

Graph and Number line Input and Exploration (GNIE) Tool Technical Report

Yee Chieh Chew¹, Brianna J. Tomlinson¹, and Bruce N. Walker^{1,2}

Georgia Institute of Technology

School of Interactive Computing¹, School of Psychology²

ycheew@gatech.edu, btomlin@gatech.edu, bruce.walker@psych.gatech.edu

1. Introduction

Mathematics education is traditionally visual, since graphs, equations, and number lines inherently present their information in a visuospatial manner; however, for students with vision impairment, math can be difficult to learn since there has been a distinctive lack of alternative methods for teaching. Interacting with and understanding graphs is one of the core skills necessary for more advanced math classes such as Geometry, Trigonometry, and Calculus. Further, low vision and blind students are at a disadvantage in other STEM classes, for which math (and especially graphs) is a critical foundation. Currently, simple graphing concepts are taught to visually impaired students through tactile graphs; while this works for the basics, it does not help the students understand all the necessary graphing concepts (Davison, 2013).

Auditory graphs (using sound attributes such as pitch to represent numerical value) present an alternative to visual graphs. Flowers found that key information from graphs can be linked to sonifications (non-speech sounds) to denote features of the graph and changes in graph values (Flowers, 2005). Math sonifications have been used to help present numerical and graphical data through different methods, including as part of a haptic system, as a game, and even as a supplementary tool for students without visual impairment (Van Scoy, McLaughlin, & Fullmer, 2005; Hetzler & Tardiff, 2006; Upson, 2001; Upson, 2002). Learning games have been an important and engaging way for visually impaired students to learn skills. While both auditory graphs and games for learning can be helpful, the designers and developers need to pay close attention to correctly identifying mental models for interactions with the software (Flores, 2004; Sánchez, 2008). There is still plenty of work to be done to develop a system that can be used in a classroom to teach graphing concepts to students with vision impairment. A complete solution needs to be accessible to blind users, but also usable by sighted students and teachers. Thus, this project involved the systematic design, development, and deployment of a software sonification system to help teachers and students in middle school learn and use math concepts via multimodal visual and auditory graphs.

2. GNIE 1.0

Our software solution began with a general graphing model, with a focus on audio and accessibility. The Accessible Graph Model (AGM) combines graph settings with graph points lines, and canvases, capturing the graph display window, audio settings, visual settings, and a variety of other components. For example, a line could now be set to have not only a certain color (e.g., blue) but also a certain timbre (e.g., grand piano), and that setting was a part of a “graph” object, which could be saved, visualized, printed, heard, or emailed to a friend. Due to

early interest in cross-platform compatibility, and interest for the AGM to act as the graph model for the Java-based Sonification Sandbox, the AGM was programmed in Java and intended for desktop purposes.

A complete multimodal graphing program was built in Java on top of the AGM. The final program was named the Graph and Number line Input and Exploration (GNIE) tool (see Figure 1). The first-generation (GNIE 1.0) user interface has three main parts. The first part is a widget area across the top with typical interaction widgets, specifically menu items, check boxes, text boxes, and buttons. These widgets inherit accessibility hooks automatically, and were enhanced with additional accessibility context for screen readers. The second part is a text area, covering the left half of the screen below the menu bar. This holds any non-graph parts of a graphing question, such as the text “Plot the point $(-2, 3)$.” The third part is a graph (or number line), covering the right half of the screen below the menu bar. All accessibility was handled explicitly by GNIE, to control for screen reader effects.

