

Determining the Best Time to Remove a Ventilator from an Apneic Baby Based on Graphs of Biophysical Signal Data vs. Time

Tom Lever

30 April 2013

A thesis submitted in partial fulfillment of the requirements for the degree of Bachelor of Science in Physics from the College of William & Mary in Virginia.

Advisor: John Delos

Senior Research Coordinator: Henry Krakauer

1 Introduction

Neo-natal apnea is the seemingly spontaneous cessation of breathing in premature infants. In hospitals these babies are monitored and supported by ventilators. For her own safety, a premature, apneic baby should be on a ventilator as long as is necessary, but no longer. We want to avoid removing a baby's ventilator before she is ready: If we remove the ventilator too early, the baby might suffocate and at best need to have the ventilator reintubated (i.e., reinserted). Intubation is painful, risks damaging sensitive tissue, and increases the risk of infection. If we keep the ventilator in too long, at best we prolong the risk of infection and at worst the baby's respiratory system might fail to develop properly. The research of John Delos, Mary Mohr, I, and others focuses on quantitatively determining when the best time to remove a ventilator from a premature, apneic baby is, to make ventilator removal more successful. To determine the best time to remove a ventilator, we study, after the fact, behaviors of the baby in the day before extubation. These behaviors include:

1. whether extubation (i.e., ventilator removal) was successful (i.e., whether the baby breathed naturally for the forty-eight hours following extubation);
2. the frequency and nature of apnea events and natural breathing; and
3. interactions of baby and ventilator.

These baby behaviors are exhibited in graphs of electronic biophysical signals vs. time for the day before extubation (e.g., heart-rate, electrocardiogram, respiration-rate, chest-impedance, oxygen pulse, and probability of apnea signals). My research focused on generating these graphs, organizing these graphs into packages for analysis, and analyzing these graphs. Using my research, we can determine that a baby is not ready for extubation (i.e., ventilator removal). Using my research, it seems we might be able to determine that a baby is ready for extubation. I'd like to explore these conclusions and how we came to these conclusions.

2 Our research conclusions

Our research team has developed a method to determine whether a baby MIGHT be ready to be taken off a ventilator. More specifically, we CAN determine that a baby is NOT ready for extubation. With more data, it seems we MIGHT be able to determine that a baby IS ready for extubation. How did we come to these conclusions?

Consider the following graph (Fig. 1). For each of ten babies, this graph plots the ratio between the time that baby was not breathing and the time the baby was breathing, vs. whether extubation succeeded or failed for that baby.

To understand this graph more clearly, consider Baby Number 2. According to the graph,

