

Data Seen through Different Lenses

Clifford Konold^a, Traci Higgins^b, Susan Jo Russell^b, and Khalimahtul Khalil^a

^aScientific Reasoning Research Institute

University of Massachusetts, Amherst

and

^bTERC

Unpublished Manuscript

February, 2004

The writing of this article was supported by the National Science Foundation under Grant Nos. REC-9725228, REC-9814898, ESI-9730683, ESI-9731064, and ESI-9818946. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Abstract

Statistical thinking focuses on properties that belong not to individual data values but to the entire aggregate. We analyze students' statements from 3 different sources to explore possible building blocks of the idea of data as aggregate and speculate on how young students go about putting these ideas together. We identify 4 general perspectives that students use in working with data, which in addition to an aggregate perspective include regarding data as pointers, as case values, and as classifiers. Some students seem inclined to view data from one of these 3 alternative perspectives, which then influences the types of questions they ask, the representations they generate or prefer, the interpretations they give to notions such as the average, and the conclusions they draw from the data.

Data Seen through Different Lenses

The graph in Figure 1 displays the bedtimes of a class of 3rd and 4th graders (from Russell, Schifter & Bastable, 2002, p. 103). Looking at this graph we might notice that:

- the average bedtime is about 9:00 p.m.
- the majority of students have bedtimes ranging from 8:30 to 9:30
- the distribution of bedtimes is a bit skewed to the right
- bedtimes on the hour and half hour are much more common than on the quarter hour.

These four particular characteristics are not properties of any of the individual students' bedtimes. Rather, they are properties of the aggregate, or collection as a whole.

Insert Figure 1 about here

Perceiving and describing aggregate features of this type is what statistics is primarily about (Fisher, 1990; Moore, 1990). Yet composing individual data values into an aggregate does not come easily to students. Cobb (1999) noted that many of the middle school students he and his colleagues worked with initially perceived graphs of data simply as "collections of points." Thus, rather than attending to features of the frequency distribution in Figure 1 — how the values cluster and spread over their range — students might describe the data by naming off the individuals with the latest or earliest bedtimes or by locating themselves within the distribution. In earlier research, Hancock, Kaput, and Goldsmith (1992) reported an intervention in which they had encouraged students (ages 8-15) to attend to aggregate qualities such as clusters and

spread in analyzing data. The students persisted, however, to hone in on “individual cases and sometimes had difficulty looking beyond the particulars of a single case to a generalized picture of the group” (p. 354). Mokros and Russell (1995) likewise observed that for many students, “a representative value had no meaning because the data set was, for them, only the values of the numbers” (p. 35). Stressing the conceptual nature of this problem, Hancock et al. (1992) concluded that “to think about the aggregate, the aggregate must be ‘constructed’” (p. 355).

These and other studies have both pointed to the importance in statistical reasoning of perceiving data as an aggregate and have documented the challenge many students face in making the transition from reasoning about cases to “constructing” and reasoning about aggregates. In this article, we analyze statements of students from three different sources to explore possible building blocks of the idea of data as aggregate and to speculate on how young students go about putting these ideas together.

Data Sources

To exemplify various types of reasoning about data, we analyzed student statements from three primary sources.

Elementary School: Teacher-Written Case Studies

The Developing Mathematical Ideas (DMI) casebook comprises 28 case studies (Russell, Schifter & Bastable, 2002). These cases describe classroom episodes from Massachusetts public school teachers in grades K-5. The teachers were participating in a teacher-development project and had completed several homework assignments that involved conducting data-analysis activities in their own classrooms. The case studies, which describe and reflect on these classroom experiences, include: a) records of student dialogue, b) copies of students’ data representations and written reports c) teacher accounts of classroom activities and discussion, d) teacher interpretations of

student reasoning, and e) teacher reflections on pedagogy. For the purposes of this analysis, we focused on a) teachers' records of classroom dialogue and b) the representations students produced.

The case studies provide a rich source of information about how younger students view and work with data. Part of this richness comes from the range of grades and activities they cover. Also, rather than being short, isolated instructional interventions, the cases describe classroom experiences embedded in sequences of instruction. Furthermore, they have been selected by teachers as particularly revealing episodes.

Middle School: Individual Interviews

Eleven eighth grade students in a public school in Nashville, Tennessee were individually interviewed at the end of a fourteen week teaching experiment. The teaching experiment, which focused on reasoning about bivariate data, was designed and conducted by Paul Cobb and associates at Vanderbilt University (see Cobb, McClain, & Gravemeijer, 2003). As part of the instruction, students used scatterplots to perceive and describe linear and non-linear trends in bivariate data. The interviews lasted about one hour and followed a structured set of questions. The interviewer posed a problem that required a student to interpret bivariate data similar to the data they had been reasoning about during the teaching experiment. The interviews were videotaped and transcribed.

High School: Pair Interviews

Konold, Pollatsek, Well, and Gagnon (1997) interviewed two pairs of high school seniors who had just completed a yearlong statistics course at Holyoke High School in Holyoke, Massachusetts. During the interview, the students explored a data set that included a variety of information obtained from 154 students at their school, producing

graphs to answer various questions. The students who were interviewed were familiar with these data as well as with the data analysis software they were using. The interviews were videotaped and transcribed.

Different Perspectives on Data

Based on an analysis of data from these three sources, we have identified four general perspectives that students use in working with data. To get a quick sense of these perspectives, consider the frequency graph in Figure 2. It displays the favorite colors of six hypothetical students, where each circle represents the color preference of a different student. Many data activities conducted in the early elementary grades involve collecting and graphing data of this sort (see, for example, DMI Casebook, case 7). Students might conduct an in-class survey using a simple question (e.g., What is your favorite color?) and then consider what they learn from the data, (e.g., What does this graph tell us?).

 Insert Figure 2 about here

Figure 3 classifies examples of four types of responses students might give in summarizing the graph in Figure 2. These include inscribing or interpreting data as:

- *pointers* to the larger event from which the data came,
- *case values* that provide information about the value of some attribute for each individual case,
- *classifiers* that give information about the frequency of cases with a particular attribute value,
- *an aggregate* that is perceived as a unity with emergent properties such as shape and center.

Insert Figure 3 about here

These perspectives attend to different questions that we can ask of data and they organize the data into different fundamental units. As stated above, a statistical perspective focuses on the entire batch of data, on the characteristics of the data set as a whole. Thus the functional unit in a statistical perspective is the aggregate. In contrast, treating data as “classifiers” involves viewing cases with similar values as a unit (e.g., all the students whose favorite color is red). Treating data as “case values” involves taking the individual data elements as the fundamental unit, focusing on characteristics of individual cases. When treating data as “pointers,” there is no obvious perceptual unit. Rather, the data serve as reminders to the larger event from which the data came.

We regard these four perspectives as forming a loose hierarchy of sorts (from data as pointer to data as aggregate) where a higher level subsumes or encapsulates lower ones. Higher levels reorganize units in the lower level into a new perceptual unit (see Figure 3). We should stress, however, that different contexts or questions may call for or cue different views of data even within the same student. And there are times during analysis when experts might regard data in ways that are akin to the case-value and classifier perspectives. For example, in inspecting a skewed distribution of income values, a statistician might quickly alternate between considering the approximate center and shape of the distribution and attending to individual data points in the tails (outliers), searching for more information about these individual cases. Note, however, that regarding an individual case as an outlier involves locating it with respect to the other values in the distribution and thus involves a coordination of aggregate and individual perspectives.

Because each of these perspectives is useful depending on the question at hand, we view them not as levels or perspectives to graduate from but rather as perspectives to coordinate and master. But as this statement implies, we frequently observe students who appear unable, for example, to perceive a batch of data as an aggregate even when a particular question or purpose seems to call for this view. Furthermore, some students seem inclined to view data from one particular perspective. This inclination then influences, and perhaps constrains, the types of questions they ask, the plots they generate or prefer, the interpretations they give to notions such as the average, and the conclusions they draw from the data.

To the extent that our descriptions of these various perspectives are valid, they offer potential insights into how students are reasoning about data. They also may provide a basis for designing interventions that encourage the development and coordination of more aggregated perspectives. In the remainder of the paper, we elaborate our description of these perspectives and illustrate them with statements and graphs made by students. In using work from a particular student as an example of reasoning from a perspective, we will not assume that that student was incapable of applying a higher-level view. The student work we had access to did not permit this type of analysis because in general it did not include examples of a particular student's responses over a range of problems. But more to the point, we view the framework we propose not as a tool for classifying students into different reasoning types, but rather as a window on how students might be reasoning when they offer a particular interpretation of data at a specific point in time. In analyzing these relatively isolated statements, we asked ourselves, "How is this student currently regarding these data such that this student's statements and actions in this context make sense to him or her?"

