

Visualization of Functional-semantic Information: Integrating the Clause and the Discourse

Hengbin Yan

Faculty of English Language and Culture, Guangdong University of Foreign Studies

Baiyun Avenue North 2#, Baiyun District, 510420, Guangzhou, China

Tel: 86-20-36209276 E-mail: yhb@gdufs.edu.cn

Yinghui Li (Corresponding author)

School of English and Education/ Center for Linguistics and Applied Linguistics, Guangdong
University of Foreign Studies

Baiyun Avenue North 2#, Baiyun District, 510420, Guangzhou, China

Tel: 86-20-36207201 E-mail: liyinghui@gdufs.edu.cn

Received: November 5, 2018 Accepted: November 19, 2018 Published: November 20, 2018

doi:10.5296/jsel.v6i1.13935 URL: <https://doi.org/10.5296/jsel.v6i1.13935>

Abstract

In this study, we present our experimental design of an interactive interface that allows users to answer linguistically sophisticated queries utilizing functional-semantic information. Building on previous visualizations of linguistic patterning and discourse structure, the proposed visualization interface presents a unified interface for interrogating the functional-semantic structure of arbitrary texts at different levels of details. To evaluate the effectiveness of the interface, we performed a comparative analysis between visualizations of manual gold-standard annotation and those automatically generated by connecting the interface to existing automatic systems, which revealed remarkable visual correspondence between the two when dealing with small to medium texts. A small-scale case study was then conducted which demonstrated the potential of the resulting tool for effective discovery of interesting patterning in large political texts.

Keywords: Linguistic Visualization, Functional Linguistics, Functional-semantic structure

1. Introduction

The study of language function from the functionalist perspective analyzes language use in context and focuses on both the content and structure of language and how they function to make meaning in different contexts. This contextual use of language is interwoven with the semantic, pragmatic, and syntactic aspects of language. Researchers have come up with different frameworks to explore such linguistic functions at two complementary levels: the discourse level, where we consider the general architecture of text and how functionally defined relations among discourse units function in the text as a whole; and the clause level, where we study the functional-semantic roles of and relations between functional constituents within the clause.

At the discourse level, one of the most influential frameworks is Rhetorical Structure Theory (RST) (Mann & Thompson, 1988), which has a number of important applications in computational tasks such as question answering, text summarization and dialogue generation (Hernault et al., 2010).

At the clause level, the mainstream approach to functional analysis is Systemic Functional Linguistics. SFL provides an ideal handle to exploring language as intentional acts of meaning, and gives a broader vision and deeper insights into the functions and meaning at the level of the discourse. The features provided by SFL have proven useful in computational tasks such as natural language generation (Teich, 1999), text classification (Whitelaw & Argamon, 2004) and sentiment analysis (Whitelaw et al., 2005). Recent advances in computational analysis with SFL has made it possible to conduct fine-grained analysis of functional-semantic information at the clause level.

Current tools for automatic functional-semantic analysis (Feng & Hirst, 2014; Yan, 2014) allows analysts to plug in arbitrary pieces of well-written texts and have the discourse- and clause-level functional-semantic structures within the texts automatically mined and stored in a generic format for subsequent retrieval. However, ordinary users (those without programming expertise) still have no means of making use of the structured data directly. The scale and complexity of the output data from the tools make it difficult to explore the intricate structure and potential patterning in the functional-semantic information. A way still needs to be provided to make the results of the automatic process easily accessible to analysts. To address this problem, visualization is a powerful tool for processing large-scale, complex data flow. However, existing visualization techniques for linguistic use are inadequate for visualizing functional-semantic structure due to their lack of linguistic sophistication and proper abstraction.

In this paper, we describe our design of a new visualization interface tailored for functional-semantic structure to help analysts quickly and accurately analyze and interpret vast quantities of retrieved data to obtain insights that would have otherwise been difficult to obtain. Our aim is to design and develop an intuitive, unified interface for visualizing functional-semantic information at both the discourse and clause levels.

2. Literature Review

2.1 Linguistic Visualization

Visualization is “a cognitive activity, facilitated by graphical external representations from which people construct internal mental representation of the world” (Mazza, 2008).

Although significant advances have been made in the field of information visualization, designing and developing visual interfaces are still difficult tasks. There remains a number of important unsolved issues such as usability, scalability and intrinsic quality measures (Keim et al., 2008; Chen, 2010) in visualization design. In addition, despite the development of increasingly sophisticated linguistic theoretical frameworks and computational models for processing them, efforts to incorporate such development into interactive linguistic analysis are still largely lacking.

Visualization has been applied in a number of linguistic applications (C. M. Collins, 2010). Linguistic information visualization involves the use of computer-based visual representation of abstract linguistic data to amplify human cognition. Visualization of complex linguistic structures can assist linguists and language learners in the comprehension and analysis of complex, ambiguous or large-scale text data flow, providing an important cognitive aid that guides researchers in illustration, organization, manipulation and analysis of linguistic data (Rohrdantz et al., 2010).

The rising importance of text visualization has given rise to newly established fields of visual analysis and digital humanities, where large-scale text data are processed using NLP techniques and visualized for analysts to gain better understanding of large bodies of text (Rockwell et al., 2010). In this area, numerous visualization techniques have emerged in recent years, for instance, Relational Links (C. Collins & Carpendale, 2007), Dependency Graphs (Nilsson & Nivre, 2008), Bubble Sets (C. Collins et al., 2009), Tag Clouds (Koh et al., 2010), Density Maps (Cao et al., 2010), Semantic Mapping (Leydesdorff & Welbers, 2011), and KWIC (Culy & Lyding, 2011). Although not developed specifically for functional-semantic analysis, they provide useful insights and inspirations for designing the visualization interface of the proposed platform.

