



Exploring Data Subsets with `vtree`

Nick Barrowman 
CHEO Research Institute
University of Ottawa

Richard J. Webster 
CHEO Research Institute

Abstract

Variable trees are a new method for the exploration of discrete multivariate data. They display nested subsets and corresponding frequencies and percentages. Manual calculation of these quantities can be laborious, especially when there are many multi-level factors and missing data. Here we introduce variable trees and their implementation in the `vtree` R package, draw comparisons with existing methods (contingency tables, mosaic plots, Venn/Euler diagrams, and UpSet), and illustrate their utility using two case studies. Variable trees can be used to (1) reveal patterns in nested subsets, (2) explore missing data, and (3) generate study flow diagrams (e.g., CONSORT diagrams) directly from data frames, to support reproducible research and open science.

Keywords: discrete, multivariate, visualization, tree, nested, subsets, R, CONSORT diagram, PRISMA diagram, study flow diagram.

1. Introduction

Data exploration is a vital step to gain insights into data sets. Raw data needs to be cleaned, merged, summarized and assessed. This process is resource intensive, accounting for 80% of time spent on data analysis, by one estimate (Hellerstein, Rattenbury, Heer, Kandel, and Carreras 2017). Furthermore, decisions made in this stage can impact scientific rigor and reproducibility. Recently, an appreciation has emerged for systematic and transparent protocols about data inspection steps to be performed prior to formal data analysis (e.g., Huebner, Vach, and le Cessie 2016). Such protocols are designed to provide structure at this key stage while preventing statistical fishing for results.

Tools for data exploration, like tables and figures, have been historically important for science. For instance, in the late 1800s Florence Nightingale used rose plots to discover patterns in data that matched her clinical intuition—that most soldiers in the Crimean War were dying from hygiene-related infections rather than on the battlefield—and subsequently used this to

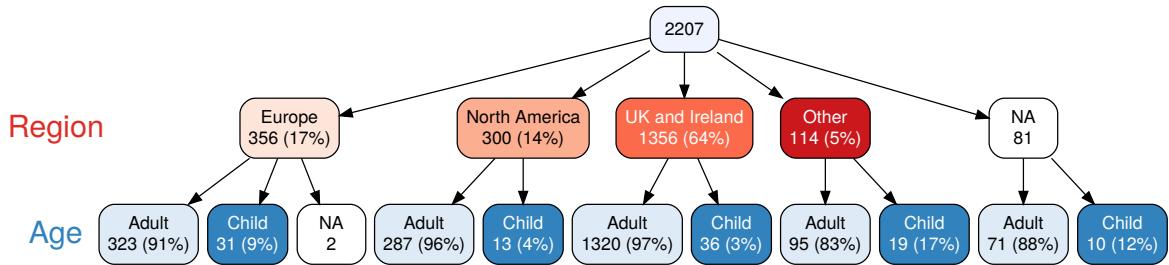


Figure 1: Variable tree for age nested within region of origin for people onboard the Titanic.

influence the British Parliament (Nelson and Rafferty 2012). This and other methods were a catalyst for the early-1900s revolution of statistical inference in many scientific fields.

Data exploration tools are more important today than ever. Data is more ubiquitous with greater volume, velocity and variety than any time in history (Katal, Wazid, and Goudar 2013). Further, these data are more accessible to analysis due to cheaper and more powerful computation (Waldrop 2016). Consequently, data literacy and intuitive data exploration tools are required for exploring and communicating findings.

In this paper we introduce variable trees as a tool for exploring subsets of data, and their implementation in the **vtree** package R (R Core Team 2025), which is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=vtree>. The objectives of this paper are (i) to compare variable trees to several established data exploration tools, (ii) to review the functionality of the **vtree** package, and (iii) to demonstrate the utility of variable trees in two case studies.

1.1. Variable trees

Subsets play an important role in almost any data analysis. Consider the variables relating to the 2207 passengers and crew members of the Titanic, represented in the the data set **titanic** from the **stablelearner** R package (Philipp, Zeileis, and Strobl 2016). Among other variables, the data set includes each person’s home country, which we have grouped into regions, and age, which we have divided into children (under age 13) and adults (see the code to perform these transformations at the beginning of Section 3). In Figure 1 each person’s home region is shown, and within each region the number and percentage of children and adults are shown. Missing values are shown as NA. We call this a *variable tree*. The **vtree** package provides a general solution for drawing variable trees and describing nested subsets.

Even in simple situations like Figure 1, it can be a chore to keep track of nested subsets and calculate the corresponding percentages. The denominator used to calculate percentages may also depend on whether the variables have any missing values, as discussed later. Finally, as the number of variables increases, the magnitude of the task balloons, because the number of nested subsets grows exponentially.

1.2. The structure of a variable tree

A variable tree consists of *nodes* connected by arrows. At the top of Figure 1, the *root node* of the tree contains all 2207 people on the Titanic. The rest of the nodes are arranged in successive layers, where each layer corresponds to a specific variable. This highlights one

	UK and Ireland	Europe	North America	Other	NA
	1356 (64%)	356 (17%)	300 (14%)	114 (5%)	81
Child	36 (3%)	31 (9%)	13 (4%)	19 (17%)	10 (12%)
Adult	1320 (97%)	323 (91%)	287 (96%)	95 (83%)	71 (88%)
NA	0	2	0	0	0

Table 1: Contingency table for `Region` → `Age` for people on the Titanic. For each combination of region and age, the frequency and column percentage for age within each region are shown. Also, directly below the name of each region, the marginal frequency and percentage are shown.

difference between variable trees and some other kinds of trees: each layer of a variable tree corresponds to just one variable. This is distinct from decision trees, where a layer may include splits based on different variables.

The nodes immediately below the root node in Figure 1 represent values of `Region` and are referred to as the *children*¹ of the root node. Inside each of the nodes, the number of people is displayed and—except for in a missing value node—the corresponding percentage is also shown. An example of a missing value node appears in Figure 1, where `Region` was missing (NA) for 81 people. Note that, by default, `vtree` displays “valid” percentages, i.e., the denominator used to calculate the percentage is the total number of *non-missing* values, in this case 2126. By default, `vtree` displays the full missing-value structure of the specified variables.

The final layer of the tree corresponds to values of `Age`. Each of these *leaf* nodes represents children and adults nested within a subset defined by a value of `Region`. Let us use the notation `Region` → `Age` to represent `Age` nested within `Region`.

A missing-value node, like any node, can have children. For example, of the 81 people for whom `Region` is missing, 10 were children and 71 were adults.

2. Methods of displaying discrete multivariate data

A variety of tools have been developed to display joint distributions of discrete variables, the most basic being the contingency table, often enhanced with row, column, or table percentages. For example, Table 1 presents the same information as Figure 1. Note that the use of column percentages encourages the reader to focus on age group nested within region.

While the contingency table above is more compact than the variable tree in Figure 1, we find the variable tree to be more intuitive. Furthermore, domain experts often respond well to such visual representations.

Now suppose we would like to examine `Region` → `Age` → `Survived` (i.e., survival within age within region of origin). Multi-way cross classifications (three or more variables) are typically displayed using several two-way tables, referred to as *layers* or *slices*. Table 2 shows two-way tables of survival within age group for each of the regions of origin. This is followed by a variable tree showing the same information (Figure 2).

¹Since the Titanic example involves actual children, we need to be careful here with this technical use of the term “children”.

UK and Ireland 1356 (64%)			
	Child	Adult	NA
	36 (3%)	1320 (97%)	0
Survived	17 (47%)	347 (26%)	0
Did not survive	19 (53%)	973 (74%)	0
North America 300 (14%)			
	Child	Adult	NA
	13 (4%)	287 (96%)	0
Survived	7 (54%)	160 (56%)	0
Did not survive	6 (46%)	127 (44%)	0
Europe 356 (17%)			
	Child	Adult	NA
	31 (9%)	323 (91%)	2
Survived	13 (42%)	91 (28%)	0 (0%)
Did not survive	18 (58%)	232 (72%)	2 (100%)
Other 114 (5%)			
	Child	Adult	NA
	19 (17%)	95 (83%)	0
Survived	16 (84%)	32 (34%)	0
Did not survive	3 (16%)	63 (66%)	0
NA 81			
	Child	Adult	NA
	10 (12%)	71 (88%)	0
Survived	5 (50%)	23 (32%)	0
Did not survive	5 (50%)	48 (68%)	0

Table 2: Contingency table layers for **Region** → **Age** → **Survived**. The name of each region is shown along with the marginal frequency and percentage, and underneath, the two-way contingency table for **Age** → **Survived** within that region. Along the top row of each table, the marginal frequency and percentage for age within that region is shown. In each table, frequency and column percentage for survival within each age and region are shown.

Note that by default, **vtree** shows percentages in each node except for the root. For example, of the 2207 people on board the Titanic, 300 (14%) were from North America, of whom 287 (96%) were adults, of whom 160 (56%) survived. In its simplest form, a contingency table only shows crosstabulated frequencies, corresponding to the frequencies shown in the leaf nodes of a variable tree. However contingency tables are often shown with additional marginal and conditional percentages.

As the number of variables increases, contingency tables for multi-way classifications become increasingly difficult to interpret. In such situations, large variable trees can also become unwieldy, however this can be mitigated by *pruning* branches of lesser interest.