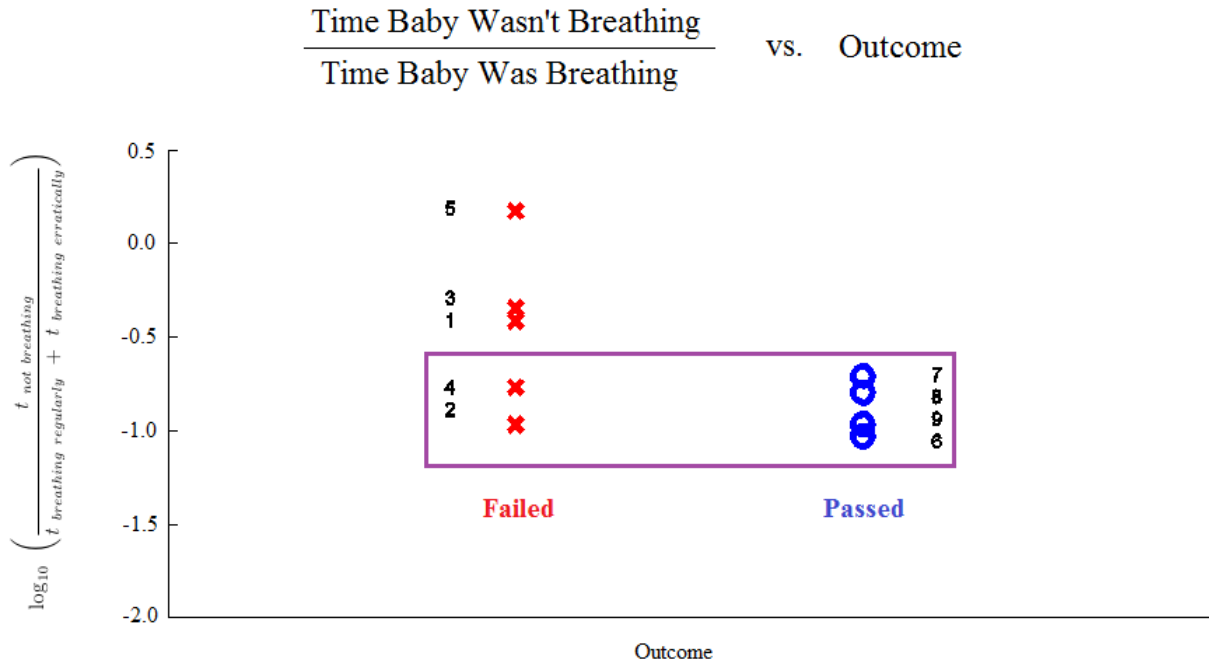


Figure 1: We CAN determine that a baby is NOT ready for extubation.

Baby 2 suffered apnea (i.e., stopped breathing) about ten percent of the time. Extubation failed for Baby 2. Consider Baby Number 6. According to the graph, Baby 6 also suffered apnea about ten percent of the time. But extubation succeeded for Baby 6. For babies whose apnea ratios lie in the purple rectangle, we don't know whether or not extubation will succeed or fail. Thus, at least for this data, we CANNOT determine that a baby IS ready for extubation. On the other hand, for a baby whose apnea ratio lies above the purple rectangle, we can say that extubation will fail for that baby. Thus, we CAN determine that a baby IS NOT ready for extubation.

To take this a step further, suppose we find a large group of babies with apnea ratios about the same and a more-or-less equal number of extubation successes and failures. Then the purple rectangle will become very precise and, presumably, our assessment of whether a baby IS NOT ready for extubation will become more accurate.

Consider another graph (Fig. 2). For each of ten babies, this graph plots the duration of that baby's longest apnea event (i.e., the longest interval the baby stopped breathing) vs. whether extubation succeeded or failed for that baby.

Consider Baby Number 4. According to the graph, Baby 4 suffered an apnea event that was about a minute and a half long. Extubation failed for Baby 4. Consider Baby 8. According to the graph, Baby 8 also suffered an apnea event that was about a minute and a half long. But extubation succeeded for Baby 8. As in the last graph, for babies whose

