

# TreeCovery : Coordinated Dual Treemap Visualization for Exploring the Recovery Act \*

Miguel Rios-Berrios<sup>†</sup>

Puneet Sharma<sup>‡</sup>

Rachel Schwartz<sup>§</sup>

Tak Yeon Lee<sup>¶</sup>

Ben Shneiderman<sup>||</sup>

Human-Computer Interaction Laboratory &  
Department of Computer Science  
University of Maryland

## ABSTRACT

The American Recovery and Reinvestment Act dedicated \$787 billion to stimulate the US economy and mandated the release of the data describing the exact distribution of that money. The dataset is a large and complex one; one of its distinguishing features is its bi-hierarchical structure, arising from the distribution of money through agencies to specific projects and the natural aggregation of awards based on location. To offer a comprehensive overview of the data, a visualization must incorporate both these hierarchies. We present TreeCovery, a tool that accomplishes this through the use of two coordinated treemaps. The tool includes a number of innovative features, including coordinated zooming and filtering and a proportional highlighting technique across the two trees. TreeCovery was designed to facilitate data exploration, and initial user studies suggest that it will be helpful in insight generation. RATB (Recovery Accountability and Transparency Board) has tested TreeCovery and considering to include the concept into their visual analytics.

**Index Terms:** H5.2 [Information Interfaces and Presentation]: User Interfaces—Miscellaneous

## 1 INTRODUCTION

In February 2009, President Obama signed an economic stimulus package into law, dedicating \$787 billion to create jobs and boost the economy, with the provision that the distribution of the money would be completely transparent. In fulfillment of this requirement, the agencies in charge of distributing the money and all recipients issued periodic reports detailing how the money they controlled was spent. These publicly available reports comprise a large amount of data, containing information about the effectiveness of the stimulus package, the general trends of distribution, and potentially interesting outliers[4].

Some effort has already been expended toward producing visualizations of this data that could assist in revealing such details. The government commissioned a website, Recovery.gov, dedicated to this purpose, and several independent journalism outlets have produced their own applications, all offering a particular take on the data. Most of the existing visualizations consist primarily of either tabular or geographical displays. In addition, while many provide some interaction, most do not allow users full control over the data visualized.

---

\*The first four authors contributed equally to this work.

<sup>†</sup>mrios@cs.umd.edu

<sup>‡</sup>puneet@cs.umd.edu

<sup>§</sup>rachel13@cs.umd.edu

<sup>¶</sup>tylee@umd.edu

<sup>||</sup>ben@cs.umd.edu

While the data lends itself well to geographical layout, given that states and counties are convenient schemas for chunking the data, exclusive use of maps cannot adequately portray alternate views of the monetary distribution. Specifically, money was distributed through 28 agencies, who assigned it to projects at their discretion; funding was placed in the charge of the prime recipient, who in turn funded sub-recipients and/or vendors as necessary for the project. Agencies naturally funded projects nationwide, and recipients for each project were not necessarily all located in the same area. This view of the data - an *agency > project > recipient* hierarchy - cannot be adequately conveyed by a geographical substrate.

Our tool, TreeCovery, offers a way to explore data both geographically and according to the monetary outlays. TreeCovery accomplishes this goal through the use of two coordinated treemaps, one drawn with a geographic hierarchy and the other one with levels corresponding to the *agency > project > recipient* money flow. While the views presented by the two treemaps differ, the underlying data remains identical at all times. Filtering is coordinated across the views and a proportional highlighting technique is used for coordination.

In addition to the coordinated treemap design, we incorporated a few other features to improve exploration techniques. From news articles about Recovery Act we found many of them using demographic statistics such as *population* or *unemployment rates*. We thus included census data for each county, and made it possible to filter by demographic attributes. We also added the ability to save *snapshots* of the current state of the treemap for later comparison. Finally, we included support for emphasizing invalid data values.

TreeCovery has been tested by RATB and received positive feedback about the concept. Hopefully features of TreeCovery will be included into next version of RATB's visual analytic platform.

The rest of the paper is organized as follows. In Section 2, we discuss related work. Section 3 provides an explanation of our analysis process, including a detailed illustration of Spotfire<sup>1</sup>'s ability to support exploration of Recovery Act data. Section 4 explains TreeCovery in detail, while Section 5 offers some sample insights found by the tool. Section 6 contains our conclusion and future work.

## 2 RELATED WORK

Because the stimulus information is both newsworthy and publicly available, many visualizations of the data are already available. First and foremost, recovery.gov offers geographical maps displaying award locality (Figure 1). The maps can be zoomed in to state and zip code level and show dots each representing a project colored by its award type - *contract, grant and loan*. The site also offers some pie and bar chart summaries, as well as tabular data. While the basic information is thus available, interaction with the visualization and customization is limited.

Many other websites offer similar tools to those of recovery.gov.

---

<sup>1</sup><http://spotfire.tibco.com>

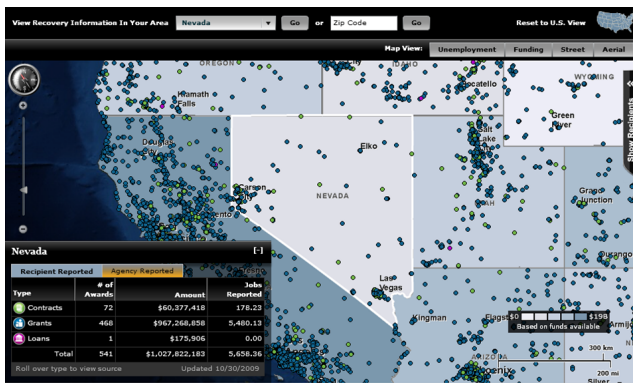


Figure 1: Recovery.gov offers geographical maps showing each award as a dot.