Substantial work has been done (Heer et al., 2007; Isenberg et al., 2011) to explore the effectiveness of applying web-based, interactive techniques in visualization. Most notably, IBM’s ManyEyes (Viegas et al., 2007) is a successful pioneer in web-based interactive visualization. It allows web users to upload textual data or structured values for visualization in several interactive formats. Its success as a visualization platform can be attributed to its pioneering use of web-based interfaces and its plug-in mechanism which makes it accessible to a wider audience. However, its lack of support for sophisticated linguistic and computational modeling renders it unsuitable for our needs for functional-semantic analysis. The visualization formats allowed are limited and are usually for raw texts or simple structured data, with no support for multilingual texts.

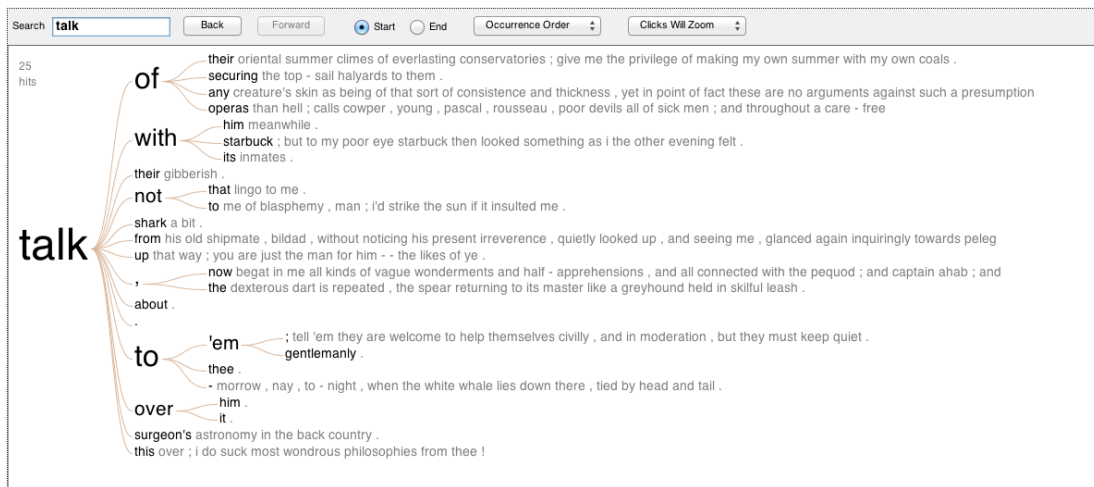


Figure 1. ManyEyes' Word Tree visualization for collocations of the verb talk.

2.2 Discourse Visualization

Discourse Visualization based on the RST theory can be naturally visualized as a hierarchical tree in accordance to RST's assumption of the discourse structure as a tree. Such structured provides a clearly organized view for detailed examination at different levels of details.

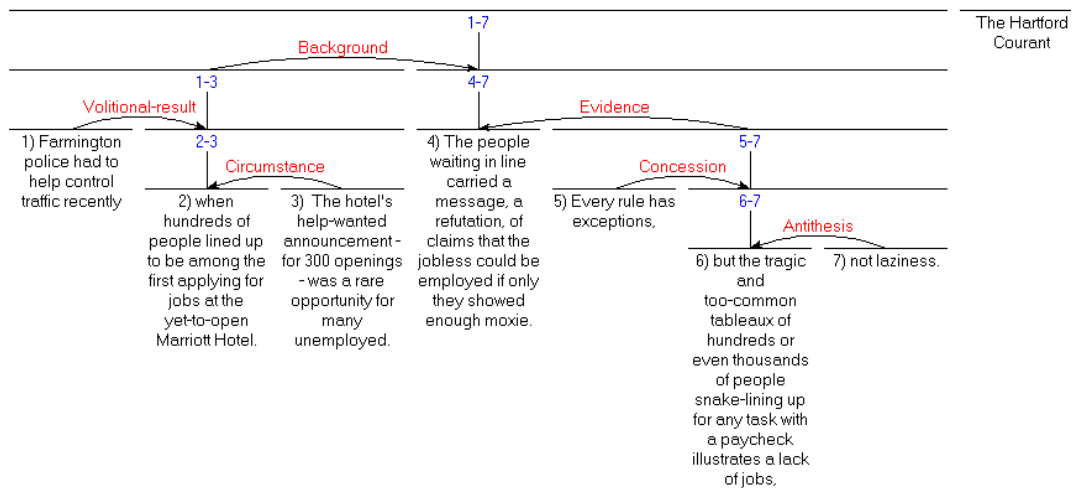


Figure 2. RST annotation view in the RSTTool.

The RST community has traditionally used the visual representation as provided in the RST tool (M. O'Donnell, 2008; Michael O'Donnell, 2000) to demonstrate small-scale RST structures. Such a visual representation, although effective for annotation purposes, is

unsuited for larger-scale interactive and exploratory analysis. First, it does not take advantage of the pre-attentive properties of effective visualization. Plain text directly embedded in the visual representation is not “pre-attentive” since people still need to consciously read the words to understand the encoded meaning. For example, analysts need to read the name of each relation to understand what it means. Secondly, apart from using lines to represent the pointed relations between textual spans, little abstraction is provided to represent the textual data. One important use of visualization is to reduce cognitive load of the user. The RST tool’s representation presents all information in one setting, making it difficult to highlight the more important parts. Users are confronted with an overwhelming amount of information they are expected to make sense of. The overloading of information also makes it difficult to provide a view of the overall structure of the text, an aim that the RST theory is specifically designed for. When we zoom out to fit the overall RST structure to a screen, the relations/textual data immediately becomes blurred and impossible to interpret. Without a proper overview of the structure, comparison between multiple texts side by side would be difficult. Finally, the visualization is only limited to analysis using the RST framework and thus unable to visualize functional-semantic structure within the clause. A better solution is needed for the visual representation of the functional-semantic structure, not only at the discourse level, but also the clausal level. In the next section, we describe our design of an innovative interface to address the issues discussed.

