

Lasagna plots: A saucy alternative to spaghetti plots

Abstract

Longitudinal repeated measures data have often been visualized with spaghetti plots for continuous outcomes. For large datasets, the use of spaghetti plots often leads to over-plotting and consequential obscuring of trends in the data. This obscuring of trends is primarily due to overlapping of trajectories. Here, we suggest a framework called lasagna plotting that constrains the subject-specific trajectories to prevent overlapping and utilizes gradients of color to depict the outcome. Dynamic sorting and visualization is demonstrated as an exploratory data analysis tool. Supplemental material in the form of sample R code and additional illustrated examples are available online.

1 Introduction

Longitudinal data are a cornerstone to causal inference in epidemiologic studies. As the number of subjects, frequency of measurements, and period of active data collection grows, so does the size of the data set. As the size of the data set grows, so does the burden and complexity of being able to graphically explore and summarize the data for comprehensible viewing without obscuring salient features.

The current gold standard for graphically displaying and exploring longitudinal data is the classical spaghetti plot, which involves plotting a subject's values for the repeated outcome measure (vertical axis) versus time (horizontal axis) and connecting the dots chronologically with lines of uniform color and intensity, for each subject. However, a number of limitations to the classical spaghetti plot present obstacles to the display of longitudinal data. Although useful for fewer subjects, trends and patterns are obscured for the larger numbers of subjects typical to modern epidemiologic studies. For example, trajectories commonly overlap in a classical spaghetti plot, as both subjects and the magnitude of the outcome measure are displayed on the vertical axis. With large datasets, the resulting plot often succumbs to "over-plotting" and can be confusing given the multiple intersecting lines with no discernible patterns. Attempts to handle large datasets have traditionally involved plotting a meaningful subset of the data based on medians or deciles, but this approach fails to utilize the whole dataset.¹ Furthermore, repeated measures data containing different enrollment times, missing data, or loss to follow-up (censoring) typically are difficult to display in the classical spaghetti plot.

Exploratory plots generally are used to reveal various structures in the data: trends

(how do most people respond over time), outliers (are some subjects different from all), clusters (are there groups of patients responding the same, maybe due to a covariate such as treatment assignment), as well as data quality checks of reasonable values and sensible collection patterns of multi-center repeated measures data. Classical spaghetti plots suffering from over-plotting lower the probability of data features being revealed.

The classical spaghetti plot can be improved. Color saturation, a method of handling over-plotting, has been implemented in parallel coordinate plots, of which the spaghetti plot is a special case.^{2;3;4} In many situations, a classical spaghetti plot or the more modern parallel coordinates plot can visualize data sufficiently. However, two data features that are common in epidemiologic longitudinal datasets limit the usefulness of parallel coordinate / spaghetti plots: categorical outcomes and missing outcome values (see Figure 19 in Supplemental Digital Content 1). Not limited to but especially in the cases of categorical outcomes and missingness we propose a procedure for visualization and exploration.

In this commentary, we advocate for the use of so-called heatmaps⁵ as an alternative, or at least complementary, graphical data exploration technique to the spaghetti plot. Heatmaps enjoy frequent use in the genomics literature and other areas with high-throughput data.^{6;7;8} However, in more standard longitudinal studies, they are less popular, as evidenced by the recommendation of spaghetti plots of the raw data^{1;9;10} or summarized data¹¹ in popular texts of longitudinal data analysis.

A lasagna plot is a heatmap depicting longitudinal data for which the aforementioned texts recommend classical spaghetti plots. Spaghetti plots are so named for each subject's trajectory over time looks like a noodle, with the simultaneous plotting of trajectories looking like spaghetti. Lasagna plots are so named for each subject's trajectory over time is a horizontal layer, with the simultaneous plotting of trajectories resulting in a stacking of layers, as in lasagna.

The proposed lasagna plot uses color or shading to depict the magnitude of the outcome measurement and fixes the vertical dimension per subject. The lasagna plot takes advantage of color to provide a third dimension and display additional information, rather than relying upon the vertical dimension to display overlapping magnitudes of change. All the information about the value of the outcome is in the color (intensity) portrayed making color choice important.

Haphazard and unprincipled color selection can result in graphics with varying appearance in different media, induce optical illusions and after-image effects, alienate those with colorblindness, and produce artifacts with respect to interpreting the data characteristics via relative aspects of color.^{12;13;14;15} One practice for dutiful color selection is that of having equally spaced hues with constant chroma and luminance across hues in a perceptually uniform colorspace, such as HCL (as opposed to RGB or HSV).¹³ For unordered categorical outcomes, such as group membership, fixed chroma and luminance for equally spaced hues within HCL provide a flexible framework for the generation of qualitative palettes. To

visualize ordered outcomes, consider the sequential palette approach of allowing at least one and possibly all three dimensions (H- hue, C - chroma, L - luminance) to vary according to some function of interestingness based upon the ordering of the categories. If the outcomes to be visualized are ordered and have an inherent neutral value from which their interestingness diverges, say 0 for correlations, then two sequential palettes of different hues can be combined into a diverging palette to achieve color gradient symmetry about the neutral value.¹⁵ Practical and immediately accessible color selection resources include <http://cran.r-project.org/web/packages/colospace/vignettes/hcl-colors.pdf> as well as <http://colorbrewer.org>.^{16;17} The colors of most of the lasagna plots in this commentary (save Figures 17 and 18) were selected from the R packages `colospace` or `RColorBrewer`.^{16;18}

To reduce after-image effects and to afford viewers more viewing time with less eye strain, it is recommended to avoid plotting fully saturated colors.¹⁶ Missing values may not be assigned a color by default and inherently be portrayed by the background color. If the background color is one of full saturation (e.g., white), then change to an off-white or directly assign missingness a color, preferably one sharing attributes of the palette in use (e.g. tantamount luminance). Epidemiologic longitudinal data may have truly continuous outcomes resulting in many unique values, which may cause some coloring procedures to categorize automatically in how values are binned and assigned color. To have fullest control, it may be best to define meaningful categories of truly continuous data for visualization, such as deciles to be visualized with a palette of 10 colors and an off-white background for missing values.

There are several advantages of the lasagna plot over the spaghetti plot. These include: (a) group, cohort, and subject level information are preserved regardless of the number of subjects or time points; (b) dynamic sorting of the data can be used to ascertain group level behavior over time; (c) the clear display of missing data and easy handling of intermittent missing data; and (d) an easier display and visualization of the distribution of onset times and ending times.

2 The Basics

In a spaghetti plot, each subject is a “noodle” across time. The unconstrained vertical dimension for each trajectory in a spaghetti plot allows the overlapping of trajectories and can lead to over-plotting, which often obscures data features. The lasagna plot displays each subject as a “layer” across time. By utilizing color and a fixed vertical position for each subject, overlapping is prevented. The framework of lasagna plots enables dynamic sorting of data. In the following section, we demonstrate how the information in spaghetti plots are represented in lasagna plots and elaborate on the dynamic sorting.

the element h_{ij} (where $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n$) is a number that corresponds to the state/outcome measurement of the i^{th} subject during the j^{th} interval. The matrix \mathbf{H} is a stack of rows, with each row depicting a subject's path over time. Each row of \mathbf{H} contains subject-specific data of repeated measurements. Each column contains cross-sectional data at the group level across subjects.

Transforming the matrix containing numbers into a graphical visualization via “painting by number” gives a snapshot of the data. A lasagna plot provides a simple image with color representing the outcome measure. Each row contains a subject's clustered measurements with each measurement taking place over the column variable, time (or location). Recognizing that the underpinning of the image is a matrix, we dynamically broaden the scope of information that can be obtained via sorting. Given \mathbf{H} , five sortings are possible: within-row, entire-row, within-column, and entire-column, and cluster.

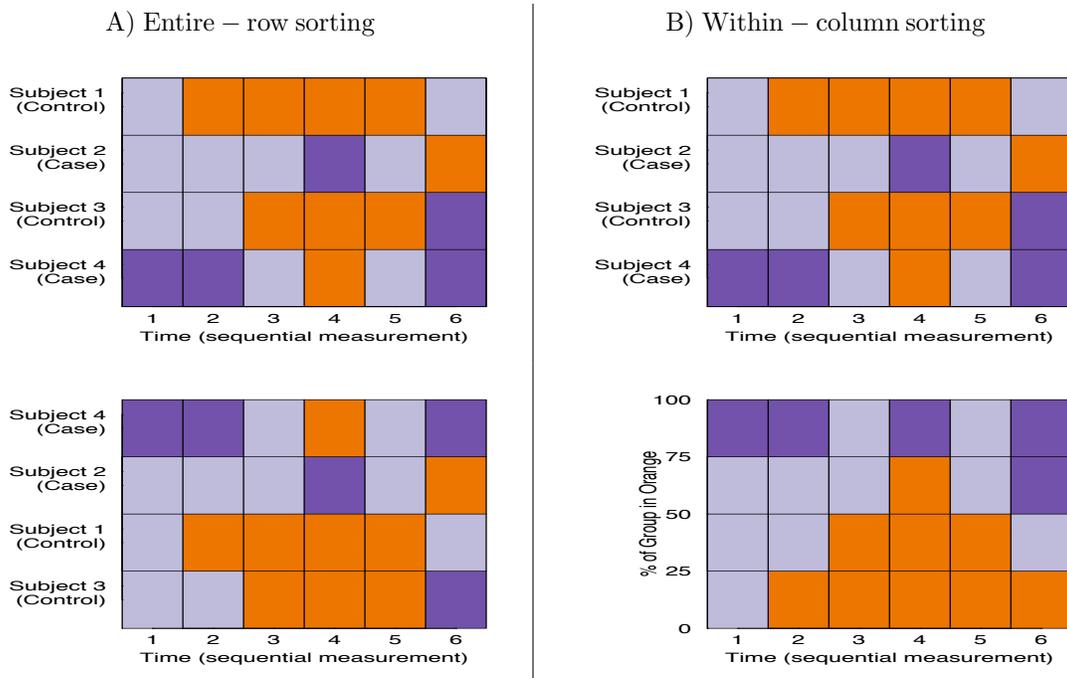


Figure 2: A) *Entire-row sorting* The order of subjects with respect to Case/Control status is initially unsorted in the top lasagna plot. Sorting on Case/Control status groups similar subjects together. Subject-specific info is retained, yet group comparisons are facilitated. For instance, it looks like the Controls have more time in the orange state. B) *Within-column sorting (vertical sorting)* Imagine the orange blocks can only be moved vertically and are “heavier” than the other colors and no longer have to stay with their subject (row) and sink to the bottom. Vertical sorting can be applied directly to the measures corresponding to the color at each measurement time across subjects. This sorting will release subject-specific info, but give group-level temporal patterns. For instance, sorting vertically reveals a symmetric distribution of an outcome centered at the fourth measurement.

2.2 Dynamic Sorting

Sortings are a way to rearrange data to better visualize potentially obscured information. The proposed sortings can be done on the original H or on a resultant sorting.

1. Within-row: sorts values of each subject in ascending order. This type of sorting preserves subject specific information, but no longer displays the temporal ordering.
2. Entire-row: sorts layers of subject in ascending order of a characteristic, which can be internal or external. An internal characteristic would be a feature in the lasagna plot, for example, the subject specific means of the outcome. An external characteristic could be baseline age of subjects. Entire-row sorting organizes cohorts together for analysis of cohort effects. This type of sorting preserves subject specific information (Figure 2).
3. Within-column: sorts values across subjects in ascending order within each epoch. This type of sorting no longer displays subject specific information, but may reveal group-level temporal patterns. This type of sorting is often called “vertical sorting” (Figure 2).
4. Entire-column: rearranges the columns of a lasagna plot in ascending order based on a characteristic, which can be internal or external. An internal characteristic is one apparent from the information displayed in the graph, such as mean outcome value across subjects at each measurement time. An external characteristic could be measurements of a second outcome variable. This type of sorting could be useful when the variable is seasonal, for example.
5. Cluster: A specific type of entire-row sorting where the characteristic is internal and a clustering algorithm such as hierarchical clustering or K-means is used so that subjects that have similar trajectories/layers are grouped together into clusters. Discrete outcomes can be ordered by strata or comparable clustering algorithms, such as correspondence analysis.

Whereas spaghetti plots are static and in the case of over-plotting stand to obscure trends, outliers, clusters, and data quality checks, a series of lasagna plots with sequential and cumulative sorting may allow the data to be clearly depicted. Subject-specific trends inherently are easier to see in a lasagna plot due to layers not overlapping as noodles do in a spaghetti plot, and with dynamic sorting trends of different cohorts (entire-row sorting on a classifying variable) and of the entire study population (within-column sorting) are possible. Outliers can be identified with various color spectra corresponding to threshold definitions as well as dynamic (entire-row) sorting. Clusters can be viewed by entire-row sorting on external variables or discovered by using cluster sorting. Data quality checks

are facilitated by the virtue that lasagna plots show each datum, unobscured by any other data. This unobscured depiction allows clear ascertainment of data collection patterns in a variety of different time scales, for example, changing the time variable from the subject-specific visit count to calendar month of visit.

3 Motivating Example

The details of the study design are described in further detail in the supplement (see Document, Supplemental Digital Content 1). Figure 3 of this commentary displays a spaghetti plots for one subject, for 118 subjects, the corresponding unordered lasagna plot for 118 subjects, and the corresponding entire-row sorted lasagna plot. To showcase the capability of lasagna plots displaying intermittent missing data, a pattern of missingness is artificially applied. To showcase the process of entire-row sorting, the outcome values between disease groups were artificially made more disparate.