The Federal Procurement Data System<sup>2</sup>, which has the raw data available for download, offers a few selected slightly interactive visualizations as well; again, interactivity is limited and the data available through the visualizations is limited as well. The website ProPublica hosts a feature with Recovery act information, Eye on the Stimulus<sup>3</sup>. Most of the site is devoted to text articles, but tabular and geographical visualizations of spending progress are offered as well. Other sites following the geographical/tabular trend include the Wall Street Journal [5] and msnbc [6].

The existing visualizations of stimulus data, while informative, do not support exploratory analysis of the data. Because the data contains dual hierarchies, the geographical one and the monetary distribution, it can be most effectively portrayed using visualizations tailored toward this structure. Several methods of achieving this have been discussed in the literature. Polyarchy Visualization allows representation of intersecting hierarchies [15]. While useful, this approach does not provide well for visualization of multi-dimensional data, as is necessary for recovery.gov data. Multitrees describe the general structure of the data: nodes that are shared by multiple ancestor trees [12]. However, they are presented as a graph theoretic rather than visual concept and, as such, are not particularly helpful in building our tool.

Treemaps are a space-filling method of visualizing large data sets which have proved very useful for displaying hierarchical structure [16]. There exist many variations of treemaps specialized for different purposes and characteristics of data sets. Burch and Diehl introduce the “Trees in a Treemap” technique to represent trees with an associated taxonomy [7]. The taxonomy is represented as a treemap, and the related tree is drawn on top. In this representation, the two (or more) hierarchies are not represented in a symmetric fashion; attribute information about the hierarchy can be easily integrated, but the nodes of the trees show only their structural position. Therefore, this approach does not suit our data. Jern, et al., suggest using a treemap visualization in combination with a cartographic one [11]. They demonstrate the value of multiple displays of the same data, but their treemap hierarchy is regional. The stimulus visualization requires a method that represents data in multiple hierarchies.

Another treemap-based hierarchical representation is that of Wood *et al.* [17, 18, 19]. A specialized ordering is used to facilitate spatial and temporal locality so that the layout of nodes is more intuitive. Wood’s approach combines two hierarchies (temporal and spatial) into a single treemap. We felt it was important for users to be able to explore each hierarchy in isolation, as well as in conjunction with the other. TreeCovey therefore uses two treemaps

to represent the data using two different hierarchies which is an instance of multiple coordinated visualizations.

Multiple coordinated visualizations help users explore complex data. There are four common types of coordinations; Brushing and linking, Overview and detail view, Drill down, and Synchronized scrolling [13]. As supporting exploration of hierarchical data sets was the goal, treecovey focuses on the coordinated drill down that allows users to navigate down successive layers of a hierarchical database [9]. Among many cases of multiple coordinated visualizations, PairTrees [12] is an interesting case as it also utilizes treemaps to support exploratory data analysis on multiple hierarchical attributes. In PairTrees, two treemaps initially show overviews of the data set using two hierarchies based on aggregation. When an element is selected in one treemap, another treemap is automatically filtered by the selected element revealing meaningful relationship between the two hierarchies. Additional flexibility is gained in this special type of dynamic query by giving no fixed role of control and representation parts.

Using two coordinated treemaps raises the design question of how to use brushing that highlights selected elements on multiple coordinated visualizations. Brushing was first used in the PRIM9 system by Tukey et al [8] and now is applied in almost every interactive visualization environment [14]. While, in most design selected elements are present on every view, the two treemaps in TreeCovey have elements partially related to each other. For example, in existing applications such as Spotfire, given an agency selected in one view another view would highlight all the states awarded any amount of money from the agency. This, however, conflicts with the space-filling rule of treemap (size of an element is the sum of its children) as the size of highlighted area does not necessarily match with the amount of money given. To the best of our knowledge, TreeCovey is the first attempt that represents actual proportion of related sub-elements when an element is highlighted.

### 3 ANALYSIS

To design our tool, we first needed to determine the chief goals of stimulus data visualization. As recipient reports of Recovery Act had just been released, it was not easy to find end-users who had already done extensive work on the data. Thus, instead of using direct interviews or a survey, we decided to analyze headlines of relevant news articles in order to understand the process of journalists analyzing it. Further, we analyzed the data with Spotfire, one of the most versatile visual analytic applications. These exploratory tasks helped us understand the data better and develop the concept of our tool.

#### 3.1 News Insights

Journalists are the most prominent group investigating Recovery Act data. From their work we could infer what people want to know and how the data is analyzed. We searched in the Google News search engine<sup>4</sup> by the keyword ‘Recovery Act’ of news reported during the first two weeks of November 2009 - right after the recipient reports had been publicized on Recovery.gov.

Most prominently, many findings focused on a **State/County-wise comparison**. Although geographical region is not the main hierarchy of the Recovery Act plan, the most frequently asked question was something like: “How much money is given to our state/county?” For example, an article compares the amounts of awards given to two states - “Idaho Gets Four times More Stimulus Money in contracts Than Louisiana” [3] It is noteworthy that state/county-wise comparisons require aggregation of projects in each state or county.

The second insight was the usage of **census data**. In order to find states/counties in similar context or to validate fairness of funding

<sup>2</sup><https://www.fpd.gov>

<sup>3</sup><http://www.propublica.org/ion/stimulus>

<sup>4</sup><http://news.google.com/news/search?aq=f&pz=1&cf=all&ned=us&hl=en&q=Recovery+Act>

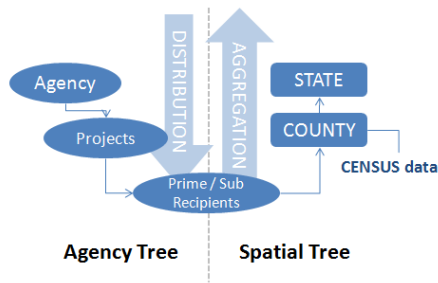


Figure 2: Recovery Act has two equally important hierarchies - agency and spatial trees.

from a specific agency, census data is quite useful. For example, an article[1] referred to *high-school graduation rates, infant mortality rates, unemployment rates, and juvenile justice incarceration* to pick the 5 worst cities for youth and compared numbers of jobs created by the Recovery Act in those places.