A more recent development of visualization of RST structures is Zhao et al. (2012), who proposed an interactive visualization system called DAView for the evaluation and annotation of results of automatic discourse parsers. DAView features an innovative interface using the dendrogram icicle, and vertical and horizontal compact views. An informal formative case study showed that the system can facilitate the comparison of different parser algorithms. DAView is specifically designed for computational linguists for the use in comparative analysis of parser results in discourse studies, and not intended for general linguists and language users whose interests lie in exploring the linguistic properties of discourse structures. In addition, like the RST tool, the system is limited to the visualization of one level of functional information – the discourse, without integrating analysis at other levels.

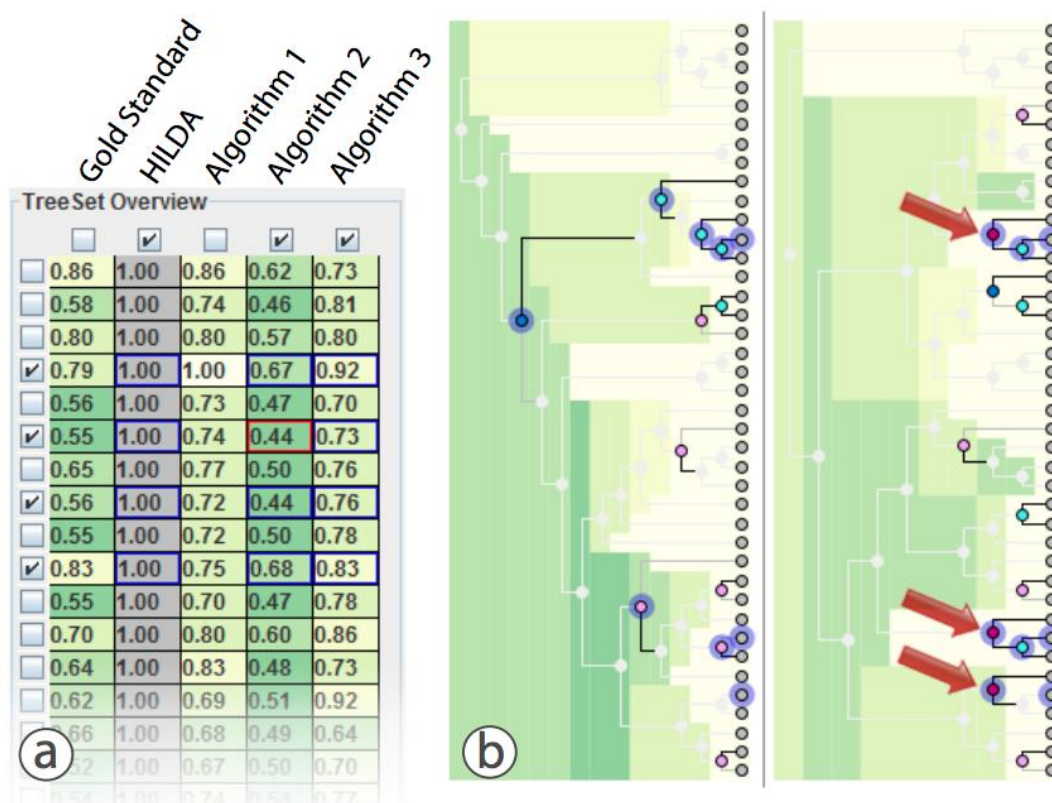


Figure 3. Interface for DAView for comparing different algorithms, adapted from (Zhao et al., 2012)

2.3 Clause Visualization

Functional-semantic patterning can appear in the form of clustering of the same functional-semantic choices at the clause level. Such patterning can take the form of repeated usage of the same lexical items, semantic or functional roles, or grammatical relations.

Traditionally, researchers use visualization to capture repeating patterns as in lexical repetition through frequency counting, regardless of contextual and positional information. For example, word clouds (e.g. Viegas et al., 2007) have been a common technique in understanding the basic theme of the text.

However, mere frequency counting seldom offers sufficient linguistic insight into the more functionally oriented aspects of textual data. Moreover, such frequency counting is often flawed and misleading, as they lack the contextual information necessary for meaningful interpretation. Without contextual information, it would be difficult to understand the meaning construed by the words in a word cloud. Other solutions such as *KWIC* (e.g. Culy & Lyding, 2011) provide some context of keywords useful in querying basic collocation patterns of words, but the lack of abstraction means that it still requires careful word-by-word reading of the actual text. These traditional word-based, linguistics-agnostic methods for pattern discovery may prove inadequate for visualizing complex structured

functional-semantic information.

2.4 Combining Two Levels of Functional-Semantic Information

In analyzing functional-semantic structure, two levels of information need to be accounted for: the architecture of text where the relations among clauses, or Elementary Discourse Unit (EDU) are studied, and the texture of text where the functional properties and their distribution within the EDUs are studied. Thus, two levels of analysis are necessary for the two facets of functional-semantic information: the discourse level for the architecture, accounted for by RST, and the clause level for the texture, accounted for by SFL.

In RST, the smallest unit for analysis is the clause (i.e. EDU). In SFL, the only recognized grammatical unit above the clause is the sentence (clause-complex). Synthesizing notations in the two frameworks, the way discourse EDUs relate and function in RST can be seen as an extension of the logical metafunction realized in structure by clause-complexing in SFL. To unify the account for the two complimentary levels of functional organization as a graph structure, we extend beyond the *rank* in functional linguistics to include *discourse* as a hierarchical parent of clause complexes/sentence. In this paper, we define “discourse level” as a unit that comprises a hierarchy of functional units of discourse and clause complexes, and “clause-level” to be ranks of clause, group, word and morpheme (Table 1).

Table 1. Different levels and ranks accounted for by RST and SFL.

Level	Hierarchy/Rank	Accounted for by
Discourse	Discourse/text	RST
	Clause-complex	RST & SFL
Clause	Clause	SFL
	Group	SFL
	Word	SFL
	Morpheme	SFL

In this way, the functions in SFL and RST correspond to each other, forming continuity from single lexical items up to the entire discourse. This functional continuity contrasts with the discontinuity of constituent analysis in formalist approaches where the integrity assumption at the sentence level no longer holds at the discourse level. With such natural functional continuity, the architecture and texture of text can now be unified and studied as a whole.