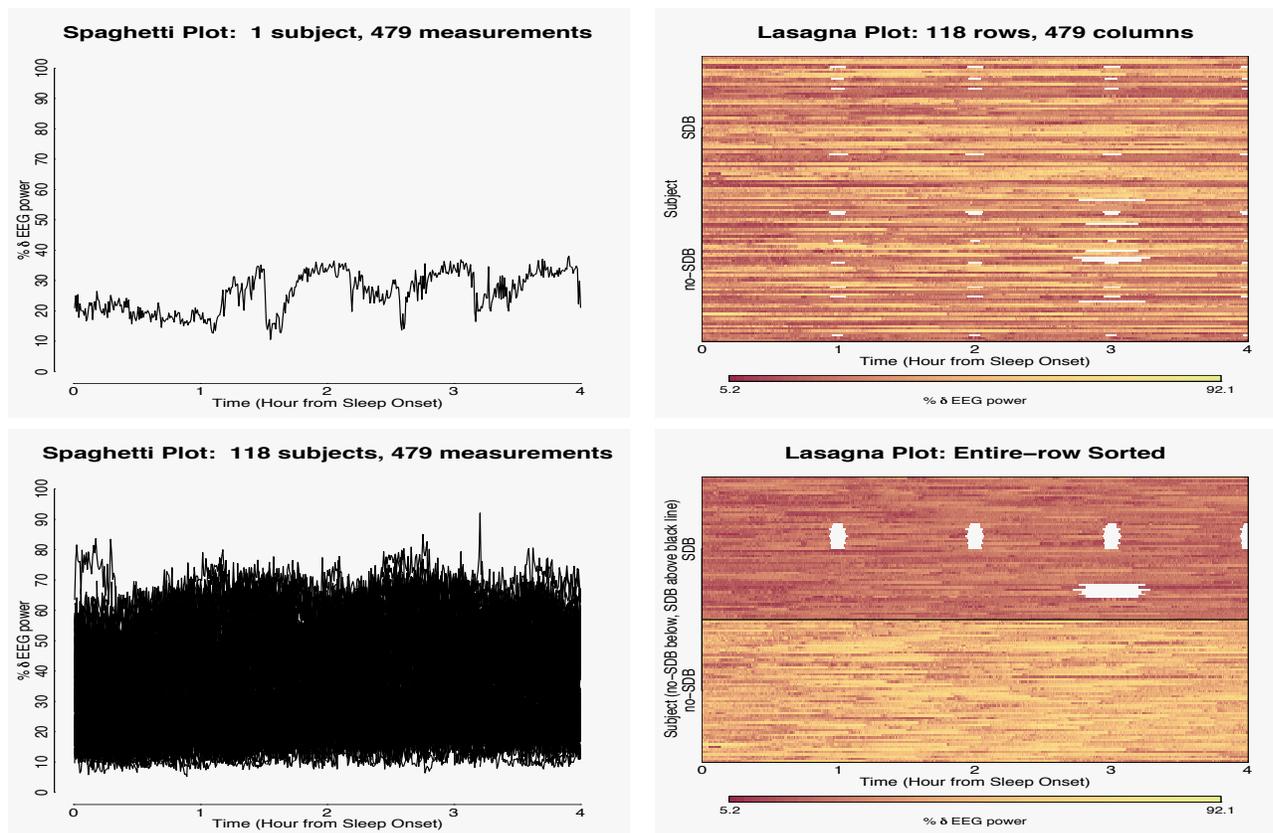


Figure 3:

Figure 3 suggests a classical spaghetti plot for one subject is informative, for 118 subjects, less so. The overlapping of multiple trajectories leads to an obscuring of trends for

a moderate number of subjects, and consequently the conveyance of intermittent missing data fails. In the lasagna plot for 118 subjects the subjects (rows) appearing in random order, but the intermittent missing data (off-white) is clearly conveyed. After an entire row sort on disease status and date of EEG recording within disease status, the intermittent missing data is not only conveyed, but the sort allows the exploration of possible trends. After the sort, the darker red region indicates that the diseased have less percent δ EEG power during sleep, it is seen that only the diseased have missing data, and that the recorder successfully recorded the first 19 SDB subjects, then malfunctioned for the next 11 recording dates in a way where it dropped measurements approximately every hour of sleep from onset. The recorder was righted and operated with full functionality for the next 14 SDB subjects, only to malfunction again by dropping measurements about three hours from sleep onset for the next six SDB subjects. The issue was addressed, and the recorder successfully recorded the rest of the SDB subjects. Comparing the two bottom panels of Figure 3, the same outcome information is in them, but lasagna plots more effectively depict the data because the non-overlapping of trajectories keeps the outcome information uncluttered and its sorting can incorporate more information. As was seen, the extra information can be intermittent missing data and cohort effects, and as outlined in subsequent examples (see Document, Supplemental Digital Content 1) it can also be in the form of informative censoring and practice effects.

In summary, we contend that lasagna plots stand to complement spaghetti plots in many aspects for graphically exploring epidemiologic longitudinal data. Coding for lasagna plots is easily done in R¹⁹ (see Documents, Supplemental Digital Content 2 and 3).

Supplemental Digital Content 1. Document with more examples and figures from Sleep Heart Health Study and the Former Lead Worker Study. pdf

Supplemental Digital Content 2. Document with R code snippets. pdf

The following document serves as an online supplement to “Lasagna plots: A saucy alternative to spaghetti plots.”

4 Additional Examples

We have used lasagna plots to aid the visualization of a number of unique disparate datasets, each presenting their own challenges to data exploration. Three examples from two epidemiologic studies are featured: the Sleep Heart Health Study (SHHS) and the Former Lead Workers Study (FLWS). The SHHS is a multicenter study on sleep-disordered breathing (SDB) and cardiovascular outcomes.²⁰ Subjects for the SHHS were recruited from ongoing cohort studies on respiratory and cardiovascular disease. Several biosignals for each of 6,414 subjects were collected in-home during sleep. Two biosignals are displayed here-in: the δ -power in the electroencephalogram (EEG) during sleep and the hypnogram. Both the δ -power and the hypnogram are stochastic processes. The former is a discrete-time continuous-outcome process representing the homeostatic drive for sleep and the latter a discrete-time discrete-outcome process depicting a subject’s trajectory through the rapid eye movement (REM), non-REM, and wake stages of sleep.

The FLWS is a study of age, lead exposure, and other predictors of cognitive decline. The study spans a decade, with up to seven study visits and three separate phases of data collection. During each visit, subjects participated in a battery of cognitive tests, resulting in a longitudinal dataset of repeated measures of test scores for each subject.

In the SHHS, we explore the data by disease status, looking for distinguishing patterns within each disease group. The disease under consideration is sleep apnea (or sleep-disordered breathing (SDB)), a condition characterized with repetitive breathing pauses during sleep. Comparing groups requires carefully selected subsamples, and thus our focus is on 59 SDB subjects and 59 subjects without SDB (no-SDB). In both the continuous and discrete outcome examples, our data assume a wide format, where the number of measurements far exceed the number of subjects. In the FLWS, all 1,110 subjects are analyzed over a maximum of 7 visits. Cluster sorting will help evaluate the presence of two common problems for longitudinal studies of cognitive function: informative censoring and “practice effects.”

4.1 Continuous Longitudinal Data with Intermittent Missingness

For continuous electroencephalogram (EEG) signals derived from the sleep studies, four distinct frequency powers are typically discerned via band-pass filters on the Fourier transform: α, β, δ , and θ . Percent δ -power is defined as $\frac{\delta}{\alpha + \beta + \delta + \theta} \times 100$. For every 30 seconds during sleep, percent δ -power was calculated for 59 SDB subjects and 59 no-SDB. In this introductory illustrative example, we look at only the first four hours of data for each sub-

ject, so that everyone has a common onset and stopping point. We also assume that the same device was used to record sleep across different subjects and thus no two subjects had sleep recorded on the same date. To showcase the capability of displaying intermittent missing data of the lasagna plot, a pattern of missingness is artificially applied. Via dynamic sorting, the pattern of missingness will be revealed, illustrating how patterns can be uncovered with this exploratory data analysis technique of sorting and visualizing. To showcase the process of entire-row sorting, the outcome values between disease groups were artificially made more disparate.

We see that a spaghetti plot is a salient display of data for one subject, but not for 118 (Figure 4). The corresponding lasagna plot for 118 subjects shows intermittent missing data, and upon entire-row sorting on the external factors of disease status and date of EEG recording reveals intriguing patterns of the missingness, as well as disease-group differences in percent δ sleep (Figure 5). It appears that only subjects with SDB have missing data and that for a period of recording dates measurements were dropped hourly. Possibly the recording device was malfunctioning, subsequently fixed, and then enjoyed a period of proper functionality only to succumb to dropping measurements 3 hours after sleep onset before being repaired again. To explore the group-level characteristics of percent δ -power during sleep evolution over the course of the night, an additional within-column sorting is conducted within disease status (Figure 6). The resulting lasagna plot from the within-column sorting highlights a temporal undulation to the signal of the no-SDB group, as well as the no-SDB group having generally higher percent δ -power during sleep than the SDB group. Delta power in the sleep EEG is thought to have an important positive association with cognition and is a maker for homeostatic sleep drive.²¹

4.2 Discrete State-Time Data with a Common Onset

The three classifications of sleep stages (Wake, REM, and Non-REM) are discretizations of several continuous physiologic acquired during sleep. The EEG signals are binned into epochs (often 30 seconds) from sleep onset and collectively used to determine the stage of sleep. To accommodate the different lengths of sleep time, an absorbing state is utilized to ensure each subject has an equal number of “measurements,” which aids visualization. This example showcases data from the SHHS, where 59 diseased subjects were matched on age, BMI, race and sex to 59 non-diseased subjects. Because the outcome is discrete (the state of sleep), the spaghetti plot is a state-time plot specifically known as the hypnogram to sleep physicians. As in the previous example, for one subject, the spaghetti plot shows the durations in states and transitions among states clearly. The subsequent spaghetti plot for all 118 subjects falls prey to over-plotting, limiting its informativeness (Figure 7). A lasagna plot shows the 1,031 outcomes of each of the 118 subjects’ trajectories in random order with respect to SDB status (Figure 8). Applying an entire-row sort on the external

characteristic of disease status and the internal characteristic of overall sleep time shows that the groups are well matched on total sleep time (even though the two groups were not explicitly matched on total sleep time). Note the degree of fragmentation and the frequency of short and long-term bouts of WAKE of those with SDB compared to controls. The difference between the two groups in the degree of fragmentation as visualized indicates there might be a difference in sleep continuity between the two groups, supporting the well-established link between SDB and fragmented sleep. It has been conjectured that sleep continuity may be important in the recuperative effects of sleep, especially in the study of sleep disordered breathing (SDB) and its impact on health outcomes.^{22;23;24;25} Applying an additional within-column sorting within disease status shows the difference in REM temporal evolutions among groups (Figure 9). The dynamic sorting of Figure 9 shows the SDB group having an overall weaker REM signal, a bimodal first peak, an absence of a peak at hour 3, the presence of a peak at ~ 7.75 hours. In addition, the peaks widen as time increases, which backs empirical findings of REM duration in state time lengthening as the overall sleep progresses.

4.3 Discretized Longitudinal Data

Lasagna plots are also useful in visualizing and detecting many of the common challenges to population-based longitudinal cohort studies in epidemiologic research. The FLWS is a study of age, lead exposure, and other predictors of cognitive decline. The study spans a decade, with up to seven study visits and three separate phases (tours) of data collection. The FLWS is a complex dataset beset with missing data, and both left and right censoring of subjects. Because subjects were enrolled over time in multiple tours, subjects could have as few as 2 or as many as 7 study visits. Study dropout is likely to be dependent upon outcome status (declines in neurobehavioral function) resulting in informative censoring. An additional challenge is the problem of a “practice effect”: scores on neurobehavioral tests of cognitive function can become better through practice, masking real declines in cognitive abilities. Lasagna plots provide a unique opportunity to visualize these complex data and detect evidence of both informative censoring and a learning effect. In order to do so, lasagna plots are made with visit as the unit of time as well as tour. Each reveals temporal patterns.

The spaghetti plot (Figure 10) for 1,110 subjects over seven visits is over-plotted, but does show a “thinning” of subjects, suggesting many had three visits, distinctly fewer had three to six, and fewer than that had all seven. In order to facilitate detection of potential informative censoring or a learning effect, neurobehavioral scores were binned based on quintiles of the first visit score distribution. The spaghetti plot of the binned quintile data (Figure 11) is over-plotted and uninformative because the number of subjects for each trajectory is not discernable. A classical spaghetti plot of discrete outcomes on the Y axis

can show possible trajectories, but no indication of how many subjects are in the study due to the exact overlapping of trajectories. The lasagna plot shows the loss to follow up for subjects over time even more clearly than the spaghetti plot (Figure 12). A cluster sort (sorting within the first column, then the second, etc.) allows us to move entire-rows so that subjects with similar trajectories are closer to one another (Figure 13). Immediately, we can identify a cluster that did not have a value reported for a first visit, but had values for subsequent visits, indicating missing data. These findings highlight the utility of lasagna plots for exploratory data analysis and data validation. The lasagna plot can also help in examining data for informative dropout and practice effects. Here it appears that if one is in the bottom (worst) quintile on the first visit, the loss to follow up is much worse than if one was in the top (best) quintile at visit 1, indicating informative dropout. A practice effect can be discerned crudely if the subjects have higher test scores on their second study visit than their first, and then scores subsequently decline over time. Overall trends in cognitive function over time are apparent as the amount of lighter colors decrease from left to right and the amount of darker colors increase, empirically confirming the overall decline in cognitive function observed with aging.