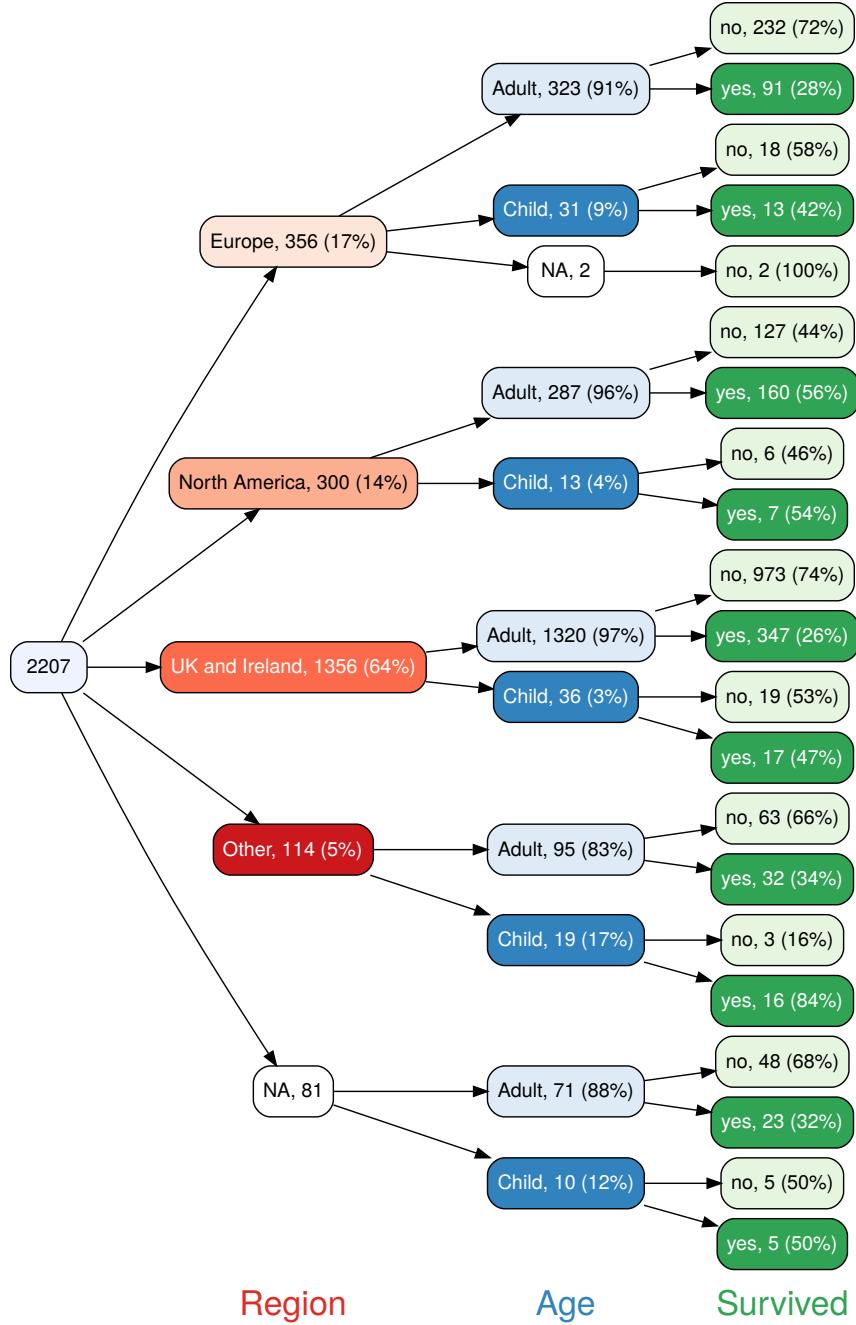


Figure 2: Variable tree for `Region → Age → Survived` with number and percent survival shown in each node. Table 2 shows the same information.

Contingency tables are not *always* more compact than variable trees. When most cells of a large contingency table are empty (in which case the table is said to be *sparse*), the corresponding variable tree may be much more compact since empty nodes are not shown. In the Titanic data set, there are two missing values of Age, and both are for individuals from Europe. This appears as a single node in Figure 2, but in Table 2 in addition to the cell showing these 2 missing values, there are 9 cells containing zero.

Like contingency tables, variable trees show numerical values (frequencies and percentages) rather than using graphical elements such as area to encode such quantities. In contrast to contingency tables, which use a tabular layout to represent subsets, variable trees use the graphical arrangement of nodes and arrows to represent the nesting structure. Put another way, *vtree visualizes relationships, not quantities*.

2.1. Visualization of discrete multivariate data

Several visualization methods have been proposed for discrete multivariate data. Here we focus on displaying raw data and descriptive statistics. Beyond the scope of this paper are various visualizations for fitted models and their residuals (Zeileis, Meyer, and Hornik 2007).

Contingency tables can be represented graphically by encoding the quantity in each cell into the size or color of an element. For example, each cell of a *balloon plot* contains a circle with area proportional to frequency (Moon 2016). Balloon plots can be produced in R using the function `ggballoonplot()` in the `ggpubr` package (Kassambara 2020). Concordance of ordinal categorical data may be visualized using *agreement charts* (Bangdiwala and Shankar 2013); these can be produced in R using the `agreementplot` function from the `vcd` package (Hornik, Zeileis, and Meyer 2006).

Barplots are another common discrete multivariate data visualization tool. They can also be produced for subsets, defined by values of another variable. A more compact representation is the stacked barplot and grouped stacked barplot. A limitation with stacked barplots is the difficulty of between-category comparisons, since there is no common baseline, except for the bottom category in a stack.

An elegant extension of the stacked barplot is the *mosaic plot* (Hartigan and Kleiner 1981). In a mosaic plot, the area of each rectangle represents the number of observations in the corresponding subset of the data. Mosaic plots are available in R through several packages e.g., `vcd`, `ggbmosaic` (Jeppson, Hofmann, Cook, and Wickham 2020), `iplots` (Urbanek and Wichtrey 2018), as well as in the base R function `mosaicplot()`. The strucplot framework implemented in the `vcd` package provides flexibility for hierarchical conditional plots (e.g., conditional mosaic plots, association plots, double decker plots, sieve plots and more). Mosaic plots can provide an intuitive visual representation of the number of observations in subsets of the data, however they tend to become visually complicated when there are more than three variables. Figure 3 is a mosaic plot for `Region` → `Age` → `Survived` for the people onboard the Titanic, as in Table 2 and Figure 2.

Visualizations like Figure 3 have advantages and disadvantages compared to text and tabular summaries. On the one hand, they represent quantitative and qualitative information in a way that is quickly decoded by our visual perceptual systems. On the other, visualizations can be unfamiliar and even perplexing compared to numerical and tabular representations. A practical advantage of tabular summaries is that they are more easily formatted and manipulated using current software. Variable trees share characteristics with both tabular representations and visualizations. To put it succinctly, variable trees visualize *relationships* not *quantities*.

2.2. Data representing set membership

A special type of discrete multivariate data is when all of the variables are binary, in which case they can be interpreted as representing set membership. Venn and Euler diagrams have

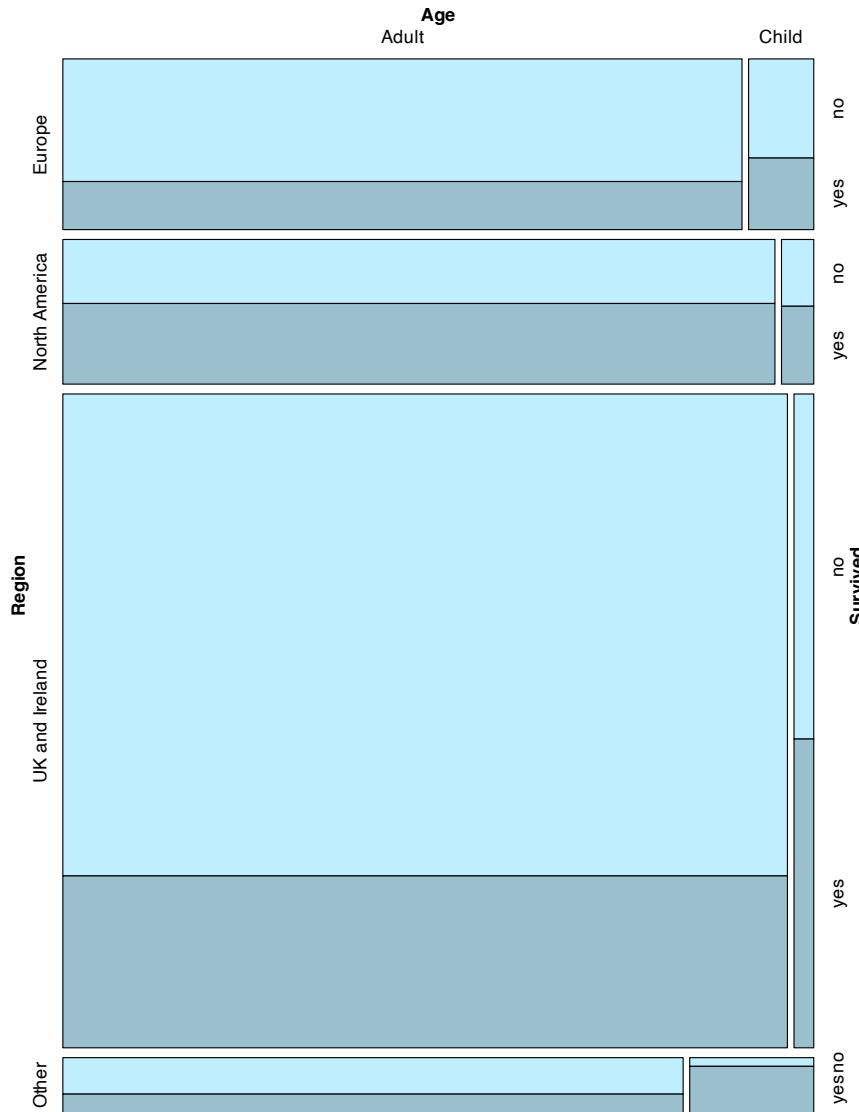


Figure 3: Mosaic plot for `Region → Age → Survived` for people on the Titanic. Each rectangle corresponds to a subset of the data and the area of the rectangle represents the relative frequency. Table 2 and Figure 2 show the same information.

long been used to represent the intersection of sets. Venn diagrams use overlapping closed curves such that all intersections between sets are represented by overlapping areas. Euler diagrams are like Venn diagrams but empty intersections need not be shown. For data sets, software is available to calculate the number of observations in each of the intersections, for example in R, the **VennDiagram** (Chen 2018) and **venneuler** (Wilkinson 2011) packages. A further elaboration of these diagrams is to make the areas of the sets and their intersections approximately proportional to the number of observations in each subset. The package **eulerr** (Larsson 2020) provides this functionality. For example, in Figure 4 a dataset of Wilkinson (2012) is represented using an approximately area-proportional Euler diagram. As the number of sets grows, Venn and Euler diagrams can become unwieldy.

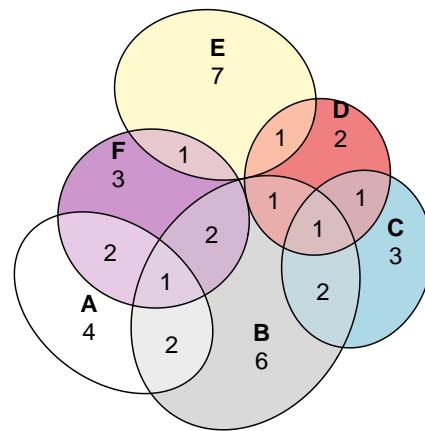


Figure 4: Euler plot for the dataset of [Wilkinson \(2012\)](#).