Data As Pointer

In this perspective, the data and the event from which they came are not clearly differentiated. We see this perspective mostly among very young students who have collected data themselves. In these instances, students treat data records much as they might a photograph of last summer's vacation — as an image that helps bring to mind whatever was salient about that event. The photo might remind one person of relaxing on the beach, while for another it might bring to mind the terrible sunburn she got. When students view data as a pointer, the data represent the whole event from which the data were generated. This orientation is reminiscent of Vygotsky's account of how young children, when asked to draw an object placed in front of them, will depict "not what they see but what they know" (1978, p. 112).

In a DMI case 7, a kindergarten class produced the plot in Figure 4 showing the frequency of the favorite colors of class members (p. 37). The students in the class came from backgrounds representing eight different language groups, and this data activity followed one in which students had learned the Chinese names for some colors.

 Insert Figure 4 about here

While collecting the data, many of the students focused on numerical information in the display, making comments such as "That's two for blue," or "Only one person likes black" as the teacher recorded new values on the board. The next day, when the teacher asked, "What does this chart tell us?" it appears that for many of the students the graph was now a more general pointer to the entire event of the previous day. Some of the students' responses included:

We know what everyone's favorite color is.

My shirt is blue.

We know everyone's name.

We learned English and Chinese colors.

When reasoning from the data-as-pointer perspective, students often mention things not explicitly represented in the display they are interpreting. In this respect, there is no specific perceptual unit they are focusing on (as indicated in Figure 3 by the question mark). Our interpretation is that these students were not regarding the display of data as a model of aspects of an event. Rather they were viewing the graph as a general rendering of that event. This orientation fits well with young students' tendencies to construct iconic data representations in which they depict considerable detail about the events they observe. For example, Figure 5 is a display made by first graders who had surveyed classmates about what recess activities they preferred (DMI case 8, p. 42). Each surveyed student is rendered in stick-figure form, and details of their preferred activities are depicted. Apparently, "Chasing games" were played on grass and the "Structures" included a slide. Not surprisingly, the class ended before these students could finish their display.

Insert Figure 5 about here

In turning observations into data, we of necessity narrow our focus, encoding only selected information. A biologist studying the health of a salmon population may record only the weight and length of each sampled fish, ignoring a host of other information about each individual that potentially could be recorded. For students viewing data as a pointer, the data records do not restrict the view to attributes of

interest but rather serve to trigger a host of potential memories or associated facts related to the observed event.

Data As Case Value

This perspective entails associating a value with an individual case. Unlike using data as a pointer to an entire event or context, this perspective acknowledges that the data encode only particular aspects of an event. The unit of analysis in this perspective is the individual case, for example, a single person in a survey of students' favorite activities (see Figure 3). This perspective is often evident in young children who tend to focus on the identity of each individual piece of data, especially data values belonging to them. For example, one of the students in interpreting the graph in Figure 4 responded, "My favorite color is red." (DMI Casebook, p. 38).

Prototypical questions based on this unit of analysis include determining the value of particular cases ("How tall is Henry?") or the case identity of salient values ("Who is tallest?" "Who is shortest?"). Figure 6 are data that a class of third and fourth graders posted on the board (DMI casebook, case 12). The graph shows how long each of their families had lived in town. To summarize the data, one pair of students wrote "The longest someone lived in our town is 37 years. The shortest time...is 0 years."

 Insert Figure 6 about here

Ordering the data, which students often do spontaneously, makes it easy to locate these extreme values. The fact that the data values were ordered in Figure 6 perhaps made the extreme values more salient.

When asked to graph numeric data, younger students frequently make the kind of plot shown in Figure 7. In this “case-value” plot, each case is represented by a bar whose length corresponds to the value on that attribute. This plot was made by fifth graders who were analyzing data from a collection of 24 cats. Note that this type of display leaves ample space for labeling each case with its case identifier, here the cat’s name.

Insert Figure 7 about here

Operating from a case-value perspective, students use a graph much as they might a phone book — to locate a particular case and read off its value. The case-value plot in Figure 7 makes this particularly easy. Indeed, case-value plots are perhaps the most common type of graph in newspapers where the bars are often ordered alphabetically by case name. Wainer (2001) pointed out the limitation of this convention for ordering the bars, noting that

Since we are almost never interested in seeing Alabama first, it is astonishing how often data displays are prepared in which alphabetical order is the organizing principle of choice. The only reason I can think of is that ...[it] is easy and obvious. p. 43

However, if one’s purpose is to locate a value of interest given the identity of a particular case (e.g., how much is Mississippi’s per capita student expenditure?), then this ordering principle makes good sense. Many students new to the study of data analysis appear to believe that associating a case with its value is mostly what graphs are used for, and thus may make or prefer data representations suited to this purpose.

Indeed, a majority of the items we find on high stakes tests that purport to assess statistical reasoning test this skill (Konold & Khalil, 2003).

We see this viewpoint about the primacy of the case-value perspective expressed by Val, an 8th grade student interviewed at the Nashville site. The interviewer had described a study designed to investigate how time spent brushing teeth affects plaque levels. After Val had read a short description of how and what data were collected, the interviewer asked:

I: How do you think the researchers organized the data?

Val: They probably organized it in a graph because most adults like graphs instead of charts. But I would probably do it in a chart just so it'd be easier to read, so you can take a specific person and know just about that person without having to know about all the other people, and you can compare them really easily.

By “chart,” we think Val meant a table of values. As Val implied at the end of her statement, this focus on individual values allowed not only reading off individual cases, but also for easy comparison of cases.

The interviewer then gave Val several representations of the data which showed the brushing time and plaque level after brushing for 48 people. The statistics unit Val had just completed had focused exclusively on how to use scatterplots to reason about these type of data, and a scatterplot was one of the options offered her (see Figure 8).

Insert Figure 8 about here

But she, along with several of the students who were interviewed, preferred the representation in Table 1. One advantage Val saw in this tabular display is that she could “see the actual numbers” and “look down and find out which person has the least amount of plaque really easily.” As she later explained, the display in Table 1 reduced the likelihood of different people reading off different values for individual cases.

Val: ... because no matter if you gave it to me or somebody else, they would always see that...[Angela] would have 21 (seconds brushing) and 61 [percent plaque]. They would know those numbers, and those numbers wouldn't change no matter who you gave it to. But if you gave somebody just like this [Figure 8], I could say that ...he had 30 something, and another person would say that he had 40 something, because there's no marks, no divisions.

In summary, students view data from the case-value perspective as a model or simplification of a real world event. The perspective focuses, however, not on features of the data set as a whole but rather on the features of individual cases, more particularly on the attribute values of individual cases. When taking this perspective, students prefer data representations that make it easy to identify and rank individual cases and to read off values accurately. This perspective is well suited to the goal of determining where a single case of interest falls in a distribution (“How do I compare with other students in my class?”). It is also a perspective that fits with the apparent objectives of many current curricula and assessment instruments (Konold & Khalil, 2003), which aim to develop and test students’ abilities to decode basic graphic elements, what Curcio (1987) has called “reading the data.”

Insert Table 1 about here

Data As Classifier

Viewing data as classifiers entails combining individual cases of the same, or similar, value into a new unit — a category or type. In our example in Figure 2, we can regard all of the students whose favorite color is red as a group, noting for example that there are 3 such students, or that in this sample red color was selected more than any other color.

In clustering cases together to quantify them, we put aside the fact that those cases differ in other respects — these three students may be different genders, ages, heights. Questions that come to the fore in this perspective concern the frequency of cases of a particular value (e.g., How many students like red best? What is the most popular favorite color?). Particularly salient in the data-as-classifier perspective is the value with the most cases — the mode. Consistent with the observations of many elementary teachers, a teacher in the DMI casebook (p. 64) described her third and fourth graders as “heading straight for the mode,” regarding it not so much as a summary of group performance but rather as the “winning” outcome: “It felt to me like they were engaged in a race-to-the-top kind of board game—whichever value has the most x’s at the end is the winner” (DMI casebook, p. 100).