Finally, if an additional within-column sort was conducted, we derive the classic stacked bar chart (Figure 14). The classic stacked bar chart removes all subject-specific trajectories and instead summarizes overall distributions of neurobehavioral test scores for each study visit. The practice effect is most visible as scores (based on quintiles of the first visit scores) appear to jump up between the first and second study visits, and then decline over time. However, from Figure 14, we cannot ascertain that the subjects in the top quintile on the first visit are in the top quintile on the second visit. Using stacked bar charts prevents statements on typical pathways, whereas a cluster sorted lasagna plot displays the trajectories for full viewing.

The time structure is complex, for subjects' visit 1 measurement may not have taken place in the same tour. Also, the amount of time lapsed between one subject's adjacent visits may not be the same as another subject's due to visit number being interlinked with what tour they enrolled. Analyzing informative censoring and practice effects is further facilitated by making a lasagna plot with tour as the time variable (Figure 15) and then sorting within each tour the subject's 1st visit quintile cognitive measure (Figure 16). Comparing those enrolled in Tour 1 of the worst and best quintile, we see that there is more dropout for those starting out in the worst quintile, possibly indicating informative censoring. The pattern of drop out being related to first visit quintile rank holds for the later tours as well. Training effects can be seen when a subject's color lightens when tracking that subject across time. For instance, a fair portion of subjects in Tour 1 were in the 2nd best quintile and advance to the best quintile in their second visit in Tour 2.

4.4 Result Tables and Covariate Selection

Simulations under different conditions often give rise to multiple tables of output. Identifying trends and comparing tables is often an arduous and obfuscating task. With lasagna plots, a quick snapshot of the tables are rendered, allowing trends within tables to be identified and compared across tables (Figure 17).

In building regression models, it is important to know what variables have high degrees of missingness. For large epidemiologic datasets modeling an outcome, a lasagna plot can be used to show the proportion of the sample covariate missingness over time (Figure 18). Here, the vertical axis is the variable, the horizontal axis is time, and the darker the plot the greater the proportion of the sample that has a reported value for that layer's variable. The plot in Figure 18 helped guide the inclusion and exclusion of covariates in the model building process.

Lasagna plots work well for data tables that have many numbers and are essentially an image of a matrix. Commonly, the layers are denoting an subject, the columns are denoting times or locations largely in common to all the subjects being visualized, and color to reflect the state occupied or magnitude/intensity of the trait. One exception to this paradigm is diary data for a subject. In diary data, the layers are days, the columns are hours, and the colors reflect activities partaken for a certain time on a particular day. The approach just described for diary data has proven useful in mapping out infant and child ideal sleep patterns.²⁶ Nutritionist colleagues are implementing lasagna plots to display caloric intake and purge cycles amongst those with eating disorders.

5 Discussion

Lasagna plots have been presented as an effective means to explore data that can be arranged into a matrix. The strengths of lasagna plotting are that it can incorporate a wide platform of data structures, ranging from longitudinal repeated measures (i.e., dominos and nonsimultaneous chains in kidney paired donation) to multidimensional temporal-spatial (i.e., FMRI) to gene expression of genes by tissue type (i.e. Barcodes).^{27;7;6} Lasagna plots are visualizations that “above all else show the data” and are more akin to the raw data than a modeling procedure.^{28;29} Row and column sorting and clustering are intuitive to a non-technical audience, and visualizations of sequential sortings and/or clustering serve as a way to engage a collaborative analysis of data. Weaknesses include the color-dependency, difficulties handling continuous time, and how growing size of the data make seeing individual layers more difficult. The first is becoming less of an issue as digital publication overtakes traditional paper publishing. The second weakness can be ameliorated by coarsening/binning time, and the third by doing sorts or making plots on subsets of the population.

Often, longitudinal data have been traditionally viewed as either a spaghetti plot or stacked bar chart, which falls prey to over-plotting and aggressive summarization, respectively. Multi-state survival (event history) data can be viewed through a longitudinal repeated measures lens, and historically were viewed with eventcharts, corresponding components of dynamic interaction and linked graphs, as well as event history graphs. These were important steps in visualizing survival data simultaneously at the group level and subject level.^{30;31;32} Limitations of the eventchart included difficulty handling multiple groups, large numbers of subjects, denoting multiple events, and the incorporation of color. These limitations are not present with the lasagna plotting of survival data. Lasagna plots work well in most trivariate and multiway data settings, conveying at least the same information as superposed level regions in color plots and multiway dot plots.³³

Genomics and pediatric sleep science are currently utilizing plots that are special cases of what we call lasagna plots. Recent graphing techniques in the statistical programming and analysis language R have come to show ingenuity in handling complex data, as evinced by the non-exhaustive list of R packages `lattice`, `ggplot2`, `seas`, `mvtpplot`, and `gplots` (see functions `heatmap.2()` and `hist2d()` for plotting similar to that of lasagna plotting).^{34;35;36;37;38} Lasagna plotting and cluster sorting is implemented by `heatmap.2()`, grouping similar genes (rows) and tissues (columns) together. The essence of lasagna plotting, within-column and entire-row sorting is captured in `mvtpplot` the best, in that it displays not only the data itself but simultaneously group level temporal trends in a smoothed curve below and subject specific summary measures on the right sidebar. Lasagna plotting encompasses `mvtpplot` and implements dynamic sorting to further explore the data. Discrete outcomes are not handled well by `mvtpplot` because the element being visualized are not necessarily numeric, thus the summarizations on the bottom and right hand panel are not useful when the data measures are nominal. Lasagna plotting and subsequent sorting handles the nominal case.

References

- [1] PJ Diggle, PJ Heagerty, KY Liang, and SL Zeger. *The analysis of longitudinal data 2nd ed.* Oxford, England: Oxford University Press, 2002.
- [2] E.J. Wegman. Hyperdimensional data analysis using parallel coordinates. *Journal of the American Statistical Association*, 85(411):664–675, 1990.
- [3] E.J. Wegman and Q. Luo. High dimensional clustering using parallel coordinates and the grand tour. In *Classification and knowledge organization: proceedings of the 20th annual conference of the Gesellschaft für Klassifikation eV University of Freiburg, March 6-8, 1996*, page 93. Springer Verlag, 1997.
- [4] E.J. Wegman. Data mining and visualization: some strategies. *Bull Int Statist Inst*, 52:223–226, 1999.
- [5] L. Wilkinson and M. Friendly. The history of the cluster heat map. *The American Statistician*, 2009.
- [6] M.J. Zilliox and R.A. Irizarry. A gene expression bar code for microarray data. *Nature Methods*, 4:911–913, 2007.
- [7] R. Baumgartner and R. Somorjai. Graphical display of fMRI data: visualizing multi-dimensional space. *Magnetic resonance imaging*, 19(2):283–286, 2001.
- [8] D. Sarkar. *Lattice: multivariate data visualization with R*. Springer Verlag, 2008.
- [9] D. Hedeker and R.D. Gibbons. *Longitudinal data analysis*. Wiley-Interscience, 2006.
- [10] G.M. Fitzmaurice, M. Davidian, G. Verbeke, and G. Molenberghs. *Longitudinal Data Analysis: A Handbook of Modern Statistical Methods*. Chapman & Hall/CRC, 2008.
- [11] G.M. Fitzmaurice, N.M. Laird, and J.H. Ware. *Applied longitudinal analysis*. ASA, 2004.
- [12] C. Poynton. Frequently asked questions about color. URL: <http://www.inforamp.net/poynton/PDFs/ColorFAQ.pdf>, 1999.
- [13] R. Ihaka. Colour for presentation graphics. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing, Vienna, Austria, ISSN*. Citeseer, 2003.
- [14] T. Lumley. Color coding and color blindness in statistical graphics. *ASA Statistical Computing & Graphics Newsletter*, 17(2):4–7, 2006.
- [15] A. Zeileis, K. Hornik, and P. Murrell. Escaping RGBland: Selecting colors for statistical graphics. *Computational Statistics and Data Analysis*, 53(9):3259–3270, 2009.

- [16] R. Ihaka, P. Murrell, K. Hornik, and A. Zeileis. *colorspace: Color Space Manipulation. R package version*, pages 1–0, 2008.
- [17] M. Harrower and CA Brewer. Colorbrewer.org: an online tool for selecting colour schemes for maps. *Cartographic Journal, The*, 40(1):27–37, 2003.
- [18] Erich Neuwirth. *RColorBrewer: ColorBrewer palettes*, 2007. R package version 1.0-2.
- [19] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. ISBN 3-900051-07-0.
- [20] SF Quan, BV Howard, C. Iber, JP Kiley, FJ Nieto, GT O’Connor, DM Rapoport, S. Redline, J. Robbins, JM Samet, et al. The Sleep Heart Health Study: design, rationale, and methods. *Sleep*, 20(12):1077–85, 1997.
- [21] L. Marshall, H. Helgad, M. Matthias, and J. Born. Boosting slow oscillations during sleep potentiates memory. *Nature*, 444(7119):610–613, 2006.
- [22] R.G. Norman, M.A. Scott, I. Ayappa, J.A. Walsleben, and D.M. Rapoport. Sleep continuity measured by survival curve analysis. *Sleep*, 29(12):1625–31, 2006.
- [23] N.M. Punjabi, D.J. O’Hearn, D.N. Neubauer, F.J. Nieto, A.R. Schwartz, P.L. Smith, and K. Bandeen-Roche. Modeling hypersomnolence in sleep-disordered breathing a novel approach using survival analysis. *American Journal of Respiratory and Critical Care Medicine*, 159(6):1703–1709, 1999.
- [24] M.H. Bonnet and D.L. Arand. Clinical effects of sleep fragmentation versus sleep deprivation. *Sleep Medicine Reviews*, 7(4):297–310, 2003.
- [25] B.J. Swihart, B. Caffo, K. Bandeen-Roche, and N.M. Punjabi. Characterizing Sleep Structure Using the Hypnogram. *Journal of Clinical Sleep Medicine: JCSM: official publication of the American Academy of Sleep Medicine*, 4(4):349, 2008.
- [26] R. Ferber. *Solve Your Child’s Sleep Problems: New, Revised, and Expanded Edition (Paperback)*. Simon & Schuster, 2006.
- [27] SE Gentry, RA Montgomery, BJ Swihart, and DL Segev. The Roles of Dominos and Nonsimultaneous Chains in Kidney Paired Donation. *American Journal of Transplantation*, 9(6):1330–1336, 2009.
- [28] E.R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2001.
- [29] T. Lumley and P. Heagerty. Graphical exploratory analysis of survival data. *Journal of Computational and Graphical Statistics*, pages 738–749, 2000.
- [30] A.I. Goldman. Eventcharts: Visualizing survival and other timed-events data. *American Statistician*, pages 13–18, 1992.

- [31] E.N. Atkinson. Interactive dynamic graphics for exploratory survival analysis. *American Statistician*, pages 77–84, 1995.
- [32] J.A. Dubin, H.G. Muller, and J.L. Wang. Event history graphs for censored survival data. *Statistics in Medicine*, 20(19):2951–2964, 2001.
- [33] W.S. Cleveland. *Visualizing data*. Hobart Press, 1993.
- [34] Deepayan Sarkar. *lattice: Lattice Graphics*, 2008. R package version 0.17-15.
- [35] Hadley Wickham. *ggplot2: An implementation of the Grammar of Graphics*, 2009. R package version 0.8.3.
- [36] Michael W. Toews, Paul H. Whitfield, and Diana M. Allen. Seasonal statistics: the 'seas' package for r. *Computers & Geosciences*, 33(7):1895, 2007.
- [37] RD Peng. A method for visualizing multivariate time series data. *Journal of Statistical Software*, 25(Code Snippet 1):687–713, 2008.
- [38] Gregory R. Warnes. Includes R source code and/or documentation contributed by Ben Bolker and Thomas Lumley. *gplots: Various R programming tools for plotting data*. R package version 2.6.0.