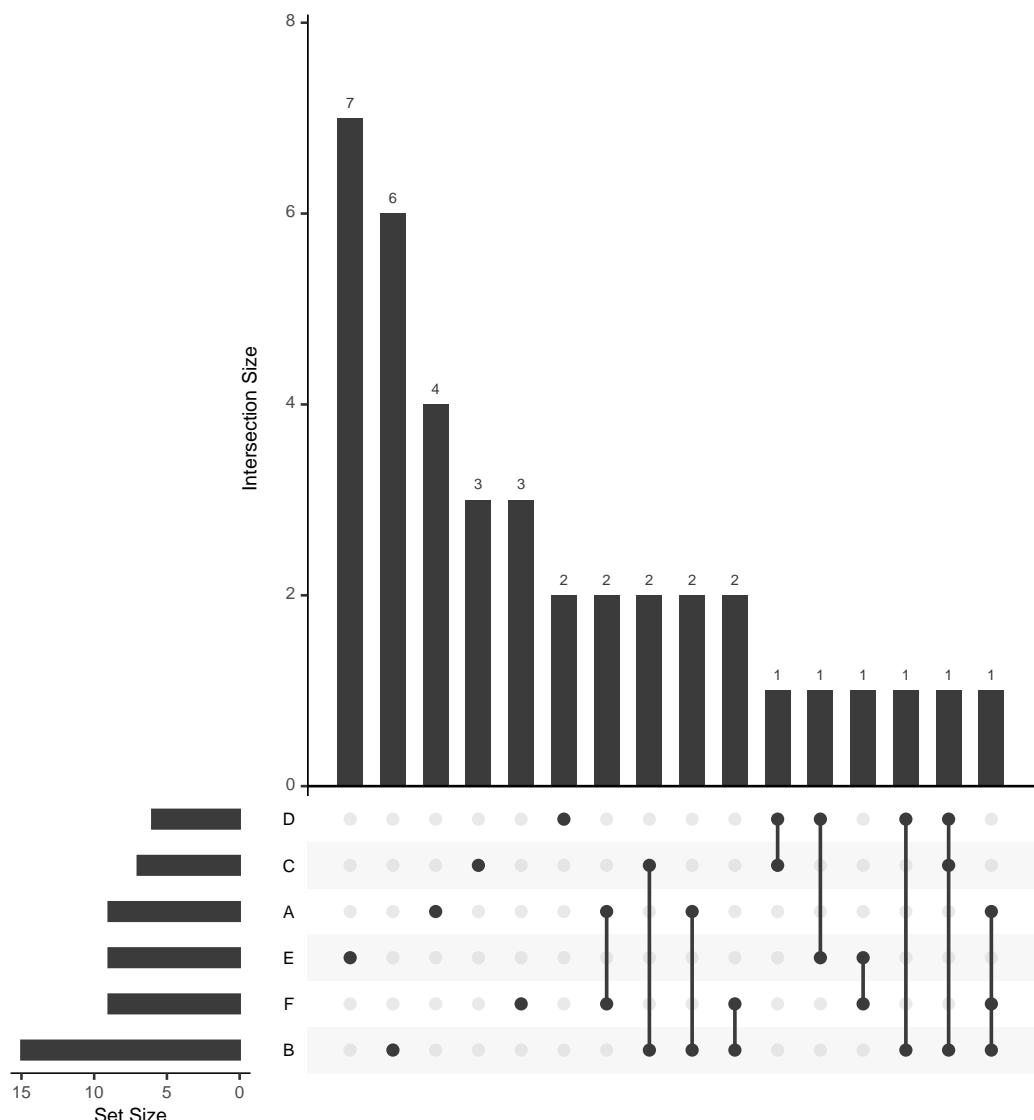
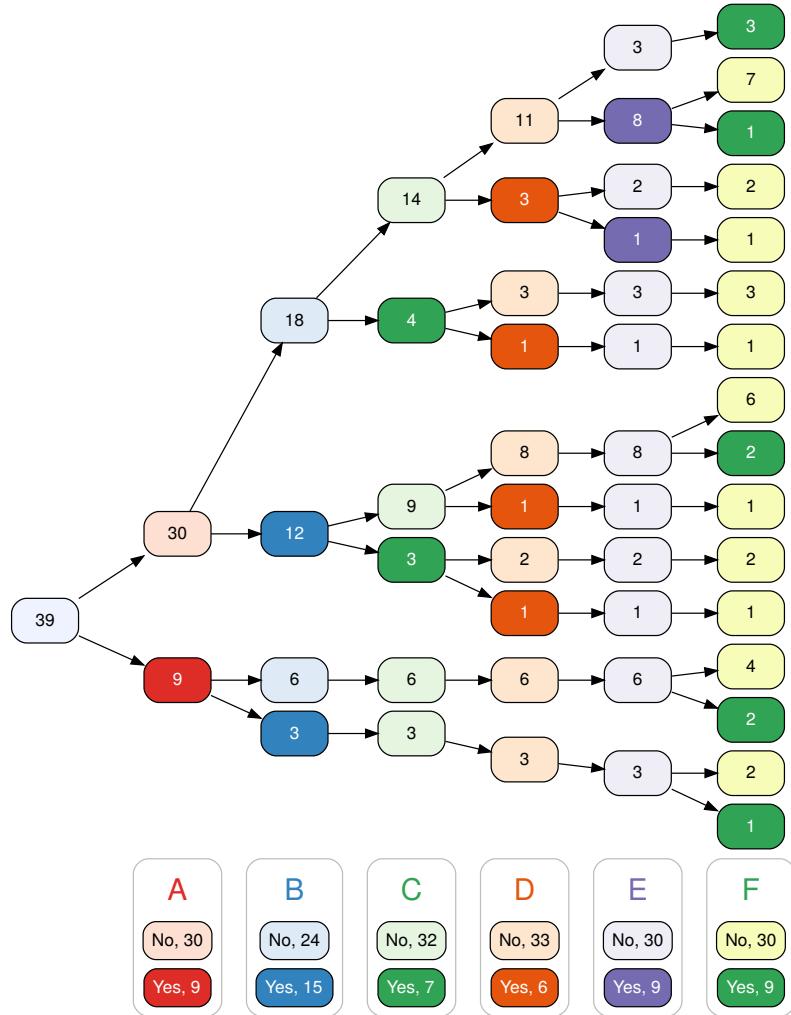


Figure 5: An UpSet plot.



An innovative way to represent the intersections of a large number of sets is the *UpSet* plot (Lex, Gehlenborg, Strobelt, Vuillemot, and Pfister 2014). The R package **UpSetR** (Conway, Lex, and Gehlenborg 2017) was used to produce Figure 5 for the dataset of Wilkinson (2012). UpSet uses a grid layout to represent the intersections (see the dots at the bottom of Figure 5), together with bar graphs on each side to represent the size of sets and intersections.

Variable trees can also represent the intersection of sets, however unlike UpSet and area-proportional Euler diagrams, they do not use graphical elements to encode quantity. Like non-proportional Venn Diagrams, variable trees graphically depict the relationships between subsets of the data, but represent quantities numerically (Figure 6). Unlike Venn, Euler, and UpSet diagrams, variable trees require a prespecified ordering. For example, Figure 6 uses the ordering A → B → C → D → E → F.

An alternative diagram that **vtree** can produce is the *pattern tree*, which depicts every intersection. Each row in Figure 7 corresponds to the combination of values represented by a terminal node in Figure 6. Since the intermediate nodes in Figure 6 are not represented,

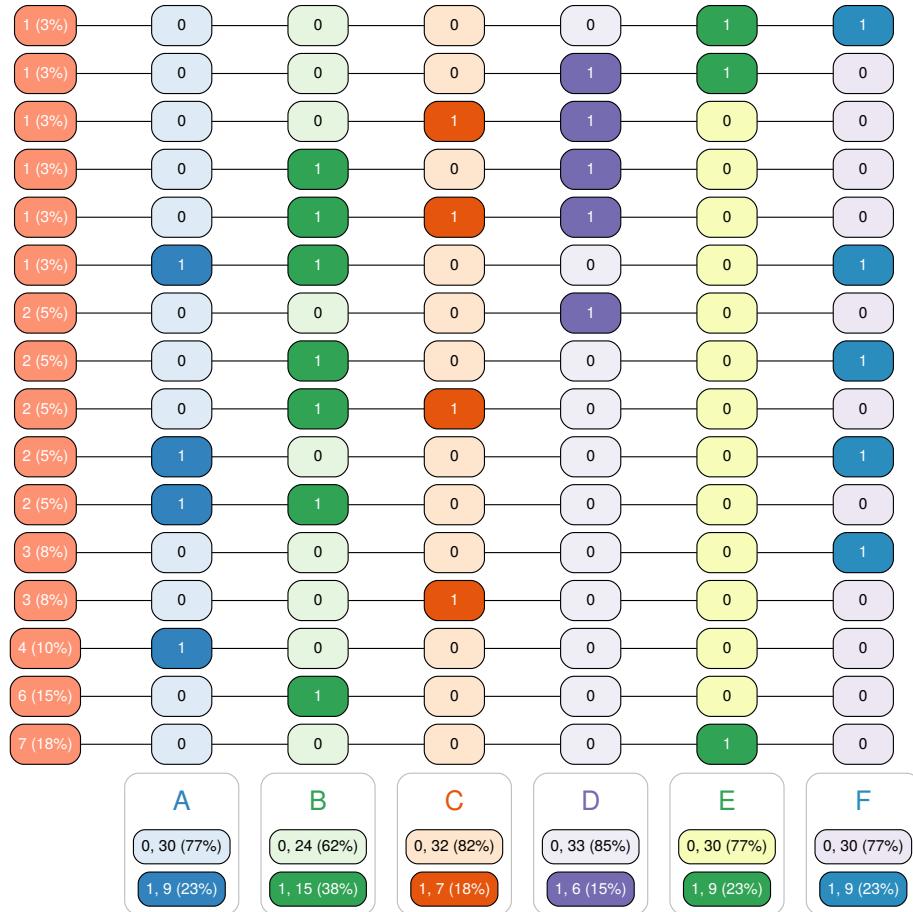


Figure 7: A pattern tree for the dataset of Wilkinson (2012). Each row represents a “pattern” (here a particular intersection), corresponding to a terminal node in Figure 6.

information is lost. The pattern tree is much easier to read, however. Pattern trees have some of the same structure as an UpSet plot, except that sizes of subsets are not represented graphically as in the bar graphs on the sides of an UpSet plot.

3. Package functionality

This section provides an overview of the features of the **vtree** package. Additional resources are available in the package [vignette](#), a [cheatsheet](#), and [video tutorials](#) on YouTube.

