Chris on 4-1-2016, and occasionally with Helen. In the beginning and ending periods of the school day, all shoes are put back together in the shoe rack. The visualization clearly captures these events too.

6.2.3 Clustered Interaction Chart

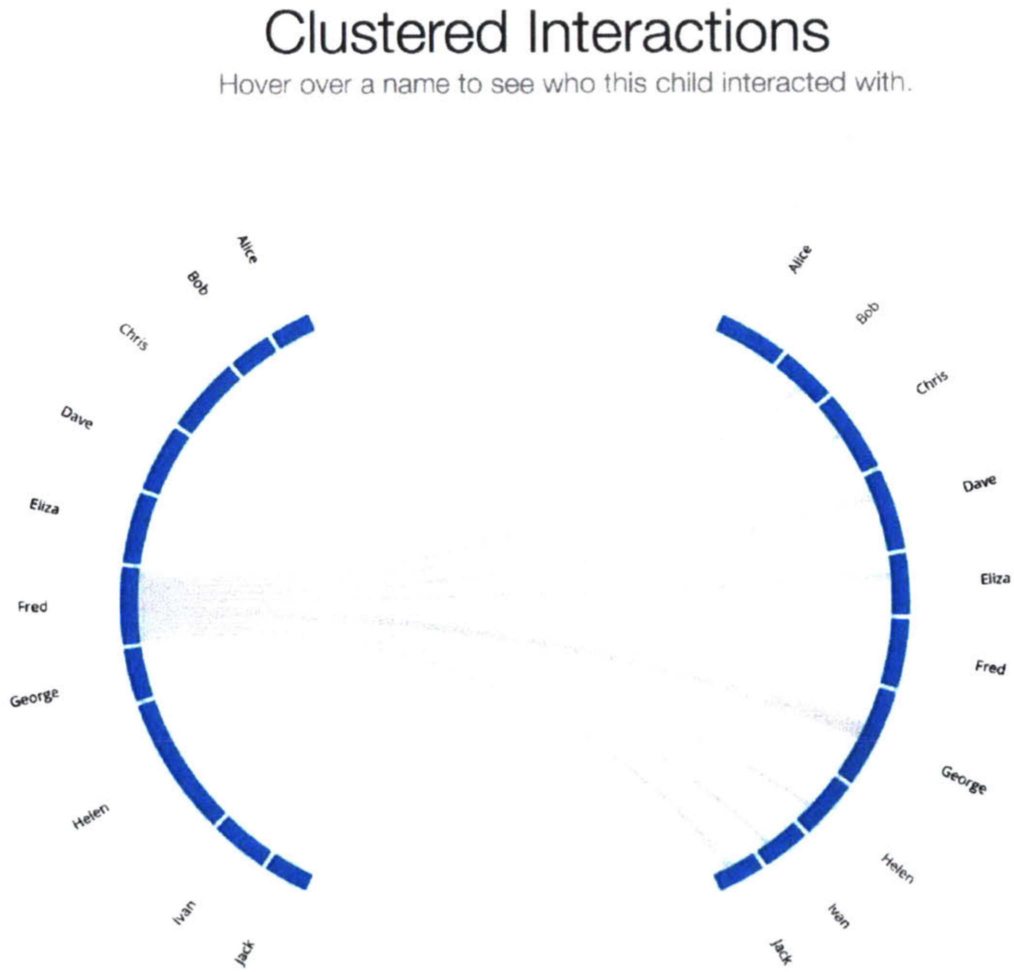

Figure 6.5. Clustered interaction aggregated from a full day of data.

This view allows the teacher to see the aggregated interaction information without the time axis. Teachers can select students by hovering over their names and see who they have worked with for the full duration of the day.

6.2.4 Teacher Interactions Radar Chart

Teachers had an immense curiosity to know how much time they spend with individual students as this information helps them optimize their interaction and lesson giving time. In order to facilitate teachers with this information, we have incorporated a radar chart in our dashboard.

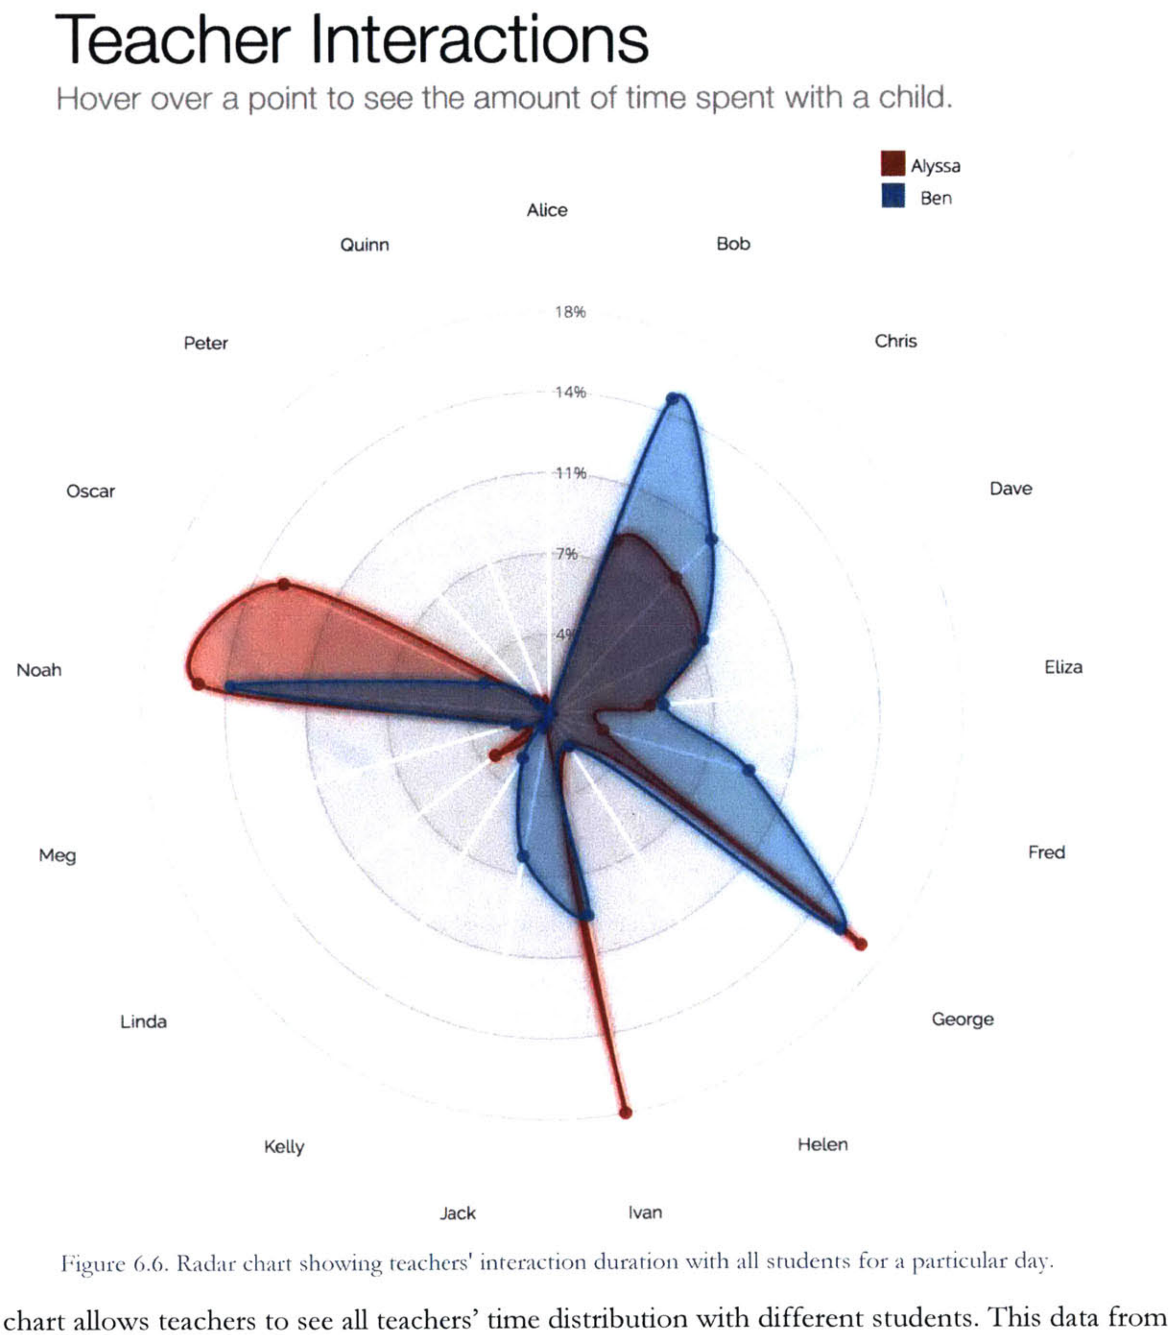

Figure 6.6. Radar chart showing teachers' interaction duration with all students for a particular day.

The chart allows teachers to see all teachers' time distribution with different students. This data from a day in the classroom shows that both teachers spent time with the same students, and other students were independently working on their own, mostly.

6.2.5 Lesson Progress Chart

The proximity mapping data from the lesson trays gives teachers a unique ability to check student's progress with different materials in the classroom. It is a simple bar chart of time spent with particular lessons for particular students.