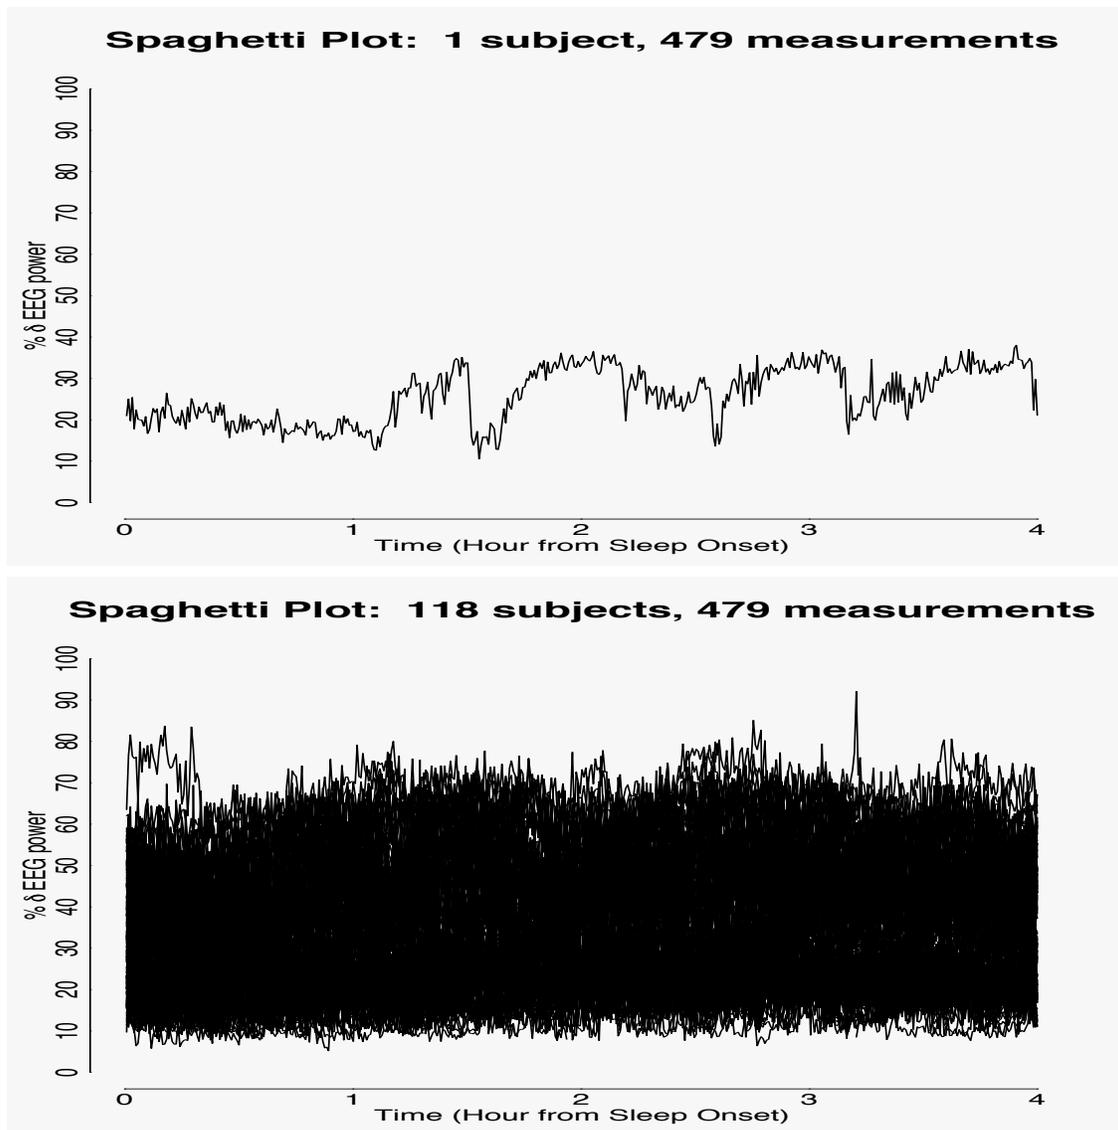


Figure 4: Top panel: Spaghetti plot for one subject. Bottom panel: Spaghetti plot for 118 subjects. The overlapping of multiple trajectories leads to an obscuring of trends for a moderate number of subjects, and consequently the conveyance of intermittent missing data fails.

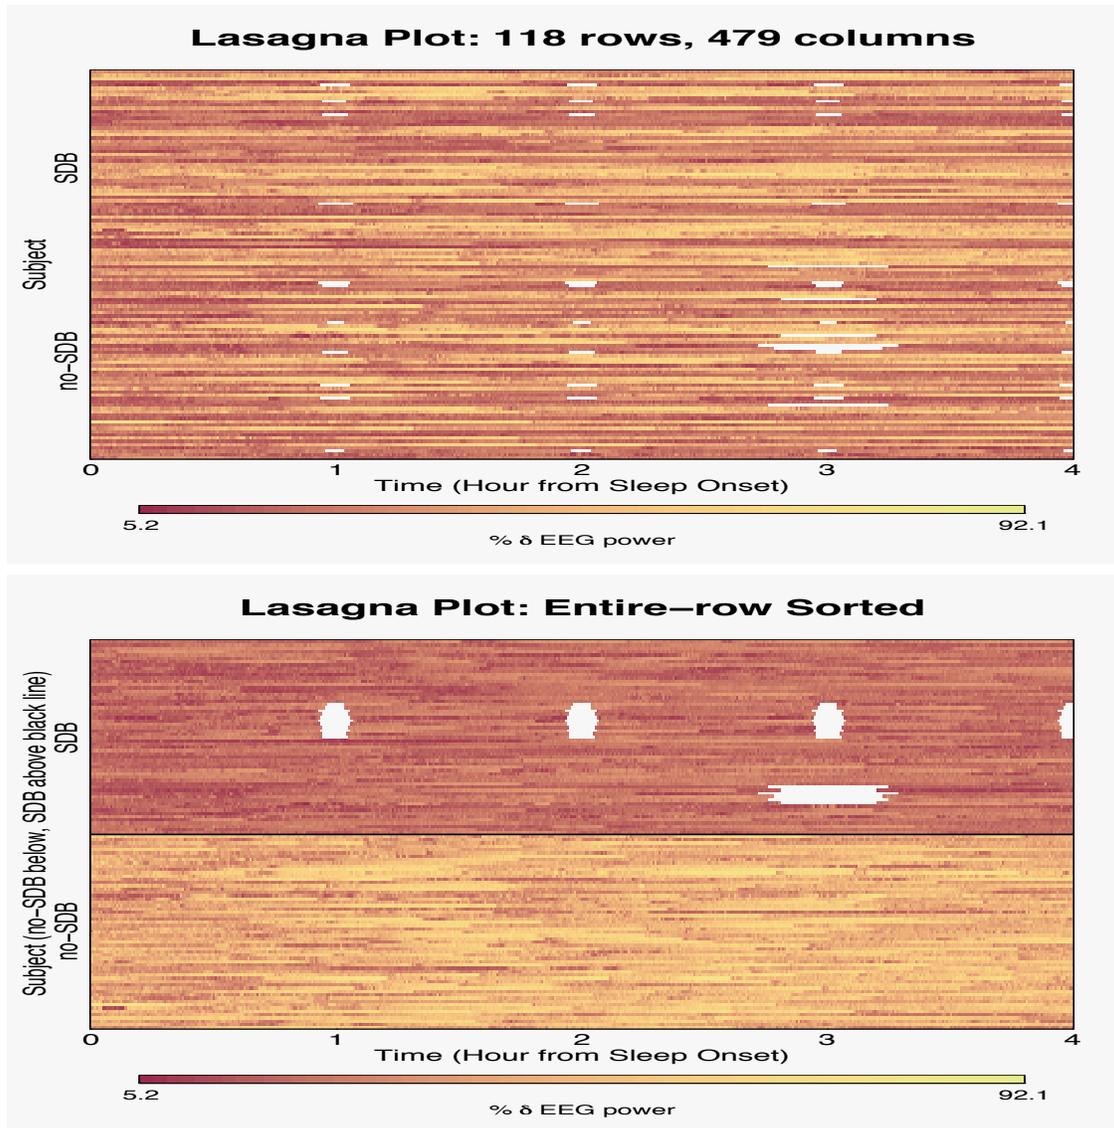


Figure 5: Top panel: Lasagna plot for 118 subjects from the bottom panel of Figure 4. The subjects (rows) appear in random order, but the intermittent missing data (white) is clearly conveyed. Bottom panel: Lasagna plot of the top panel after an entire row sort on disease status and date of EEG recording within disease status. The intermittent missing data is not only conveyed, but the sort allows the exploration of possible trends. After the sort, the darker red region indicates that the disease have less percent δ sleep, it is seen that only the diseased have missing data, and that the recorder successfully recorded the first 19 SDB subjects, then malfunctioned for the next 11 recording dates in a way where it dropped measurements approximately every hour. The recorder was righted and operated with full functionality for the next 14 SDB subjects, only to malfunction again by dropping measurements about three hours from sleep onset for the next six SDB subjects. The issue was addressed, and the recorder successfully recorded the rest of the SDB subjects. Compare the the bottom panel her to that of Figure 4. The same outcome information is contained in them, but lasagna plots more effectively depict the data because the non-overlapping of trajectories keeps the outcome information uncluttered and its sorting can incorporate more information.

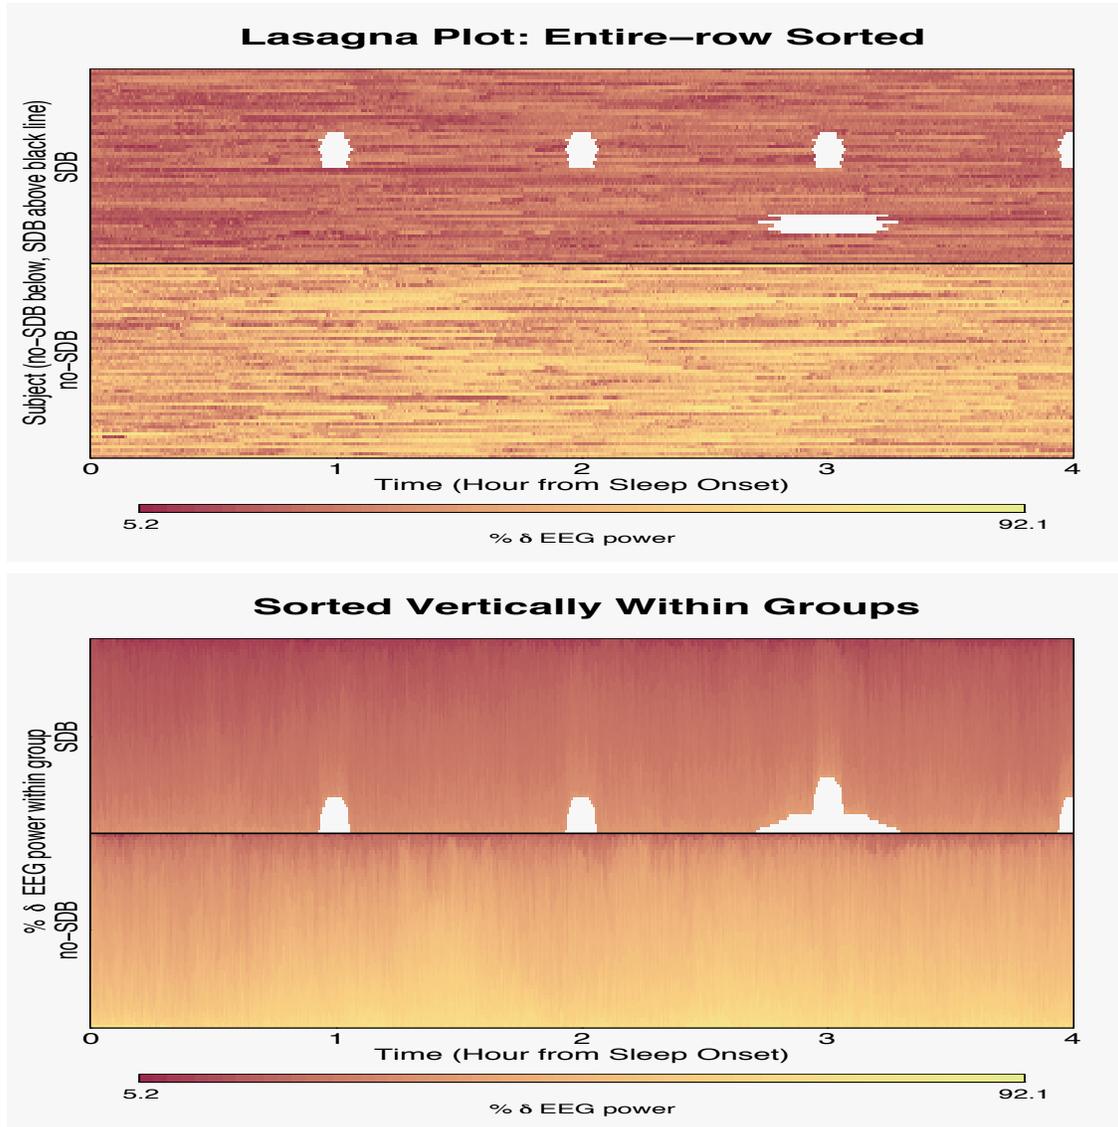


Figure 6: Top panel: The entire row sorted lasagna plot of Figure 5. Bottom panel: A within-column sort applied within disease status to the lasagna plot of the top panel. Note the waxing and waning of the yellow in the no-SDB group, depicting the group-level temporal evolution of percent δ sleep in subjects without SDB.