Baby's Longest Apnea vs. Outcome

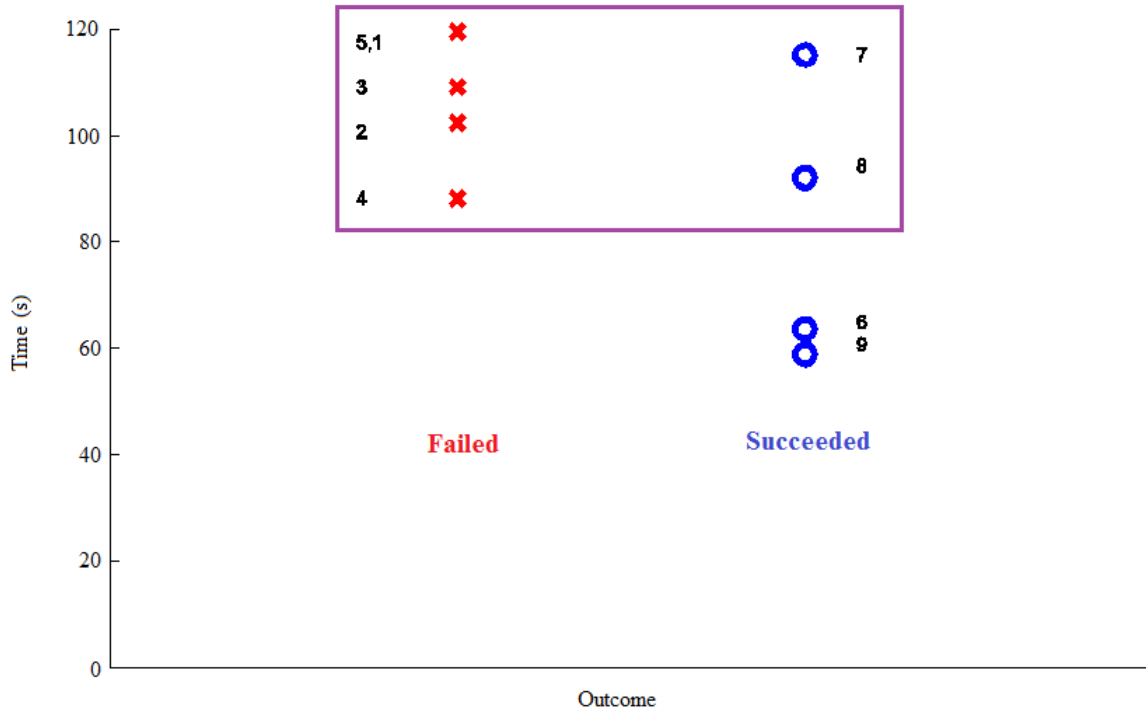


Figure 2: CAN we determine that a baby IS ready for extubation?

longest apneas lie in the purple rectangle, we don't know whether extubation will succeed or fail. On the other hand, for a baby whose longest apnea time lies below the purple rectangle, it seems that extubation MIGHT succeed for this baby. Thus it seems that with more data we MIGHT be able to determine that a baby IS ready for extubation. We should note, though, that whether or not a baby has very long apneas might be less important for predicting the result of extubation than the total number of apneas, for instance. But there is this correlation. Again, we need more data on these graphs. With more data we could define these purple rectangles more precisely and potentially predict the outcome of extubation for babies with data outside these rectangles.

Reiterating our research conclusions, we CAN determine that a baby is NOT ready for extubation. With more data, it seems we MIGHT be able to determine that a baby IS ready for extubation.

3 Collecting data for these conclusions

These conclusions depend on graphs involving five different data for each baby; namely:

1. the length of time the baby wasn't breathing;
2. the length of time the baby was breathing regularly;
3. the length of time the baby was breathing erratically;
4. the duration of the longest time interval in which the baby wasn't breathing; and
5. whether extubation succeeded or failed for the baby.

We determined the first four data from our analysis of sets of graphs of biophysical signal data vs. time for each baby. Each graph corresponded to a two-minute time interval in the day before extubation for a baby; each set of graphs corresponded to the entire day before extubation. Hospitals let us know directly whether extubation succeeded or failed for each baby.

4 My research

In August 2012 our research team had lots of biophysical signal data. I wrote code to generate 720 graphs of biophysical signal data vs. time for each of ten babies. These graphs correspond to all two minute time intervals in the day before extubation for that baby. I wrote code to organize our 7200 graphs and 1440 redundant copies into 72 packages of 120 graphs each for volunteers to analyze. I analyzed one package of graphs.

I analyzed individual graphs of biophysical signal data vs. time for each baby; Dr. Delos used our work to analyze entire sets of graphs for each baby and to come up with the five data above and the graphs from which we drew our conclusions. I'd like to talk about how I analyzed individual graphs. You can talk to Dr. Delos about how exactly he used my work to come up with our conclusions.

5 Graphical analysis for a baby

Let's study the behavior of a premature, apneic baby supported by a ventilator during a few two-minute time intervals in the day before extubation. To study the behavior of this baby I will ask and answer a series of questions.