3. Interface Design

3.1 Design Considerations for the Visualization Platform

The design of any effective visualization is the result of careful contemplation. In designing the visual representation for the functional-semantic structures, we propose the following general guiding principles:

- (1) The visualization should reflect the natural structure of the theoretical framework so that the mapping between the data and presentation becomes smoother and easier.
- (2) The visualization should be designed in such a way that the inherent patterning in the data, if any, should be made obvious.
- (3) The visualization should be able to reflect more than one stratum of information when available. In our case, for example, lexical, structural and functional information should be integrated into the visualization interface.
- (4) The visualization should provide a proper level of linguistic abstraction to avoid being linguistically naïve like many other text visualization projects. The level of abstraction should be conducive to the discovery of functional patterning.
- (5) The visualization platform should take full advantage of the output from the computational systems we have thus far developed. The mapping between the output and the visualization should be made straightforward.
- (6) The visualization should be able to provide an overview of the functional structure while remaining interactive enough for exploring details on demand. Thus, certain degrees of interactivity and collaborative analysis should be made possible. Due to the difficulties in visualization the inherent complexity functional-semantic information structure, filters should be allowed to simplify the end output to highlight the parts of interest, again with the aim of facilitating pattern discovery.
- (7) Visual variables should be unified across different data sets so that comparison between visualizations of similar texts is possible. For example, when comparing political speeches, the visual representations (e.g. the meaning of a particular structure or color) of different speeches should remain basically identical.
- (8) The visualization platform should rely on the same database structure as the annotation platform and the two should be connected in such a way that any changes in the annotation can be reflected in the visualization in real-time. This facilitates the interoperability between annotation and visualization, two steps that are traditionally separated. Such interoperability helps collaborating team members and project managers to keep track of each other's progress and changes.

It is important to note that the goal for the visualization is to augment and improve reading of the underlying data, not to completely replace it. We aim to provide a viable tool for users to explore the hidden patterns in the functional-semantic structure, as well as revealing trends and tendencies hard to discover from pure/plain reading, allowing for the analysis of

large-scale texts in a time- and cost-efficient manner.

3.2 Discovery of Linguistic Patterning through Visualization

One of the primary goals for the construction of the visualization platform is to facilitate easier discovery of functional-semantic patterning in texts. In this study, we borrow from previous corpus-driven approaches to pattern discovery such as Pattern Grammar (Hunston & Francis, 2000), which have been applied to unveiling insightful patterns in large quantities of texts. We complement the largely purely descriptive (i.e. without drawing on any theory of language) approach to grammatical patterns of lexical items/collocations and extend such a notion of patterns to include those observable not only at the lexical level, but also a broader sense of “grammar” dealing with the functional-semantic aspect of language at the clausal and discourse level, whose discovery are made possible with the help of established frameworks such as SFL and RST.

To address the drawbacks of previous methods, we propose to analyze the functional-semantic structure of a text as an interconnected whole, instead of treating each functional component as a semantically independent unit. The functionally significant components have to be investigated in context, as each functional component is configured by the context it appears in. Words are no longer the only unit that construes meaning - it is only one of the several structural levels/ranks (others may include word, group, clause, clause complex, etc.). Since there is a natural continuity from clausal semantics to discourse semantics unified under the frameworks of RST and SFL, an effective visualization of functional semantics should reflect such continuity. Therefore, the visualization should not provide a representation merely for individual words or sentences, but more importantly for the functional relations that hold among the different functional units.

Our visual representation should reflect the underlying linguistic abstraction at the two levels as afforded by RST and SFL. At the discourse level, the visualization should be able to capture the way EDUs/clauses form relations with adjacent one to serve certain structural and functional roles (e.g. nucleus/satellite) as part of a hierarchy. At the clause level, the visualization should capture how the clausal constituents function in relation to other constituents in the clause. In other words, to visualize the functional-semantic structure of a text, we have to consider the relations among the various functional components which combine to form a holistic big picture. In this way, trends and patterning in the functional structures can be captured by looking at the whole picture and keeping track of the flow of text. Interoperation between the two different levels can provide useful insights as to how they correlate and how functional patterning comes into play.

3.3 Proposed Visualization Interface

We use a sentence taken from a Wall Street Journal article to illustrate the design of the interface:

A great many federal regulations are meant for larger entities and don't really apply to small businesses, says Marian Jacob, a legislative aide to Sen. Wallop.

For our discussion, we number each of the spans in the visualization.

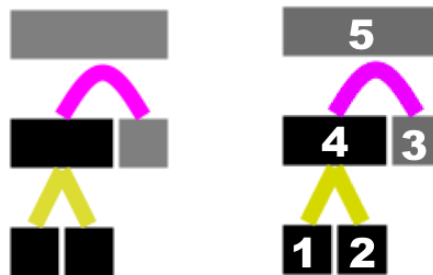


Figure 4. Basic building blocks for the visualization.

In the visualization, each span corresponds to one or more clauses in the example sentence:

(1) *A great many federal regulations are meant for larger entities*

(2) *and don't really apply to small businesses,*

(3) *says Marian Jacob, a legislative aide to Sen. Wallop.*

(4) (1) + (2)

(5) (1) + (2) + (3)

3.3.1 Discourse Hierarchy

At the discourse level, the basic unit of discourse relation, or Elemental Discourse Unit (EDU), is the clause. Each clause is represented by a rectangular block of the same width in the interface, regardless of the actual number of words in the clause. This provides abstraction away from the actual text, since we are concerned mainly with the relations between the EDUs. To represent the tree-like RST structure, a hierarchy is formed by layering the tree nodes at different levels such that the parent nodes are on top of the child nodes, with the parent nodes having a length equal to those of its children combined. In the example above, the span (1) and (2) are sibling nodes with span (4) as their parent. In turn, span (4) and (3) are siblings and both children to span (5). Such a representation affords simplicity and abstraction to the RST hierarchy.