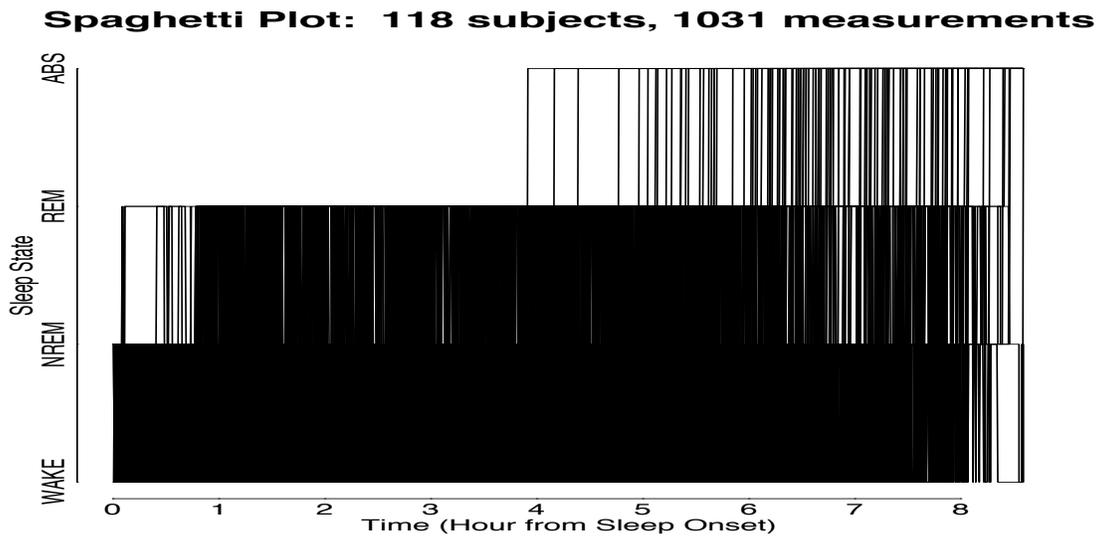
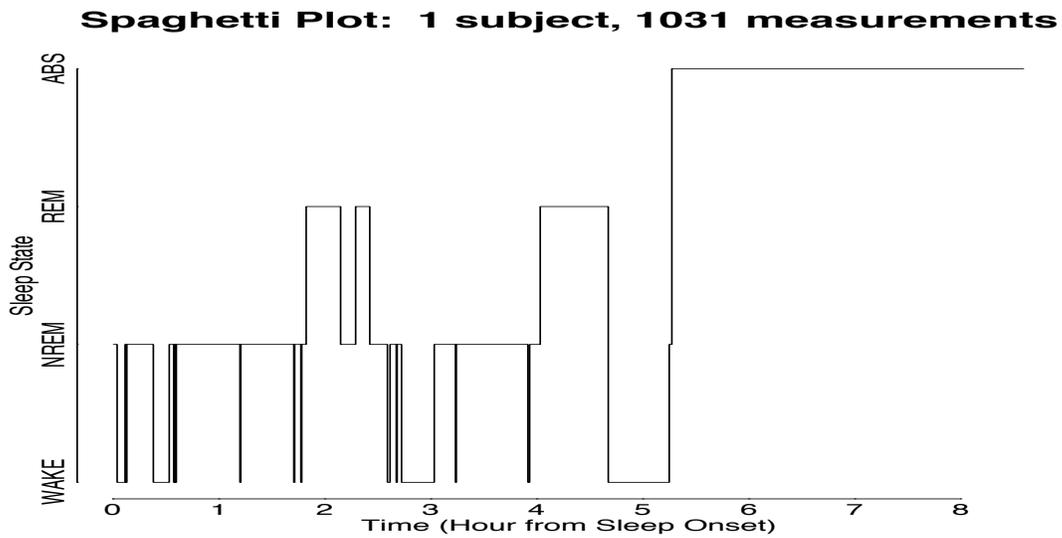


Figure 7: Top panel: a spaghetti plot for a discrete outcome for one subject. Bottom panel: a spaghetti plot for a discrete outcome for 118 subjects. Due to the discrete nature of the outcome, trajectories do not run the risk of merely crossing each other as in continuous outcome cases, but overlapping each other exactly. The informativeness of the spaghetti plot for discrete data on a moderate number of subjects is limited.

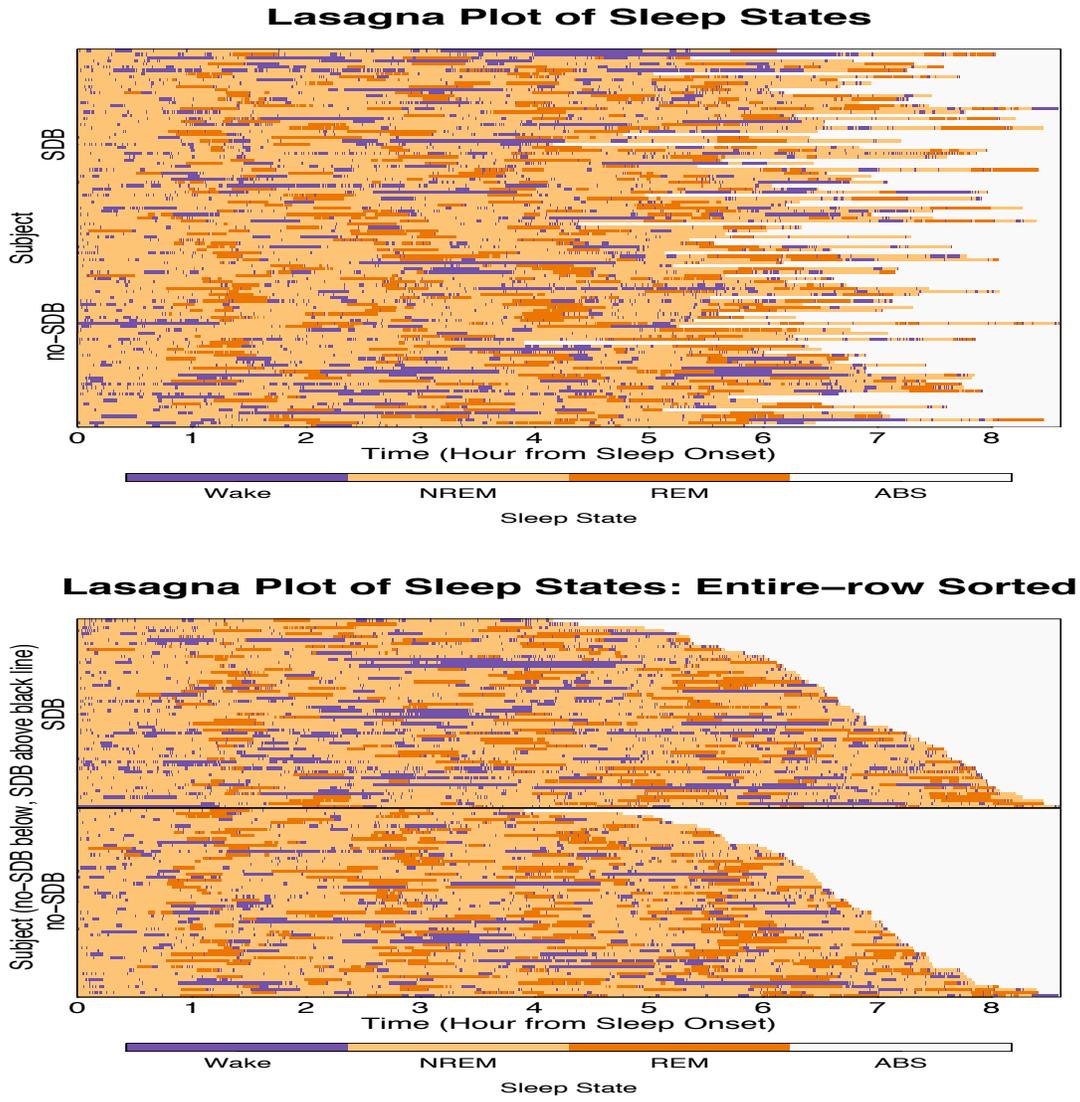


Figure 8: Top panel: corresponding lasagna plot for the spaghetti plot in the bottom panel of Figure 7, with subjects (rows) in random order. Bottom panel: the resulting lasagna plot after an entire row sort of disease status and sleep time recorded. The above data organization allows easy comparison of the absorbing state (ABS) areas showing that each group has similar distribution of sleep time recorded.

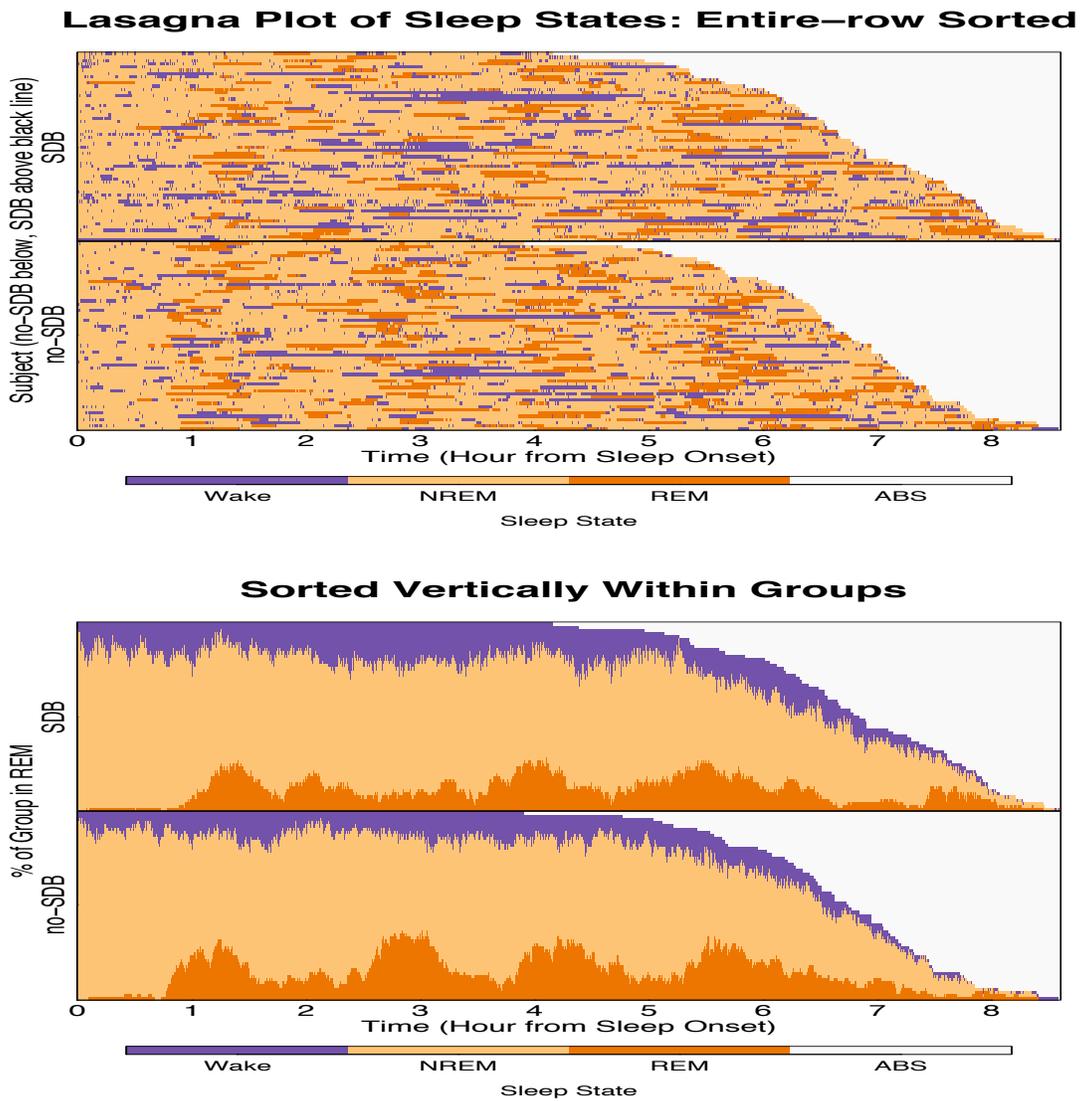


Figure 9: Top panel: same plot as the bottom panel of Figure 8. Bottom panel: the resulting lasagna plot after a within column sort applied within disease status. The above data organization shows group-level temporal evolution of REM sleep. The signal seems to be more pronounced in those without SDB.

Spaghetti Plot of Former Lead Workers Study

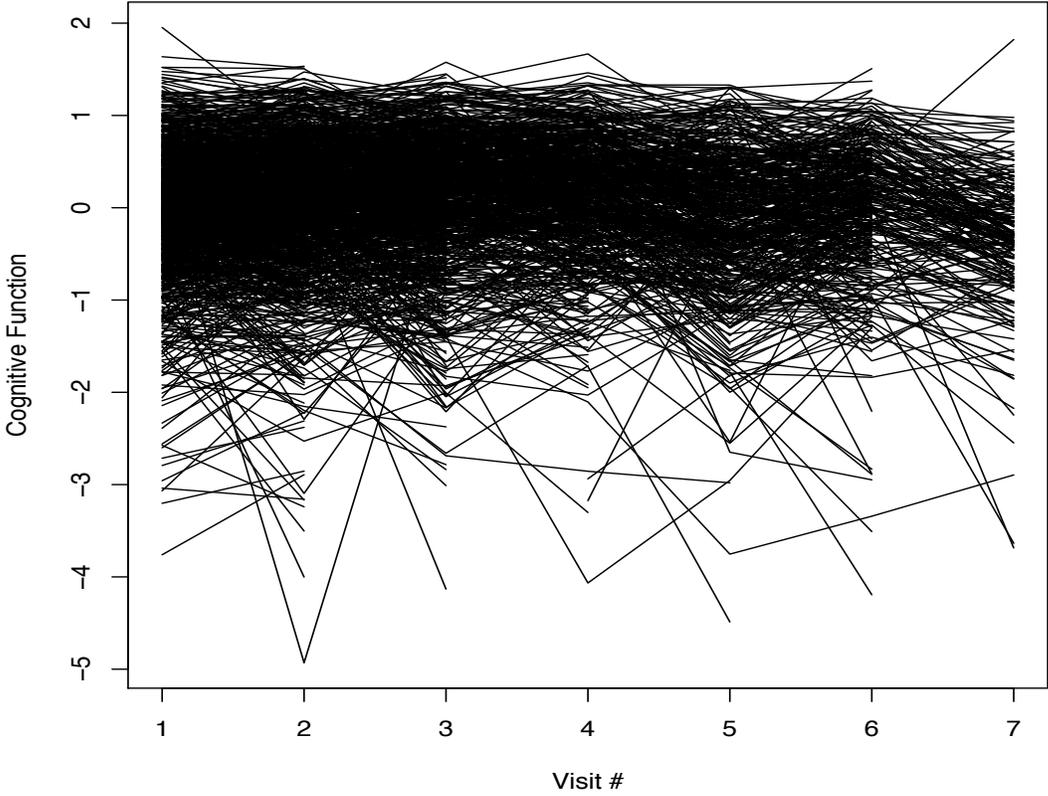


Figure 10: Spaghetti plot for a continuous cognitive measure of 1110 subjects over 7 visits.