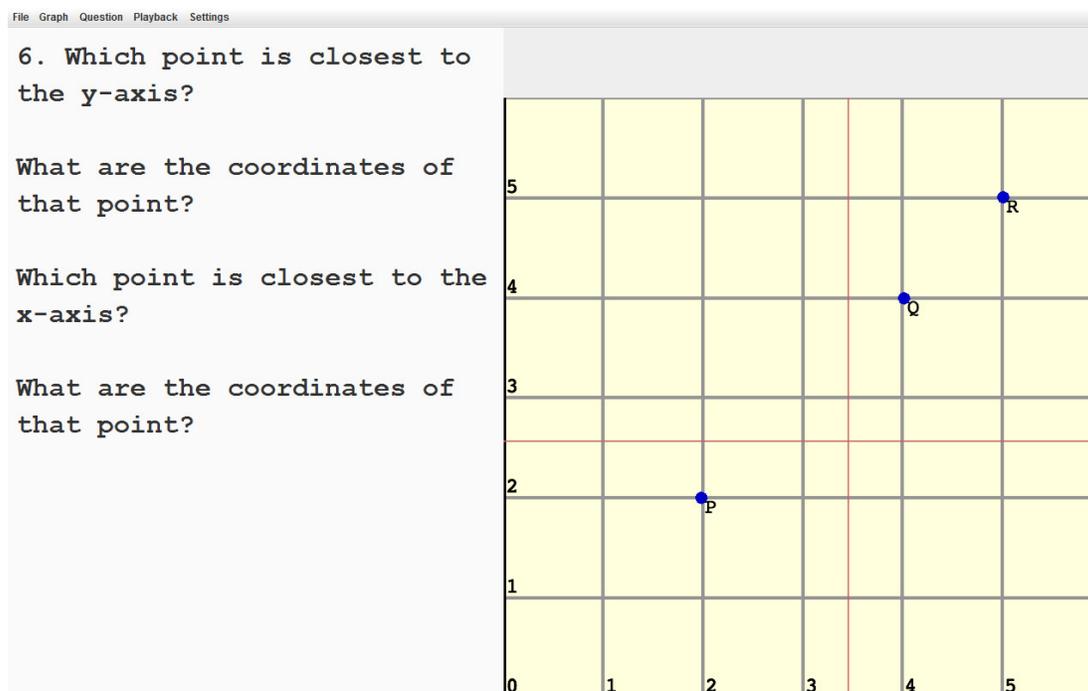


Figure 1. Screenshot of GNIE 1.0 user interface, showing menu bar, single text area (left), and graphing area (right).

Due to the dependency on existing Java systems, and the portability of Java, the system core is in Java. Later evaluations used two and three sounds in production, and up to four in development (specifically x-axis, y-axis, and either radial point distance or x and y distances each having a sound simultaneous with the others). While the latency remained low on many systems, the load was high enough to crash the system or the Java Virtual Machine. Since Java could not handle the audio processing, Max/MSP was used as an alternative. This program/language has a reputation for low-latency audio processing. In addition, using maxlink, it was possible to send messages to Max/MSP from Java. The Max/MSP program, ProcPlayer, played sounds based on messages with information about channel, instrument, pan, pitch, pitch bend, rate, and volume.

Initially, the six channels programmed in Max/MSP were mapped to context sounds for the graph (x and y), the X and Y axes, and the radial distance from a point (in the x and y directions). The sounds relating to the distance from a point begin when the cursor is within a five tick-mark radius. Each point has the same mapping, which means that every point has the same sounds associated with it, i.e. if there are two points graphed (A and B), being at a distance of 3 from point A sounds the same as being that distance from point B. The context sounds for the graph relate to the overall position of the cursor within it; this was done as a way to help give the boundaries of the graph a distinction from the middle portion. When passing over a tick-mark, the context sounds' tempos speed up, so that when the mouse is on an intersection of the marks (e.g. (4,3)) it is fastest; and when exactly between two tick marks it is slowest. These notes also utilize the mapping of pitch to provide information about the location of the cursor. Positioning the cursor near maximum for the x or y values on the graph results in a higher pitch; being near the minimum for them results in a lower pitch. Finally, the axis notes are specific noises related to crossing the x or y axis. A timbre mimicking a drum was used for the X and Y axes, to distinguish when the cursor was on the axis. In addition to sonifications alone, a text-to-speech (TTS) engine was implemented to read out the x and y tick-mark values when exploring the graph. The TTS helped provide concrete orienting information about the cursor location, to aid in exploration of the graph.

We deployed GNIE 1.0 for one year (school year 2012/13) at the Georgia School for the Blind (GAB), the only residential school for the blind in Georgia and serves students from 4-12th grade. It was used in the middle school mathematics class. GNIE was installed on four school-owned desktop computers and Georgia Tech provided two laptops so that each student would be able to work independently on their own screen. In order to allow for a more realistic usage scenario of GNIE, we left it up to the teacher to decide which lessons she would conduct using the software. Use of the software was not restricted to days when we visited the school and the teacher was encouraged to use the system as often as she wished.

3. GNIE 1.5

After being deployed for one year, we added some feature enhancements to the system to create GNIE version 1.5. The new features were based on feedback from the users of the system and observations from the researchers.