To illustrate the functionality of the **vtree** package, we use the **titanic** data set from the **stablelearner** package.

```
R> library("vtree")
R> library("stablelearner")
R> library("dplyr")
R> library("forcats")
R> data("titanic", package = "stablelearner")
```

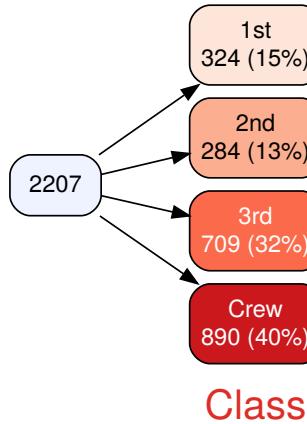


Figure 8: A simple variable tree.

```

R> td <- titanic %>%
+   rename(Survived = survived) %>%
+   mutate(
+     Age = ifelse(age < 13, "Child", "Adult"),
+     Gender = gender,
+     Class = fct_collapse(class,
+       Crew = c("deck crew", "engineering crew", "restaurant staff",
+       "victualling crew")),
+     Region = fct_collapse(country,
+       "UK and Ireland" = c("England", "Scotland", "Wales", "Ireland",
+       "Northern Ireland", "Channel Islands"),
+       "Europe" = c("Norway", "France", "Finland", "Sweden", "Latvia",
+       "Denmark", "Bulgaria", "Greece", "Hungary", "France", "Spain",
+       "Italy", "Belgium", "Germany", "Austria", "Poland",
+       "Switzerland", "Bosnia", "Croatia", "Croatia (Modern)",
+       "Yugoslavia", "Slovakia (Modern day)", "Slovenia", "Netherlands",
+       "Russia"),
+       "North America" = c("United States", "Canada", "Mexico", "Cuba"),
+       other_level = "Other"))
  
```

3.1. Calling vtree

Suppose the Titanic data are in a data frame called `td`. To display a variable tree for a single variable, say `Class`, use the following command:

```
R> vtree(td, "Class")
```

The output can be seen in Figure 8.

To produce a variable tree for `Class` → `Age`, specify `"Class Age"` (see Figure 9):

```
R> vtree(td, "Class Age", horiz = FALSE)
```

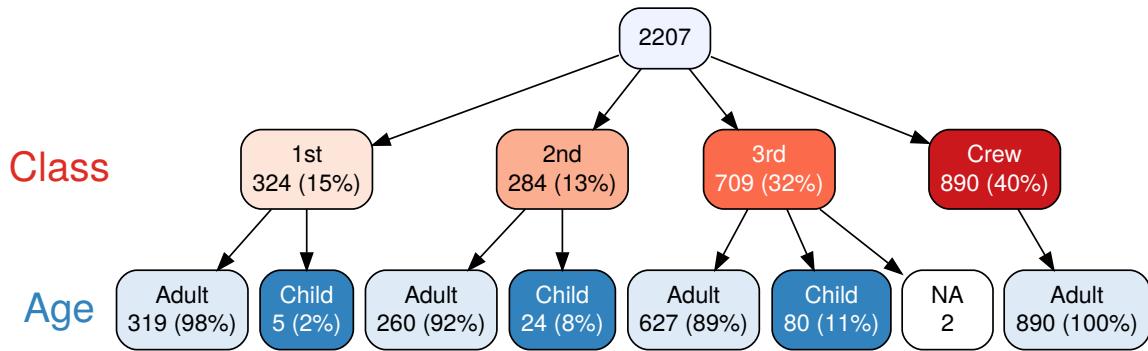


Figure 9: A two-layer vertical variable tree.

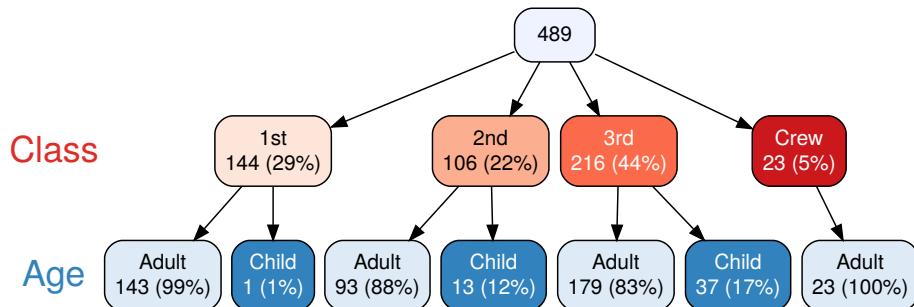


Figure 10: Class of female passengers and crew onboard the Titanic.

By default, **vtree** produces horizontal trees. The tree in Figure 9 is vertical because of the specification `horiz = FALSE`.

Note that **vtree** can also be called at the end of a **dplyr** (Wickham, François, Henry, Müller, and Vaughan 2023) pipeline. This allows variables and data to be modified for use in a variable tree. For example, the following commands produce a tree like the one in Figure 9, but *only* including female passengers and crew:

```
R> library(dplyr)
R> td %>%
+   filter(Gender == "female") %>%
+   vtree(~ Class + Age, horiz = FALSE)
```

The output is presented in Figure 10.

The example above includes the call `vtree(~ Class + Age, horiz = FALSE)`, which uses a formula to specify variables rather than a character string. (This notation is convenient, however it precludes the use of **vtree**'s built-in [variable specifications](#) as discussed in Section 3.4.)

3.2. Pruning

When a variable tree gets too big, or you are only interested in certain parts of the tree, it may be useful to remove some nodes along with their descendants. This is known as *pruning*. For convenience, there are several different ways to prune a tree, described below.

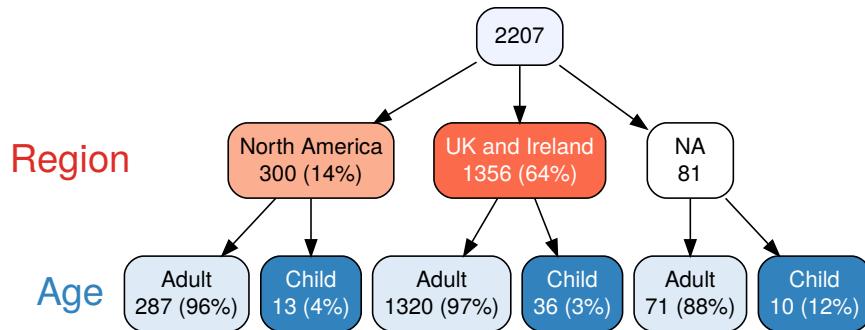


Figure 11: Using the `prune` parameter to remove “Europe” and “Other” regions.

Suppose you do not wish to show the “Europe” node or the “Other” node (which represents people from other parts of the world such as India, the Middle East, etc.). Specifying `prune=list(Region = c("Europe", "Other"))` removes those nodes, and all of their descendants:

```
R> vtree(td, "Region Age", prune = list(Region = c("Europe", "Other")),
+       horiz = FALSE)
```

The output can be visualized in Figure 11.

In general, the argument of the `prune` parameter is a *list* with an element named for each variable you wish to prune. In the example above, the list has a single element, named `Region`. In turn, that element is a vector `c("Europe", "Other")` indicating the values of `Region` to prune.

Note that once a variable tree has been pruned, it is no longer complete. This can sometimes be confusing since not all observations are represented at certain layers of the tree. For example in the tree above, only 1737 observations are shown in the `Region` nodes and their children.

Sometimes it is more convenient to specify which nodes should be *retained* rather than which ones should be discarded. The `keep` parameter is used for this purpose, and can thus be considered the complement of the `prune` parameter. Suppose we wish to create a variable tree for `Region → Class → Gender → Age`. This tree has four layers, and without any pruning it would be quite large. But suppose we are only interested in certain branches of the tree, say the “Europe” node of `Region`, the “3rd” node of `Class`, and the “male” node of `Gender`. Using the `keep` parameter, the compact tree in Figure 12 can be produced:

```
R> vtree(td, "Region Class Gender Age",
+       keep = list(Region = "Europe", Class = "3rd", Gender = "male"))
```

In the `Region` layer of Figure 12, note that along with the “Europe” node, the missing value node has also been retained. In general, whenever valid percentages are used (which is the default), missing value nodes are retained when `keep` is used. This is because valid percentages are difficult to interpret without knowing the denominator, which requires knowing the number of missing values. On the other hand, here is what happens when `vp = FALSE` (see output in Figure 13):

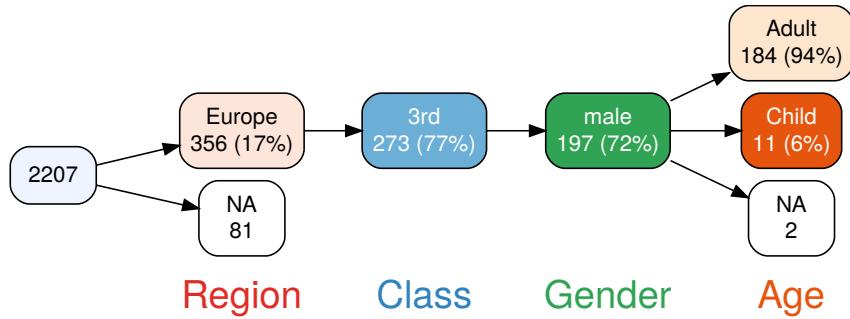
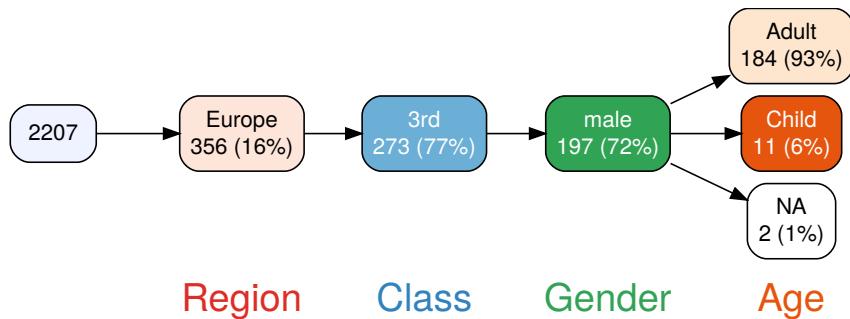
Figure 12: Using the `keep` parameter.

Figure 13: Same as Figure 12 but not using “valid” percentages.

```

R> vtree(td, "Region Class Gender Age",
+      keep = list(Region = "Europe", Class = "3rd", Gender = "male"),
+      vp = FALSE)
  
```

Note that the missing value node for `Region` is no longer present, since the percentage for the “Europe” node can be interpreted without knowing the number of missing values. Also, note that the missing value node for `Age` includes a percentage, and the percentages for the other nodes of `Age` are slightly different. (With only two missing values, the difference is slight, but as the proportion of missing data increases, the percentages become substantially different.)