3.3.2 Discourse Relation

In RST, each node in the hierarchy can form a discourse relation with one or more adjacent nodes. Although the number of adjacent nodes with which a node can form a discourse relation is unlimited, most relations are binary, with only two siblings at the same level. The discourse relation between any two nodes is represented by connecting lines and arcs. In the example above, span (3) is linked to span (4) by the relation “attribution”. Although a

hypotactic relation, the arc representing it is undirected since its direction is apparent from the nuclearity, with (4) being the nucleus and (3) the satellite. Between (1) and (2), two straight lines indicate a paratactic relation. Each relation type is represented by a distinct color. The color of a relation is consistent both within a visualization and across different visualizations. This ensures that patterns are made more obvious in a single visualization, and the patterns can be compared across different ones.

3.3.3 Nuclearity

The relation between two RST spans can be formed hypotactically or paratactically. In a hypotactic relation, only one span can serve as the nucleus while the others are satellites. For example, in arguing for a claim, evidence is often provided as support for the claim. In general, the evidence is considered a less essential part than the claim, as the claim conveys the gist of the message and can be standalone while the evidence cannot. On the other hand, the paratactic relation is multinuclear meaning all spans in the relation are nuclei with no satellites. Gray scale is employed to represent the nuclearity of each span. Nuclei are represented by dark spans, while satellites by lighter ones. Such contrast highlights the different roles played by the two different types of units. Each EDU in the visualization is connected with adjacent or parent units. In hypotactic relations, where only one nucleus is present, the nucleus is pointed to by adjacent satellites. Note that there can be more than one satellite pointing to the nucleus, although the majority of cases are binary, in pairs of one nucleus and one satellite.

3.4 *Connecting with the Transitivity Labeler*

The RST visualizer lays the cornerstone for the visualization of functional structures providing a view at the discourse level. However, to arrive at a holistic picture of the functional structure of a text, we still need to explore the clausal level functions. We represent the metafunctional structure in SFL using the same block-based structure as in the RST visualization so that each clause has two levels of analysis simultaneously.

For the example sentence, the labeler classifies the three clauses into three different process types in the ideational metafunction: cognition, happening and verbal, each represented by a distinct color. This clausal level information is then linked by dotted lines to the corresponding clauses in the RST structure. The visualization of the functional processes allows us to easily keep track of the functional strategies deployed by the author at the clausal level and makes the discovery of patterning at different stages of textual development an easier task. More importantly, this visualization allows us to explore how the two levels of functional-semantic information interweave as a whole: how different strategies at the one level may be correlated with and influenced/configured by the other.



Figure 5. Basic building blocks for the visual representation connecting the two functional levels.

3.5 Web-based Interactive Exploration

Linguistic analysis often involves subjective interpretation grounded on judgment of the analysts. The purpose of the interface is not to replace the judgment, but to provide a means to make the process smoother and more effective. The subjective nature of the analysis means that different analysts have different focus foci on what information to look at and further explore. The scale of such information often cannot be accommodated in a single visualization, necessitating further channels where analysts can explore particular areas in greater detail. In addition, previous research on interactive visualization (Stasko et al., 2008) points out that analysts often want to have direct access to the textual materials in addition to the visual representations of them. Thus, an effective visualization not only provides a high level of abstraction of the information but should also allow for ways to easily zoom in to details on demand. Such investigative activities are conducted in an interactive and explorative manner allowing analysts to freely focus on particular points of interest.

In our interface, we leverage the interactive features of the web-based platform to provide smooth interactive navigation. Users can point to a functional-semantic representation such as a RST span or process block to view additional details including the actual texts of the span and functional structure of a process.