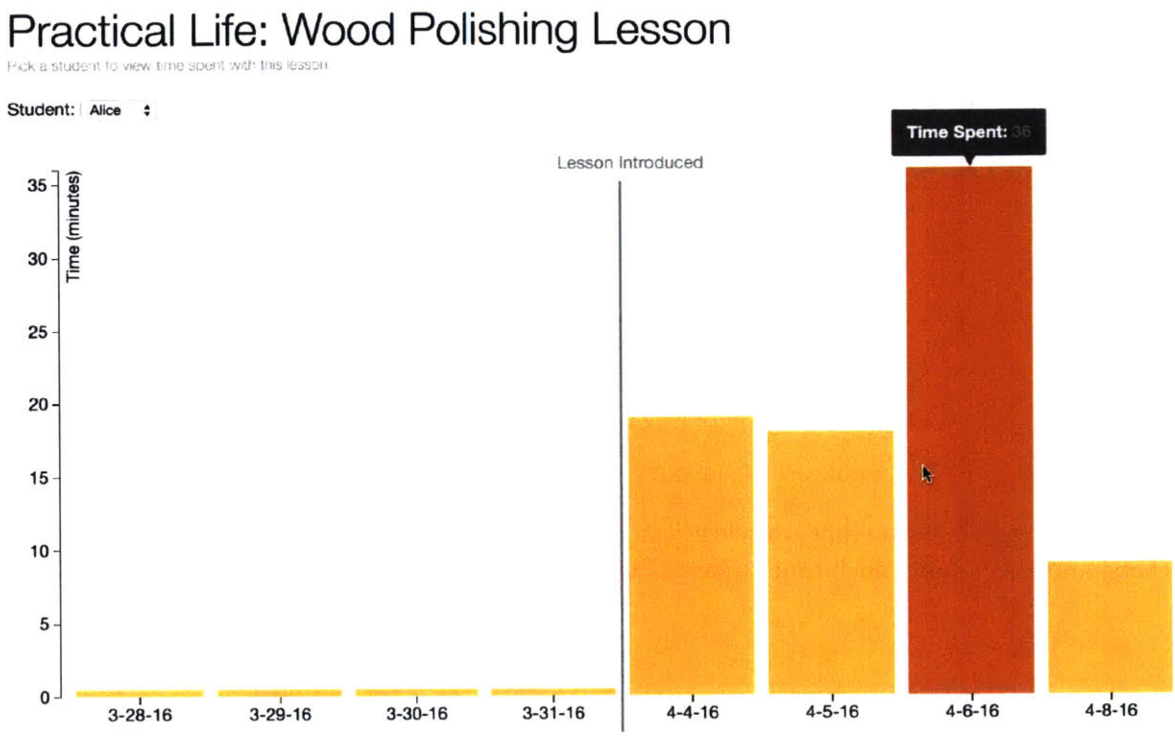

Figure 6.7. Lesson progress chart for individual students.

In this particular case, a teacher can clearly see when the wood polishing lesson was first introduced to Alice, and how long she worked with the lesson over the next few days. This visualization can also inform teachers when a student is ready to be shown a new lesson.

6.2.6 Region and Shelf Activity

Children spend their time in different areas of the classroom. We placed region trackers on different shelves to track how much time students spend near each shelf. The shelves are organized and located according to category of lessons, and students spend time near different shelves before they choose to work on a certain lesson, so this data can illuminate teachers about students' interests in different category of lessons.

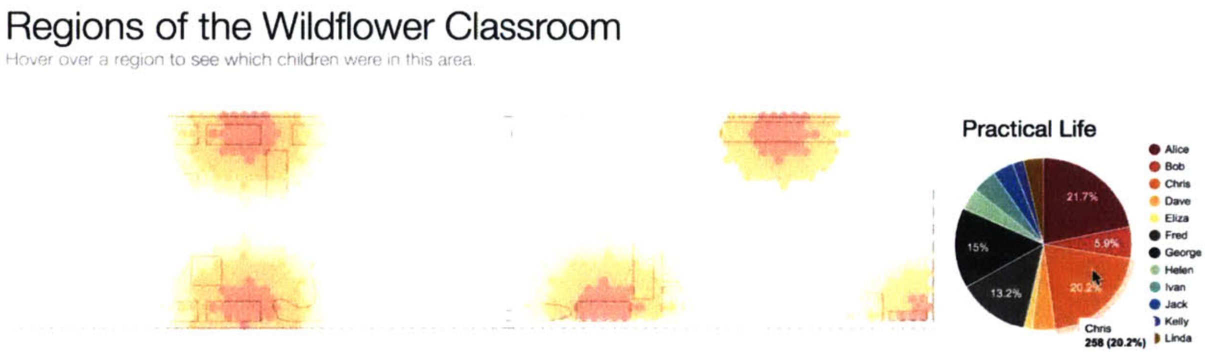

Figure 6.8. Region dynamics over the course of a full day.

This visualization allows teachers to select regions in their classrooms (where they install region trackers) and check how much time students have spent in that area over a day.

6.2.7 Motion Visualization

This visualization allows teachers to understand an aggregated motion profile of students on a particular day.

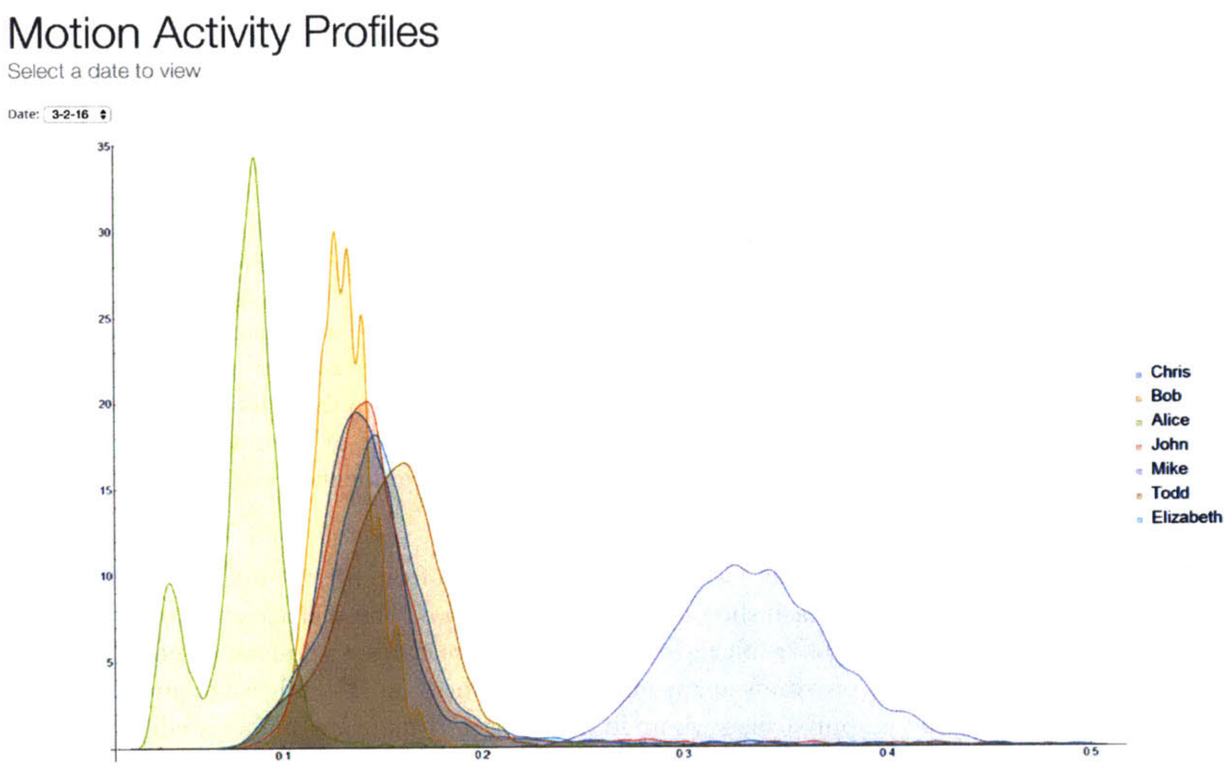

Figure 6.9. Motion activity visualization.

The motion data acquired from the accelerometer can be numerically integrated to find velocity profiles for each student. The above is a histogram of velocity values sampled every 10 seconds for the whole day, filtering stationary values and keeping only the motion values over a small threshold to get rid of noise and inaccuracies due to the numerical integration.

7 Network Dynamics and Evolution

One of the most interesting aspects of Sensei is that it enables researchers, educators, and teachers to analyze longer term learning behavior at scale. Early childhood classrooms have rich learning dynamics, both from social and lesson progress perspectives. The ability to collect fine grained social proximity information with high time resolution (10 seconds) provides us ways to test hypotheses about individual learning, deploy different interventions in the classroom and research their effectiveness, and many other similar research opportunities. In this section, we demonstrate some analysis methods that can be used as tools to research and understand early childhood social proximity data.

7.1 Temporal Contact Network Dynamics

As students and teachers come into the social proximity range of each other, they form a temporal network [14] between themselves. These networks are highly dynamic in nature, are relatively short lived compared to the span of the whole day, and are difficult to observe without Sensei proximity mapping in the classroom. There are several kind of analysis that can be done on such networks to understand community structure, dominant figures and their followers in the classroom, and how the social structure change over time (both during a day and over a scale of days, weeks, and months). These can help researchers understand how social factors influence learning.

7.1.1 Network Formation

In this demonstration, we take one day's worth real data collected from the classroom, and chop the proximity mapping time series of each shoe sensor into hourly bins. This will allow us to form temporal networks on an hourly basis. Next, for a given hour, we create an undirected edge between two shoes if there is a trace of proximity at any instant during the hour. The edges are undirected as there is no concept of who is approaching whom in this dataset. If there are duplicate edges because both shoes saw each other at the same time, we merge both the observations into one edge in the temporal network.

7.1.2 Network Structure Evolution and PageRank

The PageRank algorithm [15, 16] can be used to assign a centrality measure to students and teachers at different hours of the day. There are other centrality measures that can be interesting in this context [17, 18]. We use PageRank (a version of eigenvector centrality) for the purpose of demonstration.

The teachers in this dataset are numbered 10 and 11. The rest are children in a Wildflower Montessori classroom.