The DMI casebook contains many examples of students employing this perspective. For example, students in a fourth grade class, who had little previous experiencing working with data, conducted a survey about the number of people in their families (case 3, p. 12). Three students made the stacked dotplot in Figure 9.

Insert Figure 9 about here

On their plot they wrote, “Most people in 5 and 6 have a lot. Most people in 9, 11, 12, 18, 7, 4 have a less.” In making these observations, these students attended not to the values of single cases, but rather to the frequencies of occurrence of *types* of values. Furthermore, their treatment of these categories (families with 11 members, families with 7 members) suggests that they were regarding them as nominal rather than numeric in nature as indicated by their unordered list of values that have “less.” When the teacher asked them to explain their graph, Jacob said, “Well three kids in the class have 5 people in their family and five kids have 6. That’s more than the other numbers.” The teacher then asked, “What about the numbers that you say have less?” Tyrone responded, “At all of those places, there are only one or two kids with that many people in their family.” All of their responses indicated that they were focused on the frequencies of same-valued data types. They gave no evidence of attending to aggregate features such as the overall shape of the distribution or to relative densities of values (e.g., most of the families are bunched up on the lower end).

The Holyoke High School students who were interviewed in the study by Konold et al. (1997) were remarkably consistent in viewing data from the classifier perspective. For example, to provide an overview of the data they were examining, R and P wanted to produce a display that summarized grade levels of students in the data set. Working with the data analysis software they had been using during the course, they first produced a table of descriptive statistics, which included the median and mean grades. They were not satisfied with this table.

- R: No, that doesn't give us what we want. Isn't there one [a display type] that give us...
- I: What do you want to get?
- R: We want to separate them, you know like how many are of 12, how many of 8, how many of 11? How many... [the display in Table 2 appeared]. Yeah, there we go.

On this as well as most occasions during the interview, these students searched among the program's menu options for a workable display. Regardless of their question or the type of variables involved, the display they finally settled on was almost always a frequency table. Note that the moment Table 2 appeared on screen, R sensed they had what they wanted.

 Insert Table 2 about here

Asked to summarize the information in Table 2, R read off all of the frequencies from the table.

- R: That says we have one person from 8th grade, 10 from 9th grade, 6 from 10th grade, 31 from 11th grade and most of them are from 12th grade, which is 106 people.

Their attention seemed clearly focused on the frequency of various data types, in this case grade levels. At this point, the interviewer asked if there were other ways to display the data that might be helpful.

R: Yeah, we could use a plot ...the bar, bar graph. [Produced display in Figure 10]. That would do it. And probably the box plot.

I: You've got a bar graph [on the screen] and a moment ago you had a table with frequencies. Which do you think is the best for sort of giving a...

R: The table of frequencies. [Laughter]

I: Why is it better?

P: Cause you have to estimate on this [pointing to bar graph show in Figure 10].

Insert Figure 10 about here

We see here, as we did with the case-value perspective, a preference for displays that permit reading off exact values (in this case frequencies) rather than having to estimate them by, for example, sighting to an axis.

As we mentioned, during the interview these students used frequency tables to answer nearly every question they investigated. During the yearlong statistics course, they had practiced making and interpreting a number of other displays, including histograms, box plots, and scatterplots. One reason they may have come to prefer frequency tables over these other displays is that frequency tables allow for easy identification and counting of responses of a particular type — treating data as classifiers.

In addition to preferring displays that allow them to accurately read off frequencies of particular response types, students often attempt to extract this same information from displays in which it is not available. For example, in the process of exploring their question about curfew and study time, R and P produced the box plot display shown in Figure 11. As soon as it appeared, P responded:

 Insert Figure 11 about here

P: That doesn't help.

R: I know [Laughing].

I: [...] Why doesn't that help?

P: Because, it's, like, confusing. We don't know what the hours are, how many hours...

R: [...] Like how many students studied 10 hours on "no" [curfew], and how many students studied 10 hours on "yes" [curfew].

They tried to get this frequency information from the box plot using a point-ID feature of the tool that displays the value of particular cases. When they could not get frequency information from boxplots using this feature, they produced the histogram display shown in Figure 12.

 Insert Figure 12 about here

Again they were stymied because the number of students at each grade level was not directly available.

R: No. This is confusing.

P: You can use the point ID though, can't you?

I: Now, now what is that you want to find out?

R: [...] The number of students like who were studying 5 hours, who studied 5 hours and “yes,” and who will have a curfew and who don't.

In summary, when viewing data as classifiers, students focus on the frequency of like-valued cases to answer such questions as what outcome type is the most frequent or to compare frequencies of various outcome types. For students taking this perspective, it is still important that data displays make it easy for them to accurately read both case values and their frequencies. When taking this perspective, students tend to view different response categories as independent of one another and thus do not attend to how frequency may be changing systematically over outcome type (i.e., to how frequency distributions are shaped). And because they do not think about the collection of categories as constituting a unity, they tend to reason about frequencies rather than relative frequencies of category types.

Coordinating Case-value and Classifier Perspectives

As they share and discuss their work in the classroom, we frequently see students working to distinguish between the case-value and classifier perspectives. In the excerpt below (from DMI case 9, p. 47), two kindergarten students were looking at a

representation of their responses to the question, “Do you like to work on a computer?” (See Figure 13).

Insert Figure 13 about here

The chart had emerged as each student placed a marker on one of the possible responses to the question. The teacher asked:

Teacher: So do you think someone else could tell something about us from our survey?

Rhea: Yes, they know most of us like computers a lot.

Amanda: Not me, I said no.

Melinda: Me either, I said I never played before.

Rhea: I said most of us!

By her response we conclude that Rhea was attending to the large number of “yes’s.” Many younger students use the term “most,” to mean “more than any other.” Thus we conclude from her response that she was working from the perspective of data as classifier. Amanda and Melinda protested that Rhea’s summary did not take into account their particular responses. Taking a view of data as case-value, they were attending to their individual data values. From this perspective, they apparently interpreted Rhea’s assertion as ignoring their contributions to the plot. Rhea responded with some frustration, emphasizing that she had said “most,” which from her perspective takes into account that fact that not all the students like computers a lot.

With attributes coded as integers (e.g., family size), students can find it particularly challenging to translate between the case-value and classifier perspectives. Because in this instance both value and frequency are integers, it becomes easy to confuse the value of a case (a family of 4) with the number of cases of that type (5 families with 4 members). We see this confusion at play in a discussion among three fourth grade students who were describing to one another the graphs they each had made of the size of the families of their classmates (DMI case 14, p. 83). Kenny had made a case value plot like the one shown in Figure 14 using a bar (composed of individual x's) to represent the size of each of the 12 students' families. Thus the length of each bar corresponded to the number of members in a particular student's family. His partners, Cara and Tim, had both plotted the same data using stacked dot plots. A stack of x's in each of their graphs represented a given number of families of a particular size. Kenny tried to explain his plot to his partners.

 Insert Figure 14 about here

Kenny: The first family had 12 [members], the second family had 8, and the third family had 5. I put X's for the number of people in the family.

Tim: That looks like too many.

Cara: It is. You have to put the amount of people [family size] under the line.

Tim: How many families goes on top of the line.

Kenny: Huh?

As their teacher observed, Cara and Tim saw Kenny's plot not as an alternative way to display the data, but as a mistake to correct. Interpreting the x's in Kenny's plot as standing for each student in the class, Tim saw "too many" students (families). Likewise Kenny was unable during this exchange to understand the objections of his partners.

The potential for confusing case-value and frequency graphs of integer data is heightened by the fact that traditionally both have been referred to as "bar graphs" (see Konold & Higgins, 2003, p. 200). Making the difference between these two plot types an explicit part of instruction could help students distinguish them. Furthermore, Cobb (1999) claims that case-value plots are an easier graph for students to interpret and that by first introducing students to them we can provide a suitable grounding for frequency plots (see Konold & Higgins, 2003, p. 200-201 for an illustration of the sequence of graphs recommended by Cobb, 1999).

Data As Aggregate

When viewing data as an aggregate, the perceptual unit is the entire distribution of values. In focusing on the distribution, one attends to emergent features not evident in any of the individual data values. These features include the general shape of the frequency distribution, how spread out the cases are, and where in the distribution cases tend to be concentrated. Using aggregate reasoning, one might describe the relative number of cases in various parts of the distribution, using either percentages or quantitative descriptors such as "most" or "majority." With numeric variables, one might summarize group features with measures of center, of spread, or of shape. For students applying this perspective, questions that come to the fore include whether two

or more groups differ on some aggregate measure or whether two variables are related to one another.