Did extubation succeed or fail for this baby? For this baby extubation succeeded.

What signals vs. time do the graphs for this baby present? All graphs present, vs. time, from top to bottom: heart rate, the "best" electrocardiogram for this baby, respiration rate, raw chest impedance, oxygen pulse, singly filtered chest impedance, doubly filtered chest impedance, and probability of apnea. There are 720 graphs for this baby for the day before extubation: each graph spans two minutes in the day before extubation.

What are all the signals graphed? A baby's heart rate is the number of heart beats per unit time (beats per minute). An electrocardiogram is a graph of the electrical potential vs. time between two points on a baby's chest (volts). The electrocardiogram signal varies with the electrical activity of the heart. A baby's respiration rate is the number of times a baby inhales per unit time (inhalations per minute). A baby's chest impedance is the electrical resistance between two points on the baby's chest (ohms). A baby's oxygen pulse, while not directly measurable, is the volume of oxygen consumed by the body per heartbeat (liters per heartbeat). We monitor the color of a baby's blood to determine this signal. The probability of apnea for a baby at any given time is a signal calculated by an algorithm of Lee et al. It is used in hospitals to determine whether a baby is suffering apnea and to sound an alarm.

How do we interpret a graph? To the right of all graphs is a list of possible interpretations, including:

1. "Uninterpretable" (graph exhibits chaotic behavior);
2. "Baby and vent together" (baby and ventilator breathing are in phase and periodic);
3. "Unsynchronized" (baby and ventilator breathing are out of phase or are chaotic and show no phase relationship yet seem to have an overall envelope);
4. "Apnea" (ventilator breathes for baby or there is no breathing).

Let's consider a graph (Fig. 3). Which of the 720 graphs for this baby are we looking at? We are looking at graph 251 of 720, beginning fifteen hours and twenty-eight minutes before extubation. How do we interpret this graph? The baby's filtered chest impedance signals vary directly with natural and ventilator breathing. Reading these signals, we conclude that most of the time, the baby's breathing is erratic and out of sync with her ventilator. Sometimes the baby's breathing becomes regular and the ventilator supports her breathing.

Let's consider another graph (Fig. 4). Which of the 720 graphs are we looking at? We are looking at graph 593 of 720, beginning four hours and fourteen minutes before extubation. How do we interpret this graph? In this case the baby's breathing is regular, and the ventilator supports her breathing most of the time. Yet the baby stops breathing (i.e., suffers apnea events) four times over this two-minute time period. Note that the probability of apnea signal is zero for these apnea events: the algorithm of Lee et al. fails to filter ventilator signals and identify flat lines in our chest impedance signals.