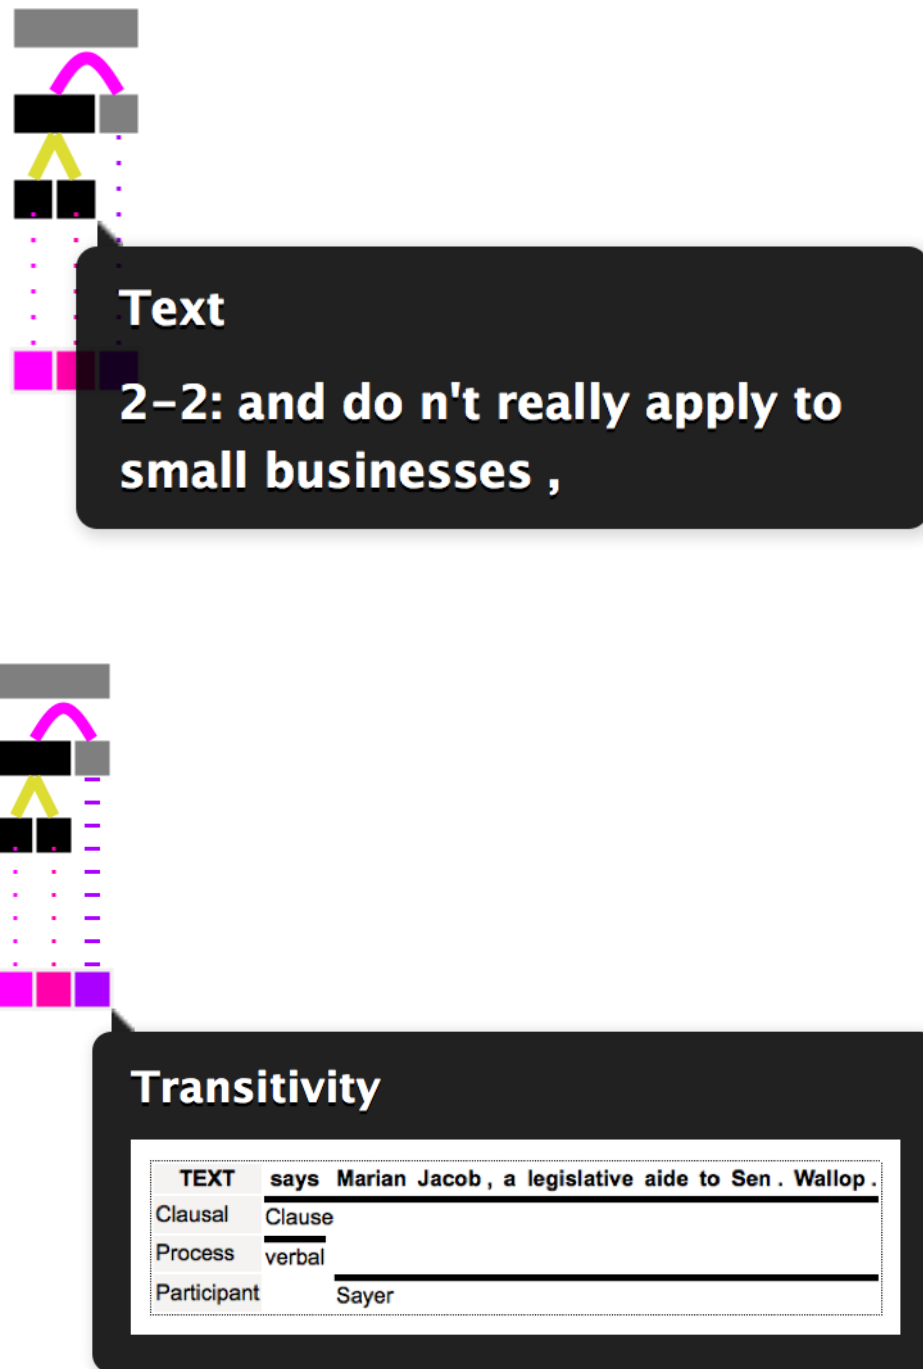


Figure 6. Interactive details-on-demand navigation.

We leverage state-of-the-art technologies for web-based interactive visualizations. The Data-Driven Document (D3) (Bostock et al., 2011) JavaScript library is employed to bring dynamic, interactive visualization on the client side. Compared to traditional visualization toolkits which hide the underlying scenograph and impose restrictions on the interoperability between different resources, D3 exposes the Document Object Model (DOM) to the user and

allow for direct manipulation of HTML (Hypertext Markup Language) and SVG (Scalable Vector Graphics) elements. Users bind data directly to DOM objects before generating or transforming the content to produce the desired visual effects. It is argued that with visualization applications increasingly leveraging not one but varied technologies such as HTML, CSS, JavaScript, SVG for handling different aspects of the visualization, such representational transparency improves expressiveness and efficiency. The data transform model in D3 allows clear separation of data and algorithms that manipulate them. Linguistic data are fetched from the unified storage system (Yan & Webster, 2013) and preprocessed at the backend server using a server-side language (in our case Python) before sending them to the client-side to be visualized by users. Users then interactively explore on the visual interface and manipulate with the document elements, resulting in task-specific queries automatically sent to the server in real-time using AJAX. The server handles the query by querying the database and performing the necessary computations before sending the results back to the client which updates the visual view accordingly. This cycle repeats until the user's needs are met and the applications are terminated.

3.6 Manual versus Automatic Analysis

One of the primary goals of the proposed work in this paper is to free humans from time-consuming repetitive tasks that can be automated through data-driven techniques. We aim to maximize automation so that humans can focus on what is really important and difficult requiring human experience and intuition. But to what degree can we achieve this? The question can be answered by comparing analyses that have been conducted through our automated pipeline with those by human experts. For this purpose, we perform a small-scale comparative analysis using texts annotated with RST structure from the RST Discourse Treebank (RST-DT). As our goal is to evaluate the efficacy of the visualization of unseen text samples, the texts used in the analysis are randomly selected from the test set of the RST-DT to ensure that they have not been used by the F & H discourse parser (Feng & Hirst, 2014) during the training phase.

In the following, we first present the visualizations of the texts from the RST-DT using our visualization interface. The samples are ordered by the size of the texts, from the smallest to the largest. The visualizations are presented in pairs. The first visualization in the pair is based on RST annotations from the RST-DT, which we create the visualization at the discourse level and also automatically generate its transitivity structure at the clause level. The second visualization is generated in a completely automatic manner, the discourse-level structure parsed by the F & H parser and the clause-level structure created from the transitivity labeler. After presenting the visualizations (in the Appendix), we discuss our observations in a following section.

4. Discussions

From the comparative visualizations presented (in the Appendix) we are able to make a few observations.

First, there seems to be a general correlation between the length of a visualized text and the

similarity in the two versions of its hierarchical discourse structure. In very small texts with about 10 clauses/100 tokens, the parser version is structurally similar to the corpus version, with some of the samples being nearly identical. In these samples, we can easily align individual clauses with their counterparts within the pair. In slightly larger texts of about 20 clauses/200 tokens, the discourse structure in some of the parser samples begins to deviate from the corpus ones, although the basic “architectures” of the text pair are still roughly the same and we can align some of the structural counterparts. Starting from 40 clauses/400 tokens, the differences between the two versions are widening quickly. While the overall contours of a pair can still be roughly identified, such as an upward/downward trend at a particular point, it is difficult to align the individual clauses in the two versions. Finally, with texts of over 90 clauses/800 tokens, even the overall contours/trending become unrecognizable, although small degrees of correspondence can still be found in parts of the versions.

On the other hand, while the differences in the overall discourse structure become more significant with the increase of text length, the clausal analysis have remained relatively stable and recognizable irrespective of text length, to the degree that it could be used to pinpoint and align individual clauses in the overall discourse structure.

However, the visualizations also reveal some problems in the automatically parsed version. It is apparent from the visualizations that more diversity of discourse relations is observed in the manual version than the automatic version, as indicated by the numbers of colors representing discourse relations. The dominant type of discourse relation in both versions is *Elaboration* (in dark blue color). However, the percentage of *Elaboration* in the parser version is much higher than its corpus counterpart, as reflected by the much richer varieties of colors representing discourse relations (further programmatic analysis confirms the impression, showing that the figures are 67.59% for the parser version and 29.17% for the corpus version). This suggests that the parser seems to suffer from a priori bias whereby the most frequent class in the training samples of the parser tends to be skewed. The visualizations further reveal that the same is true for multinuclear/paratactic relations, which, due to their relatively low frequency in the training samples, is underrepresented in the parser output.