Spaghetti Plot of Former Lead Workers Study -- Discretized

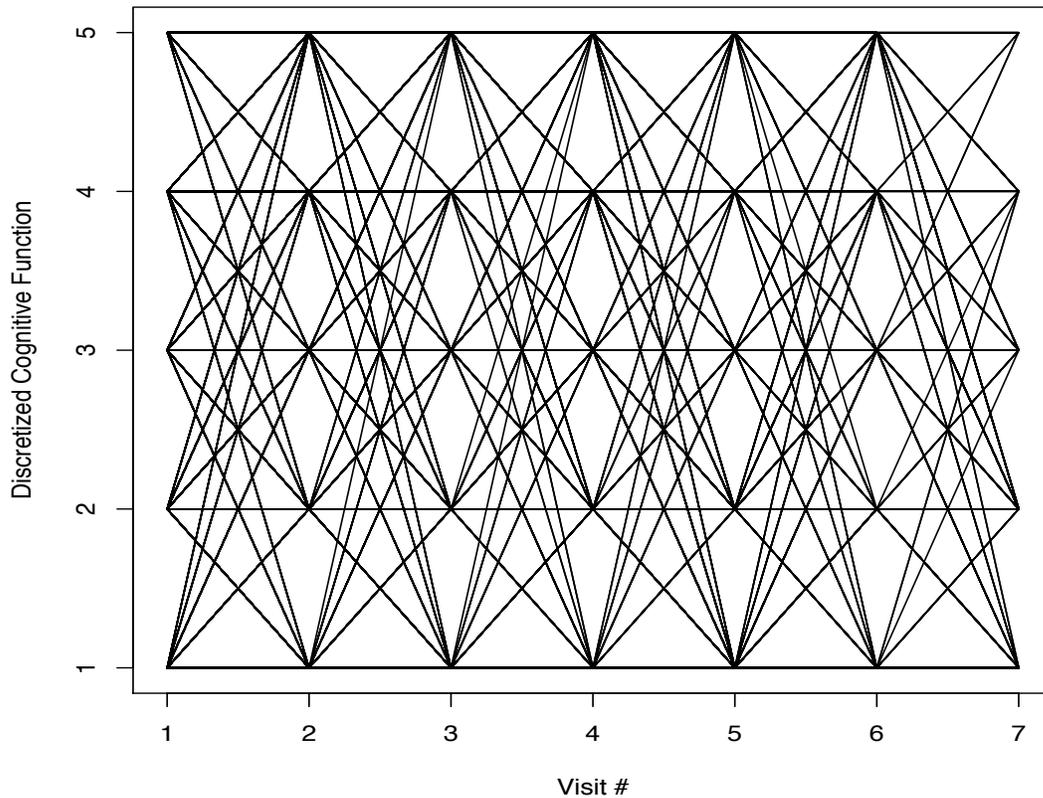


Figure 11: Spaghetti plot for the discretized cognitive measure of Figure 10 for 1110 subjects over 7 visits. The discretization was based on quintiles of the 1st visit outcome measures. The only information this plot can guarantee is that if a line exists between quintile nodes for adjacent visits, at least one subject made that move - it does not show how many subjects took the path, and it cannot show specific paths over multiple nodes for one subject. For instance, notice the absence of a line connecting the 1st visit 5th quintile node to the 1st quintile node of visit 7 (there is no line going from the upper left of the graph to the lower right). The absence of a line means no subject was recorded on only the first visit and the last visit with no visits between with the measurements recorded having her start out in the top quintile and declining to the bottom quintile. The absence of a line means the path was not taken. However, in a classical spaghetti plot of discretized data, the presence of a line over multiple nodes does not indicate that the path was taken by a subject. For instance, no subject only had two measurements taken on visit 1 and visit 5 and went from the top quintile to the bottom quintile, yet there is a line between 1st visit 5th quintile node to the 5th visit 5th quintile node. One cannot tell from the spaghetti plot alone if a path is made of one subject between two non-adjacent nodes, or several subjects making the pairwise adjacent transitions. For instance, the line between 1st visit 5th quintile node to the 5th visit 5th quintile node could comprise four subjects: one going from the 5th quintile to the 4th from visit 1 to visit 2, one going from the 4th quintile to the 3rd from visit 2 to visit 3, one going from the 3rd quintile to the 2nd from visit 3 to visit 4, and one going from the 2nd quintile to the 1st from visit 4 to visit 5.

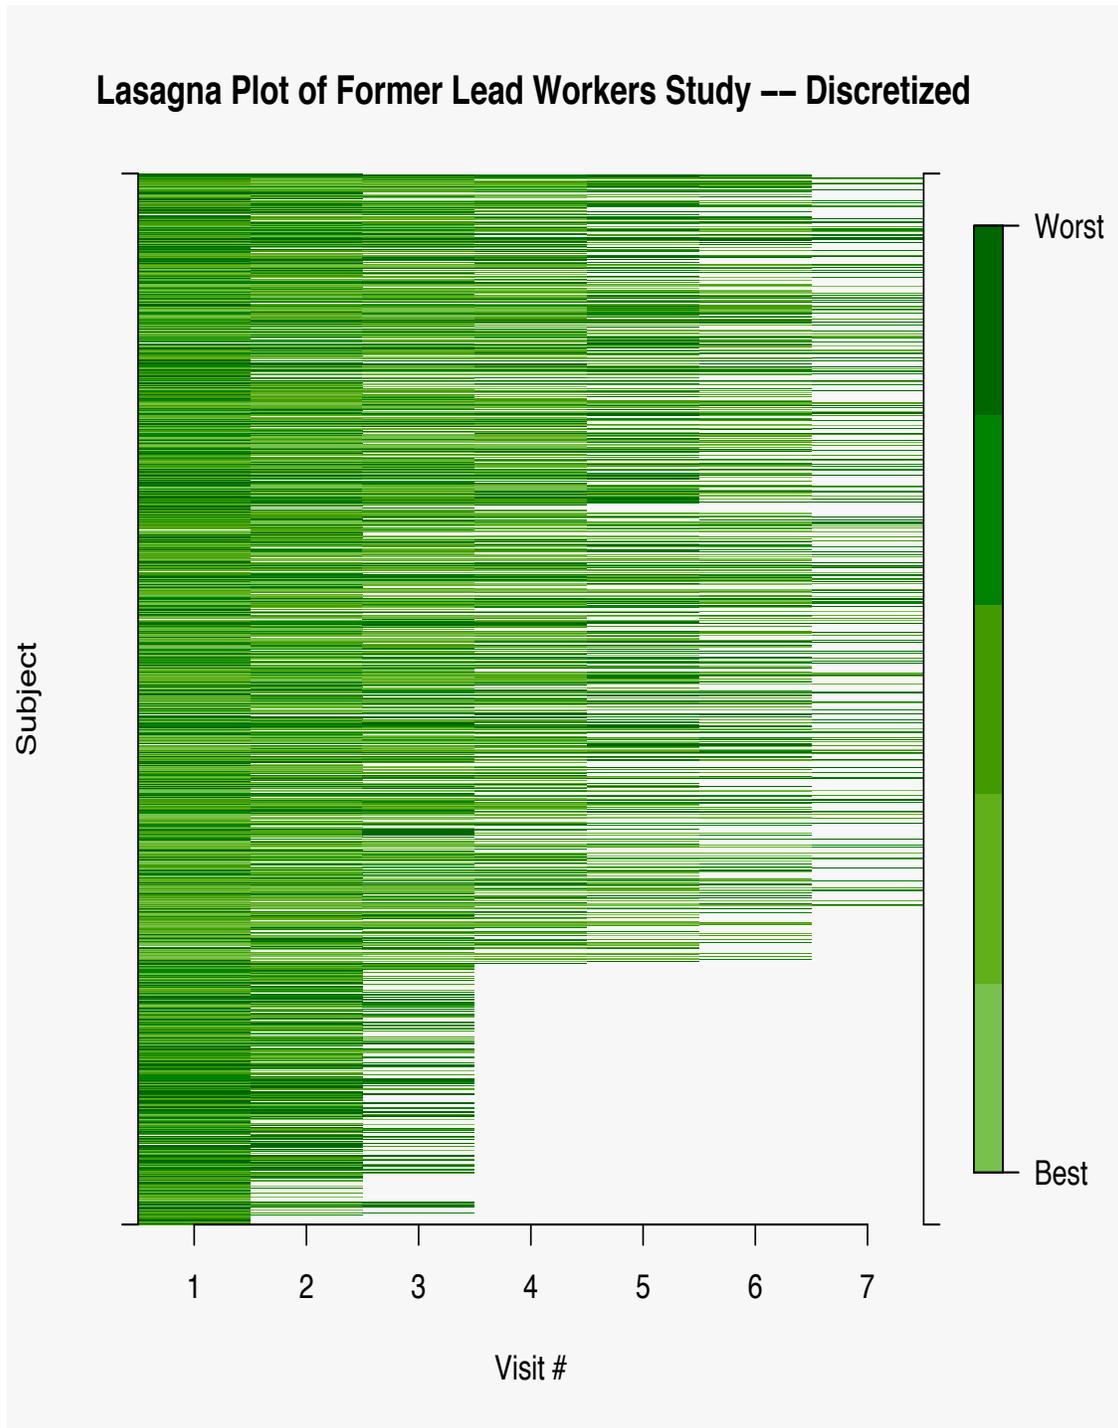


Figure 12: Lasagna plot for 1110 subjects over 7 visits, from Figure 11. The above image depicts paths taken by subjects more clearly than the discretized spaghetti plot.