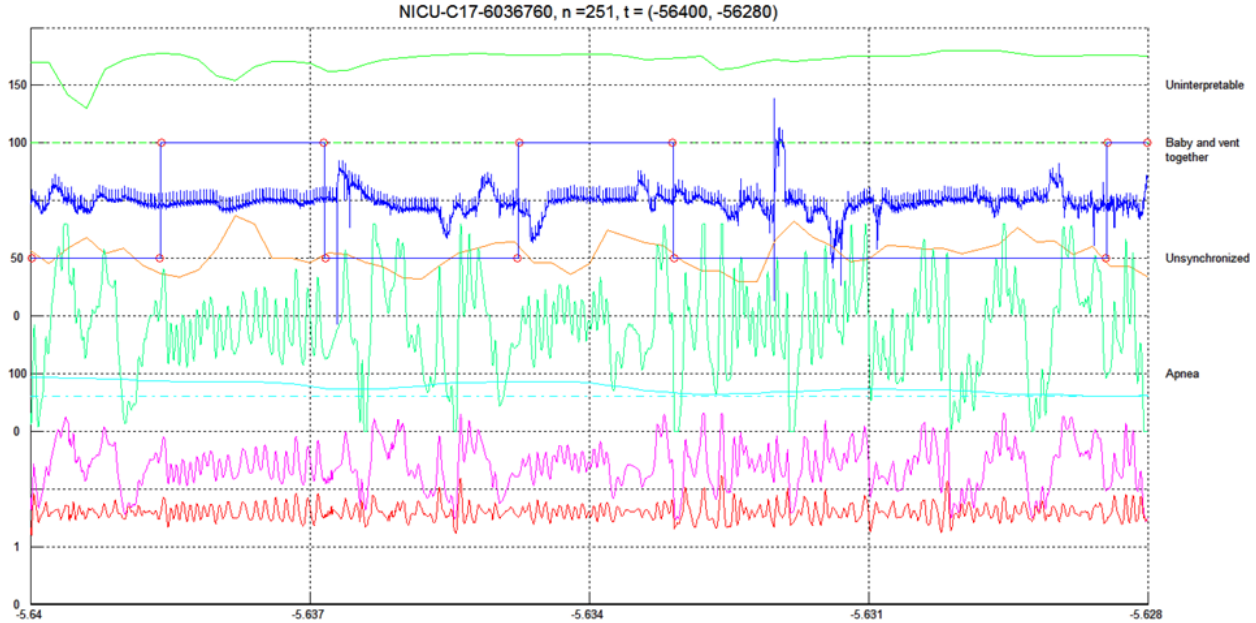


Figure 3: Baby breaths erratically and regularly.

Let's consider another graph (Fig. 5). Which of the 720 graphs are we looking at? We are looking at graph 400 of 720, beginning ten hours and forty minutes before extubation. How do we interpret this graph? For the first minute or so the baby's breathing is erratic and out of sync with her ventilator. After a minute or so the baby's breathing becomes regular and the ventilator supports her breathing.

Let's consider another graph (Fig. 6). Which of the 720 graphs are we looking at? We are looking at graph 566 of 720, beginning five hours and eight minutes before extubation. How do we interpret this graph? Most of the time the baby's breathing is regular and the ventilator supports her breathing. For about ten seconds the baby's breathing is erratic and out of sync with her ventilator. In this case, it seems we can distinguish baby and ventilator breathing: the ventilator breathes abruptly and at about one third the frequency of the baby's natural breathing.

Let's consider a final graph (Fig. 7). Which of the 720 graphs are we looking at? We are looking at graph 68 of 720, beginning twenty-one hours and forty-four minutes before extubation. How do we interpret this graph? For the first minute or so the baby suffers frequent apnea events. After a minute or so the baby's breathing becomes regular and the ventilator supports her breathing. I present this graph in part to demonstrate that our probability-of-apnea algorithm sometimes correctly recognizes apnea. Of course, in the other graphs I showed you, our algorithm failed to recognize apnea. We need to develop a method to filter ventilator signals and identify flat lines in our chest impedance signals.




Figure 4: Baby breaths regularly and suffers apnea.

Reiterating my contribution to our research, I wrote code to generate 720 graphs of biophysical signal data vs. time for each of ten babies. These graphs correspond to all two minute intervals in the day before extubation for that baby. I wrote code to organize our 7200 graphs and 1440 redundant copies into 72 packages of 120 graphs each for volunteers to analyze. I analyzed one package of graphs.

I analyzed individual graphs of biophysical signal data vs. time for each baby; Dr. Delos used our work to analyze entire sets of graphs for each baby and to come up with the five data above and the graphs from which we drew our conclusions. Reiterating our research conclusions, we CAN determine that a baby is NOT ready for extubation. With more data, it seems we MIGHT be able to determine that a baby IS ready for extubation.

6 Research log

I'd like to give a history of my research.

On Monday, 27 August 2012, Dr. Delos re-introduced me to neo-natal apnea and assigned me to develop a method to collect, read, and analyze data in real time.