First, we introduced two new modes for accessing GNIE – GNIE Teacher and GNIE Student. The mode is selected via a dialog box when the program starts up. GNIE Student mode (see Figure 2) is an expansion/enhancement of the existing GNIE interface, as follows. During our classroom observations we noticed that some students had difficulty determining where they should type their answer. Not only that, but there were instances where the teacher had problems finding the location of student answers. The main cause of this problem was that the system only had one text area (basically, the left half of the screen). The question and any answers were all located in that same pane, and there were no cues for the students or teacher to differentiate question text and answer text. This led to answers being typed into unexpected places (such as before the question or at the end of a multiple choice selection) instead of at the end of the question. To address this issue, we subdivided the left pane into two text areas: the top pane remained the question area, not editable by the student; and the bottom pane was designated as the answer area where the student enters her response (or enters a message to refer the grader to the graph area). Separating the question and response text areas also had the

secondary benefit of reducing text-to-speech time for reading questions or answers. Before this change, students were forced to listen to the entire left pane even if they only wanted to check their answer, which might be located at the end of a lot of text. Now students are able to tab between question and answer easily. The teacher is also able to quickly determine whether or not the student answered the question correctly.

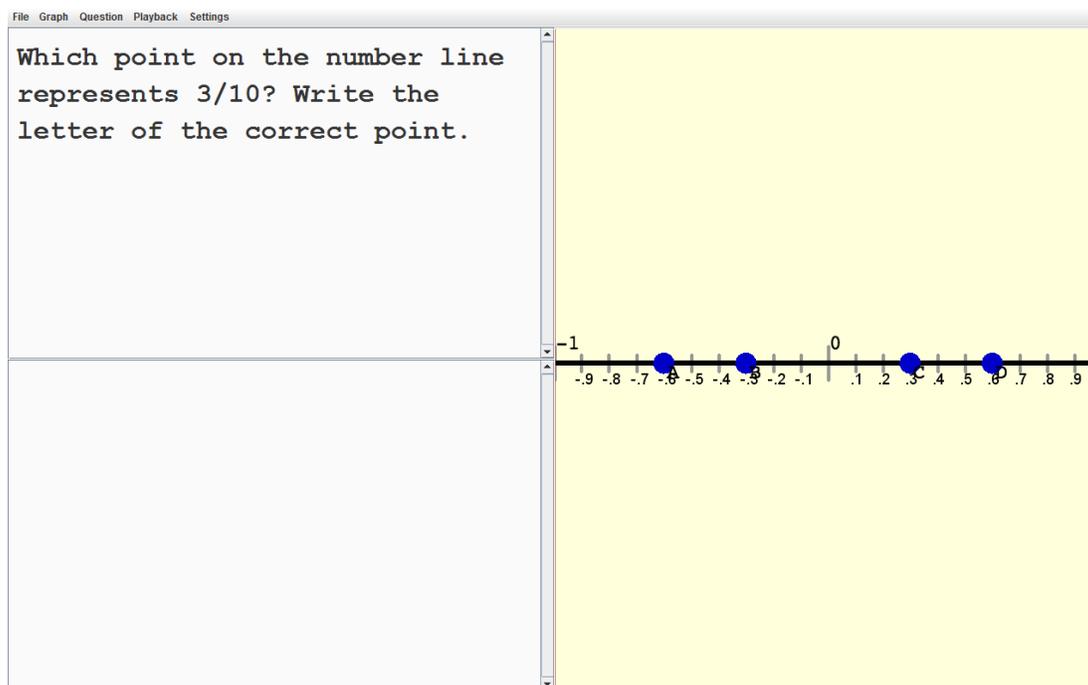


Figure 2. Screenshot of GNIE 1.5 user interface in Student mode. The text area on the left is split into Question (upper) and Answer (lower) areas. The graph supports larger values and decimals.

GNIE Teacher mode allows the teacher to create lessons for her students and does not include the answer area (thus it resembles GNIE 1.0, with basically the split screen layout containing text on the left and graph on the right). This interface does not have an answer area since the teacher does not need it. This view provides more space for question text and for reviewing questions created for the lesson.