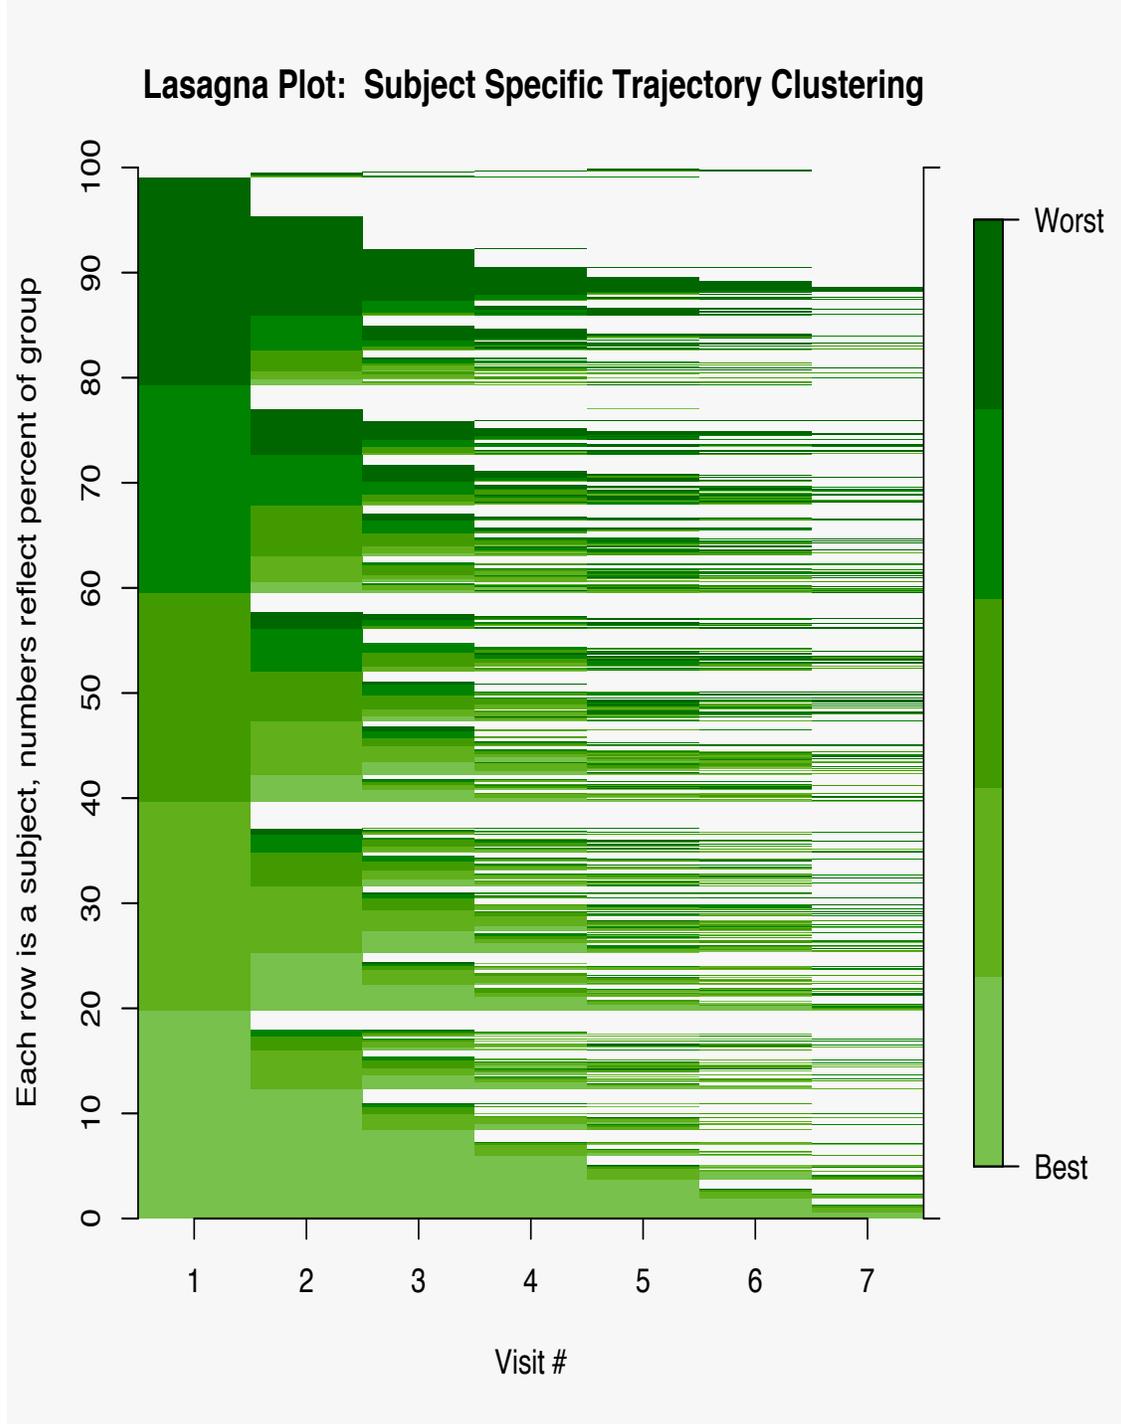


Figure 13: Cluster sort of Figure 12. Similar trajectories are grouped together and subject-level information is maintained and the association of cognitive ability by the metric of baseline quintiles across the visit structure can be analyzed.

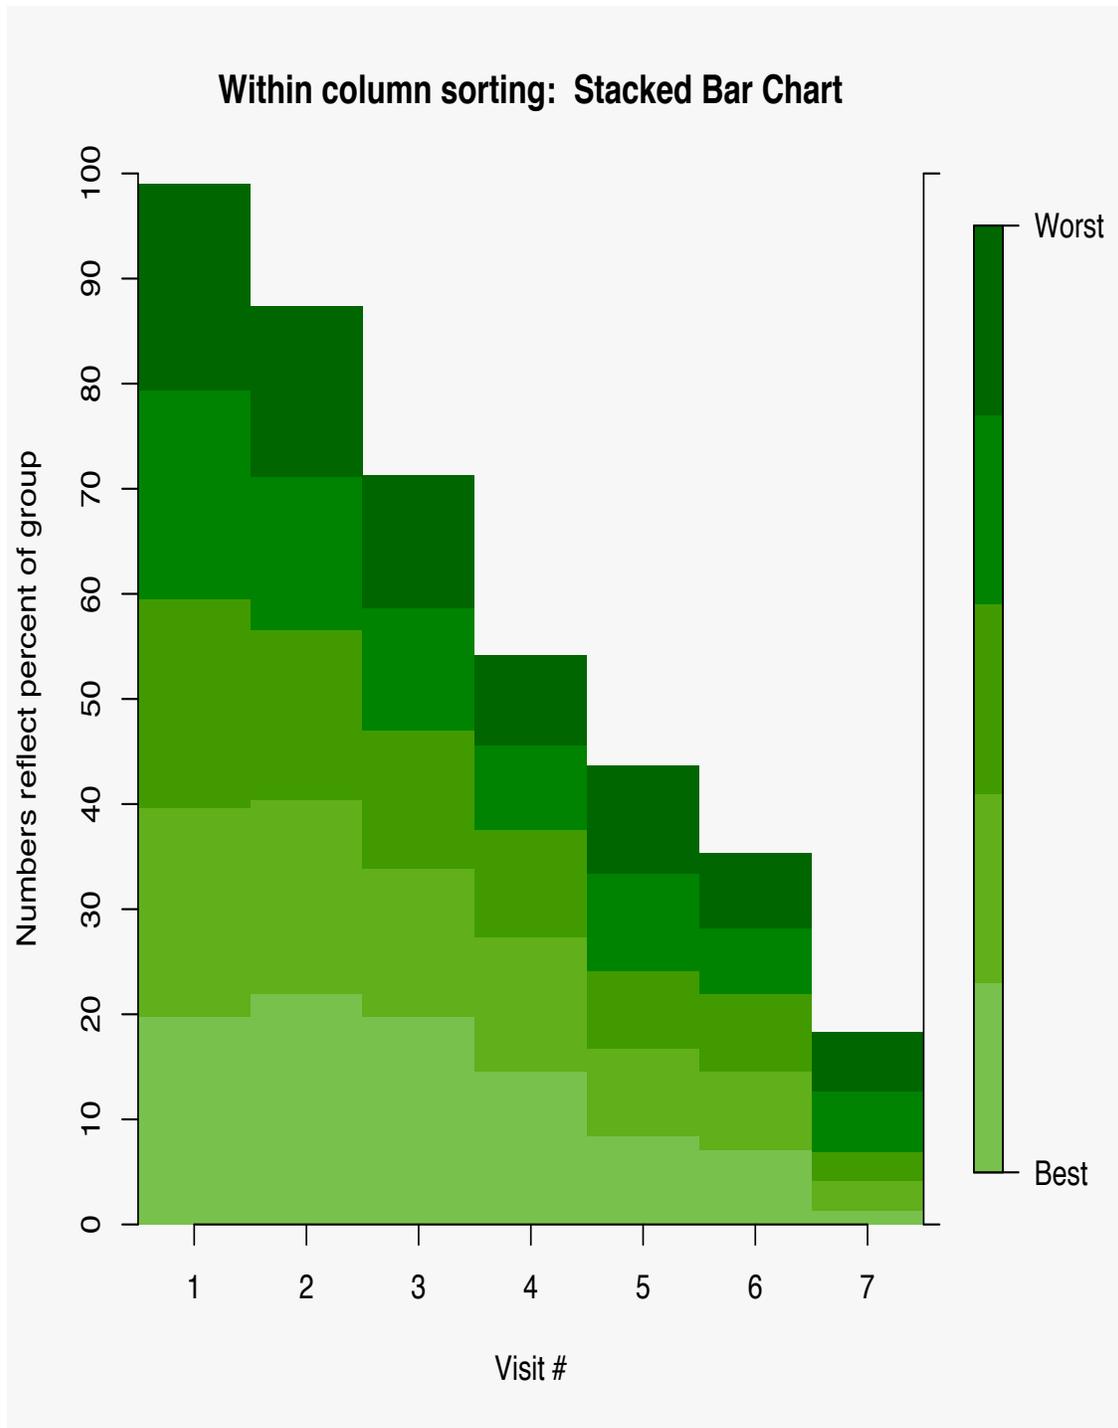


Figure 14: Within-column sorting of Figure 13, which shows the derivation of the classic stacked bar chart. The within-column sorting severs the connection of repeated measures within subject completely and is a strong summarization of the data in that it discards a lot of information. In the above graphic, we cannot see the distributions of 2nd visit best quintile score conditional on quintile of the score of visit 1. However, we can see a bump of best scores from visit 1 to visit 2, indicating a possible training effect.

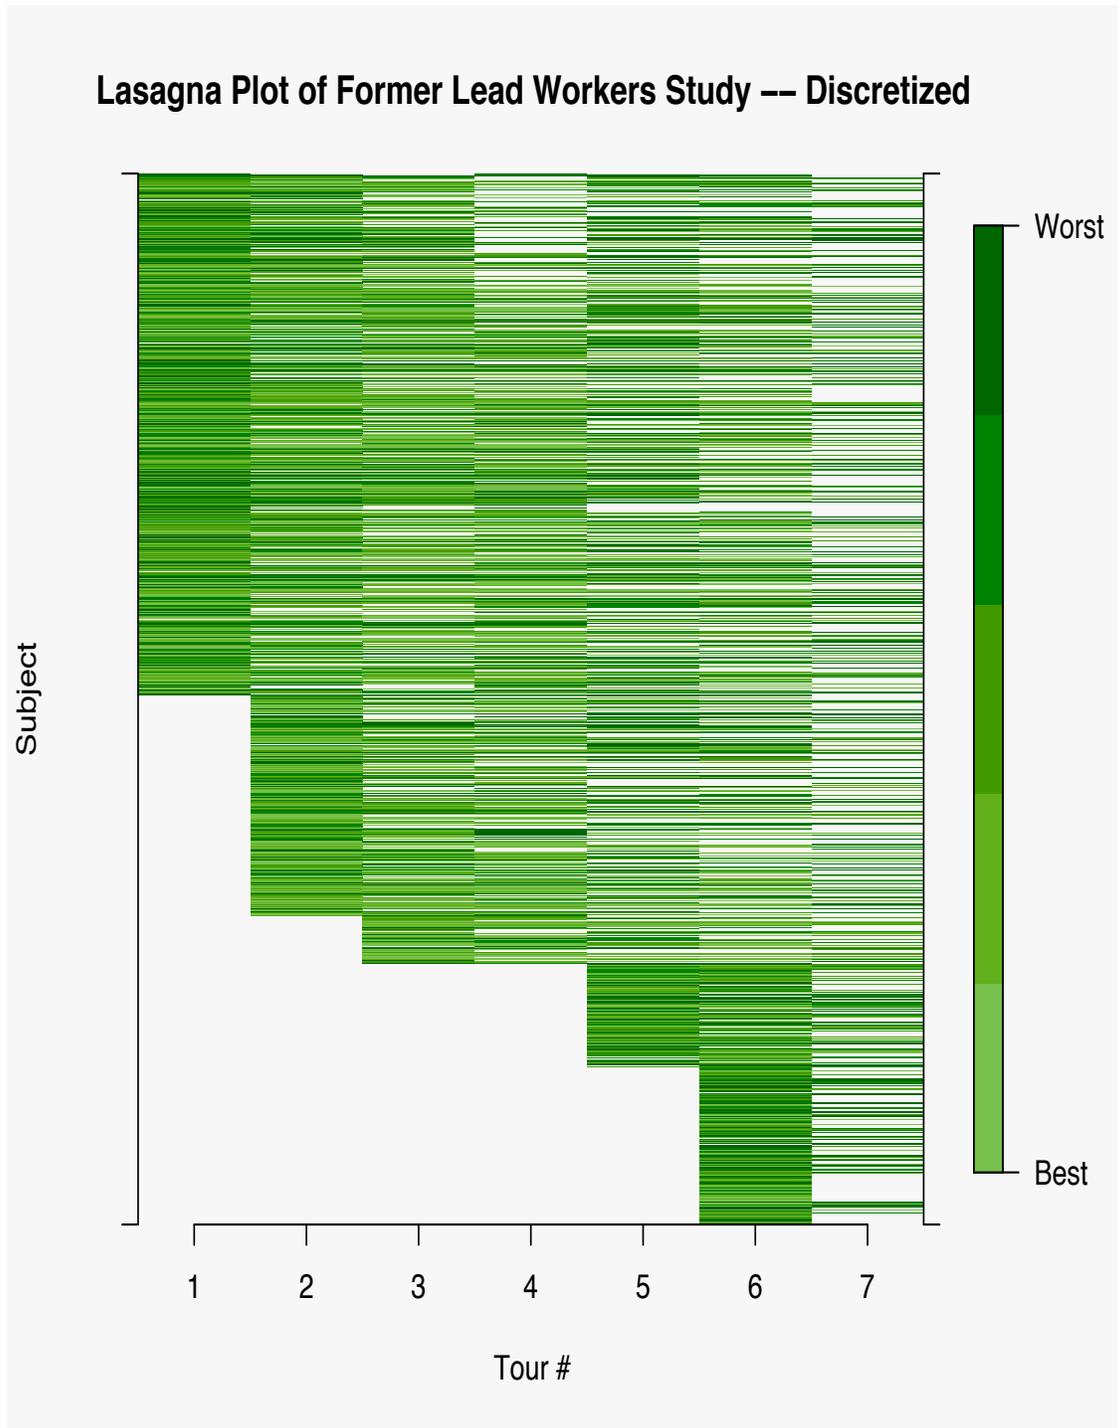


Figure 15: Lasagna plot for 1110 subjects over 7 tours, compare to the same data plotted by visit (Figure 12), with subjects in random order within their tour of enrollment.