The third insight we found was **validity checking** to reveal unlikely numbers and non-existing categorical values. As each recipient report was submitted through an online form by the recipient, it is natural for the reports to have some errors due to simple mistakes. Non-existing congressional district codes are a typical case of simple mistakes criticized by CNN[2], however, other more complicated cases can exist too. For instance, *Number of jobs created* which is an important gauge of success, can be interpreted in many different ways. Usually, invalid values are either resolved or ignored in information visualization, although they have significant importance especially for a federal government website.

Summing up the findings above, we came up with the idea of bi-hierarchical data exploration in Figure 2. While the Recovery Act funding is distributed to recipients along the agency tree, the information of recipients is also aggregated by county and state in combination with census data.

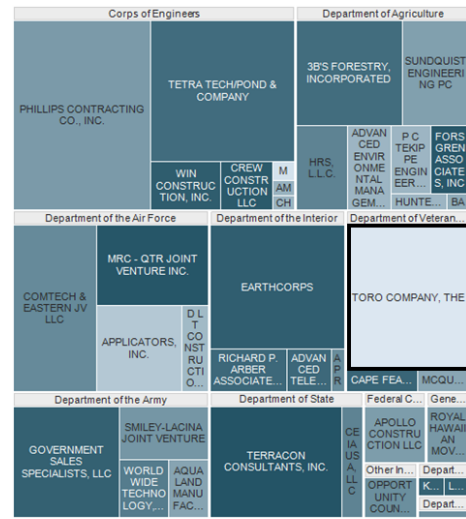
### 3.2 Spotfire Insight

Prior to building our own tool, we analyzed the recipient data<sup>5</sup> with an existing visual analytic toolkit. Spotfire was chosen because of its wide set of features for dealing with multivariate data. This analysis had three purposes. Firstly, analyzing the actual data provided us a deeper understanding. Secondly, we understood the capabilities and limitations of current visualization techniques. Finally, the exploration was helpful in suggesting opportunities for improvements on existing visualization techniques.

In addition to the recipient report, we also incorporated the State and County QuickFacts dataset<sup>6</sup> from the US Census Bureau. The census data consists of a wide variety of demographic profiles for each county, including Population, Infant deaths, Housing unit, Household income and Unemployment rate.

Paragraphs below are examples of the headlines we found by Spotfire analysis.

The most effective job creators are suspect. Knowing that there are projects with an improbably high number of jobs created, we drew a treemap visualization (Figure 3) showing which departments or states were related to those projects. Colors of elements represented *money per job*, which means how much money was spent for creating each job in a project. We had to filter projects by *money per job* in order that small important outliers be visible in the treemap. We found that treemap visualization was suitable for displaying job creation anomalies.



Agency > Project hierarchy treemap. SIZE by amount of money, COLOR by money per job

Figure 3: Treemaps can show the distribution of Recovery funds effectively with other attributes as well. TORO COMPANY is an exceptional recipient created a number of jobs with very small money.

Florida, the most senior state in the US, gets the most money from military sources. We tried to incorporate census data with recipient reports in this example. In Figure 4, the scatter plot(bottom right) shows the distribution of counties with percentage of population over 65 along the x axis. By selecting counties with a high percentage of senior citizens, all the related counties and projects were highlighted, which is called the *brushing and highlighting* technique [12]. Although very useful, the highlighting technique is often misleading because it always fills up the entire area of corresponding elements. Thus even when one state contains only a tiny single project related to the *brushed* selection, it would look as if the state has a lot of related projects.

## 4 TREECOVERY

We designed TreeCovey to be useful for investigative journalists and citizen watchdogs who have some domain knowledge and experience in data analysis. It streamlines the exploration process available through existing visualization techniques and adds more features for data analysis. This section elaborates on the development platform, data and UI components of TreeCovey.

### 4.1 Software Architecture

TreeCovey is implemented as a desktop application written in the Java programming language using the Prefuse visualization toolkit[10]. We decided to use the Treemap implementation in the Prefuse toolkit due to its strength in the visualization area as well as its data management. Although the implementation was not robust enough, it served as a perfect base in which we could build our application.

TreeCovey's implementation has three main components: the container of the user interface with the controls and settings; the treemaps, and a data processing tool. When the data file (in CSV format) is loaded into the application, the data processing tool converts the tables into trees, by using Prefuse's tree implementation. The trees are fed into the treemap layouts, which are initialized with several internal settings as colors, fonts and others, and then they are displayed in the container of the main user interface.

The filtering process in the treemaps is done using Prefuse predicates, which filter the underlying data table fast and efficiently by

<sup>5</sup><http://www.recovery.gov/FAQ/Pages/DownloadCenter.aspx>

<sup>6</sup>[http://quickfacts.census.gov/qfd/download\\_data.html](http://quickfacts.census.gov/qfd/download_data.html)

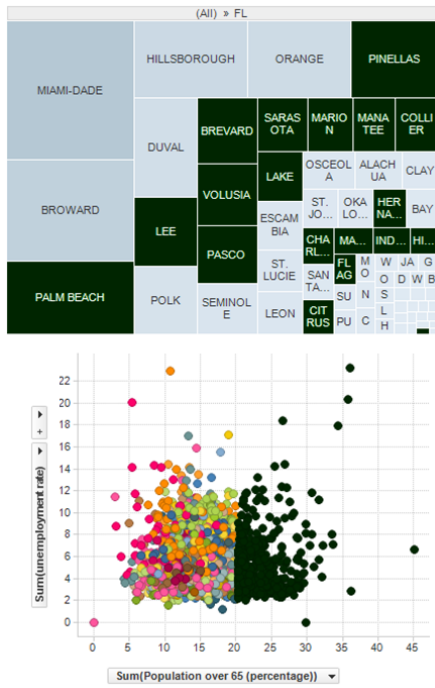


Figure 4: Brushing and highlighting technique is useful for connecting multiple visualizations.

binding directly the visual sliders to the data that is stored in the memory. Prefuse provides in-memory querying of data, what was essential in our case because we are using real-time dynamic filters using several criteria. This created a limitation because the amounts of data we are handling are huge, and using this approach all the data had to fit in memory. We developed a version that stored the data in a SQL-based database and performed informal benchmarks. As TreeCovary relies in intensive real-time filtering and is well-known that database access is slower than in-memory access, we decided to continue using the former and mitigated memory bottlenecks by increasing the Java memory heap size in our application.