Network Structure **PageRank Chart**

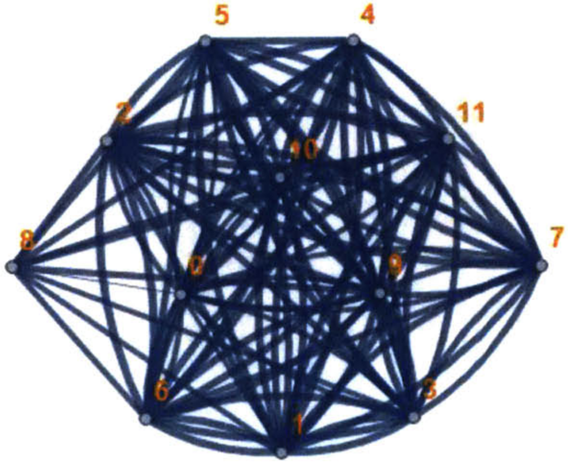

10 am – 11 am

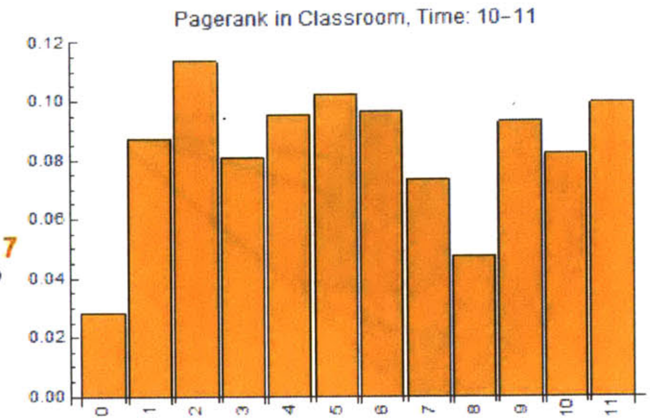

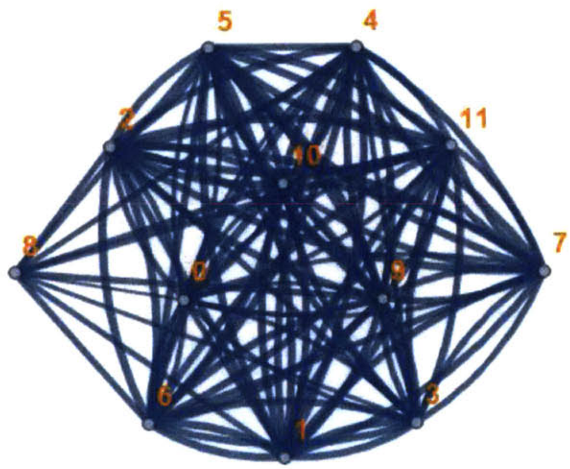

11 am – 12 pm

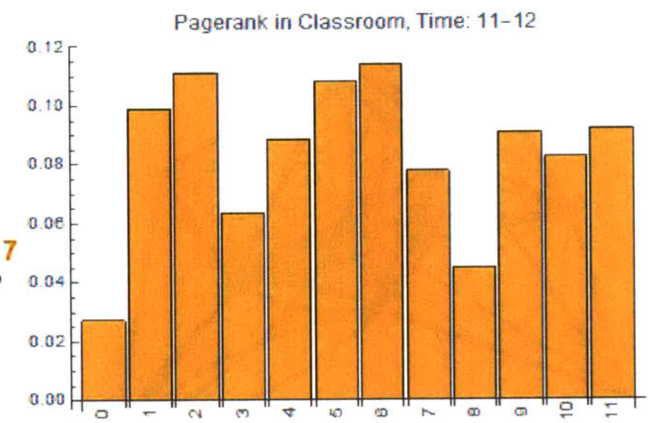

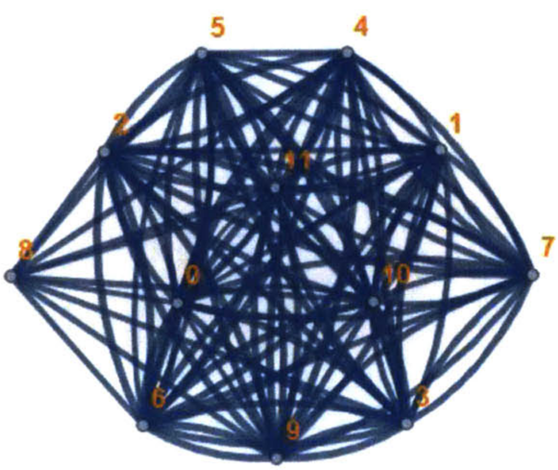

12 pm – 1 pm

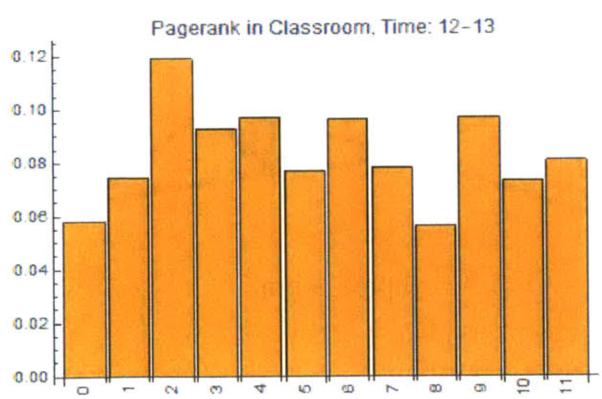

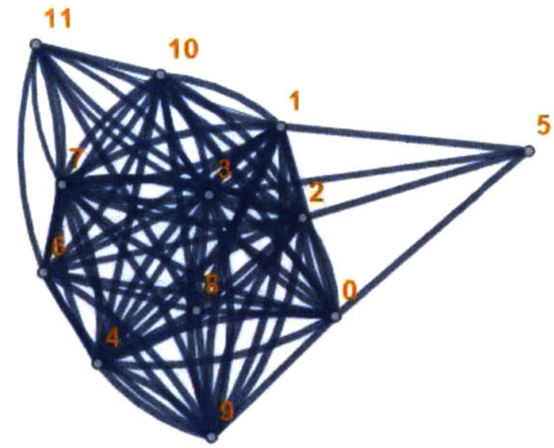

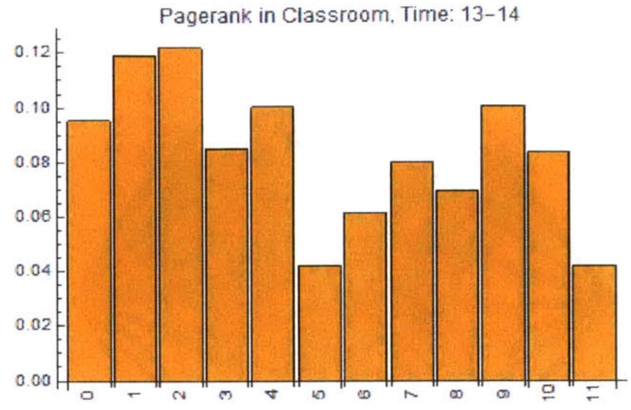

1 pm – 2 pm

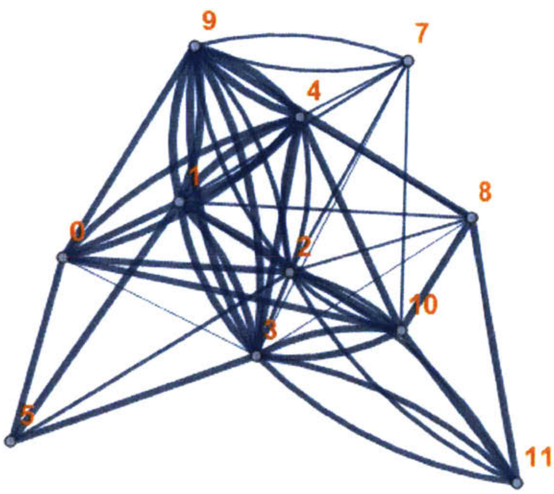

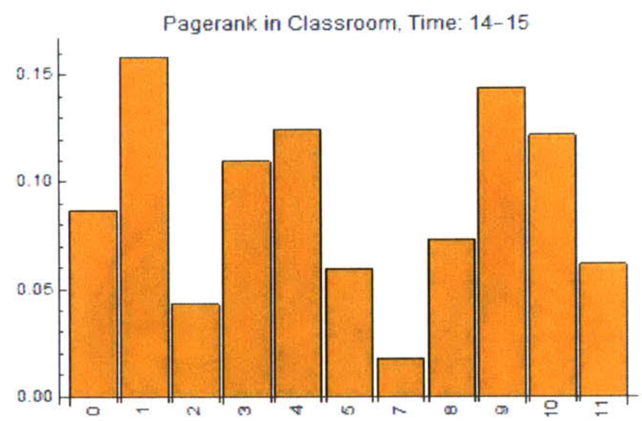

2 pm – 3 pm

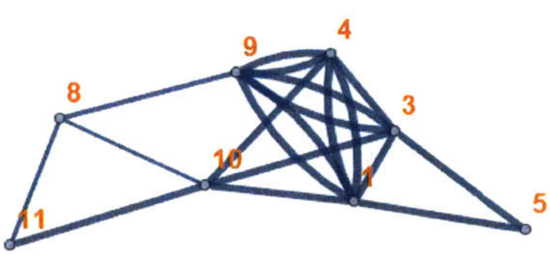

3 pm – 4 pm

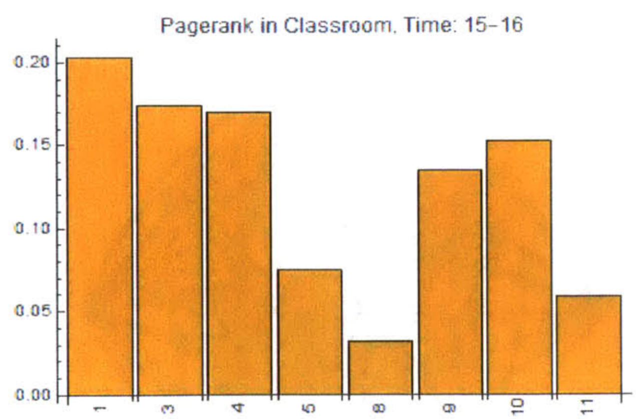

Table 1. Network structure evolution on an hourly basis, with PageRank charts for all students and teachers.

The initial hours' networks are dense, a lot of interactions are happening in those times. Teachers have dominant role in the social network as lesson giver and caregivers too. A few students do have more sense of independence, and are surrounded by students who spend time with both of these influencers and teachers. As the day goes on, students are dispersed in the classroom and the PageRank chart reveals a shift in the usual dynamics. Students 5 and 6 are were playing a dominant role by spending time with both teachers and other students. However, their roles diminish as the day goes on, probably they work by themselves for the rest of the day. At the end of the day, students start leaving, so the number of nodes are decreasing in the hour-based temporal networks. We still see two dominant clusters in the network based on the PageRank score.