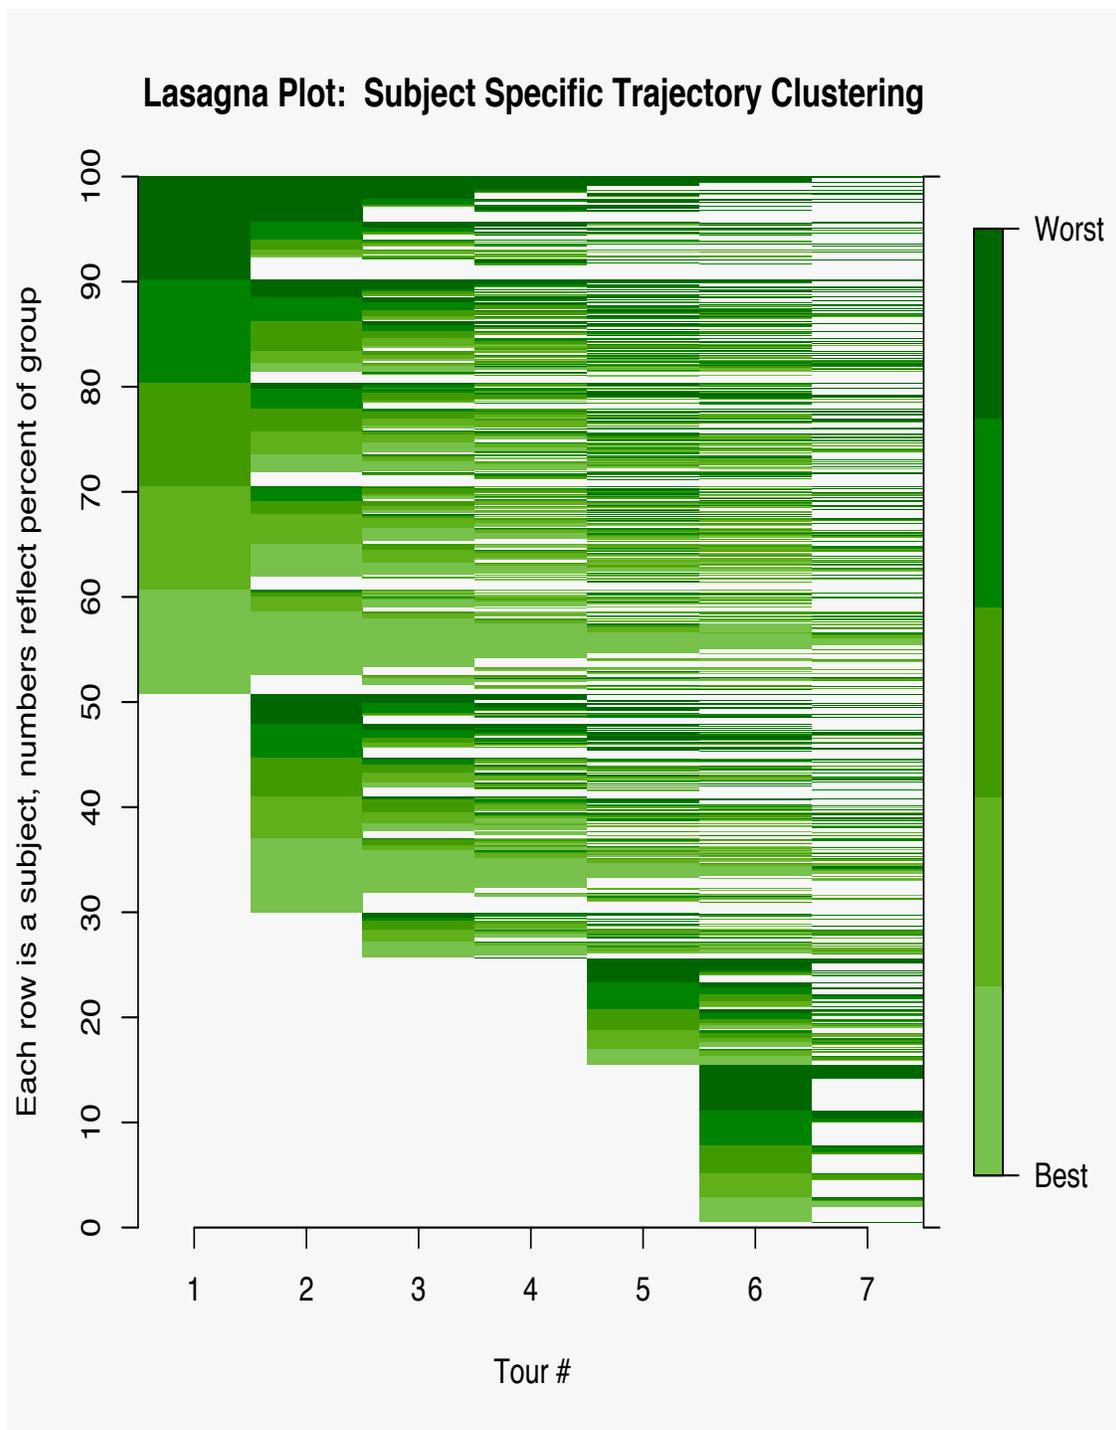


Figure 16: Cluster sort of Figure 15. Similar trajectories are grouped together and subject-level information is maintained and the association of cognitive ability by the metric of baseline quintiles over time can be explored. Comparing those enrolled in Tour 1 of the worst and best quintile, we see that there is more dropout for those starting out in the worst quintile, possibly indicating informative censoring. The above pattern of drop out being related to first visit quintile rank holds for the later tours as well. Training effects can be seen when a subject's color lightens tracking that subject across time. For instance, a fair proportion of subjects in Tour 1 were in the 2nd best quintile and advance to the best quintile in their second visit in Tour 2.

Table 1: Error rate of the confidence interval - simulation

n_1/n_2	σ_1^2/σ_2^2						
	1/8	1/4	1/2	1	2	4	8
1/2	0.012	0.022	0.032	0.056	0.085	0.114	0.148
1	0.059	0.055	0.050	0.050	0.058	0.056	0.058
2	0.149	0.115	0.086	0.053	0.033	0.021	0.014
4	0.252	0.193	0.119	0.058	0.024	0.005	0.003
8	0.364	0.258	0.150	0.056	0.017	0.003	0.000

Table 2: Error rate of the confidence interval - normal approximation

n_1/n_2	σ_1^2/σ_2^2						
	1/8	1/4	1/2	1	2	4	8
1/2	0.011	0.016	0.028	0.050	0.080	0.11	0.133
1	0.050	0.050	0.050	0.050	0.050	0.050	0.050
2	0.133	0.110	0.080	0.050	0.028	0.016	0.011
4	0.237	0.179	0.110	0.050	0.016	0.004	0.001
8	0.331	0.237	0.133	0.050	0.011	0.001	0.000

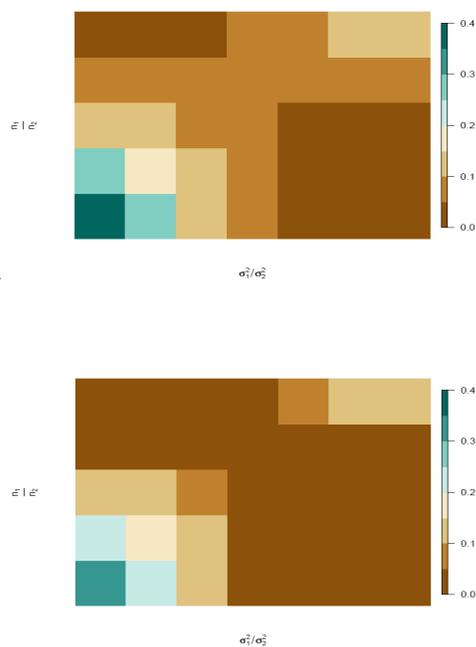


Figure 17: Simulation tables often have outcomes as result of different permutations of parameters. Making a lasagna plot of such a table gives a heatmap that might convey trends more clearly than the numbers themselves.

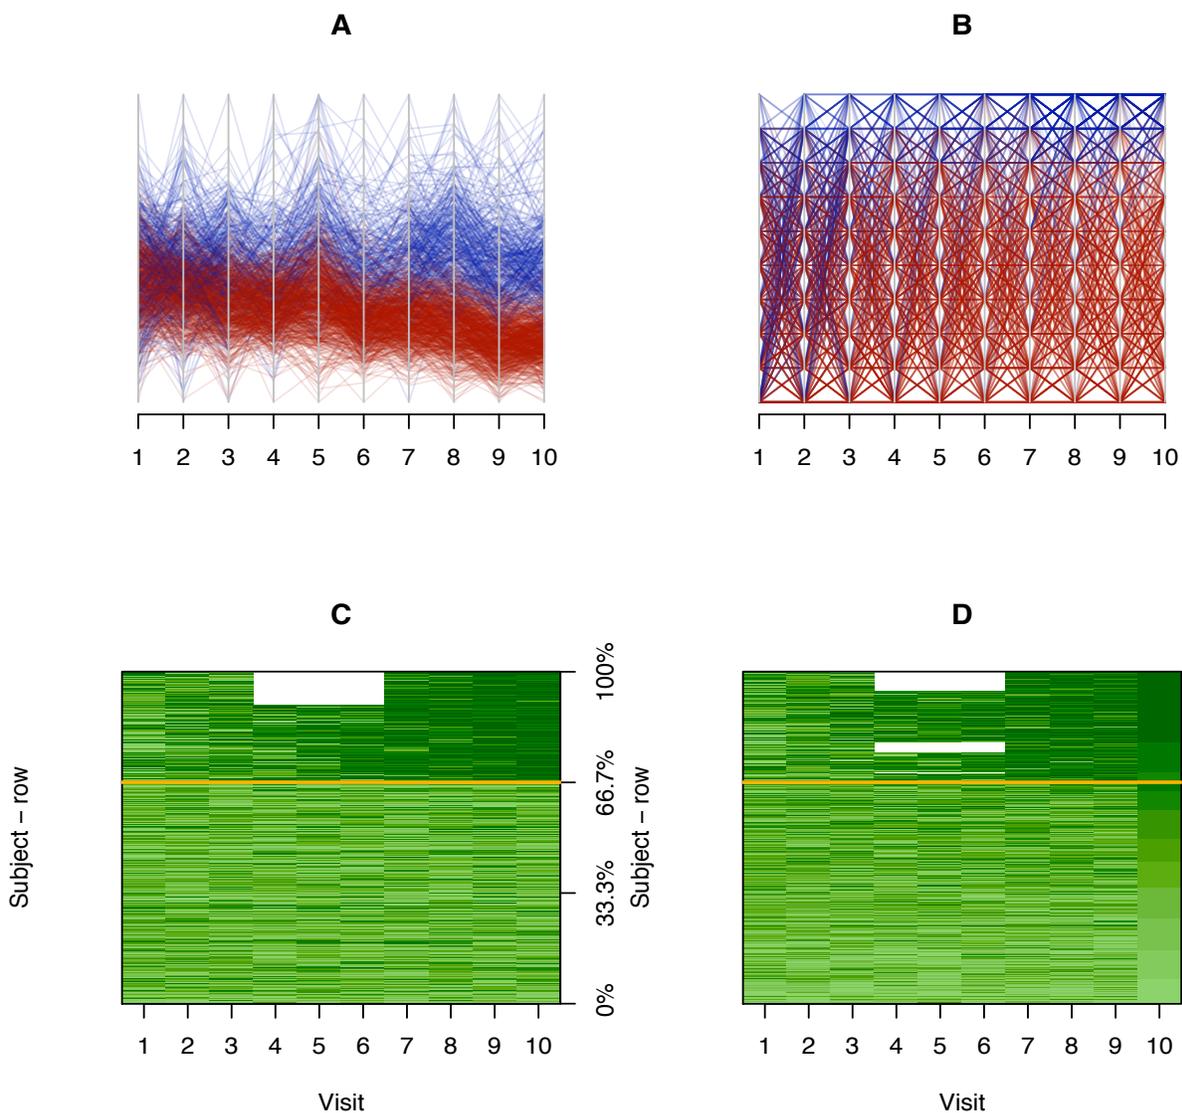


Figure 19: Panel A: A parallel coordinate plot with 10 measurements per subject, 1000 subjects. The color (red, blue) corresponds to the cluster to which the subject belongs, with the density of the plotting reflecting how many subjects trajectories are being plotted. Panel B: The discretized data of Panel A, where the outcomes were binned by decile. Panel C: Lasagna plot of Panel B, with those of the blue cluster appearing above the orange line, those of the red cluster below it, showing 2/3 the subjects are in cluster 1 (below the orange line). The darker the plot the greater the outcome magnitude, showing that cluster 2 as a group had greater values than cluster 1. White is missing, implying that cluster 2 alone and a decent proportion of which missed visits 4, 5, and 6. Panel D: The lasagna plot of Panel C entire row sorted within cluster based on the outcome value of visit 10. In either Panel A or Panel B, the substantial amount of missing data is not conveyed, as well as it may be difficult to ascertain relative number of subjects among clusters. In Panel B, the color density of the parallel coordinate plot of discretized data is constrained to exact segments of trajectories, potentially not fully conveying data features. In Panel C and D, missingness and relative number of subjects among clusters is more clearly conveyed, suggesting that in cases of discretized epidemiologic longitudinal data with missingness, lasagna plots may facilitate exploratory data analysis.