Other components as the proportional highlighting visualization were developed from scratch. In this case the overlaid component is drawn according to the dimensions retrieved from a process that retrieves and resizes all the components of the original treemap accordingly to the computed proportion. Other main component of the application that was built from scratch is the synchronization between treemaps which was achieved by setting listeners that respond to events in any of the treemaps. They are monitoring each change, and when appropriately, they synchronize and show back the filtered or zoomed visualizations.

## 4.2 Data Pre-processing

TreeCovary relies on the recipient report data from [recovery.gov](http://www.recovery.gov)<sup>7</sup>. Information is available for each recipient of any stimulus dollars, including those attributes that are necessary to build the hierarchies (department, project, state, and county) as well as attributes (*number of jobs created*, *award amount*) that are used for filtering, size-by, and color-by, and other attributes that can be viewed in the details-on-demand panel. Our dataset includes a few interesting values for each record generated during the preprocess. For instance, we included *money per job* for each recipient, as that proportion may be more germane to the questions of stimulus effectiveness than either

<sup>7</sup><http://www.recovery.gov/FAQ/Pages/DownloadCenter.aspx>

job creation or award amount alone. Besides the standard recipient data, TreeCovary also integrates census data fields such as *population*, *education level*, and *unemployment rate*. We found that many of the headlines discussing recovery data also used census information so we felt it would be very useful to integrate that information directly into the tool.

In addition to this standard preprocess, three specific types of invalid values are searched and marked: invalid zip codes, invalid congressional districts and projects that had no primary recipient, only sub-recipients. Invalid zip codes were simply those that do not exist in the standard zip codes table (253 recipients found). Invalid congressional districts were determined on a state-by-state basis, so that if a district did not exist in that recipient's state, it was flagged as invalid (38 project found). Finally, there were 561 projects with no prime recipient reported at all. All these errors were flagged and integrated into the TreeCovary display to be presented to RATB.

## 4.3 Visualization

The chief component of the tool is, of course, the **dual treemap display** (Figure 5). The left treemap displays the agency (*department > project*) hierarchy, while the right one is spatial (*state > county*). Although the two share identical underlying data (the recipients) as their leaves, this level is never visible on the treemap. The shared leaf level data does, however, make it possible to coordinate the displays. Essentially, filtering occurs simultaneously (Figure 6); zooming in on one treemap, which amounts to filtering on the recipient leaves of the zoomed node, causes the other treemap to be filtered on those same recipients. Although Spotfire allows users to both zoom and filter, the two actions are independent, so that side by side treemaps will not stay coordinated automatically. In TreeCovary, users can zoom in and out on both treemaps in any arbitrary order, and the recipient leaves included in the layout will remain coordinated.

In addition to the filtering capabilities provided through zooming, TreeCovary also allows **customization of the treemap displays** using a set of controls. Double-sided sliders allow filtering on various attributes of the recipients and the census information for their areas. One innovative feature we added was the ability to filter out invalid values found in the preprocessing stage. Each treemap can be sized and colored on a chosen attribute as well.

TreeCovary's main innovation lies in its **proportional highlighting capability**. This feature completes the coordination of the two treemaps. Although Spotfire highlights child nodes in a treemap when the parent is selected in another one, our highlighting technique is much more finely tuned. When a node is clicked in one treemap, a highlighting square is placed inside all nodes in the other treemap that share any child recipients with the selected node. The square's size is proportional to the ratio of the shared children's areas to the total area of the node. For example, if the selected node has a child with an area of 10, and its parent in the other tree has an area of 100, then the highlighted square will take up 1/10 of the parent in the other tree's area. In this way, the exact distribution of the size by attribute for the selected attribute across the other hierarchy becomes easily apparent. Using the default proportionality constant of 1, the total highlighted areas will, in fact, equal the area of the selected node. Users can change the proportionality constant for sizing the squares, in case they want to emphasize the highlighted areas. They can also control the opacity of the squares, so that the underlying labels will be visible. This highly flexible highlighting capability displays extensive information about attribute distribution across hierarchies in a powerful and intuitive manner.

Besides the main treemap functionality, TreeCovary provides additional features for data exploration. Firstly, details of the recipients that comprise each node are available when selected. A Details-on-Demand table displays the data for each recipient, including values for all attribute fields in our data set. Each row of