An alternative is to prune *below* the specified nodes (i.e., to prune their descendants), so that the counts always add up. In the present example, this means that the other nodes will be shown, but not their descendants. The `prunebelow` parameter is used to do this (see output in Figure 14):

```

R> vtree(td, "Region Age",
+      prunebelow = list(Region = c("UK and Ireland", "North America",
+                                "Other")))
  
```

The complement of the `prunebelow` parameter is the `follow` parameter. Instead of specifying which nodes should be pruned below, this allows you to specify which nodes should be *followed* (that is, *not* pruned below).

As a variable trees grow, it can become difficult to see the forest for the tree, as it were. For example, the following variable tree is very wide, which makes it difficult to read. (One

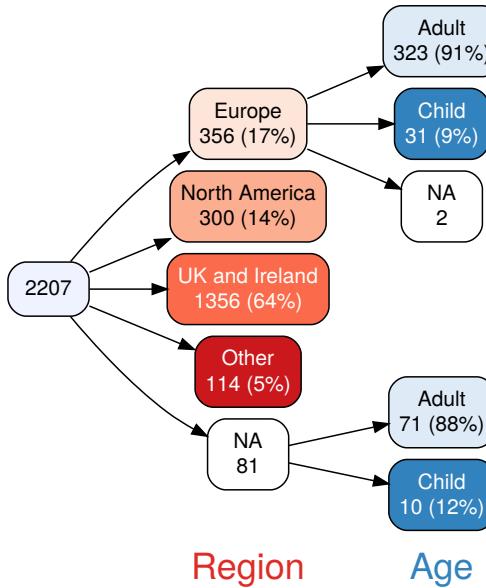
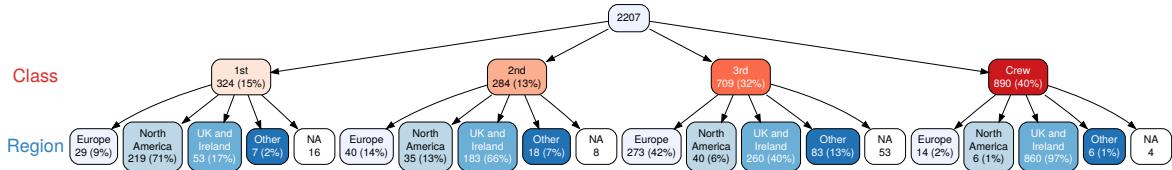
Figure 14: Using the `prunebelow` parameter.

Figure 15: A variable tree that is too wide to read easily.

small modification to make it a little narrower is the use of the `splitwidth` parameter, which specifies the number of characters of text before it gets split onto another line.)

```
R> vtree(td, "Class Region", horiz = FALSE, splitwidth = 5)
```

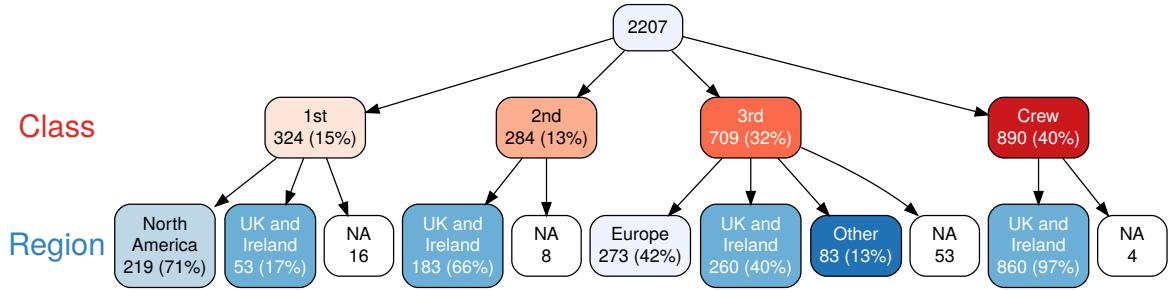
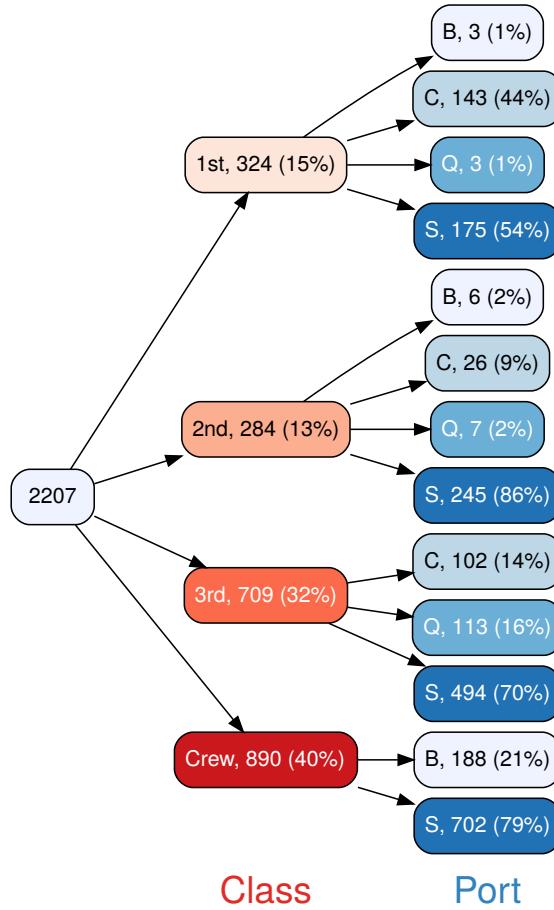
The output can be visualized in Figure 15. One solution is to prune nodes that contain small numbers of observations. For example if you want to only see nodes with at least 50 observations, you can specify `prunesmaller` = 50, as in this example (see Figure 16):

```
R> vtree(td, "Class Region", horiz = FALSE, prunesmaller = 50,
+         splitwidth = 5)
```

Similar to the `keep` parameter, when valid percentages are used (`vp = TRUE`, which is the default), nodes represent missing values will not be pruned. As noted previously, this is because percentages are confusing when missing values are not shown. On the other hand, when `vp = FALSE`, missing nodes can be pruned.

3.3. Labels for variables and nodes

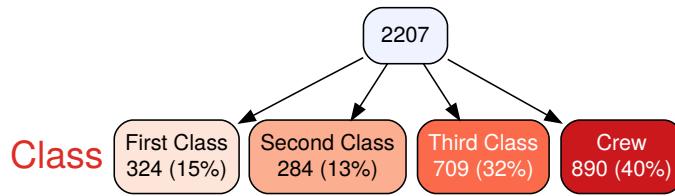
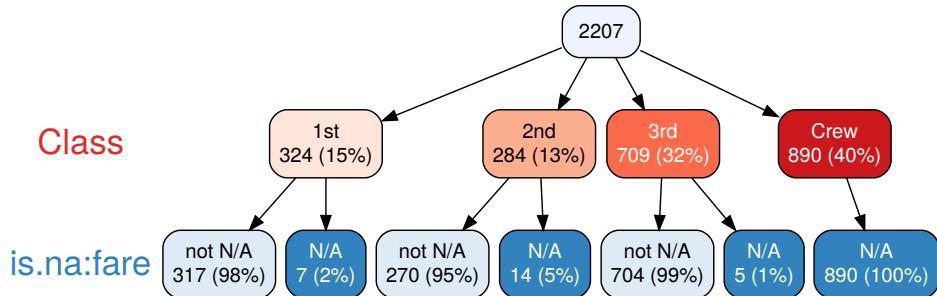
Readability of a variable tree can be improved by customizing the variable and node names using the `labelvar` and `labelnode` parameters. By default, `vtree` labels variables and nodes exactly as they appear in the data frame.

Figure 16: Using the `prunesmaller` parameter.Figure 17: Using the `labelvar` parameter.

For example, the `embarked` variable indicates the port where a passenger or crew member boarded the Titanic. Suppose we wish this variable to appear as *Port* in the variable tree. The `labelvar` parameter specifies this.

```
R> vtree(td, "Class embarked", labelvar = c("embarked" = "Port"),
+         sameline = TRUE)
```

The output is presented in Figure 17. By default, `vtree` labels nodes (except for the root

Figure 18: Using the `labelnode` parameter.Figure 19: Using the `is.na:` variable specification.

node) using the values of the variable in question. (If the variable is a factor, the levels of the factor are used). Sometimes it is convenient to instead specify custom labels for nodes. The `labelnode` argument can be used to relabel the values. For example, to relabel the classes as “First Class”, “Second Class”, and “Third Class” we can specify (see output in Figure 18):

```
R> vtree(td, "Class", horiz = FALSE, labelnode = list(Class = c(
+   "First Class" = "1st", "Second Class" = "2nd", "Third Class" = "3rd")))
```

3.4. Specification of variables

For convenience, in the call to the `vtree` function, you can specify variable names (separated by whitespace) in a single character string. (If, however, any of the variable names have internal spaces, the variable names must be specified as a vector of character strings.) Additionally, several modifiers can be used, as detailed below.

If an individual variable name is preceded by `is.na:`, that variable will be replaced by a missing value indicator in the variable tree. This facilitates exploration of missing data, for example:

```
R> vtree(td, "Class is.na:fare", horiz = FALSE)
```

The output is presented in Figure 19.

A variety of other specifications are available. For example `<`, `=`, and `>` can be used to dichotomize numeric variables. While this is a powerful tool for data exploration, a word of caution is needed. To ensure scientific rigor, it is essential that this functionality not be used to explore a variety of dichotomizations of a predictor variable in relation to the outcome variable. There is a large literature on the misuse of dichotomization and its detrimental effect

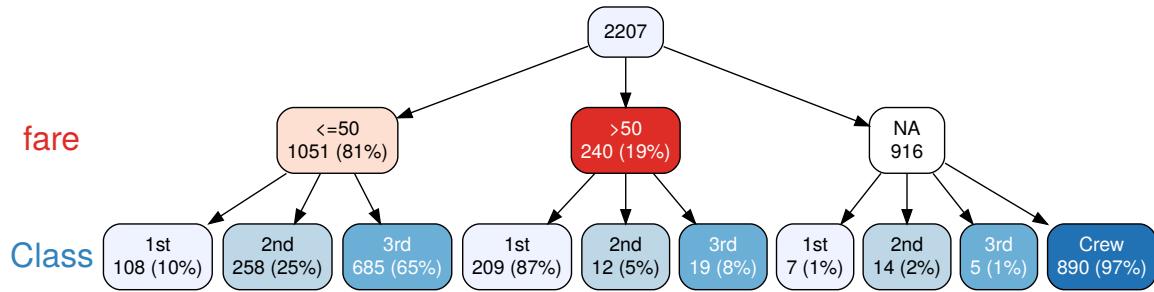
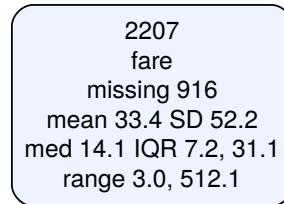
Figure 20: Using the `>` variable specification.