7.1.3 Community Detection in Temporal Networks

We use a modularity maximization based community detection algorithm [19] to find community structure in these temporal networks. The results illuminate the PageRank evolution results further, and provide basis for future work for social structure analysis in early-childhood learning domains.

Community Structure Evolution

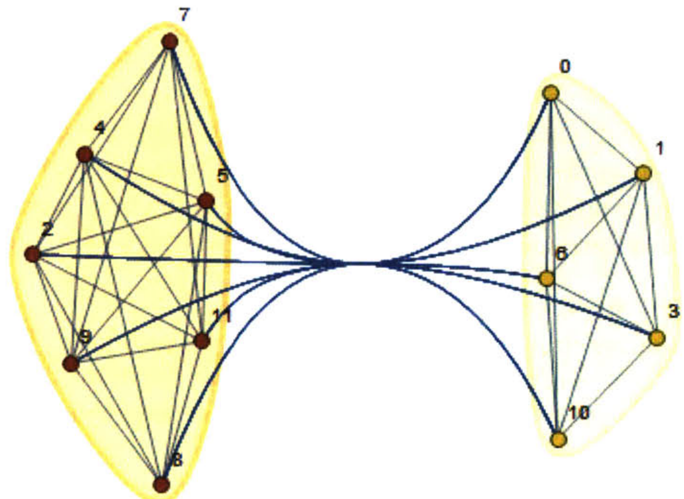

10 am – 11 am

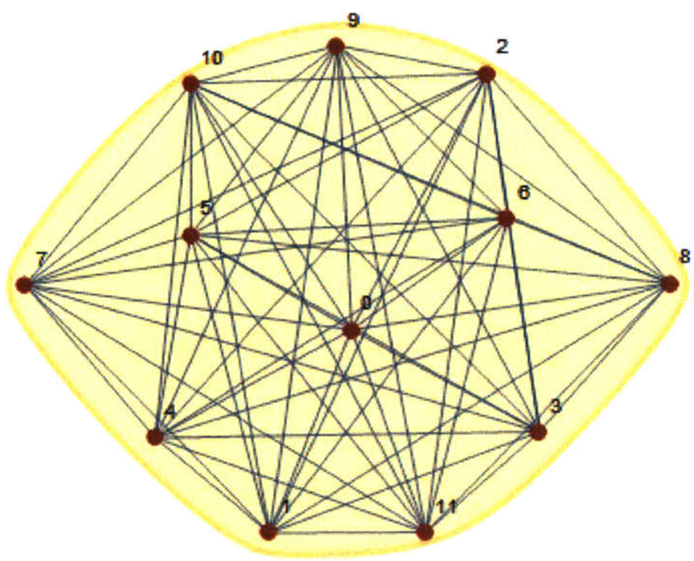

11 am – 12 pm

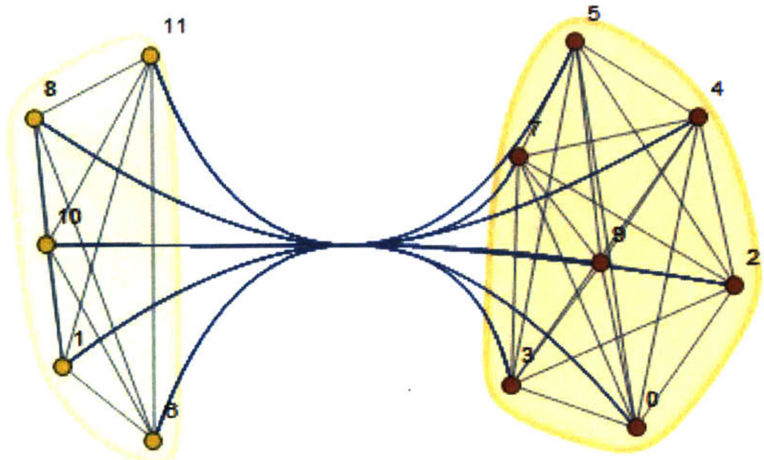

12 pm – 1 pm

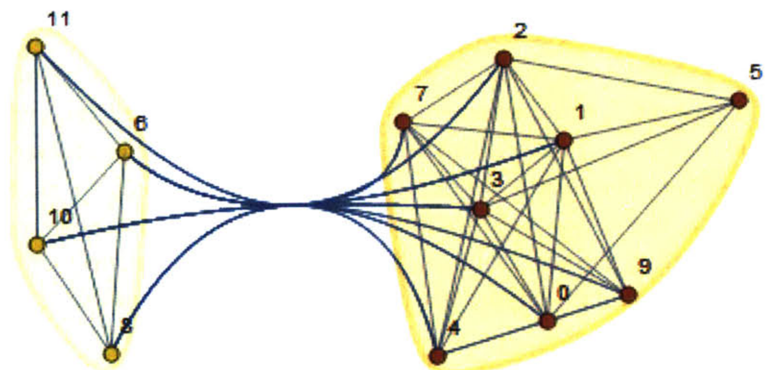

1 pm – 2 pm

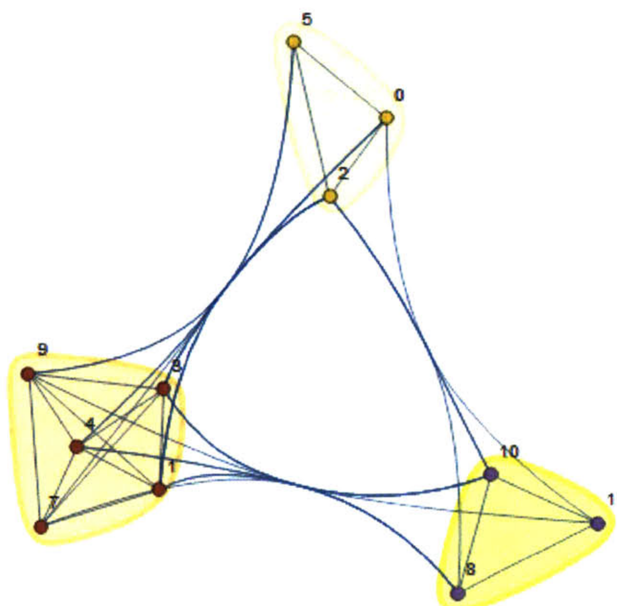

2 pm – 3 pm

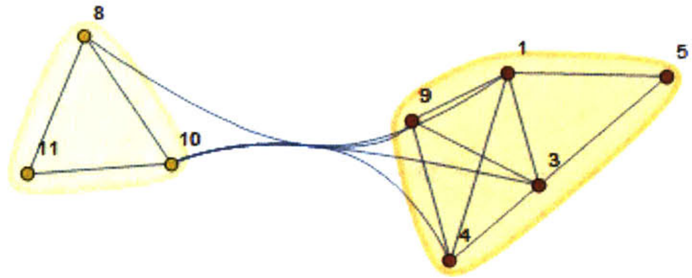

3 pm – 4 pm

Table 2. Community evolution in temporal networks in the classroom.

Community detection methods reveal finer details of student-teacher and student-student dynamics that unfold during a typical day. We have verified this community structure by manual observation during that particular day. Initially, teachers (10 and 11) have their own clusters in the first hour, when both of them are giving lessons to students, separately. From 11 am – 12 pm, the students and teachers come together for circle time, when social activities happen involving everyone. For the rest of the day, some students work on lessons by themselves, whereas teachers spend time together with a few students who need more attention. Students start to leave at the end of the day, and there are more clustered interaction happening. Community detection, in such cases, can give us a very good picture of social clusters and dynamics, along with social learning and group work. Such analysis done over months of data can reveal unique patterns in social learning and effectiveness of teaching lessons to students.

7.2 Region Preference Dynamics

As discussed before, the regions in the classroom are important markers of interest among students. Our visualization of region occupation over time provides a way for teachers to understand their classroom design better. Additionally, some statistical analysis techniques can group students according to their preferred locations in the classroom.

7.2.1 Student-Location Frequency Matrix

The region data collected over a day can be used to form a matrix M of student-location frequency matrix, similar to term-document frequency matrix. Rows are individual students, and columns are marked shelf locations where region trackers were placed. As each student visits different shelves in the classroom, his/her proximity data are logged in the region tracker. Every 10 seconds, the occupation count increases once in this way. The symmetric matrix M.M' can be used to find co-occurrence of students with each other. A principal components analysis would also reveal clusters of students based on common locations (the two most varying dimensions of PCA can be plotted to find these clusters). Thus data collected over months can show unique patterns in learning preference.

For demonstration purpose, we show the analysis a few sample days spread over a month: a few days before and a continuous week after spring break in a classroom.

7.2.2 Principal Components Analysis (PCA)

The PCA results on the timeline stated above reveal social groups clustered according to region preference in the classroom. This gives a relevant spatial context to the community detection methods.