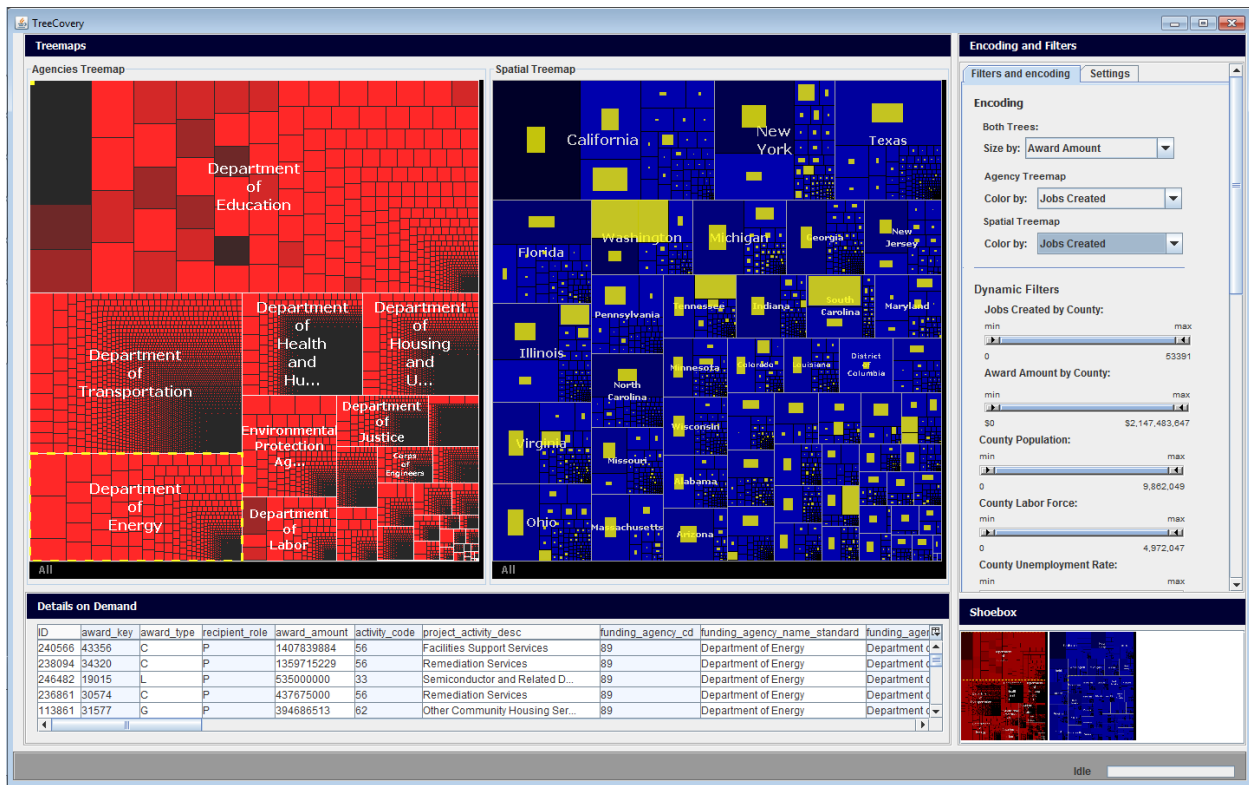


Figure 5: TreeCovary User Interface with the Recovery Act data. The pair of treemaps is at the upper left, the control panels for Filters and Settings are at the upper right, the tabular data is at the lower left, and the shoebox with user selected screenshots is at the lower right.

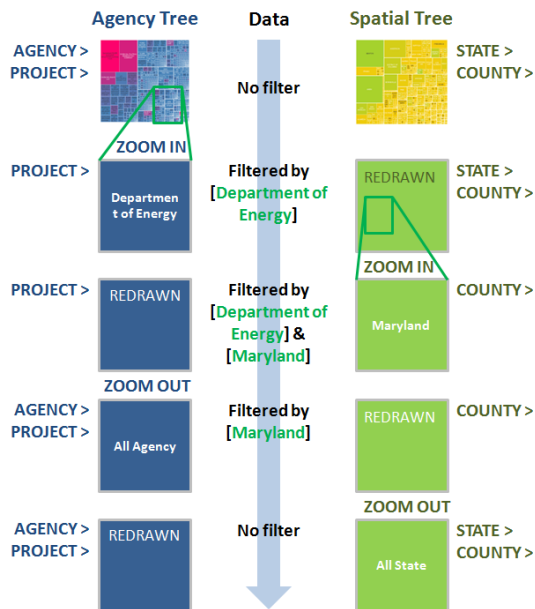


Figure 6: Two treemaps share the identical data whose filter is controlled by zooming activity on either treemap.

the table can be double-clicked to bring up the recovery.gov page on that project. TreeCovary also allows users to save the current view of a treemap for later viewing. The image of a single treemap is saved, along with the current filter, zoom, size-by, and color-by settings. All saved images are shown as thumbnails in the shoebox area. Users can select and view any number of the saved images side by side in a separate window.

## 5 EVALUATION

### 5.1 Insights

To demonstrate the utility of TreeCovary, we had a short exploratory session. Below, we present three insights that illustrate the use of TreeCovary in data analysis.

The major source of awards assigned to California (CA) is a few biggest projects from the Department of Education (Figure 7). TreeCovary was run with the contract, grant and loan data. Both treemaps were sized by the amount of money, and colored by the number of jobs created. California received the highest amount of awards, which was easily visible in the spatial treemap (right). By selecting California, it quickly highlights the related awards in the agency treemap (left). The distribution of high-lighted awards illustrates whether California received a few bigger awards or many smaller awards. From the department of education, California received the two biggest awards in a dark color that means very few jobs have been created by those projects. It is worthwhile to look into more details of them. The coordinated highlighting feature enables users to quickly examine the overview and motivates further exploration.

District of Columbia (DC), the federal hub, received the highest amount of contract money from the General Service Administration (Figure 8). To illustrate TreeCovary's zooming