Figure 21: Summary information over the entire data frame.

on statistical inference (Altman 1994). It is therefore recommended that any dichotomization using **vtree** be conducted according to a pre-specified protocol (Huebner *et al.* 2016).

For example, consider the fares paid by passengers (in pounds sterling). To examine the class of passengers split according to whether they paid more than £50 or not:

```
R> vtree(td, "fare>50 Class", horiz = FALSE)
```

The output is presented in Figure 20.

3.5. Displaying summary statistics in nodes

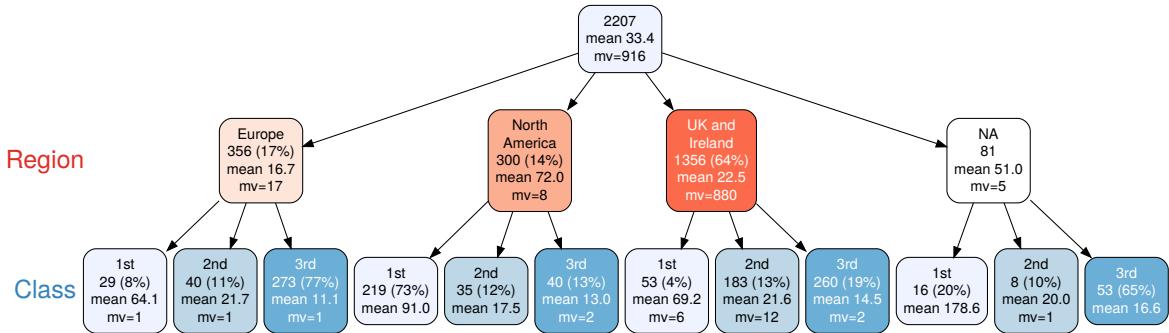
It is often useful to display information about *other* variables (apart from those that define the tree) in the nodes of a variable tree. This is particularly useful for numeric variables, which generally would not be used to build the tree since they have too many distinct values. The **summary** parameter allows you to show information (for example, the mean of a numeric variable) within each subset of the data frame.

To obtain summary information about fares paid by all of the passengers on board the Titanic (i.e., in the root node), you do not need to specify any variables for the tree itself (see Figure 21):

```
R> vtree(td, summary = "fare", horiz = FALSE)
```

Using a *summary code*, `%mean%`, the mean of this variable can be displayed within, for example, levels of **Region** and **Class**, as shown in Figure 22. (To reduce the size of this tree we'll hide the crew and the “Other” region.)

```
R> vtree(td, "Region Class", summary = "fare \nmean %mean%", horiz = FALSE,
+     prune = list(Region = "Other", Class = "Crew"), splitwidth = 5)
```

Figure 22: Using the `summary` parameter to produce customized summaries.

Summary code	Variant	Result
<code>%mean%</code>	<code>%meanx%*</code>	Mean
<code>%SD%</code>	<code>%SDx%*</code>	Standard deviation
<code>%sum%</code>	<code>%sumx%*</code>	Sum
<code>%min%</code>	<code>%minx%*</code>	Minimum
<code>%max%</code>	<code>%maxxx%*</code>	Maximum
<code>%range%</code>	<code>%rangex%*</code>	Range
<code>%median%</code>	<code>%median%*</code>	Median
<code>%IQR%</code>	<code>%IQRx%*</code>	Interquartile range
<code>%freq%</code>	<code>%freq_%**</code>	Frequency of values of a variable
<code>%npct%</code>		Frequency and percentage
<code>%pct%</code>		Same as <code>%npct%</code> but percentage only
<code>%list%</code>	<code>%list_%**</code>	List of individual values, separated by commas

Table 3: Summary codes. * = Missing values are suppressed. Caution is recommended.
** = Shows each value on a separate line.

Code	Summary information restricted to
<code>%noroot%</code>	All nodes <i>except</i> the root
<code>%leafonly%</code>	Leaf nodes
<code>%var=v%</code>	Nodes of variable v
<code>%node=n%</code>	Nodes named n

Table 4: Control codes.

Summary codes always start and end with %. A list is shown in Table 3. Sometimes, you might want to only show summary information in particular nodes. Table 4 lists codes to control where summary information is shown.

3.6. Pattern trees

Each node in a variable tree provides the frequency of a particular combination of values of the variables. The leaf nodes represent the observed combinations of values of *all* of the variables.

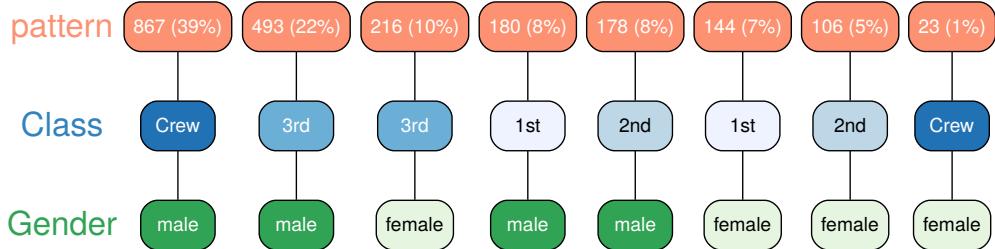
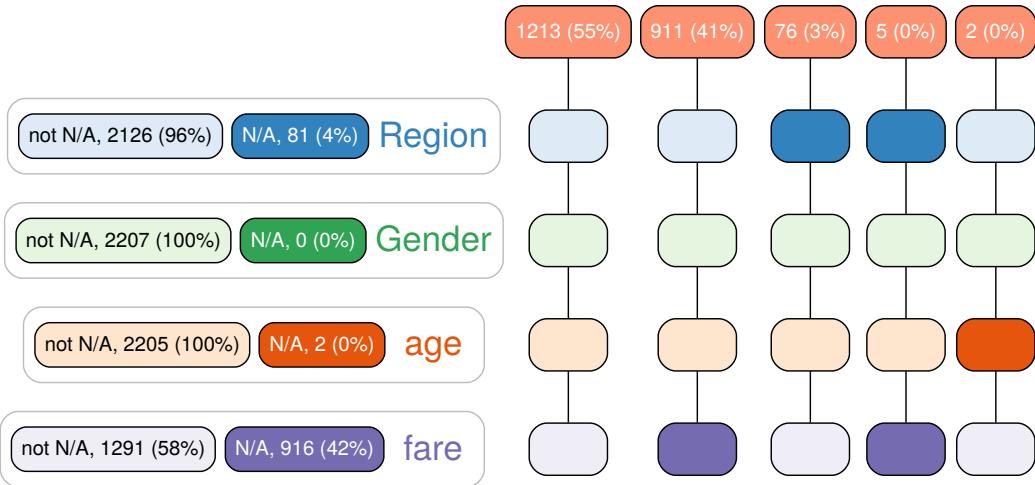


Figure 23: A pattern tree.

Figure 24: Using the `check.is.na` parameter.

For example, in a variable tree of `Gender` nested within `Class`, the leaf nodes correspond to Male and Female. These combinations, or *patterns*, can be treated as an additional variable. And if this new pattern variable is used as the first variable in a tree, then the branches of the tree will be simplified: each branch will represent a unique pattern, with no sub-branches. A *pattern tree* can be easily produced by specifying `pattern = TRUE`. For example Figure 23 is produced by:

```
R> vtree(td, "Class Gender", horiz = FALSE, pattern = TRUE)
```

A special pattern tree with missing values shown as darker nodes can be produced using the `check.is.na` parameter. Additionally the `showlegend` parameter produces marginal frequencies for each variable (see Figure 24).

```
R> vtree(td, "Region Gender age fare", check.is.na = TRUE,
+       horiz = FALSE, showlegend = TRUE)
```

4. Case study: A study flow diagram

Study flow diagrams provide a visual representation of how participants (or study units) meet or do not meet a sequence of inclusion criteria. These diagrams provide critical information

to the reader of published study. Medical research in particular has embraced these data visualizations as part of recommended reporting guidelines. Randomized clinical trials use CONSORT diagrams to show the flow of participants through a single study (Schulz, Altman, and Moher 2010). Systematic reviews use PRISMA flow diagrams to depict study screening (Page *et al.* 2020; Stovold, Beecher, Foxlee, and Noel-Storr 2014). While presenting study flow diagrams is widely considered to be best practice, preparing these diagrams has traditionally been a slow, resource intensive, manual process, which has to be repeated when small changes are made to the data.

vtree uses an R data frame to make a data-driven study flow diagram. This automates the production of study flow diagrams. As more data arrives, data cleaning changes the existing data and the analysis plan is modified after initial assessment of the data (Huebner *et al.* 2016), the study flow diagram is easily kept up to date. Not only does this increase efficiency, it minimizes the risk of introducing human error.

Consider, for example, the Remdesivir trial of Spinner *et al.* (2020), in which 612 patients with confirmed severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) infection and moderate COVID-19 pneumonia were screened for inclusion. Although, in this case, the full data set is not publicly available, the variables required for the flow diagram can be reconstructed from Figure 1 of the published paper. The `build.data.frame` function built into the **vtree** package makes it easy to construct a data frame indicating which participants were screened, included (and of these, who was eligible, and who consented), the group participants were randomized to, and who started the intervention. (Additional details have been omitted for the sake of brevity.)

```
R> rem <- build.data.frame(
+   c("included", "elig", "consent", "randgrp", "started"),
+   list(0, 0, 1, 0, 0, 13),
+   list(0, 1, 0, 0, 0, 3),
+   list(1, 1, 1, 1, 1, 193),
+   list(1, 1, 1, 1, 0, 4),
+   list(1, 1, 1, 2, 1, 191),
+   list(1, 1, 1, 2, 0, 8),
+   list(1, 1, 1, 3, 1, 200))
```

Next, let us define node labels:

```
R> nodelabels <- list(
+   included = c("Randomized" = "1", "Excluded" = "0"),
+   randgrp = c(
+     "Randomized to receive 10 d of remdesivir" = "1",
+     "Randomized to receive 5 d of remdesivir" = "2",
+     "Randomized to continue standard care" = "3"),
+   started = c(
+     "Did not start remdesivir" = "0",
+     "Started remdesivir" = "1"))
```

Having set up these objects, the code to produce a CONSORT-style diagram is fairly straightforward. In particular, the `follow` parameter makes it easy to specify which branches of the tree should be retained. The following code produces Figure 25.