Above all, the visualization interface has proven to be an effective tool in helping us conduct comparative/contrastive studies whereby the patterns/trends in the visualizations are easily identified. Discourse-level analysis based on the RST framework is known to be a difficult task with many variables, with a human inter-annotator agreement of about 65.3% in the RST-DT, and state-of-the-art automated parsers achieving about 84% of the human level. We should also note that unlike syntactic parsing, the interpretation of the discourse structure of a text is often subjective and there might not be a single gold standard acceptable to all analysts. While parser accuracy can still be improved with further development of discourse parsing technologies, the current visualizations have demonstrated remarkable efficacy of the platform in producing analysis that is often comparable to human experts. The correspondence between manual and automated analysis is high in small-scale texts, while in medium to larger texts, we are able to observe similar trends and patterns of varying degrees.

5. Case Study

The visualization interface presented in this study exposes many of the functional-semantic components both at the discourse and the clause level, making it possible for discovery of interesting linguistic patterns in texts. In this section, we briefly discuss a case study using the visualization interface in linguistic analysis of real-life texts.

During the development process, one expert analyst in Systemic Functional Linguistics and Discourse Analysis from our research team used the interface to help him theorize about the underlying patterning in the texts of two speeches as delivered by former U.S President Barack Obama: Obama's victory speech and the inaugural speech, both from Obama's first term. The expert notes that the two speeches, while supposedly written by the same author over approximately the same period of time, receive quite different response from critics: the victory speech is acclaimed for its eloquence, while the inaugural speech has been criticized for its lack of eloquence and logical progression. Hypothesizing that the difference might have been due to differences in the architectural consideration of the two texts at the discourse level, the expert turns to the visualization interface for further exploration. Using the interface, the expert was able to confirm that in the victory speech (Figure 7), there is a pattern for high-level, larger spans to tend to be connected more hypotactically than paratactically, with an incremental building up of successive spans culminating in the powerful resounding message: "Yes, we can". Contrastively, the high-level discourse spans in the Inaugural speech (Figure 8) tend to be more paratactically related, leading to critics' judgment that the speech is more diffuse without a cohesive central message.

Analysis of the transitivity patterns can also benefit from such an interface. Repeated occurrences of certain patterns of transitivity as uniquely distinguished from their neighbors can be easily seen in the visual representation. For instance, when we look at the visual representation at the clausal level, one interesting pattern emerges near the end of the inaugural speech: nearly all the processes in a block of 13 clauses are relational processes which is unlikely to be due to mere coincidence. Looking at the RST structure corresponding to each of the clauses, we discover that each of them is linked paratactically to each other. Interactive exploration of the texts in the clauses reveals that there exists an interesting number of parallel constructions within clause 200 to clause 212 (clause numbers are indicated in angle brackets):

<200> **It is** ultimately the faith and determination of the American people upon which this nation relies.

<201> **It is** the kindness to take in a stranger when the levees break; the selflessness of workers who would rather cut their hours than see a friend lose their job which sees us through our darkest hours.

<202> **It is** the firefighter's courage to storm a stairway filled with smoke, but also a parent's willingness to nurture a child, that finally decides our fate.

<203> Our challenges **may be new**,

<204> the instruments with which we meet them **may be new**,

<205> but those values upon which our success depends, honesty and hard work, courage and fair play, tolerance and curiosity, loyalty and patriotism -- **these things are old**.

<206> **These things are true**.

<207> They **have been** the quiet force of progress throughout our history.

<208> **What is demanded** then **is** a return to these truths.

<209> **What is required of us** now **is** a new era of responsibility -- a recognition, on the part of every American, that we have duties to ourselves, our nation and the world, duties that we do not grudgingly accept but rather seize gladly, firm in the knowledge that there is nothing so satisfying to the spirit, so defining of our character than giving our all to a difficult task.

<210> **This is the price** and the promise of citizenship.

<211> **This is the source** of our confidence: the knowledge that God calls on us to shape an uncertain destiny.

<212> **This is the meaning** of our liberty and our creed, why men and women and children of every race and every faith can join in celebration across this magnificent mall.