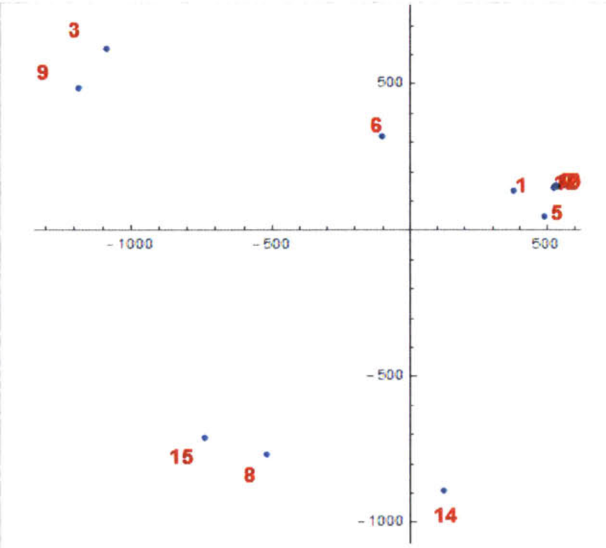

03-10-16

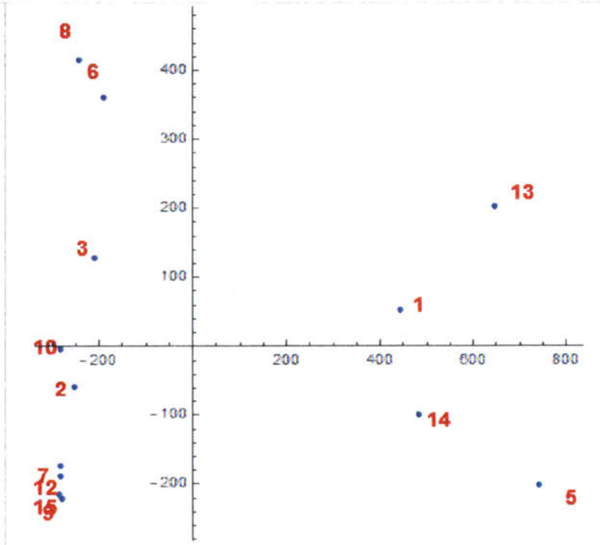

03-28-16

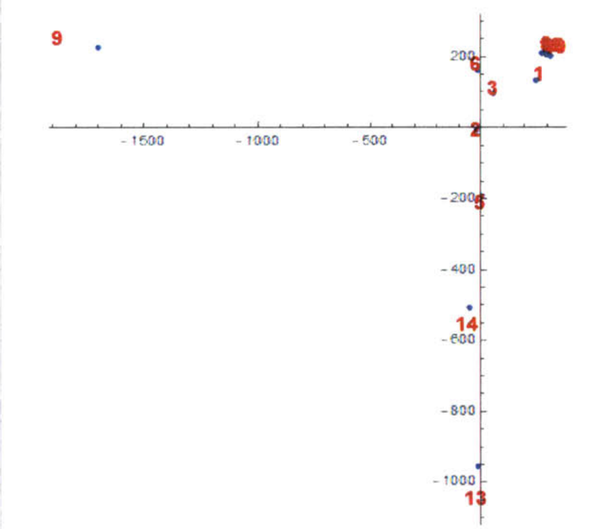

04-04-16

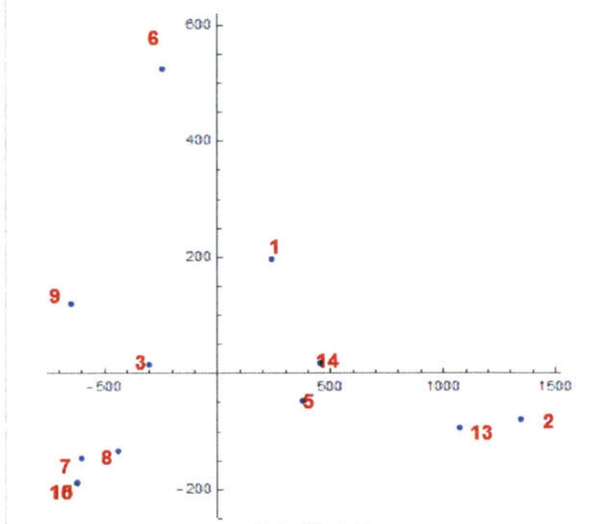

04-05-16

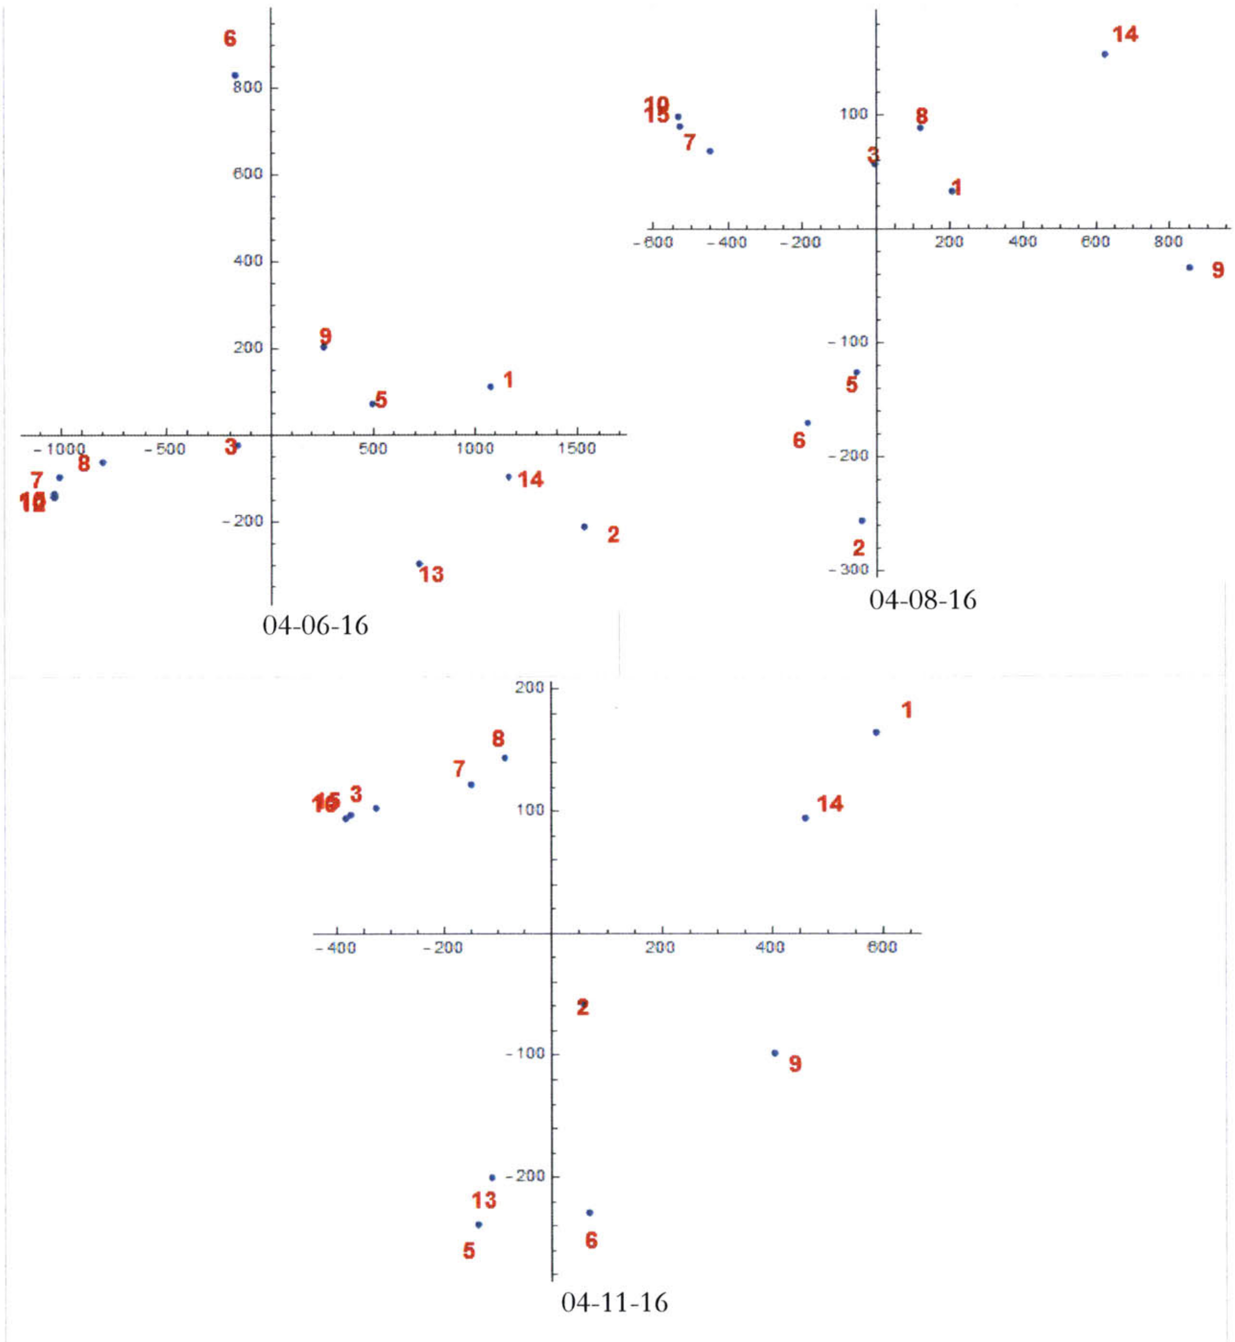

Table 3. The first two dominant PCA components of Student-Location matrix.

On some days, there are a few students who cannot be assigned to a tight cluster, but the rest are usually very much clustered that show preference towards certain locations. Before and after the spring break (3/29 to 4/3), there is a significant change in shelf preferences, but the preferences stabilize soon enough to previous patterns. As we learned, this can be due to certain classroom designs teachers changed over the break.

7.2.3 Co-occurrence Matrix

Another way to look at the same data is by calculating the co-occurrence matrix for the same days. The student-location matrix columns are normalized before we calculate the co-occurrence.