One of the major hurdles the teacher mentioned during our interviews is the limited range of values (e.g., -10 to +10) that the number line and coordinate plane could represent in our system. GNIE 1.0 could only count by ones and could not handle decimals. Thus, for GNIE 1.5 we included the option to specify the interval of the number line/axis, thereby allowing for both decimal and larger interval representations.

We also changed the architecture of the system so that it behaved analogously to software the teacher was already familiar with. GNIE 1.0 did not recognize hierarchical folder structure, so we added support for different folders for lessons. This change allows the teacher to create multiple folders for lessons and for classes. This feature is useful when planning future lessons and also for keeping track of old lessons (e.g., creating a review lesson, or preparing for the next school year). Additionally, we created a “student” folder where students’ work is saved. When students select “GNIE Student” mode a dialog prompts them to enter their name and select their

lesson folder for the day. All progress is saved in the student folder under the student's name, which provides quick retrieval and review of lessons.

Finally, we changed the way GNIE saves and orders questions for presentation. In the first version, questions were assigned a random number ID by default. This number was used to randomize question presentation, which may be helpful in a testing situation. However, when GNIE is used for class, it is important that questions are ordered sequentially, since some questions may have multiple parts and concepts may build on each other. The teacher was frustrated when she could not understand why her questions were being displayed out of order. She said, "I mean I get the feeling even still that I put a question into GNIE and I'm really not sure where it is. You know.... I quit numbering them because it's like 'what for?'" In GNIE 1.5 we eliminated the random number assignment and incorporated a naming scheme that made more sense – a timestamp marked by when the question was created.

The changes built into GNIE 1.5 are part of the iterative process of creating a system of features that are useful in the real-world. These changes will help streamline the lesson preparation process and save time for the teacher. The design decisions were guided by our experience during the deployment and feedback from our users. GNIE 1.5 was deployed at GAB, starting Fall 2013.

4. GNIE 2.0

During Fall 2013, the development team began implementing a major new feature for the system - support for lines. GNIE 2.0 allows for a line to be drawn between two points in the graphing area (see Figure 3.0). While basic line functionality did exist in previous versions of GNIE, it was never fully available to be used by the teacher or students. The lines that existed previously were made up of a series of points plotted within the range of the two endpoints; however, this might have led to confusion with the students because the sonifications for a line would sound exactly the same as the sonification for the point, but with more of them present. We determined that it was important to differentiate between the sonifications mapped to the end points versus other values on the line segment between them.

GNIE 2.0 also has updated the sound mappings from version 1.0's Max/MSP programming. The pitch-specific context sounds depicting the location of the cursor on the graph were removed, since through deployment and observation we determined that the TTS X and Y values were more useful in displaying the location in the graph. Using tempo to demonstrate when the mouse was on an intersecting tick-mark is still implemented. Additionally, the sounds mapped to the plotted points are the same as well, except that the distance threshold for the radius was changed to four instead of five. Also, the TTS rate was slowed down to 115 words per minute, for better understanding of the question when the question area is selected.

A majority of the sonification changes were in the way lines were represented. It was important to show the continuity of the line segments, therefore the sounds initially mapped to the X and Y axes were changed to map onto each line. The dimension of pitch was used to represent the x and y values, and to mitigate any confusion which may arise from hearing both values' pitches at the same time (both the X and Y mapping), the graph must be locked on either the X or the Y axis. When locked on the X axis, the y values of the line segment may be explored (and vice versa). For all of the values, a higher number is represented by a higher pitch. Therefore, when exploring in the Y direction from $y = 5$ to $y = -5$, the pitch at $y = 5$ will be the highest and the pitch at $y = -5$ will be the lowest. The same concept applies to the x values, as

well, when the Y axis is locked. To determine the location of the cursor when locked on either the X or Y axis, a distance algorithm from linear algebra was applied to calculate whether or not the mouse was within the two end points (this is to keep the graph from playing line specific sonifications when the cursor is not within the length of the line segment).