In Figure 3 (on page xx), we depict the data structure of an aggregate perspective as not only bounded, but also as internally structured. Each different data type occupies a well-defined space in the mental representation. Effectively, the different color types are joined together in a unity, a dimension called color which comprises a number of mutually exclusive values. In this sense, a classification system is more than a device for placing name tags on cases (Bowker & Star, 1999). It is a system for dimensionalizing them — placing values in relation to one another. It is this structure that allows us to quickly add a case we have not yet observed (e.g., a favorite color of purple) and to make quantitative comparisons between various case types, to say for example that half of the students have red as a favorite color, or that “the majority of families have lived in town less than 10 years” (see Figure 6). To put these objects into such a relation to one another requires that we regard them as part of a homogenous group despite the fact that they differ not only with regard to other attributes that we are ignoring (boys that are all 6 feet tall nevertheless differ in other respects), but also with respect to the common attribute we are currently examining (individuals vary on height). Stigler (1999) points out that this step of regarding different observations as belonging to a single group proved to be a major barrier to the wide-scale application of statistical methods to both the social and physical sciences:

The first conceptual barrier in the application of probability and statistical methods in the physical sciences had been the combination of observations; so it was with the social sciences. Before a set of observations, be they sightings of a star, readings on a pressure gauge, or price ratios, could be combined to produce a single number, they

had to be grouped together as homogeneous, or their individual identities could not be submerged in the overall result without loss of information. This proved to be particularly difficult in the social sciences, where each observation brought with it a distinctive case history, and individuality that set it apart in a way that star sightings or pressure readings were not. ... If it were felt necessary to take all (or even many) of these [distinct case histories] into account, the reliability of the combined result collapsed and it became a mere curiosity, carrying no weight in intellectual discourse. Others had combined ...[the individual cases], but they had not succeeded in investing the result with authority. (pp. 73-74).

A few students in the Vanderbilt interviews appeared to use aggregate reasoning. Susanne, for example, said she preferred a scatterplot display in analyzing the calories-energy data to the other displays because it offered a “general view” as compared to the “hardcore facts” of the ordered table of values. Her interpretation of the scatterplot included a description of the trend for the full range of the data (an aggregate feature) and a statement relating that trend to the problem context:

Susanne: I would use ...[Figure 15] because, you know, it shows around the middle part, where the calories are, like, steady at that spot. It shows that the people who are in the area are like up [at high energy] while the people who are extremely low in calories or extremely high vary, and were like, decreasing [in energy level].

Insert Figure 15 about here

She not only described the general trend, but noted differences she perceived in variability around that trend, describing the “middle part” of the scatterplot as “steady” compared to the extremes.

On the whole, instances of aggregate reasoning are relatively rare in the three data sources of this study. This observation accords with reports from numerous researchers concerning the challenges students face in learning to reason about statistical aggregates (e.g., Ben-Zvi & Arcavi, 2001; Hancock et al. 1992; Konold et al. 1997; Konold & Higgins, 2003; Mokros & Russell, 1995). In the DMI casebook, we do often see teachers prompting the students in an attempt to support aggregate reasoning. For example, in the 4th grade discussion of family size referred to earlier (DMI, case 3), one group created the representation in Figure 16.

Insert Figure 16 about here

The teacher reflected, “Their line plot looked very similar to that of the previous group [see Figure 8] ... However, I noticed that they ... had used terms such as typical ... and range.” Accordingly, she asked:

Teacher: I noticed that you wrote 5 or 6 were typical numbers of family members for this class. Can someone say something about why you chose these numbers?

Inez: There were more people with 5 or 6 in their families than the other numbers.

This last statement may have meant “More than half of the families have 5-6,” which would be an aggregate statement about the collection as a whole. But consistent with a classifier perspective, she may have meant, “Families with 5 or 6 are more frequent than families of other specific sizes.” After making sure students were aware of how many total data values there were, the teacher probed in a way that led another student to offer an aggregate summary:

Teacher: How many students in this class had either 5 or 6 people in their family? [I was met with many responses of 7.] If 7 out of 14 students have 5 or 6 family members, how else can we say that? Is there a fraction we could use to describe this?

Denise: I know! One half of the class has 5 or 6.

In a similar episode (DMI case 12, p. 63-68), a teacher helped her combined third/fourth grade class to take an aggregate perspective. She did this by building on students’ tendencies to identify modal clumps (Konold et al. 2002) and by encouraging one student to consider what she meant by “a lot.” The teacher had arranged the students in pairs and asked them to write a description of the data displayed in Figure 6.

Teacher [Reflecting]. Several other papers referred to the four x’s at 3 as being “the most” or “a lot,” and I thought, here we are again—mode rules. However, while the kids were writing, I had a chance to question one group... I heard Anna, a fourth grader, saying right off the bat that 3 [years] had “a lot.”

Teacher: What do you mean by “a lot”?

Anna: It's the only one that has four, and all the rest have less than that.

Teacher: Does that mean that it has a lot?

Teacher: [Reflecting]. This seemed to make her stop and think, and I decided I wanted to push her a bit.

Teacher: How many x's are up there altogether?

Anna: 23.

Teacher: How many of them are at 3?

Anna: Oh, it's only four. That's not really a lot, is it?

In the class discussion that followed, Anna explained to the class how even though 3 years did have the most families, that was not "a lot" when compared to the total number of families. The following exchange occurred when the teacher again asked students to summarize the data. Anna was now able to apply the same reasoning she had used in evaluating "a lot" to provide a more precise meaning for "most."

Kevin: Most of the x's are between 0 and 6.

Teacher: How did you decide that?

Kevin: It's the biggest clump.

Teacher: How many x's are in that clump?

Anna: There's eleven. That's almost half.

Teacher: Almost half of what?

Anna: Well there's 23 altogether, and it's almost half of that....

In both of these instances, the teacher saw an opportunity to support students in taking what they noticed about individual values and relating it to the data set as a whole. Part of what prompted these learning episodes was the teachers' being sensitive to, and inquisitive about, the meanings that their students attached to words such as "most."

Statisticians use averages such as the mean and median as one way to characterize a distribution of values, to locate the approximate center of the distribution. Even before they take a unit or course in statistics, young students have already encountered statistical terms such as "average" and have incorporated them into their everyday discourse (Gal, Rothschild, & Wagner, 1990; Watson & Moritz, 2000). It is important that the teacher not assume, however, that because students are using the terminology of aggregates, that they are perceiving and describing aggregate features. As an example, we show below how students apply the terminology of averages in ways that are more consistent with case-value or classifier perspectives than they are to use of averages to summarize an entire data set.

Averages From the Viewpoint of the Case-value and Classifier Perspectives

In DMI case 22, a third grade teacher asked her class, "What would you say is the average height of kids in our room?" Brita answered, "It's me. I think I am average." To further explore this question, the teacher had the students line up according to their height. Looking at this physical graph, other students used the term *average* in the same way Brita had— as a feature of particular individuals in the line-up.

Phoebe: I think I'm taller than average because I notice that on the playground.

Brita: I was right. Sam is average, and I'm average too. We are the same.

Tiffany: I'm average too.

Katie: I'm not average. I'm shorter.

We see this adjectival usage of average as consistent with the case-value perspective. To say that Sam is of "average" height is to characterize this single case, and not necessarily the group the case is part of. Certainly to make this attribution, one must attend to where Sam is located in the distribution of heights. In the same way, to identify a particular case as the "smallest" or "largest" requires locating that case with respect to the other cases in the group. Having the students line up in order makes it relatively easy to locate these salient cases. However, it seems clear that the aim of making these attributions — "of average height," "shorter than average" — is not to characterize the whole data set, but rather to assign another feature, albeit it a comparative one, to a single case. To Sam's list of features ("male," "3^h grader," "42 inches tall") is added yet another — "of average height." In the same way, Katie applies to herself the attribute "shorter."

Watson and Moritz (2000) interviewed Australian students in grades 3 to 9 to explore their understanding of averages. Many of these students seemed to hold this perspective on what averages were. For example, a grade 3 student who was asked "Where have you heard "average?" responded, "My stepbrother plays cricket, and he usually gets an average amount of runs." A grade seven student who was asked the

same question responded “Average is like ... the middle standard.... So, if you said the television program was average, it means that it wasn’t very good and it wasn’t very bad.”