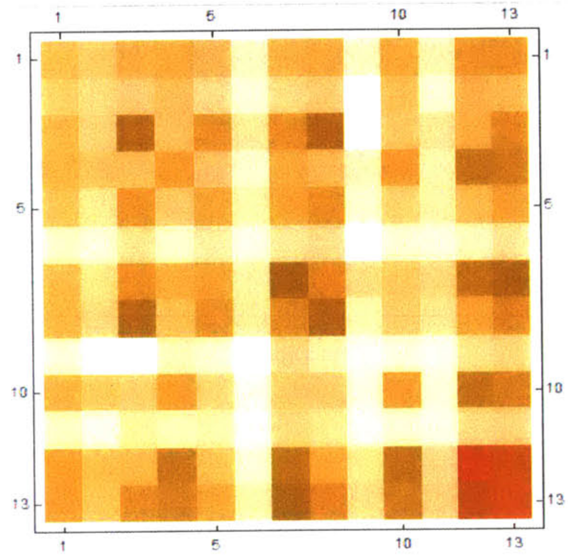

03-10-16

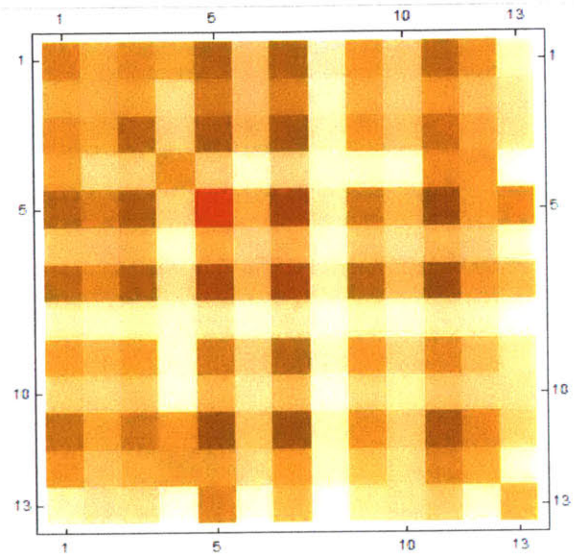

03-28-16

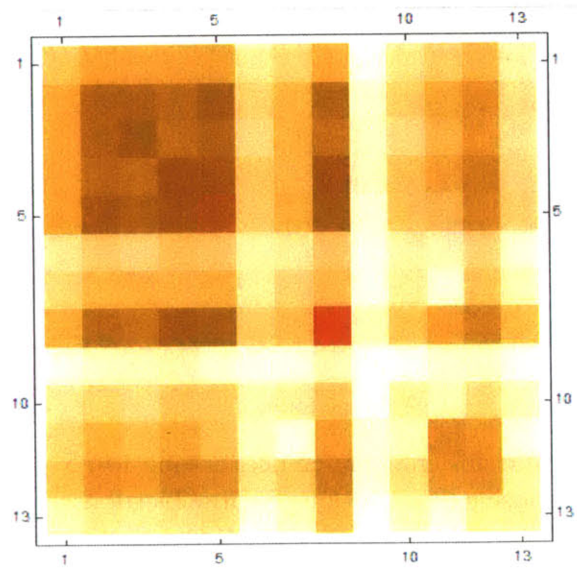

04-04-16

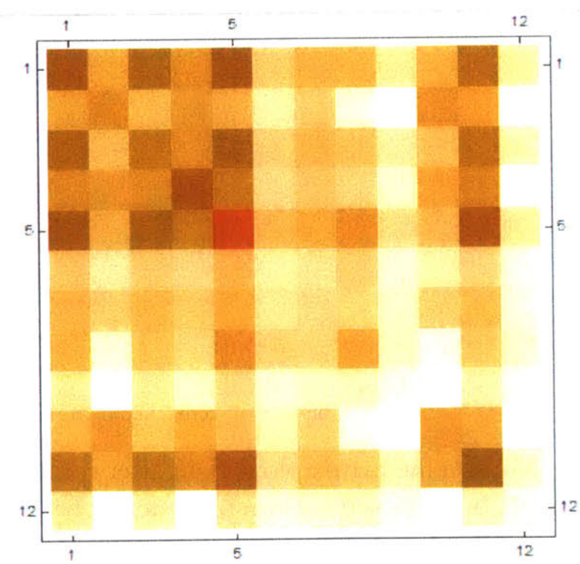

04-05-16

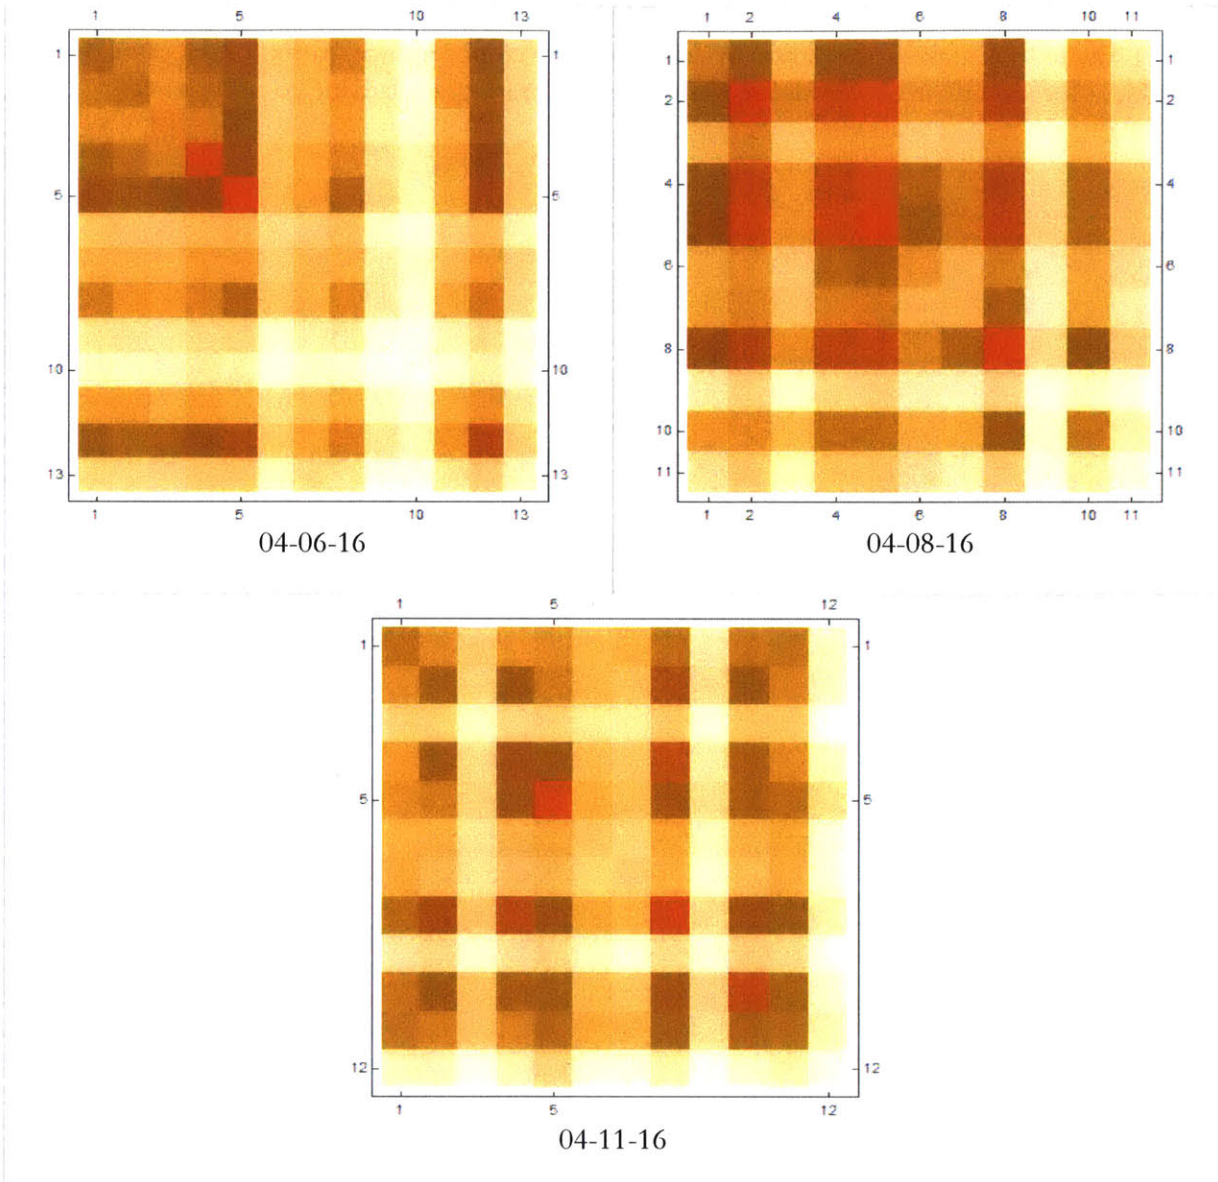

Table 4. Co-occurrence matrices for the same days PCA was performed.

As seen from the matrix plots, right after spring break, the clustering preferences among 1 to 5 changed, but they settle down onto similar habits of choosing their previous favorite locations before spring break.

Doing a co-occurrence analysis based on the spatial context can provide such insights over months of data, which has the potential to reveal preferential bias among students to adhere to certain category of lessons and shelves.

8 Use Cases Evaluation

After deploying Sensei for four weeks in three different Montessori classrooms, we interviewed teachers about the intended use cases of the system. They were given access to the web dashboard to look at their own classroom data, and described what they thought about the usefulness of Sensei in their classroom.

Ten certified Montessori teachers were recruited from classrooms with three different types of student demographics: ages 1.5-3, ages 3-6, and bilingual schools. This variety of schools gave us a diversity of opinions and how they thought Sensei would be useful in the context of their classrooms. Out of the ten teachers, six of them were in classrooms where we deployed our final pilot study with Sensei sensors. We conducted a 30 minute long qualitative study with them, and found that they 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 in the web dashboard and we recorded videos of the 30-minute interviews to transcribe their comments.

8.1 Perception about Sensei Deployment

In our initial interviews, Montessori teachers indicated that they were hesitant to introduce technology in the classroom, especially in the form of screens. However, teachers seemed quite enthusiastic to deploy Sensei in their classrooms. They found the visualizations useful in different ways according to their own classroom settings. Teachers who were from the classrooms where we already deployed our system mentioned that they did not notice any change in natural classroom interactions because of Sensei; students could barely recognize the changes we made in their classrooms. In other words, Sensei was successful in blending in with the seamless classroom experience.

8.2 Augmenting Manual Observation with Sensei

All in all, the teachers most appreciated the aspect of having "specific quantifying elements" to augment their observations. Almost all of them remarked that Sensei is an excellent tool for capturing insights that may otherwise have been lost.

One teacher commented that she can now have more meaningful conversations about her students with their parents, so this system "would be especially meaningful around parent-teacher check-in time". Another teacher from a 3-6 year old classroom divided with walls said that this data would be particularly useful for her, as she does not have a full view of her classroom to observe properly.

Another teacher from a 1.5-3 year old classroom commented that these classrooms are more difficult to observe as younger children prefer to rapidly move around the classroom and socialize, rather than spend too much time on lessons. These teachers need to be more active and hence their observation records are usually minimal compared to other Montessori classrooms.

8.3 Needs for Increased Interaction

When shown the three visualizations using social interaction data, nearly every teacher drew comparisons between the visualizations and their own classroom. Moreover, teachers were able to identify clusters of children based on the data presented to them.

Teachers who wore the sensors during our final round of pilot study were shown their level of interaction with the student body. They expressed concern about children who spent less time with them and more time alone, and thus the teachers were able to make informed decisions about to whom they would need to reach out.