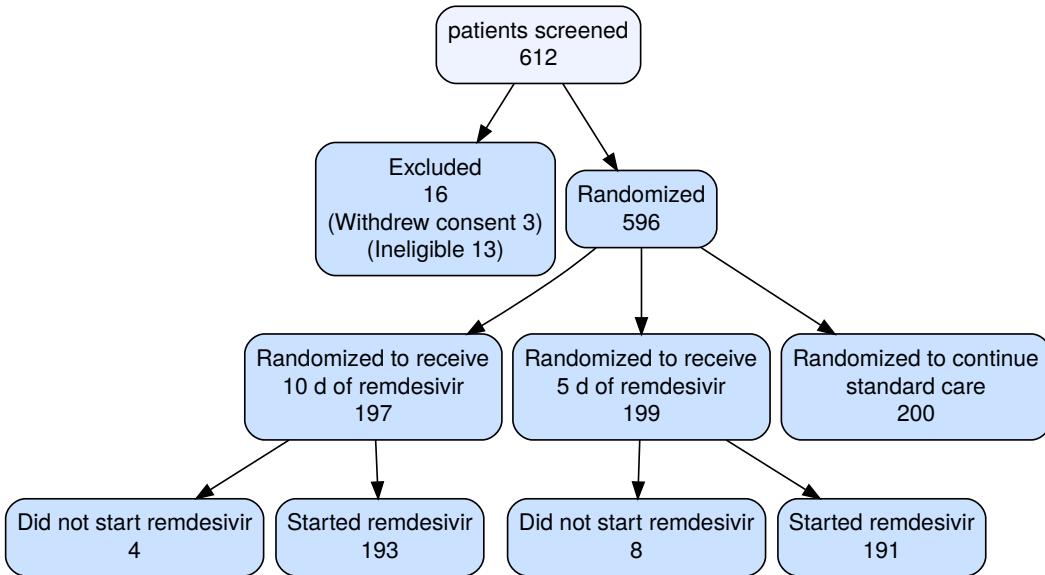


Figure 25: A variable tree providing a CONSORT-style diagram for the Remdesivir trial.

```

R> vtree(rem, "included randgrp started",
+   labelnode = nodelabels,
+   follow = list(included = "1", randgrp = c("1", "2")),
+   summary = c(
+     "consent=0 \n(Withdrew consent %sum%%var=included%%node=0%)",
+     "elig=0 \n(Ineligible %sum%%var=included%%node=0%)",
+     cdigits = 0, showvarnames = FALSE, title = "patients screened",
+     horiz = FALSE, fillcolor = "lightsteelblue1", showpct = FALSE)
  
```

5. Case Study: Ottawa Police Service Traffic Stops Data

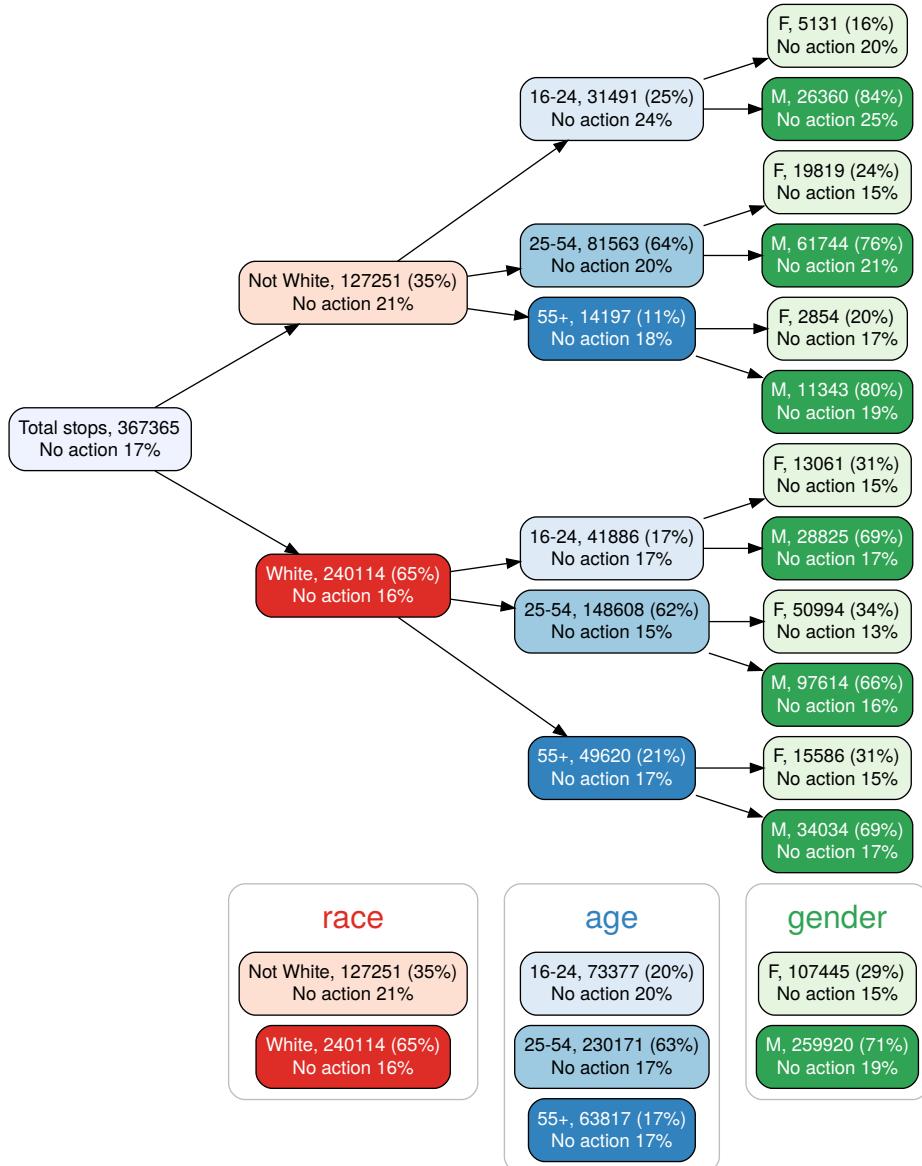
Following a 2005 racial profiling complaint to the Ontario Human Rights Commission, the Ottawa Police Service agreed to collect race data in traffic stops, known as the Traffic Stop Race Data Collection Project (TSRDCP). The TSRDCP required police officers to record *their perception* of the driver's race, for traffic stops from 2013 to 2023. A data set representing these traffic stops was made public² (licensed under the Open Government Licence – Ottawa Police Service³). After downloading the file from the Ottawa Police Service website, we read it into R (note that the file is also available in the replication files accompanying this paper):

```

R> data <- readr::read_csv("TSRDCP_TrafficStops_v1_164883526439542163.csv",
+   col_select = c(`How Cleared`, `Driver Gender`, `Driver Race`,
+   `Driver Age`),
+   col_types = cols())
  
```

²https://data.ottawapolice.ca/datasets/76a54b7f7ef44ab1b78dd96f3061c7cb_0/explore

³<https://data.ottawapolice.ca/pages/open-data-licence>



Important questions concern whether some racialized or ethnic groups are stopped at a rate disproportionate to overall makeup of the population. This would require external data, not presented here. See the report by researchers at York University, dated October 2016, for a comprehensive analysis of data from 2013 to 2016: https://www.researchgate.net/publication/344906617_FinalOPSOTTAWA_REPORT_-_2016EN. In the York University report, some records from the raw data were removed due to errors. Additionally, since some drivers were stopped more than once, only a single report per driver was included. It was not possible to replicate this last step because driver identifiers were not included in the publicly available data set. We remove missing data and anomalous values by:

```
R> z <- data |>
+   rename(race = `Driver Race`, age = `Driver Age`,
+         gender = `Driver Gender`, how_cleared = `How Cleared`) |>
+   filter(!is.na(race) & !is.na(age) & !is.na(gender)) |>
+   filter(gender != "D", age != "-15")
```

The data set includes a number of variables including age group and gender of the driver. One important variable is the outcome (renamed `how_cleared`) of the traffic stop: *charged*, *warning*, or *final (no action)*. This last outcome is of particular interest, because it means that the driver was neither charged nor given a warning, which may raise the question of whether the stop was actually necessary. Figure 26 shows the percentage of stops with this outcome (via the `summary` parameter) in each node of a tree for $\{\text{race}=\text{White}\} \rightarrow \text{age} \rightarrow \text{gender}$. Here race has been dichotomized as white or non-white, denoted $\{\text{race}=\text{White}\}$ (where the braces are for clarity in the arrow notation). Legend nodes are shown for each variable. Additionally the percentage of stops for which no action was taken is also shown in the legend nodes because `showlegendsum = TRUE` has been specified.

A number of interesting patterns emerge when looking at Figure 26, produced by the code below. The following drivers were more likely to receive neither a charge nor a warning: (1) male drivers, within all combinations of race and age; (2) younger drivers, within each race; and (3) non-white drivers.

```
R> vtree(z, "race=White age gender", splitwidth = Inf, sameline = TRUE,
+       summary = "how_cleared=Final \nNo action %pct%",
+       title = "Total stops", showlegend = TRUE, showlegendsum = TRUE)
```

6. Concluding remarks

Variable trees are an intuitive way to represent discrete multivariate data. The **vtree** package in R provides an implementation of variable trees along with a number of convenient extensions. There are a variety of other methods for displaying discrete multivariate data, and depending on the context, one of these methods may be preferable. However, the simple structure of variable trees provides not only ease of interpretation but also considerable generality. We have found that variable trees facilitate iterative data exploration when a statistician is working together with a domain expert.

A key characteristic of variable trees is that the order of variables is important. Sometimes the ordering of variables is natural (e.g., school board \rightarrow school \rightarrow teacher), in other cases it is dictated by the research question, and in still other cases the choice of ordering is up to the analyst. Depending on the situation, the variable-order dependence may be a strength or a weakness.

While **vtree** can be used to explore data, it can also be used to generate study flow diagrams. In recent years there has been growing concern about the “reproducibility crisis” in science (Baker 2016). In order to produce study flow diagrams using **vtree**, all of the variables and the corresponding set of inclusion/exclusion steps must be in a single data frame, which encourages a reproducible workflow. The design of **vtree** was influenced by the tidyverse philosophy (Wickham *et al.* 2019), with its emphasis on reproducible workflows, and **vtree**

works well with tidyverse tools. A key barrier to the wider adoption of study flow diagrams has been the difficulty required to produce them. **vtree** facilitates reproducible research by making it easy to produce accurate study flow diagrams directly from the study data.