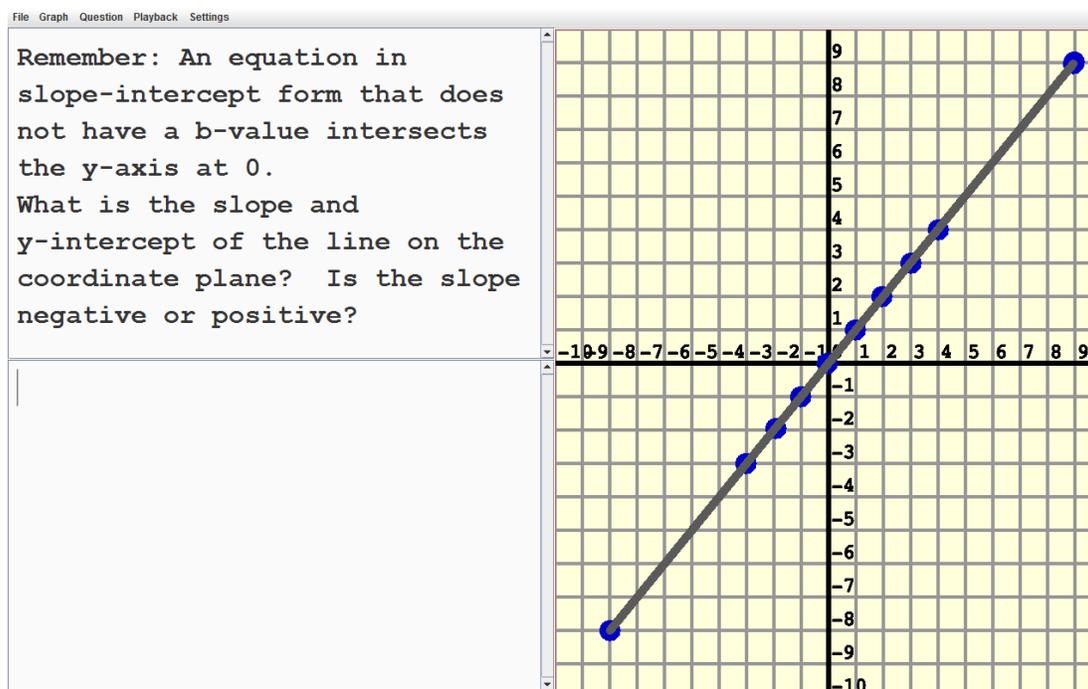


Figure 3. Screenshot of GNIE 2.0 in Student mode. The graph supports drawing of lines between points.

Multiple lines may also be created in the same graph. A line can be drawn from a series of two previously plotted, consecutive points, or when an endpoint of a line is selected, a new line can be drawn at the location of the cursor. This means that although a single line may be plotted, one may draw a rectangle, or any number of polygons within the graph area, as long as points are plotted (or selected) in pairs to form the end points.

5. Continued Deployment

GNIE 2.0 remains deployed and is in frequent use in middle school math classes at GAB. We continue to conduct observations and evaluations, and will be rolling new features and fixes into GNIE. Next steps include an expansion of the use of GNIE on homework and tests; and improvement of GNIE's math rendering capabilities.

6. Download

Information about the project, and downloads of the latest version of GNIE are available at <http://sonify.psych.gatech.edu>.

7. References

Davison, B. K. (2013). *Understanding Graph Literacy: Understanding How Blind and Low*

Vision Students Can Satisfy the Core Standards with Accessible Auditory Graphs.
Georgia Institute of Technology, unpublished doctoral dissertation.

- Flores, H. E. (2004). AudioMath: blind children learning mathematics through audio, 183–189.
- Flowers, J. H. (2005). Thirteen Years of Reflection on Auditory Graphing: Promises, Pitfalls, and Potential New Directions, 1–5.
- Hetzler, S. M., & Tardiff, R. M. (2006). Two Tools for Integrating Sonification Instruction into Calculus Instruction. Department of Mathematics and Computer Science , Salisbury University 1101 Camden Ave ., Salisbury , MD , 21801 , USA. *International Conference on Auditory Display*, 281–284.
- Sánchez, J. (2008). User-Centered Technologies for Blind Children. *Human Technology*, 4(2), 96–122.
- Upson, R. (2001). Sonifications as Mathematics Teaching Tools. Center for Science Education Portland State University PO Box 751. *International Conference on Auditory Display*, 218–221.
- Upson, R. (2002). Educational Sonification Exercises: Pathways for Mathematics and Musical Achievement. Center for Science Education, Portland State University. *International Conference on Auditory Display*, 1–6.
- Van Scoy, F., McLaughlin, D., & Fullmer, A. (2005). Auditory Augmentation of Haptic Graphs: Developing A Graphic Tool for Teaching PreCalculus Skill to Blind Students, 1-5.