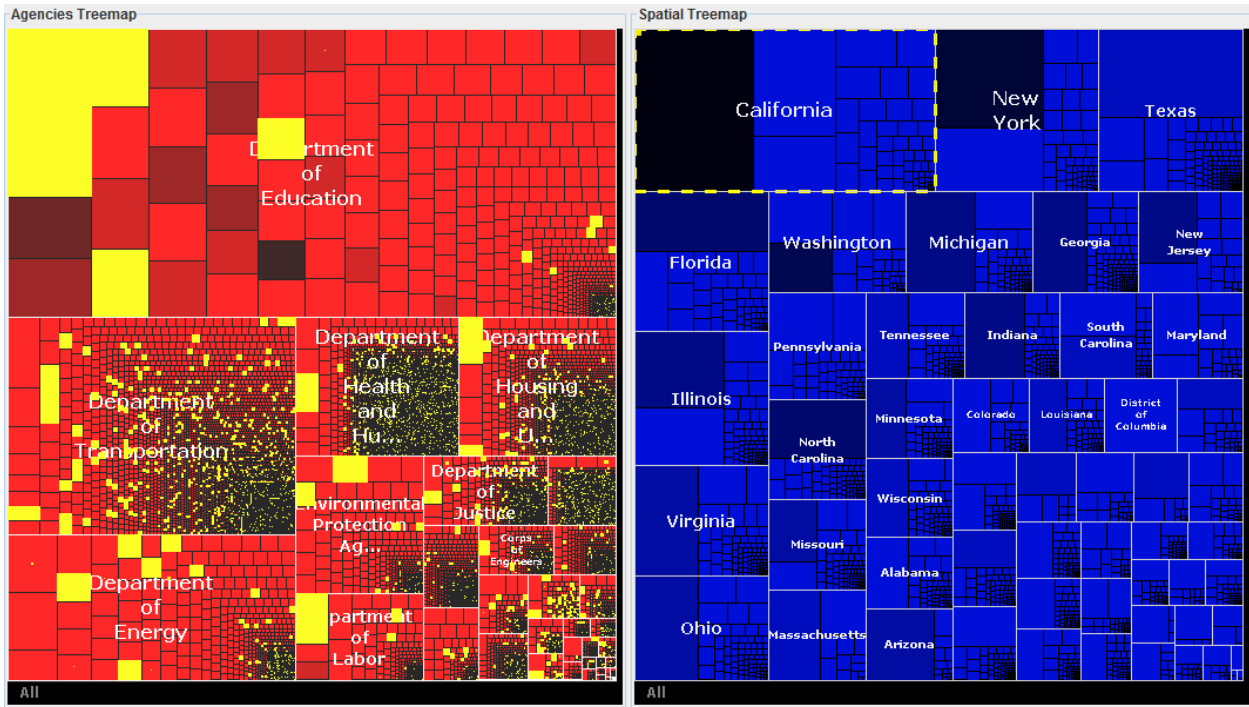


Figure 7: The major source of awards assigned to California(highlighted) is a few biggest projects from the Department of Education. While the size of each treemap represents the amount of money of the award, the color shows the number of jobs created.

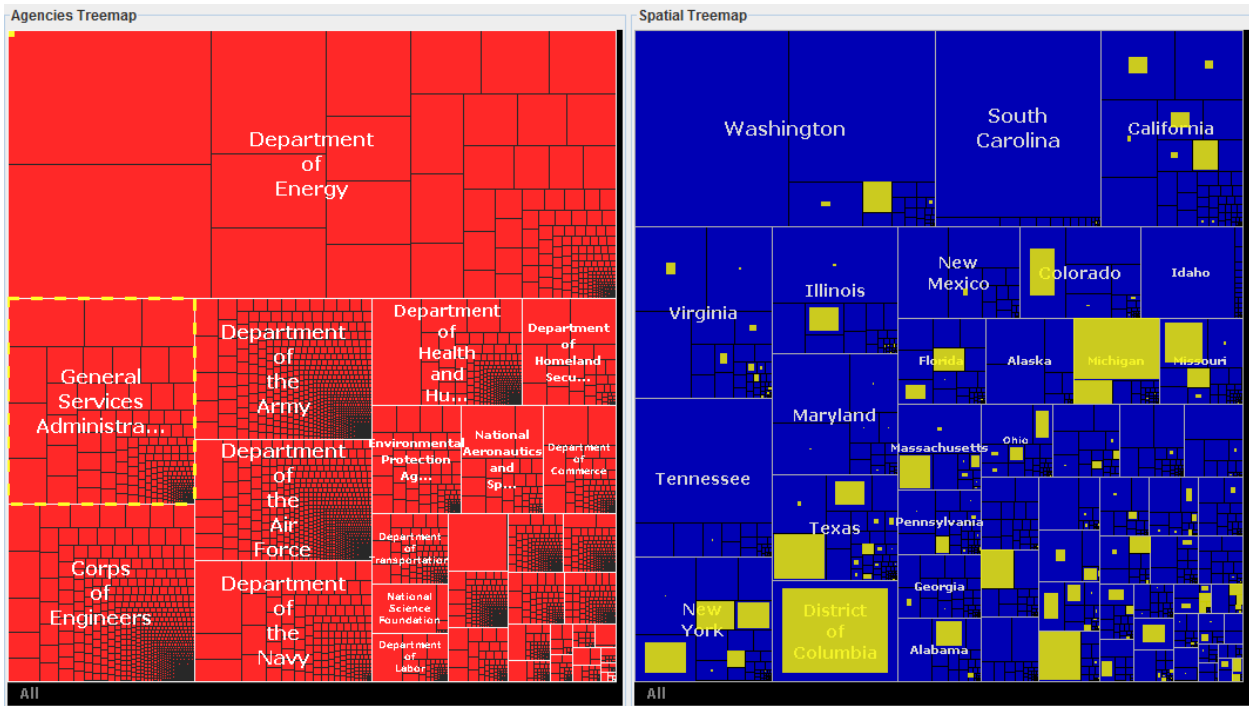


Figure 8: The District of Columbia(highlighted) received the highest contract amount from the General Service Administration Department. The size of each treemap represents the amount of money of the award.