All of the teachers were drawn to the data about their own interactions, which reflected our findings in our initial teacher interviews. One teacher reflected on how much time was spent with children: "I need to get these two kids more independent from us. One of the key Montessori principles is independence." In the bilingual classroom, teachers were interested in seeing which students spent more time with the Spanish-speaking teacher over the English-speaking teacher. With this visualization, teachers can self-reflect on their own methods in the classroom.

8.4 Tracking Learning Progress

Teachers from 3-6 year old classrooms were very interested in the data gathered from sensors located around the room and on lesson trays. According to them, this information can assist in optimizing learning outcomes in the classroom by determining what lessons to introduce to children.

When shown the visualization using lesson interaction data, teachers who taught the 1.5-3 year old classrooms were not very interested, because younger children spend less time working on specific lessons. However, teachers who taught older classrooms were very interested in using this to inform their weekly lesson planning. They were also interested in this data to inform them when they should introduce a new lesson into the classroom: "If these lessons aren't being used at all, then my class is done with them".

Teachers were also interested in identifying a child's specific interest. One teacher described a scenario where a child chooses a lesson, removes it from the shelf, but then returns it when it is too advanced for them. By seeing proximity data from this scenario, they can recognize these often quick interactions and choose to present the lesson to the child instead of having the child attempting it by himself/herself.

When shown the visualization of region tracker data, teachers, especially those who taught the 3-6 year old classrooms, were drawn to the pie chart depicting a child's time spent in particular areas of the classroom. They remarked on how this information would help them determine when a child should be encouraged to explore a new subject area.

Teachers also raised questions on how Sensei can measure a student's actual focus on a lesson. Although a student might be in close proximity to the lesson, their attention might be divided.

8.5 Proximity as a Proxy for Interaction

Proximity does not always represent social interaction, which is an assumption on which we designed the system. This was an issue raised by teachers too. To determine how often children are actually interacting, we studied how well proximity maps to social activity by quantifying how often these interactions happen on a typical day.

We recorded four hours of video in one school that has 3 - 6 year old children. Three observers manually annotated the video, recording timestamps when they thought social interactions occurred between children. To imitate our sensor network configuration, less than ten seconds of proximity were ignored in the manual annotation. They also recorded when the children were in close proximity but were not interacting.

Using aggregate values of social interaction and proximity durations, we calculated that children were actually engaged in interaction 84.9% of the time when they were in proximity. By tracking patterns over time and gathering more longitudinal data, we can help alleviate this limitation of our system.

9 Related Works

9.1 Wearable Electronics and Software for Contact Networks

The Sociometric Badge and Sociopatterns proximity sensor log proximity data between groups of people [20, 21, 22]. The Sociometric Badge is a wearable badge that was designed to measure proximity during meetings and conferences. Bluetooth modules search for similar modules in their proximity and use the RSSI value of the incoming data packets to measure that proximity. Smartphones have also been used as proximity sensors in some studies. The authors of [23] used Bluetooth RSSI signal strength between phones to understand proximity and create a social interaction network of students in the Copenhagen Network Study.

A challenge of using these sensors in an early childhood classroom setting is their high power requirement and the consequent effect on the size of the circuitry. The batteries tend to be larger than coin cells, and using a big badge in early education classrooms is discouraged by Montessori educators and teachers. Screens are also discouraged, and children cannot really carry a smartphone around. In contrast, our sensors are small and discreet in nature, so they do not disrupt the classroom experience.

9.2 Computer Vision in the Classroom

A review of the literature on people tracking is well beyond the scope of this thesis, we will only mention a few examples of the related work here. Authors of [24, 25] demonstrated, among many, that it is possible to track people in a mildly cluttered scene when individual motion tracks are intersecting in a video. In the recent years, multilayer neural networks have become popular in recognizing and segmenting objects in a video. A survey of several methods inspired by this technique is available in [26]. Depth sensing imaging devices like Kinect have been used widely to detect and track multiple human figures [27]. Thermal imaging was used in [28] to improve human figure detection, and infrared motion tracking systems exist [29] that occasionally require people to wear tracking tags.

There are certain caveats to using computer vision in classrooms. Other than privacy concerns, most computer vision algorithms require good lighting conditions and are challenged by occlusion limits. It is difficult to track 15 - 20 children of small heights in a cluttered classroom scene where tables and furniture act as major occlusion. Kinect and similar stereo based camera systems have a maximum range within which they are accurate (5m for Kinect) [30], and Kinect can track only up to a finite number of human figures simultaneously using its skeleton tracker features. Thermal and infrared camera systems also require multiple cameras installed in the classroom.

Moreover, they are typically very expensive and are not suitable for a small school budget. Our sensors are low-cost, uniquely identifiable, and do not require changing the existing classroom design to install a camera system.

9.3 Technology in Montessori Classrooms

Both early childhood classrooms and Montessori classrooms have started to employ technology in different ways to augment lesson presentation and curriculum planning [31, 32]. There are existing tools to help teachers record their observation notes [33]. These tools are useful for comparing notes later, but they are not sensor enabled and do not present any quantitative data on lesson progress, social interaction or region based activities.

9.4 Proximity Data Analysis and Visualization

Visualization platforms for proximity data are limited to understanding social and organizational behavior [34]. Authors in [35] demonstrated how social network analysis could be used with qualitative evaluation data to understand the context behind classroom interactions. Data from Sociometric badges have been used to identify important social factors in organizational design and management. The authors of [36] used the badges to study and visualize longitudinal social interaction patterns over weeks. The visualizations were designed for research purpose only. Sociopattern proximity sensors have been used in a high school [37, 38] to demonstrate the superiority of proximity data in understanding social interactions, compared to observation diaries and friendship surveys. However, to our knowledge, no proximity visualization framework exists that is accessible to a person outside the domain of research and engineering. Additionally, no system currently enables a teacher or a parent to understand changes in social behavior in the classroom and lesson progress of children.

10 Future Work

Sensei will be enhanced both in the hardware and firmware fronts. The region tracker solution can log proximity data for three weeks at maximum. Better sleep schedules and smarter protocols can enhance the battery life further. We plan to develop rigid-flex PCBs [39] of Sensei hardware so we can deploy the sensors in student's work rugs and carpets.

We would like to run more pilots and deploy Sensei in different classrooms of Wildflower Schools to evaluate its usefulness in individualized curriculum design. Classroom design changes inspired by Sensei will also be evaluated, as this is an important and impactful use case of our system.

The Sensei API will facilitate early childhood learning research in many different ways. Sensei deployed in Montessori classrooms for a period of months will provide a unique dataset about early childhood education. This rich dataset can be used, among other things, to create models of social and learning interaction. For example, using Hidden Markov Models [40], we can classify "hidden states" underlying children's learning patterns based on lesson activity. Using Probabilistic Latent Semantic Indexing (PLSI) [41] on student-location or student-lesson time series, we can group students according to their interest in different topics and track how the groups evolve.

Maria Montessori and other early childhood researchers have observed that children have "sensitive periods" in which they are particularly open to learning a certain subject [7]. HMMs can help to understand and quantify these sensitive periods.

Even though the social network produced from our proximity data is an undirected graph, we can treat it as a directed graph based on who is speaking in a group (captured from microphone data). This creates a better ground for PageRank-like algorithms to analyze the network, helping to understand the influence of different children in the classroom.

11 Bibliography

[1] J. Dunn, "The Emanuel Miller Memorial Lecture 1995 children's relationships: Bridging the divide between cognitive and social development," *Journal of Child Psychology and Psychiatry*, vol. 37, no. 5, pp. 507-518, 1996.

[2] A. S. Lillard, O. N. Saracho and B. Spodek, "Playing with a theory of mind," *Multiple perspectives on play in early childhood education*, pp. 11-33, 1998.

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

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

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

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

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

[8] P. B. Uhrmacher, "Uncommon schooling: A historical look at Rudolf Steiner, anthroposophy, and Waldorf education," *Curriculum Inquiry*, vol. 25, no. 4, pp. 381-406, 1995.

[9] A. Lillard and N. Else-Quest, "Evaluating montessori education," *SCIENCE-NEW YORK THEN WASHINGTON-*, vol. 311, no. 5795, p. 1893, 2006.

[10] R. F. Digital, *RFduino Datasheet,* 2014.

[11] R. F. Digital, *Simblee Datasheet,* 2015.

[12] "Physicaloid Android Library," 2014. [Online]. Available: https://github.com/ksksue/PhysicaloidLibrary.

[13] M. Bostock, "D3. js," *Data Driven Documents,* 2012.

[14] P. Holme and J. Saram{\"a}ki, Temporal Networks, Springer Berlin Heidelberg, 2013.

[15] T. Haveliwala, S. Kamvar and G. Jeh, "An analytical comparison of approaches to personalizing pagerank," 2003.

[16] L. Page, S. Brin, R. Motwani and T. Winograd, "The PageRank citation ranking: bringing order to the web.," 1999.

[17] M. Barthelemy, "Betweenness centrality in large complex networks," *The European Physical Journal B-Condensed Matter and Complex Systems,* vol. 38, no. 2, pp. 163-168, 2004.

[18] P. Grindrod, M. C. Parsons, D. J. Higham and E. Estrada, "Communicability across evolving networks," *Physical Review E,* vol. 83, no. 4, p. 046120, 2011.

[19] S. Fortunato, "Community detection in graphs," *Physics reports,* vol. 486, no. 3, pp. 75-174, 2010.

[20] "Sociometric Badge," [Online]. Available: http://hd.media.mit.edu/badges/.

[21] "Sociopattern Proximity Sensor," [Online]. Available: http://www.sociopatterns.org/.

[22] A. Barrat, C. Cattuto, V. Colizza, J.-F. Pinton, W. Van den Broeck and A. Vespignani, "High Resolution Dynamical Mapping of Social Interactions With Active RFID," *ArXiv e-prints,* #nov# 2008.

[23] A. Stopczynski, V. Sekara, P. Sapiezynski, A. Cuttone, M. M. Madsen, J. E. Larsen and S. Lehmann, "Measuring large-scale social networks with high resolution," *PloS one,* vol. 9, no. 4, p. e95978, 2014.