On Monday, 10 September 2012, Dr. Delos re-assigned me to developing a graphical user interface to help volunteer data analyzers when apneatic babies were breathing naturally

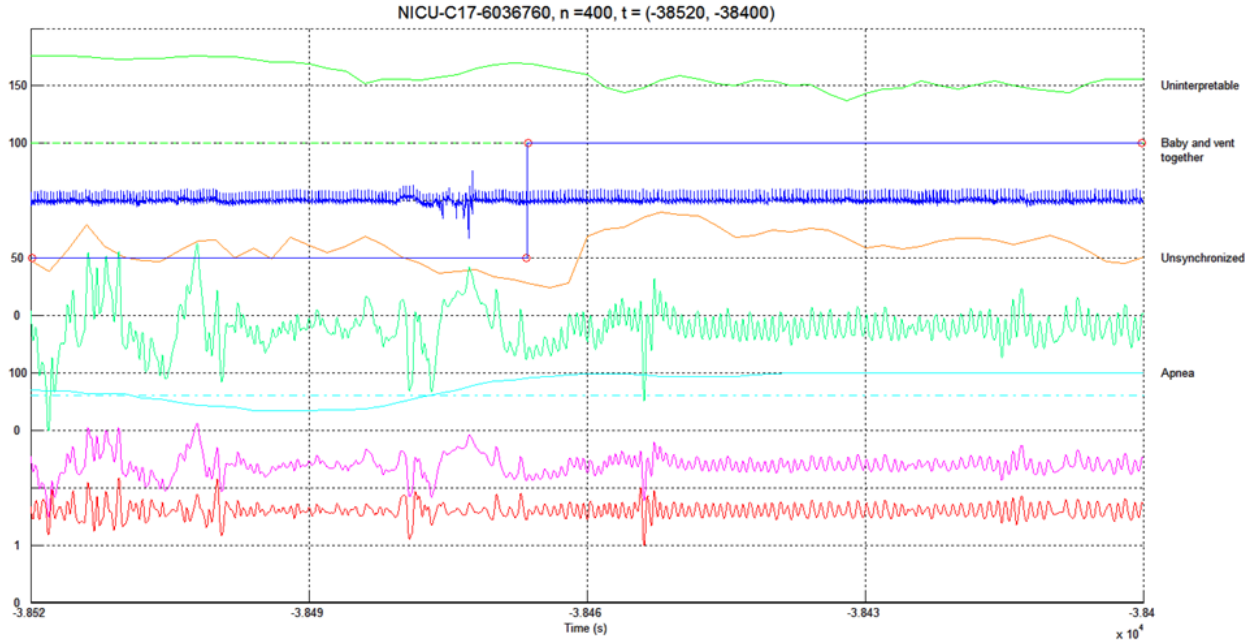


Figure 5: Baby breathes erratically and regularly.

and when their respiration was assisted by a ventilator.

On Wednesday, 12 September 2012, I met Mary Mohr, a graduate student at William & Mary and a leader in data analysis for this research. Dr. Delos and Mary agreed that I would create graphs of biophysical-signal vs. time data for many babies over the twenty-four hours before extubation (i.e., ventilator removal). These graphs would represent two-minute intervals, to zoom in on signals and make analyzing graphs more manageable.

On Friday, 14 September 2012, Dr. Delos and I met to find me access to SciClone's LittleBaby database. Up to this point I was unable to contact Mike Tighe, the custodian of my password to SciClone. I also did not know how to download large sets of data from LittleBaby or how to run Mary's analyzing and graphing programs to process data on SciClone.

On Monday, 17 September 2012, I met with Mary to learn to navigate SciClone.

On Tuesday, 18 September 2012, Dr. Delos granted me permission to access the LittleBaby database. I downloaded WinSCP, a file-transfer protocol, to download binary data files from SciClone. The LittleBaby database is located at `\sciclone\data20\LittleBaby\`. The biophysical signal data I'm interested in is located in multiple directories within the LittleBaby database. Three electrocardiograms and a raw chest impedance signal are located in the binary files

`\BedmasterData\NICU_XXX\NICU_XXX_XXXXXXXX_wave.bin`, with associated index files




Figure 6: Baby breathes regularly mostly. Ventilator supports at one third the frequency.