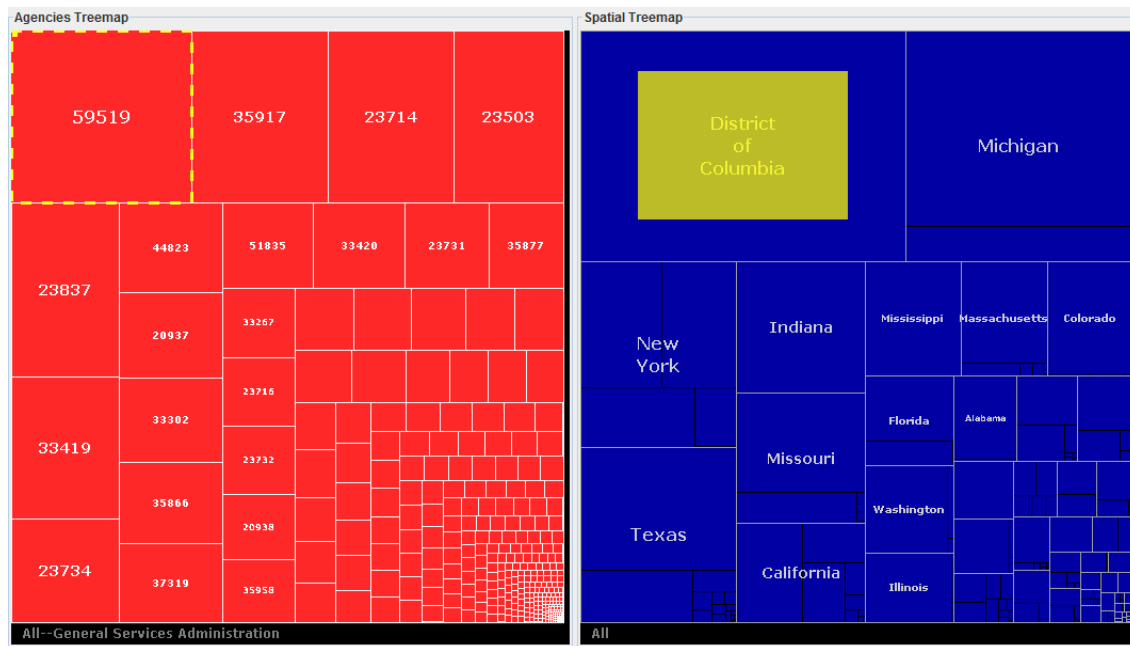


Figure 9: The District of Columbia received the biggest projects from the GSA department. The size of each treemap represents the amount of money of the award.

feature, both treemaps were sized by the award amount and the General Service Administration (GSA) node was selected in the agency treemap. Various counties were highlighted in the spatial Treemap, and it was clear that DC had received the highest amount of contract money. We zoomed in on the GSA node (Figure 9) and found that one of the biggest projects of GSA was allocated to DC and all of its recipients (Prime and Sub-Prime) were in DC.

Georgia is creating more jobs while getting less contract award money (Figure 10). To see the utility of the shoebox in analysis, the spatial treemap was sized by award amount and a snapshot was taken. Another snapshot was taken after sizing the spatial treemap by jobs created. Both images were selected in the shoebox and opened in the comparison window. In the figure, the left treemap shows the money distribution and the right treemap shows *jobs creation*. Usually, states getting more money create more jobs but Georgia (GA) stood out as an outlier, with more jobs created and a comparatively small contract award amount.

## 5.2 Usability Evaluation

To find unexpected usability problems and insights for revision, we conducted a usability evaluation with four graduate students in Computer Science department at University of Maryland. Each participant spent 10 minutes to learn the bi-hierarchical structure of Recovery Act data and the usage of TreeCover. They then explored the data freely for 20 minutes, looking for interesting patterns and outliers. They were also asked to use think-aloud protocol. At the end of the evaluation, they gave us general remarks of the tool. Below we present a few usability issues and remarks.

The coordinated dual treemap interface proved to be usable. All the participants understood the dual treemap representation in a few seconds with brief explanation. They found the proportional highlighting very helpful to overview the hierarchical distribution on both axis. It also frequently raised further questions such as “Why does maryland receive a lot from the Department of Health and Human Services? Is it because of the National Institutes of Health?”. By the way, it was interesting to see that most

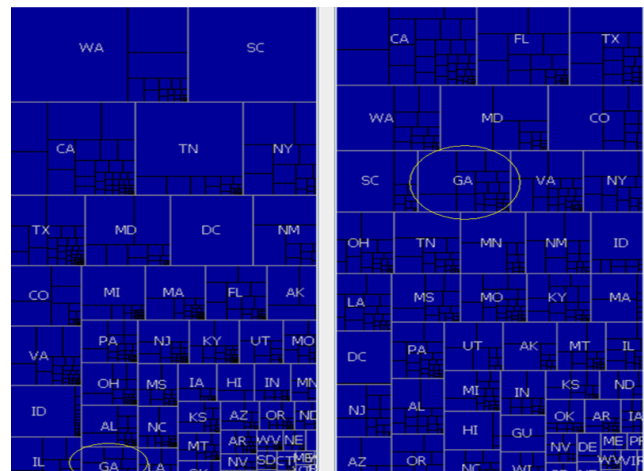


Figure 10: Comparison window showing two spatial treemaps, one sized by the award amount (left) and another sized by the jobs created (right)

participants tried to control the left treemap only for a while. This is probably due to the familiarity the users have with layouts that have controls in the left and representations in the right side of the screen.

Some users get lost when the entire treemap is redrawn. During the evaluation participants often tried to apply dynamic filters on highlighted selections, however, as the current version of TreeCover does not keep the selection when the entire treemap is redrawn, some participants got lost. Although we told them that the *Shoobox* can be used for comparing treemaps under multiple settings, they seem to need visual reference that relates the previous and the current status of the treemap.

Users need visual reference of the proportion of the current views to the entire data. Participants often wanted to know how big the size of the current treemap is within the entire data. One quick remedy can be a thumbnail window of the entire treemaps with a smaller region representing the portion of current view.

### 5.3 Expert Evaluation

TreeCovary was presented to RATB and the recovery.gov development team at Synteractive<sup>8</sup>. Their feedback was very positive especially about the proportional highlight and coordinated exploration features. Both groups received the source code of TreeCovary, and are now evaluating the possibility of two practical applications; First, coordinated multiple treemaps can be incorporated into recovery.gov website by Synteractive as an exploration tool for citizens. Second, RATB can use TreeCovary application for their internal visual analytic tool.