Given this case-value perspective of averages, it make sense that some students reject average values that do not correspond to the value of at least one case. For example, in DMI case 25 (p. 148) third grader Robbie had made two attempts at blowing a Styrofoam cylinder as far as he could, recording the values 152 cm and 186 cm. The class was discussing whether he should use the average of these two values (169 cm) to represent how far he could blow the cylinder. Robbie would have none of this. “But I didn’t get 169 as one of my distances,” he objected. “It’s a lie!”

Another commonly accepted view of the average is consistent with the classifier perspective. This is the view of the average as a characteristic of either the most frequently occurring cases (the mode) or of a subset of cases in the middle of a distribution, what Konold et al. (2002) have termed a “modal clump.” We think this is the perspective that predisposes many young students to focus in particular on the mode, considering it “the end-all way to describe what’s typical in a set of data” (case 12, case 19). Here, the attribution of “average” or “typical” is applied to a subset of cases with the same, or close to the same, value. Several researchers have reported students’ tendencies to regard numeric data as comprising three groups: middle, low, and high. A student in the Bakker and Gravemeijer (in press) study, for example, summarized a graph of student heights with the observation that “You have smaller ones, taller ones, and about average.” Having partitioned a distribution in this way sets the stage for applying these attributes “low” “high,” “average” to the cases in the corresponding partitions.

In both the classifier and case-value interpretations of average, the term “average” is used as an adjective to describe a characteristic of a particular case or subset of cases. “This person (group of people) who is (are) 65 inches tall has the characteristic ‘average.’” The people are average, of course, due to their central location in the distribution and/or because their value is the most commonly occurring. In both instances, students are using “average” as a label which they apply to one or perhaps several of the cases in the group, but certainly not to the entire group: some people are “average,” others are below or above average or even “outliers.”

In the aggregate sense, average is used as a noun. The claim is “This group has an average of x .” In this sense, average is not a label applied to a single case or subset, but rather is a measure that applies to the entire group. The whole group is not average, but rather *has* a particular average value. Because an average in this sense applies to the whole group, we can use it to represent the group and to compare one group to another. It is interesting in this regard to note that many researchers have reported that even though many students know how to compute and use averages with single groups, that few of them use averages to compare two groups (Bright & Friel, 1998; Hancock, Kaput & Goldsmith, 1992; Jones et al. 1999; Konold et al. 1997; Watson & Moritz, 1999). Our analysis suggests a possible explanation for this result, which is similar to an argument made in Konold and Pollatsek (2002). If it is indeed the case that the averages students are using are used as descriptors of only one or a subset of cases in a group, then it seems reasonable that they would not think of using these averages to make a comparison between two groups.

Summary

Given that statistics is fundamentally about the behavior of aggregates, the overarching question in statistics education is how we can foster students’ abilities to

perceive and reason about aggregates. The question of how to do this with younger students has recently been taken up by a number of researchers including Bakker (in press), Bakker and Gravemeijer (in press), Cobb (1999), Cobb, McClain, and Gravemeijer (2003), Konold and Pollatsek (2002), and Lehrer, Schauble, Carpenter, and Penner (2000). In this article, we have not focused on this question per se but rather have proposed a number of perspectives, in addition to viewing data as an aggregate, that we see students using as they reason about data. Our hope is that by differentiating these different perspectives, and suggesting what purposes they serve, that we will be better able to understand and help shape statistical reasoning of our students.

The fact we so rarely see elementary students using the aggregate perspective may suggest that treating data as an aggregate is beyond the abilities of these students. It is important to keep in mind, however, that we are only beginning sustained efforts to teach statistics in the elementary grades, and that we still have much to learn about how to engender statistical reasoning. However, we should also reiterate that we do not regard an aggregate perspective as THE way to look at data. It would be a serious mistake if educators saw as the major instructional goal to quickly move students past interpreting data as pointer, case values and classifiers. As we mentioned, many of the questions young students have about data involve locating themselves within distributions, identifying the highest and lowest values and the most frequently occurring outcome. Thus, the case-value and classifier perspectives are well suited to the interests of many young students. They serve well the purposes they have in mind when they collect and analyze data. And we maintain that choices about how to collect and analyze data should be based on one's questions and purposes. The perspectives one takes on data should service one's questions rather than the other way around. Furthermore, the case-value and classifier perspectives highlight important features of

data that are essential components of an aggregate perspective. A case-value perspective, for example, makes salient the fact that individuals vary, an awareness that too often gets neglected in both our curriculum (Biehler, 1994; Shaughnessy, Watson, Moritz, & Reading, 1999) and our public discourse (Gould, 1996).

By focusing on case values, students stay connected to the meaning of the data, to what the graph is about and to the information it encodes. Many of the difficulties students can experience in interpreting graphs in fact results from them losing this meaning. Focusing on the values of individual cases is a way to establish and maintain this connection.

References

- Bakker, A. (2003). Reasoning about shape as a pattern in variability. Paper presented at the Third International Research Forum on Statistical Reasoning, Thinking and Literacy. Lincoln, Nebraska.
- Bakker, A., & Gravemeijer, K. P. E. (in press). Learning to reason about distribution. In D. Ben-Zvi & J. Garfield (Eds.), *The challenge of developing statistical literacy, reasoning, and thinking*. Dordrecht, the Netherlands: Kluwer Academic Press.
- Batanero, C., Estepa, A., Godino, J. D. (1997). Evolution of students' understanding of statistical association in a computer-based teaching environment. In J. B. Garfield & G. Burrill (Eds.), *Research on the Role of Technology in Teaching and Learning Statistics: Proceedings of the 1996 IASE Round Table Conference* (pp. 191-205). Voorburg, The Netherlands: International Statistical Institute.
- Ben-Zvi, D. & Arcavi, A. (2001). Junior high school students' construction of global views of data and data representations. *Educational Studies in Mathematics*, 45, 35-65.
- Biehler, R. (1994). Probabilistic thinking, statistical reasoning, and the search for causes—Do we need a probabilistic revolution after we have taught data analysis? In J. Garfield (Ed.), *Research papers from ICOTS 4* (pp. 20-37) Minneapolis: University of Minnesota.
- Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. Cambridge, MA: MIT press.
- Cobb, P. (1999). Individual and collective mathematical development: The case of statistical data analysis. *Mathematical Thinking and Learning*, 1(1), 5-43.
- Cobb, P., McClain, K., & Gravemeijer, K. (2003). Learning about statistical covariation. *Cognition and Instruction*, 21, 1-78.

- Curcio, F. R. (1987). Comprehension of mathematical relationships expressed in graphs. *Journal for Research in Mathematics Education*, 18, 382-393.
- Fisher, R. A. (1990). *Statistical methods experimental design and scientific inference*. Oxford University Press, Oxford. (First published in 1925).
- Gal, I., Rothschild, K., & Wagner, D. A. (1990). Statistical concepts and statistical reasoning in school children: Convergence or divergence? Paper presented at the Annual Meeting of the American Educational Research Association, Boston, MA.
- Gould, S. J. (1996). *Full house*. New York: Harmony Books.
- Hancock, C., Kaput, J. J., & Goldsmith, L. T. (1992). Authentic inquiry with data: Critical barriers to classroom implementation. *Educational Psychologist* 27(3), 337–364.
- Jones, G. A., Thornton, C. A., Langrall, C. W., Mooney, E. S., Perry, B., & Putt, I. J. (1999). *A framework for assessing students' statistical thinking*. Paper presented at the Annual Meeting of the Research Pre-session of the National Council of Teachers of Mathematics, San Francisco, CA.
- Konold, C. & Higgins, T. L. (2003). Reasoning about data. In J. Kilpatrick, W. G. Martin, & D. Schifter (Eds.), *A research companion to Principles and Standards for School Mathematics* (pp. 193-215). Reston, VA: National Council of Teachers of Mathematics.
- Konold, C. & Khalil, K. (2003). If u can graff these numbers — 2, 15, 6 — your stat literit. Paper presented at the annual meeting of the American Educational Research Association Annual Convention, Chicago.
- Konold, C. & Pollatsek, A. (2002). Data analysis as the search for signals in noisy processes. *Journal for Research in Mathematics Education*, 33(4), 259-289.
- Konold, C., Pollatsek, A., Well, A., & Gagnon, A. (1997). Students analyzing data: Research of critical barriers. In J. B. Garfield & G. Burrill (Eds.), *Research on the role of technology in teaching and learning statistics: 1996 Proceedings of the 1996 IASE Round*