`wave.idx`. Heart rate, oxygen pulse, and respiration rate signals are located in the binary files

`\BedmasterData\NICU_XXX\NICU_XXX_XXXXXXXX_vital.bin`, with associated index files `vital.idx`. A singly-filtered chest impedance signal, a doubly-filtered chest impedance signal, a probability of apnea curve, and a signal used to choose the best electrocardiogram signal are located in the binary files

`\Rusin_Server\Batch_37\NICU_XXX_XXXXXXXX_aprob.bin`, with associated index files `aprob.idx`. Not only are signal and time data for various signals dispersed among multiple directories, but also the data are disjointed: the vectors that store the data have different lengths and correspond to different time periods. Signal data don't always match up with measured points in time and often are missing entirely.

In late September I worked to develop various MATLAB programs to neatly parcel disjointed signal and time data into column vectors corresponding one-to-one and to arrange these vectors in a single structure for each baby. I developed a program to plot data from this structure corresponding to the day before extubation and in two-minute segments. I revised this program to format plots and add graphical and textual annotations, including a detailed title and hints for volunteers.

Mary sent me two spreadsheets of the ID's, filenames, and extubation times of many babies. One spreadsheet listed babies for whom extubation failed and the other listed babies for whom extubation succeeded.




Figure 7: Baby suffers apnea and breathes regularly. Algorithm correctly identifies apnea.

In late October Mary and I worked to iron out my code to correctly read and organize data from the binary files in the LittleBaby database. Mary rewrote our code to correctly enter signal and time data into each baby’s MATLAB structure. I kept my structure-compiling and plotting programs.

Throughout November I turned to Mary’s spreadsheets. I began processing and checking signal data for babies for whom extubation failed and succeeded. To process data, I download binary and index files for each baby from SciClone, run my reading, compiling, and plotting programs, and check whether all signals are present. Most of the time data files are missing from the LittleBaby database or signals are missing from the binary files. By the beginning of the Spring 2013 semester I generated analyzable plots for five babies for whom extubation failed and five babies for whom extubation succeeded.

In late January Dr. Delos wrote the program `analyze.m` (and assisting programs) for volunteers to use to record their analyses of a series of graphs on the graphs themselves. A volunteer calls `analyze` with reference to a particular plot. MATLAB opens this plot. Using a mouse, a volunteer can delimit intervals in each two-minute period in which the signal was uninterpretable, indicated that baby and ventilator were breathing together, indicated that baby and ventilator breathing were unsynchronized, or indicated that the baby was suffering apnea. After analyzing a plot, the volunteer may move on to the next plot, go back to the previous plot, or edit the present plot. Upon moving on, `analyze` saves the annotated plot. I present an annotated plot below (Fig. 6).




Figure 8: This graph indicates that a baby and ventilator breathed together, that baby and ventilator breathing were unsynchronized, and that the baby suffered apnea.

On 18 and 19 February 2013 I wrote `PackagePlots.m`, `CheckListOfPlotNamesForRedundancy.m`, and `ListPlotNamesInVolunteerPackages.m` to provide ten volunteers each with a series of graphs to analyze using Dr. Delos's `analyze`. First, for each of ten babies I organized the data structure and folder of 720 plots for the day before extubation in a folder corresponding to that baby. I organized these ten baby folders in two parent folders: `BabiesForWhomExtubationFailed` and `BabiesForWhomExtubationSucceeded`. In `PackagePlots`, I assembled a list of all 7,200 plots in these two folders. I randomized this list. I made a smaller list of the first 1,440 plots on this randomized master list. I randomized this smaller list. I randomly distributed these 7,200 plots and 1,440 redundant copies into 72 packages of 120 plots each for volunteers to analyze. I organized these 72 packages into one large folder.