### 6 FUTURE WORK

While TreeCovary provides some innovative features and encompasses many exploration aids, it can, of course, be greatly improved. As observed during usability evaluation, we can make exploration easier by adding visual references of the size of current between previous and current status. The shoebox feature could potentially allow more extensive comparison among saved treemaps if the saved views were more interactive. The future version of TreeCovary will allow entire treemaps to be saved and loaded for viewing, rather than just a screenshot. Future versions will also include support for data manipulation, including for user-defined columns, in the manner of Spotfire. This will allow greater flexibility in the way users build treemaps. We would also like to incorporate an advanced color scheme, where the color gradient follows the distribution of the data, rather than staying linear. This will allow close but not identical values to be easily differentiable in color. These extra features will enhance the TreeCovary exploration experience, and hopefully lead to more insight generation.

Even though our application of the coordinated dual treemaps was the visualization of the money spent in Recovery Act, the concept can be generalized to any multi-hierarchical data. One of such examples is the representation of the money spent by the United States Agencies in different industry. The hierarchies in this case could be agencies and projects in one treemap, while the category of industry (e.g. manufacturing, technology, etc.) and the companies that work in each category will be represented in another treemap.

### 7 CONCLUSION

The American Recovery and Reinvestment Act provided for a substantial sum of money, \$787 billion, to be distributed with the goal of economic stimulus. Tracking that distribution involves a large, multi-attribute set that can be organized as a dual hierarchy of money flow and geographical allocation. Many visualizations of the stimulus data have already been developed, but none of them adequately portray this dual hierarchy or offer flexible exploration capabilities. Our tool, TreeCovary uses coordinated treemaps to accomplish exactly that task. We use coordinated zooming and filtering and finely tuned highlighting to streamline exploration across the two hierarchies. The tool incorporates a number of other features to aid in customization and flexibility of the display. Insights that would be difficult or impossible to see with previously available tools become readily apparent when TreeCovary is used to visualize the data.

### ACKNOWLEDGEMENTS

Special thanks to Synteractive and Recovery Accountability and Transparency Board for their continuous support. We are grateful to Angela Noh and researchers in Human-Computer Interaction

Laboratory at University of Maryland for the invaluable insights and feedback.

### REFERENCES

- [1] The 5 worst cities for urban youth - ABC news. Website, 2009. <http://abcnews.go.com/Politics/Politics/worst-cities-urban-youth/story?id=9083935>.
- [2] Feds: False congressional districts removed from stimulus site - CNN.com. Website, 2009. <http://www.cnn.com/2009/POLITICS/11/18/stimulus.district.errors/index.html>.
- [3] Idaho gets four times more stimulus money in contracts than louisiana - NAM. Website, 2009. [http://news.newamericamedia.org/news/view\\_article.html?article\\_id=e538921e01538f6f6fcc5a2ecfcfba9c](http://news.newamericamedia.org/news/view_article.html?article_id=e538921e01538f6f6fcc5a2ecfcfba9c).
- [4] Recovery.gov. Website, 2009. [http://www.recovery.gov/About/Pages/The\\_Act.aspx](http://www.recovery.gov/About/Pages/The_Act.aspx).
- [5] Stimulus spending, breakdown by states - The Wall Street Journal Online. Website, 2009. <http://online.wsj.com/public/resources/documents/info-STIMULUS0903.html>.
- [6] The stimulus tracker. Website, 2009. <http://www.msnbc.msn.com/id/33498869/#/all/all/us/all/>.
- [7] M. Burch and S. Diehl. Trees in a treemap: visualizing multiple hierarchies. In *Proceedings of SPIE*, volume 6060, pages 224–235, 2006.
- [8] M. A. Fisher, J. H. Friedman, and J. W. Tukey. *PRIM9: An Interactive Multidimensional Data Display and Analysis System*, pages 91–109. Wadsworth Brooks/Cole, 1988.
- [9] A. Fredrikson, C. North, C. Plaisant, and B. Shneiderman. Temporal, geographical and categorical aggregations viewed through coordinated displays: a case study with highway incident data. In *Proceedings of the 1999 workshop on new paradigms in information visualization and manipulation in conjunction with the eighth ACM international conference on Information and knowledge management*, pages 26–34, 1999.
- [10] L. J. A. HEER J., CARD S. K. prefuse: a toolkit for interactive information visualization. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, 2005.
- [11] M. Jern, J. Rogstadius, and T. Astrom. Treemaps and choropleth maps applied to regional hierarchical statistical data. In *2009 13th International Conference Information Visualisation*, pages 403–410, Barcelona, Spain, 2009.
- [12] B. Kules, B. Shneiderman, and C. Plaisant. Data exploration with paired hierarchical visualizations: initial designs of PairTrees. In *Proceedings of the 2003 annual national conference on Digital government research*, page 6, 2003.
- [13] C. North and B. Shneiderman. Snap-together visualization: can users construct and operate coordinated visualizations? *International Journal of Human Computer Studies*, 53(5):715–740, 2000.
- [14] J. C. Roberts and M. Wright. Towards ubiquitous brushing for information visualization. *IEEE Proceedings of the Information Visualization (IV06)*, 2006.
- [15] G. Robertson, K. Cameron, M. Czerwinski, and D. Robbins. Polychrome visualization: visualizing multiple intersecting hierarchies. In *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*, pages 423–430, 2002.
- [16] B. Shneiderman. Tree visualization with treemaps: A 2-d space-filling approach. In *ACM Transactions on Graphics*, 11(1), pages 92–99.
- [17] A. Slingsby, J. Dykes, J. Wood, and A. Crooks. The role of layout and order in treemaps for showing spatial and temporal variation in house prices. In *Abstract for the GeoViz Hamburg Workshop, Hamburg, Germany*, 2009.
- [18] J. Wood and J. Dykes. From slice and dice to hierarchical cartograms: Spatial referencing of treemaps. In *GISRUK 2008 the 15th national GIS research conference for the UK, Manchester*.
- [19] J. Wood, A. Slingsby, and J. Dykes. Using treemaps for variable selection in Spatio-Temporal visualization. *Information Visualization*, 7(3), 2008.

<sup>8</sup><http://www.synteractive.com/>