[24] H. Ben Shitrit, J. Berclaz, F. Fleuret and P. Fua, "Multi-commodity network flow for tracking multiple people," *Pattern Analysis and Machine Intelligence, IEEE Transactions on,* vol. 36, no. 8, pp. 1614-1627, 2014.

[25] M. Xu, J. Orwell and G. Jones, "Tracking football players with multiple cameras," in *Image Processing, 2004. ICIP'04. 2004 International Conference on*, 2004.

[26] P. Agrawal, R. Girshick and J. Malik, "Analyzing the performance of multilayer neural networks for object recognition," in *Computer Vision--ECCV 2014*, Springer, 2014, pp. 329-344.

[27] J. Han, L. Shao, D. Xu and J. Shotton, "Enhanced computer vision with microsoft kinect sensor: A review," *IEEE transactions on cybernetics,* vol. 43, no. 5, pp. 1318-1334, 2013.

[28] S. Kumar, T. Marks and M. Jones, "Improving person tracking using an inexpensive thermal infrared sensor," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2014.

[29] P. Zappi, E. Farella and L. Benini, "Tracking motion direction and distance with pyroelectric IR sensors," *Sensors Journal, IEEE,* vol. 10, no. 9, pp. 1486-1494, 2010.

[30] K. Khoshelham and S. O. Elberink, "Accuracy and resolution of kinect depth data for indoor mapping applications," *Sensors,* vol. 12, no. 2, pp. 1437-1454, 2012.

[31] M. Srivastava, R. Muntz and M. Potkonjak, "Smart kindergarten: sensor-based wireless networks for smart developmental problem-solving environments," in *Proceedings of the 7th annual international conference on Mobile computing and networking*, 2001.

[32] E. R. Hubbell, " Authenticity and Technology in Montessori Education," *Montessori Life* , vol. 18 , no. 2 , pp. 16-20, 2006 .

[33] "Montessori Compass," [Online]. Available: http://montessoricompass.com/.

[34] K. Stein, R. Wegener and C. Schlieder, "Pixel-Oriented Visualization of Change in Social Networks.," *ASONAM,* vol. 10, pp. 233-240, 2010.

[35] A. Martínez, Y. Dimitriadis, B. Rubia, E. Gómez and P. de la Fuente, "Combining qualitative evaluation and social network analysis for the study of classroom social interactions," *Computers & Education* , vol. 41, no. 4, pp. 353-368, 2003.

[36] D. Olguin-Olguin and A. Pentland, "Sensor-based organisational design and engineering," *International Journal of Organisational Design and Engineering,* vol. 1, no. 1-2, pp. 69-97, 2010.

[37] R. a. F. J. a. B. A. Mastrandrea, "Contact Patterns in a High School: A Comparison between Data Collected Using Wearable Sensors, Contact Diaries and Friendship Surveys," *PLoS ONE,* vol. 10, 09 2015.

[38] J. Stehl{\'e}, N. Voirin, A. Barrat, C. Cattuto, L. Isella, J.-F. Pinton, M. Quaggiotto, W. Van den Broeck, C. R{\'e}gis, B. Lina and others, "High-resolution measurements of face-to-face contact patterns in a primary school," *PloS one,* vol. 6, no. 8, p. e23176, 2011.

[39] J. Isaac, "Rigid-flex technology: mainstream use but more complex designs," *CIRCUITREE-CAMPBELL-,* vol. 20, no. 10, p. 30, 2007.

[40] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE,* vol. 77, no. 2, pp. 257-286, Feb 1989.

[41] T. Hofmann, "Probabilistic latent semantic indexing," in *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval,* 1999.

Appendix A

Field interview samples.

1. What events lead/motivate you to an observation at different times of the day?

2. What motivates you to write down these observations? How do you use these notes later?

3. I want you to walk me through a day in the school. What do you observe at different times in relation to the activities in the school?

4. What are some key lessons/categories that are housed in lesson trays?

   Materials are categorized by shelves, always. Complexity of lessons are placed from top to bottom. Sorted.

5. What leads you to interact with a child?

6. How do you define focus or engagement for a child? For example, is this the length of time spent on the lesson, or the intensity?

7. What are the different kind of social interaction that happens in the classroom?

8. Tell us about specific patterns that seem to be a recurring activity in the classroom, like a rhythm. Is it important for you to observe this rhythm/pattern?

   Mary Movement rhythms on days [illegible] repetition.

   Erin: Dependence. do they always depend on a teacher to make the first [illegible]

* 9. How can we help you observe? Or enable you to observe certain things.

   What are some key questions?

   Mary. is there a correlation between teacher and student's stress?

   Mary is [illegible] [illegible] [illegible] for children.

   Missy Parent Fridays How does stress [illegible] [illegible] or [illegible]

10. If you see a child is not interacting much with a category of lessons, do you intervene or try to encourage them to use it?

all the steps involved? [illegible] response. How they use language? complex phrases] pronouns. How the language for [illegible] what is in the [illegible] that might come [illegible]. How much [illegible] do they [illegible].

1. **What events lead/motivate you to an observation at different times of the day?**

completion of the lesson, social [illegible], [illegible] materials
How they are using [illegible], how they are walking, engaged? focused?
anything that interrupts [illegible] piece. Do they need to intervene [illegible]?

[illegible] take [illegible] notes [illegible] child

2. **What motivates you to write down these observations? How do you use these notes later?**

How do we design the environment? [illegible] design of environment
lesson design. Comm. to parent.
At the end of the day, [illegible] you remember and write down

3. **I want you to walk me through a day in the school. What do you observe at different times in relation to the activities in the school?**

4. **What are some key lessons/categories that are housed in lesson trays?**

5. **What leads you to interact with a child?**

When child not working in area of lesson [illegible] to [illegible] attention. Once you identify a pattern, then intervention is more effective.

6. **How do you define focus or engagement for a child? For example, is this the length of time spent on the lesson, or the intensity?**

✓ 7. **What are the different kind of social interaction that happens in the classroom?**

[illegible], [illegible], aggression, greeting, [illegible], invitation to work on lessons. (depends on age level)

8. **Tell us about specific patterns that seem to be a recurring activity in the classroom, like a rhythm. Is it important for you to observe this rhythm/pattern?**

9. **How can we help you observe? Or enable you to observe certain things.**

Appendix B

Observation cards filled by teachers during our ideation phase.

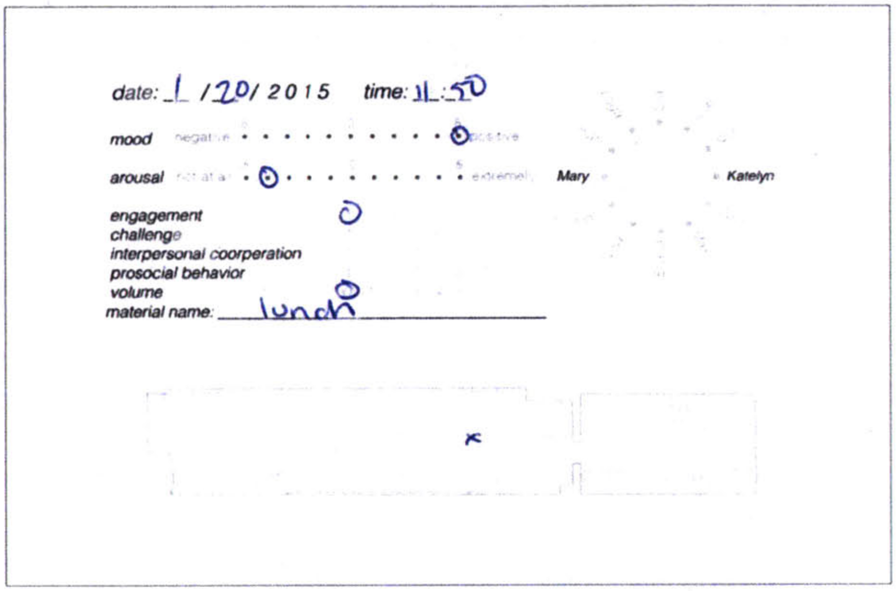

date: 1 /20/ 2015 time: 11:50

mood negative . . . . positive

arousal not at all . . . . extremely Mary Katelyn

engagement
challenge
interpersonal coorperation
prosocial behavior
volume
material name: lunch

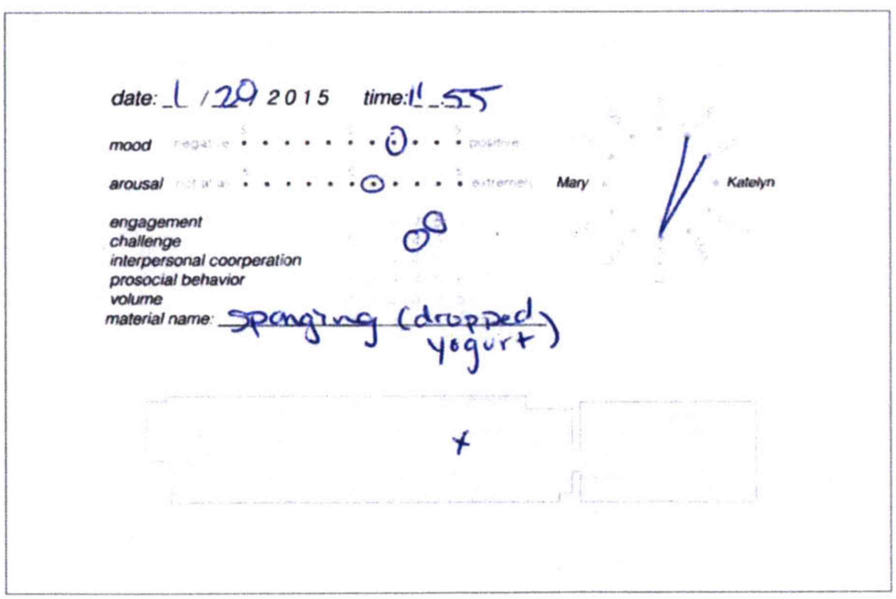

date: 1 /29 2015 time: 11:55

mood negative . . . . positive

arousal not at all . . . . extremely Mary Katelyn

engagement
challenge
interpersonal coorperation
prosocial behavior
volume
material name: sponging (dropped yogurt)