In

`CheckListOfPlotNamesForRedundancy`, I checked to see whether a plot and its redundant copy were in the same package; if so, I manually swapped the copy out for a plot in another package. In `ListPlotNamesInVolunteerPackages`, I saved lists of the plots in each package and a master list to `.mat` files.

In late February, Dr. Delos, I, and three volunteers met to begin analysis. We discovered that all the plots from one baby were too noisy to analyze. I submitted about 108 plots that I had analyzed to Dr. Delos in late-March.

In March, Dr. Delos began to share data with medical doctors in the University of Virginia Health System. By graphing on a logarithmic scale the ratio of the total apnea time to the

total breathing time for all ten babies vs. whether extubation succeeded or failed, Dr. Delos determined that the babies with the highest apnea ratios consistently failed extubation. Dr. Delos discovered a correlation between babies having short apnea times and passing extubation.

7 Mechanics of processing data

What do I do to process data? I open WinSCP, connect to SciClone, and download binary and index files from the above directories into my local directory

`C:\NeoNatalApnea\BabiesForWhomExtubationSucceeded\babyXXXX\`. I run my programs in this directory. These programs generate `\Plots_babyXXXX\` containing 720 plots (e.g., `NICU_XXX_XXXXXXX_1.fig`), each presenting signal data for two minutes in the day before extubation. These programs also generate `babyXXXX.mat`, a MATLAB data file containing the structure containing all the signal and time data for babyXXXX. I examine this structure and the plots to determine whether sufficient data is presented for volunteers to analyze. I submit satisfactory plots to the office PC of our research team.

How do my programs work? I present them all below. `ProcessData` is a managing program for my other programs. It takes the names of the binary and index files I downloaded, using them for reading the files and titling the plots. It takes `start_time` and `end_time`, two integers from Mary's spreadsheets delimiting the day before extubation for which to process data, corresponding to one day before extubation and extubation, respectively. `ProcessData` outputs the folder of plots, structure of signal and time data, and `.mat` file, all mentioned above.

`GetData` reads data from one set of binary files and combines it into a structure of signal and time data. Mary wrote it: I updated it to provide the structure and contents I wanted.

The data for some babies listed in Mary's spreadsheets is spread out over multiple sets of `wave`, `vital`, and `aprob` binary files. `GetData` reads data from one set of binary files and generates one structure of signal and time data. `CompileData` calls `GetData` multiple times and combines the single structures into one large structure. It chooses the "best" electrocardiogram signal to use, looking at the whole, combined structure. It saves the compiled structure to a `.mat` file.

`Read_Data_Segment_From_BIN_File` is used by `GetData`. I don't know who wrote it or how it works. I can't even read the file.

`PlotData` plots signal vs. time data for all two-minute intervals in the twenty-four hours before extubation. It formats, annotates, and organizes the plots.

To process data I call, for example,
`baby6991 = ProcessData({'NICU_A4-6350663'}, 9104400-86400, 9104400);`

8 Summary

One more summary. I wrote code to generate 720 graphs of biophysical signal data vs. time for each of ten babies. These graphs correspond to all two minute intervals in the day before extubation for that baby. I wrote code to organize our 7200 graphs and 1440 redundant copies into 72 packages of 120 graphs each for volunteers to analyze. I analyzed one package of graphs.

I analyzed individual graphs of biophysical signal data vs. time for each baby; Dr. Delos used our work to analyze entire sets of graphs for each baby and to come up with the five data above and the graphs from which we drew our conclusions. Reiterating our research conclusions, we CAN determine that a baby is NOT ready for extubation. With more data, it seems we MIGHT be able to determine that a baby IS ready for extubation. We have many more babies to study to make these conclusions more rigorous. We should consider improving our probability of apnea algorithm, baby breathing categories, and human analyses.