- Table Conference* (pp. 151-167). Voorburg, The Netherlands: International Statistical Institute.
- Konold, C., Robinson, A., Khalil, K., Pollatsek, A., Well, A., Wing, R., & Mayr, S. (2002). Students' use of modal clumps to summarize data. In Proceedings of the Sixth International Conference on Teaching Statistics. Cape Town, South Africa.
- Lehrer, R., Schauble, L., Carpenter, S., & Penner, D. (2000). The inter-related development of inscriptions and conceptual understanding. In P. Cobb, E. Yackel, & K. McClain (Eds.), *Symbolizing and communicating in mathematics classrooms: Perspectives on discourse, tools and instructional design* (pp. 325-360). Mahwah, NJ: Lawrence Erlbaum Associates.
- Mokros, J. & Russell, S. (1995). Children's concepts of average and representativeness. *Journal for Research in Mathematics Education*, 26(1), 20-39.
- Moore, D. S. (1990). Uncertainty. In L. A. Steen (Ed.), *On the shoulders of giants* (pp. 95-137). Washington, D.C.: National Academy Press.
- Russell, S. J., Schifter, D., & Bastable, V. (2002). *Developing mathematical ideas: Working with data*. Parsippany, NJ: Dale Seymour Publications.
- Shaughnessy, J. M., Watson, J., Moritz, J., & Reading, C. (1999). School mathematics students' acknowledgment of statistical variation. Paper presented at the 77th annual meeting of the National Council of Teachers of Mathematics, San Francisco.
- Stigler, S. M. (1999). *Statistics on the table: The history of statistical concepts and methods*. Cambridge, MA: Harvard University Press.
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.
- Wainer, H. (2001). Order in the court. *Chance*, 14(1), 43-46.

Watson, J. M. & Moritz, J.B. (1999). The beginning of statistical inference: Comparing two sets of data. *Educational Studies in Mathematics*, 37, 145-168.

Watson, J. M., & Moritz, J. B. (2000). The longitudinal development of understanding of average. *Mathematical Thinking and Learning*, 2(1), 9-48.

Table 1.

Part of the table of values of the tooth brushing data used in the Vanderbilt interviews. The table orders cases by ascending value of brushing time. To save space, we have included here only the first and last five values.

Name	Brushing time	Plaque
Angela	21	61
Peter	33	52
Eric	34	42
Wendy	35	18
Tina	38	48
.	.	.
.	.	.
.	.	.
Joseph	169	30
Kelly	170	18
Daniel	181	31
Andrew	202	18
John	221	27

Table 2.

Frequency table showing the number (and proportion) of surveyed students at grades 8-12.

Grade	count	proportion
8	1	0.01
9	10	0.06
10	6	0.04
11	31	0.20
12	106	0.69
total	154	1.00

Figure Captions

Figure 1. Bedtime frequencies for a class of third and fourth graders.

Figure 2. Graph of the favorite colors of six hypothetical students. The gray scale is meant to suggest the different colors.

Figure 3. Depiction of four different perspectives students might take in summarizing information displayed in Figure 2. A gray scale is intended to suggest the three different colors, red (darkest), green (lightest), and blue (middle value). The figures in the column labeled “Data Structure” depict how data from the point of view of each perspective might be mentally represented. With the data-as-pointer perspective, there is no a clear distinction made between the data and the real-world event (hence the porous boundary). In an aggregate perspective, each type of data value occupies a well-defined space in the mental representation such that new data types (e.g., a favorite color of purple) can be quickly incorporated (blank divisions within the rectangle) and existing types can be combined to form composite events (“green” and “blue” into “not red”).

Figure 4. Favorite colors of students in a kindergarten class.

Figure 5. First graders’ iconic representation of results from a survey of students’ favorite recess activities. The figure shows eight students who like “chasing games,” and one student who likes the “structures,” which apparently includes a ramp or slide.

Figure 6. Stacked “dotplot” of the number of years 23 students’ families had lived in their town. Focusing on individual cases, we may note that the longest a family has lived in town is 37 years (data as case value). Observing that four families have lived in town for 3 years, we shift the perceptual unit to cases of

the same value (data as classifier). To perceive that about half of these families have lived in town less than 10 years entails viewing the entire data set as the unit (data as aggregate).

Figure 7. A portion of a graph made by two fifth graders at Fort River Elementary School, Amherst, MA. The entire graph spanned 3 pages. The students represented each of 24 cats with a bar whose height corresponded to the length (in inches) of the cat's tail. The bars are labeled with abbreviated forms of the cats' names and include tail length at the top. Students were graphing different attributes of the cats, including body length, weight, and tail length, to use as a basis for making recommendations to a make-believe business that wanted to design one-size-fits-all cat clothing. The students who made this graph later commented that it was not very good for showing what a "typical" tail length was.

Figure 8. Scatterplot showing the relationship between time (in seconds) spent brushing and percent of plaque remaining on teeth for a sample of 48 people. These hypothetical data were used in a post instruction interview of eighth grade students who were participating in the teaching experiment by Cobb and his colleagues (Cobb, McClain, & Gravemeijer, 2003).

Figure 9. Graph showing the family sizes of each student in the class, with written summaries below the graph, DMI Casebook, p. 15. The fourth graders who made this graph used stick figures to indicate cases (a student/family) and a zero to clearly indicate values with zero frequencies. Gender is indicated with long hair (girls) and baseball caps (boys). What the student who wrote the summaries meant by the phrase "Most people" is unclear.

Figure 10. Frequency bar graph showing the number of surveyed students at grades 8-12.

Figure 11. Boxplots of weekly homework hours for a sample of Holyoke High School students with (yes) and without (no) curfews. This type of display proves difficult to interpret when viewing data as either case values or classifiers because information about particular cases or the frequency of case-types have been forfeited to offer a more holistic view.

Figure 12. Histograms showing the number of weekly homework hours for surveyed students with (yes) and without (no) curfews. Even though these histograms depict frequencies, the fact that they are frequencies of aggregated values makes them trickier to interpret when focusing on data as either classifiers or case values.

Figure 13. Results of a survey in a kindergarten class to the question "Do you like to work on a computer?" Each student responded to the question by attaching a marker to the appropriate chart area. The elaborated responses were invented by students who felt that their view was not adequately captured by a "yes" or "no."

Figure 14. Kenny's case-value plot showing the size of 12 families (family #3 has 5 members). The x axis probably indicates the order in which Kenny obtained the data.

Figure 15. Scatterplot showing the relationship between calories consumed and self reports of energy level for 50 football players. These hypothetical data were used in a post instruction interview of eighth grade students who were participating in the teaching experiment by Cobb and his colleagues (Cobb, McClain, & Gravemeijer, 2003).

Figure 16. Stacked dotplot of the number of family members of students in a 4th grade class.

Figure 1.

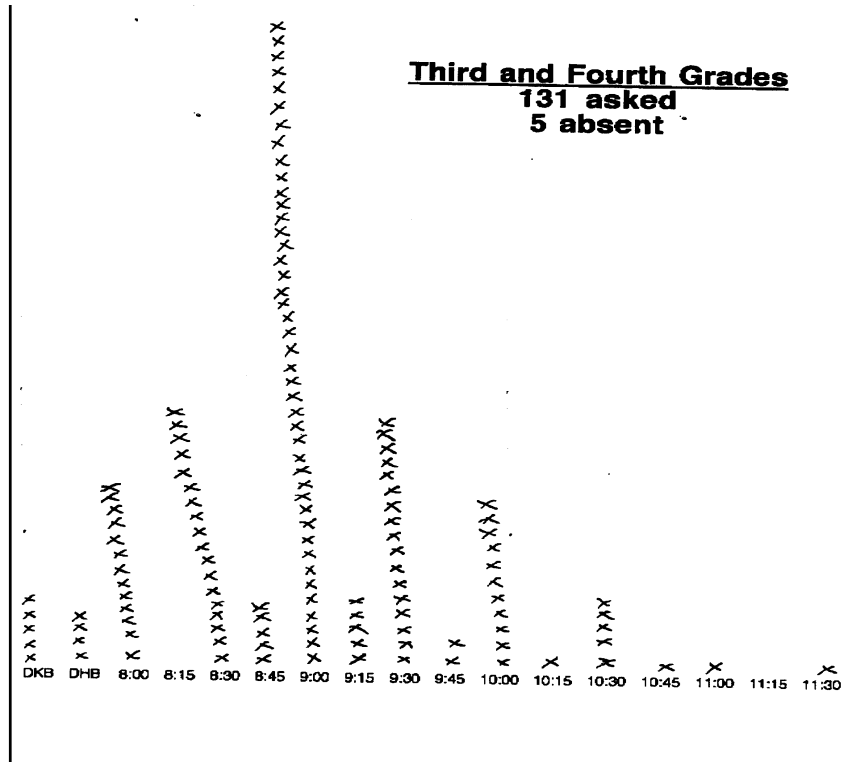


Figure 2.