To conclude, variable trees are an intuitive new data exploration tool for visualizing nested subsets. Applications of variable trees include revealing patterns in data, understanding missingness and producing study flow diagrams for reproducible research.

Acknowledgments

The **vtree** package builds on the **DiagrammeR** package (Iannone 2020), which in turn is based on the *Graphviz* graph visualization software (Gansner and North 2000). Sebastian Gatscha also contributed code to **vtree**. Development of **vtree** was partially supported by the Clinical Research Unit (CRU) at the Children’s Hospital of Eastern Ontario Research Institute. Members of the CRU contributed helpful suggestions and endless patience.

References

Altman DG (1994). “Problems in Dichotomizing Continuous Variables.” *American Journal of Epidemiology*, **139**(4), 442–442. [doi:10.1093/oxfordjournals.aje.a117020](https://doi.org/10.1093/oxfordjournals.aje.a117020).

Baker M (2016). “1,500 Scientists Lift the Lid on Reproducibility.” *Nature*, **533**(7604), 452–454. [doi:10.1038/533452a](https://doi.org/10.1038/533452a).

Bangdiwala SI, Shankar V (2013). “The Agreement Chart.” *BMC Medical Research Methodology*, **13**(1), 1–7. [doi:10.1186/1471-2288-13-97](https://doi.org/10.1186/1471-2288-13-97).

Chen H (2018). **VennDiagram**: Generate High-Resolution Venn and Euler Plots. [doi:10.32614/CRAN.package.venndiagram](https://doi.org/10.32614/CRAN.package.venndiagram). R package version 1.6.20.

Conway J, Lex A, Gehlenborg N (2017). “**UpSetR**: An R Package for the Visualization of Intersecting Sets and Their Properties.” *Bioinformatics (Oxford, England)*, **33**. [doi:10.1093/bioinformatics/btx364](https://doi.org/10.1093/bioinformatics/btx364).

Gansner ER, North SC (2000). “An Open Graph Visualization System and Its Applications to Software Engineering.” *Software: Practice and Experience*, **30**(11), 1203–1233. [doi:10.1002/1097-024X\(200009\)30:11<1203::AID-SPE338>3.0.CO;2-N](https://doi.org/10.1002/1097-024X(200009)30:11<1203::AID-SPE338>3.0.CO;2-N).

Hartigan J, Kleiner B (1981). “Mosaics for Contingency Tables.” *Computer Science and Statistics: Proceedings of the 13th Symposium on the Interface*. [doi:10.1007/978-1-4613-9464-8_37](https://doi.org/10.1007/978-1-4613-9464-8_37).

Hellerstein JM, Rattenbury T, Heer J, Kandel S, Carreras C (2017). *Principles of Data Wrangling: Practical Techniques for Data Preparation*. O’Reilly Media, Inc., Sebastopol.

Hornik K, Zeileis A, Meyer D (2006). “The Strucplot Framework: Visualizing Multi-Way Contingency Tables with **vcd**.” *Journal of Statistical Software*, **17**(3), 1–48. [doi:10.18637/jss.v017.i03](https://doi.org/10.18637/jss.v017.i03).

Huebner M, Vach W, le Cessie S (2016). “A Systematic Approach to Initial Data Analysis Is Good Research Practice.” *The Journal of Thoracic and Cardiovascular Surgery*, **151**(1), 25–27. [doi:10.1016/j.jtcvs.2015.09.085](https://doi.org/10.1016/j.jtcvs.2015.09.085).

Iannone R (2020). **DiagrammeR**: *Graph/Network Visualization*. [doi:10.32614/CRAN.package.diagrammer](https://doi.org/10.32614/CRAN.package.diagrammer). R package version 1.0.6.1.

Jeppson H, Hofmann H, Cook D, Wickham H (2020). **ggbmosaic**: *Mosaic Plots in the ggplot2 Framework*. [doi:10.32614/CRAN.package.ggbmosaic](https://doi.org/10.32614/CRAN.package.ggbmosaic). R package version 0.3.3.

Kassambara A (2020). **ggpubr**: *ggplot2 Based Publication Ready Plots*. [doi:10.32614/CRAN.package.ggpublisher](https://doi.org/10.32614/CRAN.package.ggpublisher). R package version 0.4.0.

Katal A, Wazid M, Goudar RH (2013). “Big Data: Issues, Challenges, Tools and Good Practices.” In *2013 Sixth International Conference on Contemporary Computing (IC3)*, pp. 404–409. [doi:10.1109/ic3.2013.6612229](https://doi.org/10.1109/ic3.2013.6612229).

Larsson J (2020). **eulerr**: *Area-Proportional Euler and Venn Diagrams with Ellipses*. [doi:10.32614/CRAN.package.eulerr](https://doi.org/10.32614/CRAN.package.eulerr). R package version 6.1.0.

Lex A, Gehlenborg N, Strobelt H, Vuillemot R, Pfister H (2014). “UpSet: Visualization of Intersecting Sets.” *IEEE Transactions on Visualization and Computer Graphics*, **20**, 1983–1992. [doi:10.1109/tvcg.2014.2346248](https://doi.org/10.1109/tvcg.2014.2346248).

Moon KW (2016). *Learn ggplot2 Using Shiny App*, chapter Balloon Plot, pp. 85–90. Springer-Verlag, Cham. [doi:10.1007/978-3-319-53019-2_10](https://doi.org/10.1007/978-3-319-53019-2_10).

Nelson S, Rafferty AM (2012). *Notes on Nightingale: The Influence and Legacy of a Nursing Icon*. Cornell University Press, Ithaca. [doi:10.7591/9780801460241](https://doi.org/10.7591/9780801460241).

Page MJ, McKenzie J, Bossuyt P, Boutron I, Hoffmann T, Mulrow C, Shamseer L, Tetzlaff J, Akl E, Brennan SE, Chou R, Glanville J, Grimshaw J, Hróbjartsson A, Lalu M, Li T, Loder E, Mayo-Wilson E, McDonald S, McGuinness LA, Stewart L, Thomas J, Tricco A, Welch V, Whiting P, Moher D (2020). “The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews.” [doi:10.31222/osf.io/v7gm2](https://doi.org/10.31222/osf.io/v7gm2).

Philipp M, Zeileis A, Strobl C (2016). “A Toolkit for Stability Assessment of Tree-Based Learners.” In A Colubi, A Blanco, C Gatu (eds.), *Proceedings of COMPSTAT 2016 – 22nd International Conference on Computational Statistics*, pp. 315–325. The International Statistical Institute/International Association for Statistical Computing. Preprint available at <http://EconPapers.RePEc.org/RePEc:inn:wpaper:2016-11>.

R Core Team (2025). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. [doi:10.32614/R.manuals](https://doi.org/10.32614/R.manuals). URL <https://www.R-project.org/>.

Schulz KF, Altman DG, Moher D (2010). “CONSORT 2010 Statement: Updated Guidelines for Reporting Parallel Group Randomised Trials.” *BMJ*, **340**. [doi:10.1136/bmj.c332](https://doi.org/10.1136/bmj.c332).

Spinner CD, Gottlieb RL, Criner GJ, López JRA, Cattelan AM, Viladomiu AS, Ogbuagu O, Malhotra P, Mullane KM, Castagna A, Chai LYA, Roestenberg M, Tsang OTY, Bernasconi

E, Le Turnier P, Chang SC, SenGupta D, Hyland RH, Osinusi AO, Cao H, Blair C, Wang H, Gaggar A, Brainard DM, McPhail MJ, Bhagani S, Ahn MY, Sanyal AJ, Huhn G, Marty FM (2020). “Effect of Remdesivir vs Standard Care on Clinical Status at 11 Days in Patients with Moderate COVID-19: A Randomized Clinical Trial.” *JAMA*, **324**(11), 1048–1057. [doi:10.1001/jama.2020.16349](https://doi.org/10.1001/jama.2020.16349).

Stovold E, Beecher D, Foxlee R, Noel-Storr A (2014). “Study Flow Diagrams in Cochrane Systematic Review Updates: An Adapted Prisma Flow Diagram.” *Systematic Reviews*, **3**(54). [doi:10.1186/2046-4053-3-54](https://doi.org/10.1186/2046-4053-3-54).

Urbanek S, Wichtrey T (2018). *ipplots: Interactive Graphics for R*. [doi:10.32614/CRAN.package.ipplots](https://doi.org/10.32614/CRAN.package.ipplots). R package version 1.1-7.1.

Waldrop MM (2016). “The Chips Are Down for Moore’s Law.” *Nature News*, **530**(7589), 144. [doi:10.1038/530144a](https://doi.org/10.1038/530144a).

Wickham H, Averick M, Bryan J, Chang W, D’Agostino McGowan L, François R, Grolemund G, Hayes A, Henry L, Hester J, Kuhn M, Pedersen TL, Miller E, Bache SM, Müller K, Ooms J, Robinson D, Seidel DP, Spinu V, Takahashi K, Vaughan D, Wilke C, Woo K, Yutani H (2019). “Welcome to the Tidyverse.” *Journal of Open Source Software*, **4**(43), 1686. [doi:10.21105/joss.01686](https://doi.org/10.21105/joss.01686).

Wickham H, François R, Henry L, Müller K, Vaughan D (2023). *dplyr: A Grammar of Data Manipulation*. [doi:10.32614/CRAN.package.dplyr](https://doi.org/10.32614/CRAN.package.dplyr). R package version 1.1.4.

Wilkinson L (2011). *venneuler: Venn and Euler Diagrams*. [doi:10.32614/CRAN.package.venneuler](https://doi.org/10.32614/CRAN.package.venneuler). R package version 1.1-0.

Wilkinson L (2012). “Exact and Approximate Area-Proportional Circular Venn and Euler Diagrams.” *IEEE Transactions on Visualization and Computer Graphics*, **18**(2), 321–331. [doi:10.1109/tvcg.2011.56](https://doi.org/10.1109/tvcg.2011.56).

Zeileis A, Meyer D, Hornik K (2007). “Residual-Based Shadings for Visualizing (Conditional) Independence.” *Journal of Computational and Graphical Statistics*, **16**(3), 507–525. [doi:10.1198/106186007x237856](https://doi.org/10.1198/106186007x237856).

Affiliation:

Nick Barrowman
CHEO Research Institute
401 Smyth Road
Ottawa, Ontario, K1H 8L1, Canada
and
Department of Pediatrics
Faculty of Medicine
University of Ottawa
Ottawa, Ontario, K1N 6N5, Canada
E-mail: nick.barrowman@gmail.com
URL: <https://nbarrowman.github.io/vtree>

Richard J. Webster
CHEO Research Institute
401 Smyth Road
Ottawa, Ontario, K1H 8L1, Canada