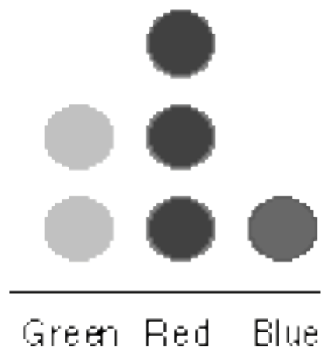


Figure 3.






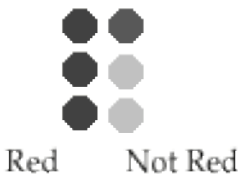
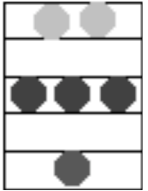
View of Data	Student Observation	Perceptual Unit	Data Structure
Pointer	We said our favorite colors	?	
Case Value	Juan likes red		
Classifier	Three like red		
Aggregate	Half like red		

Figure 4.

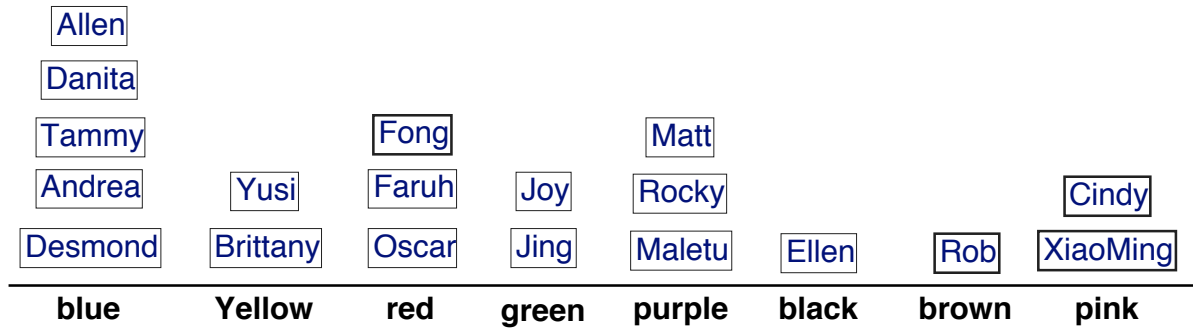


Figure 5.

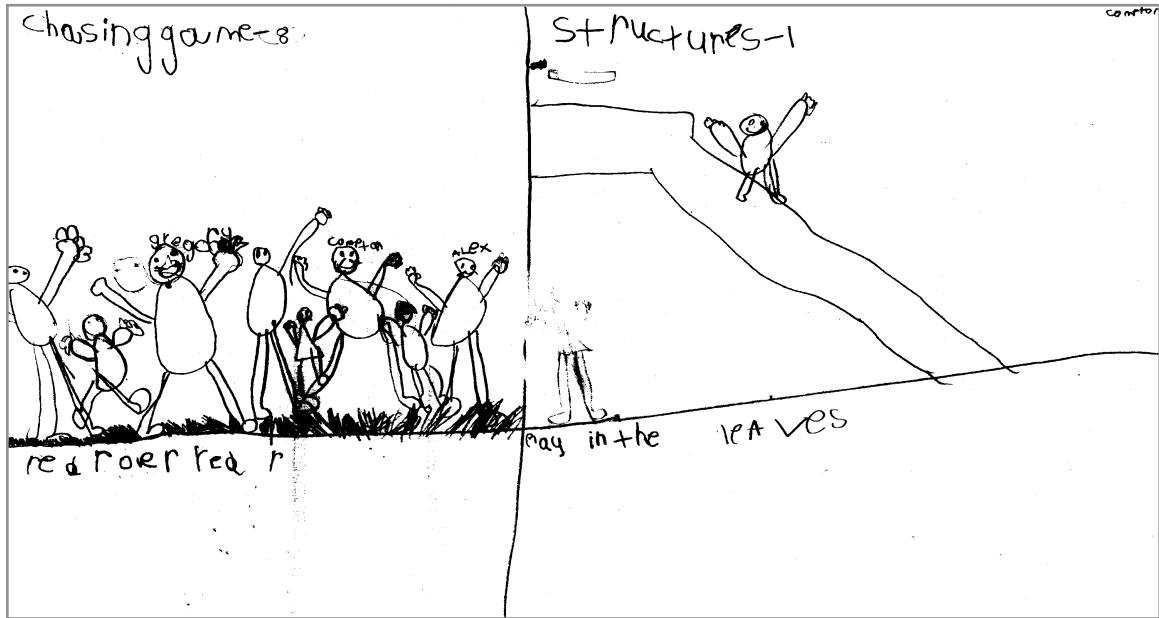


Figure 6.

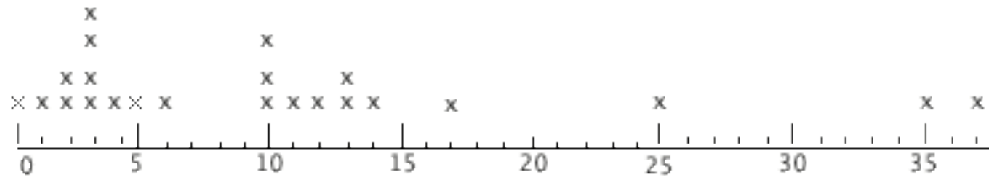


Figure 7.

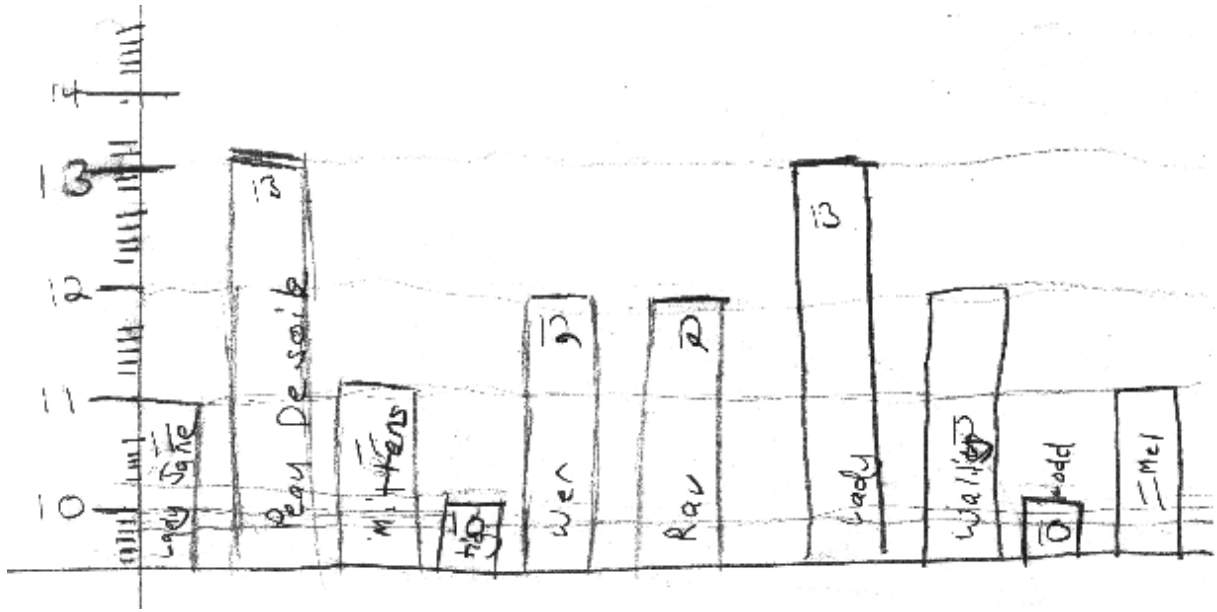


Figure 8.

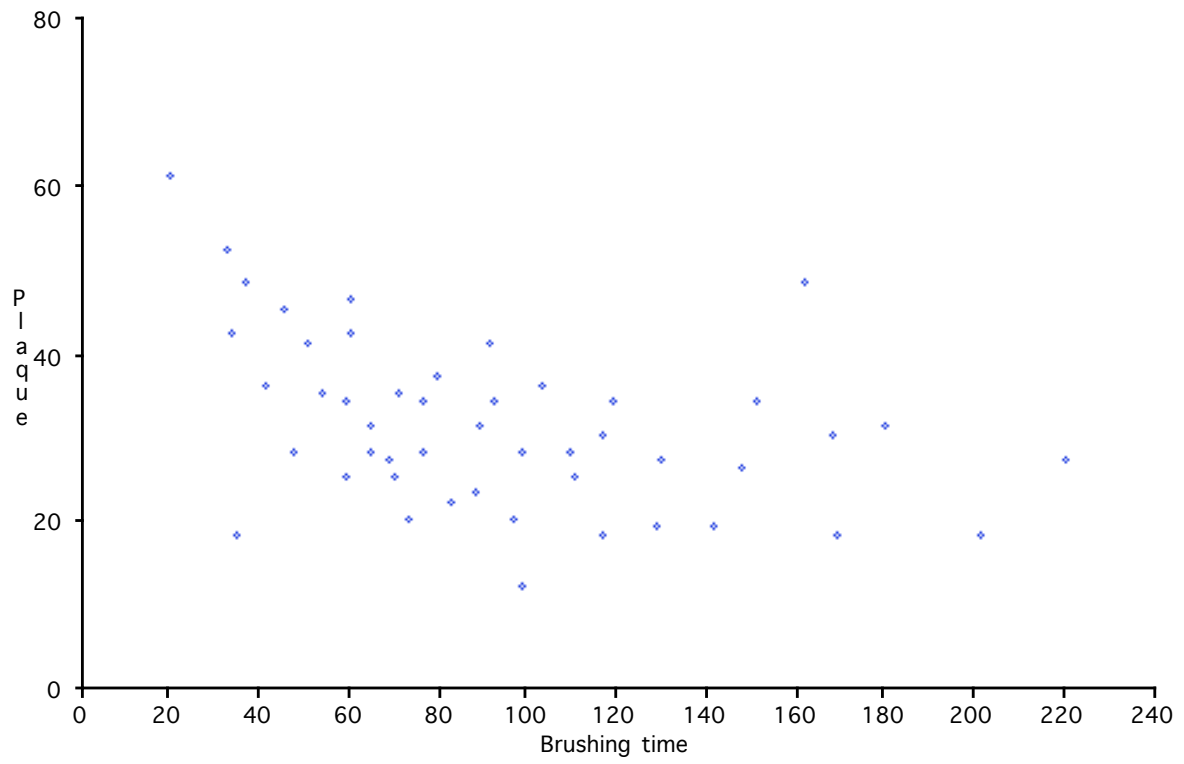


Figure 9.

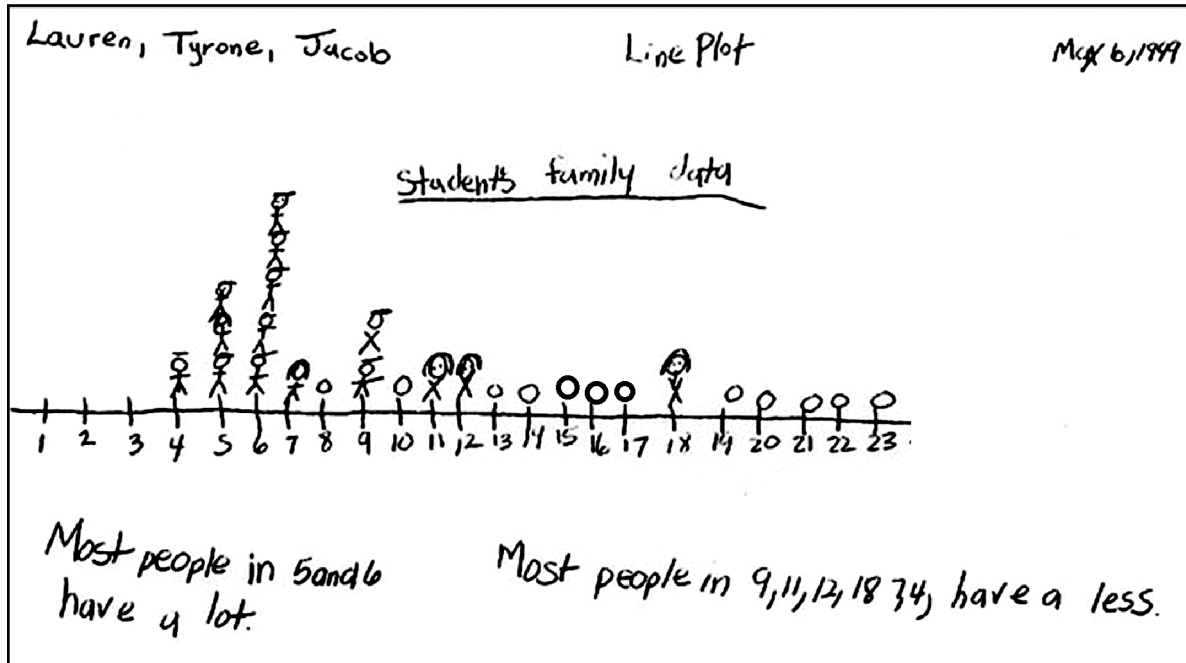


Figure 10.

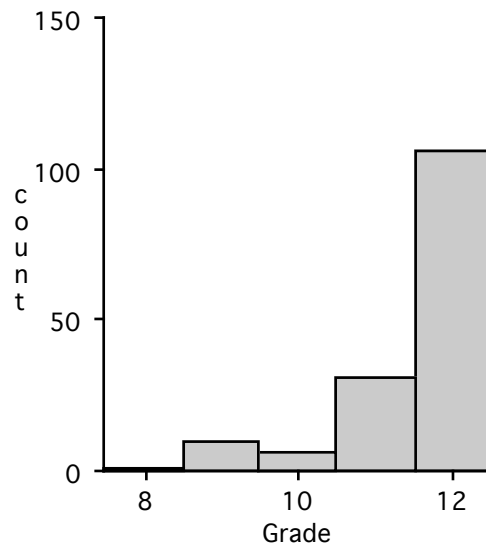


Figure 11.




Figure 12.





Figure 13.




Figure 14.




Figure 15.

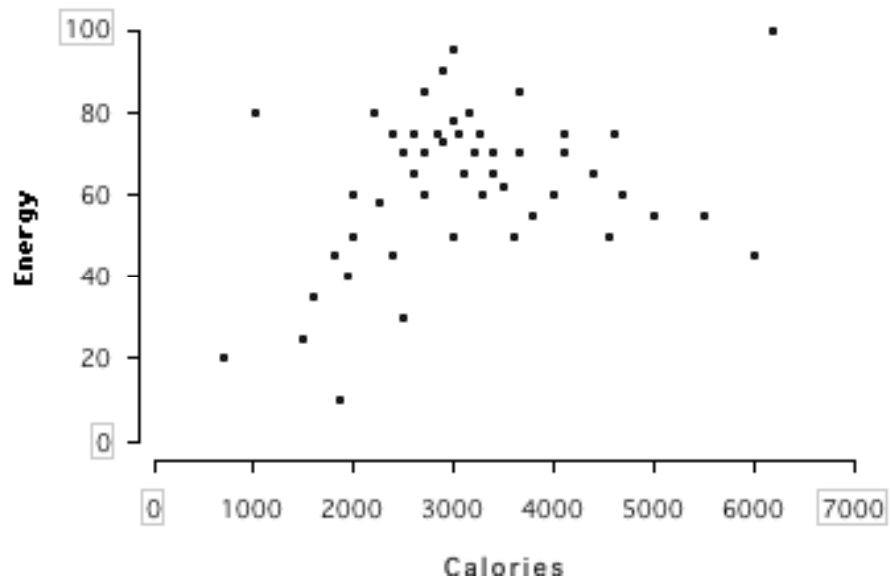


Figure 16.