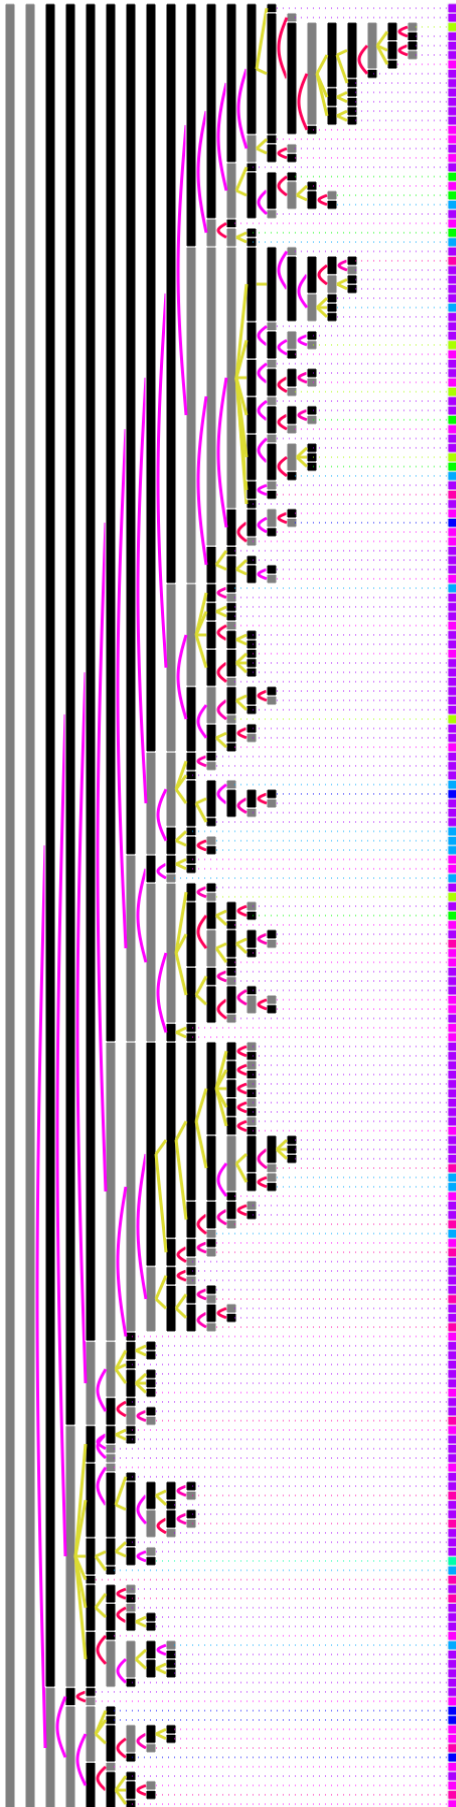


Figure 8. Complete visualization for Obama's Inaugural Speech.

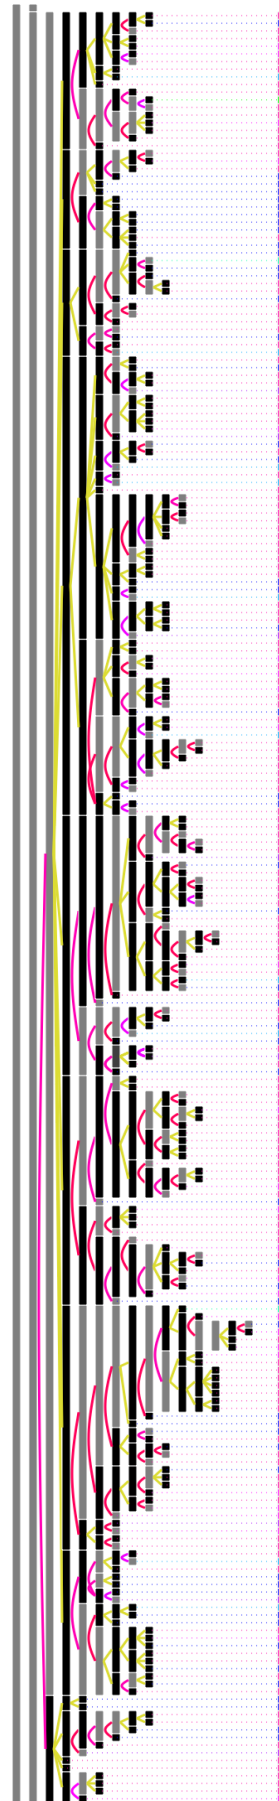


Figure 7. Complete visualization for Obama's Victory Speech.

These constructions in the speech are realized all in the same process type culminating in a long series of parallel patterns that are connected practically to achieve powerful rhetoric effects typical of Obama’s speeches. Even without careful word-by-word reading of the original text, the visualization makes it straightforward to discover such functional-semantic patterning at both the clausal and discourse levels.

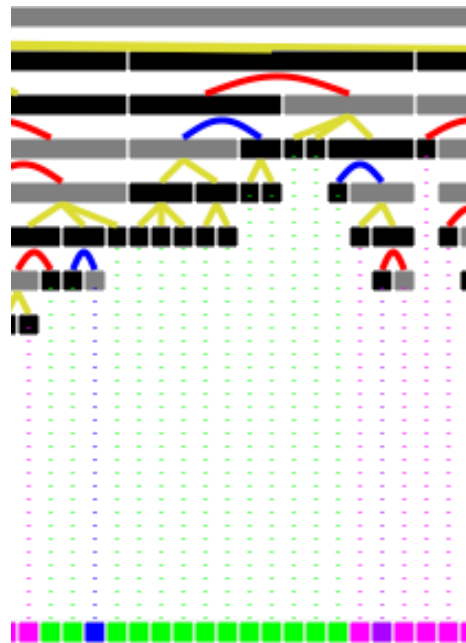


Figure 7. Patterning in clausal visualization.

6. Conclusion

In this paper, we have presented our initial design of a visual representation of the functional-semantic information. We first reviewed some of the recent development in interactive linguistic visualization of textual data. We then presented our design of an innovative visualization interface aimed to help analysts explore and understand functional-semantic components more efficiently and effectively by providing a high-level architecture of the overall functional-semantic structure of texts and an abstraction from the lower-level texture at the clause level. The visualization interface highlights the relationships and connections among different functional-semantic components for analysts to browse and explore.

To test the efficacy of the interface, we first performed a comparative study by analyzing the manually annotated and automatically generated versions of the same texts from the test set of the RST-DT corpus. The results demonstrated the feasibility of the platform for fully automated discourse analysis both at the discourse and clause level. We then conducted a small case study by an expert analyst using the interface, whose initial discovery demonstrates the feasibility of the proposed visualization design to help analysts to more easily uncover functional-semantic patterns in text. While the results are promising, the discussion presented here is by no means an exhaustive survey of the possible applications to

which the visualization interface can be put. Rather, it serves as a preliminary demonstration that showcases the usefulness of the proposed visualization system.

Acknowledgement

This study was supported by Guangdong Planning Office of Philosophy and Social Science (grant number: GD18YWW02), the Foundation for Distinguished Young Talents in Higher Education of Guangdong (grant number: 2015WQNXCX030), Department of Education of Guangdong Province, P. R. China, and by Innovative School Project in Education of Guangdong, China (project number: 312-GK18CQ70).

References

- Bostock, M., Ogievetsky, V., & Heer, J. (2011). D3 Data-Driven Documents. *Visualization and Computer Graphics, IEEE Transactions On*, 17(12), 2301–2309.
- Cao, N., Sun, J., Lin, Y.-R., Gotz, D., Liu, S., & Qu, H. (2010). Facetatlas: Multifaceted visualization for rich text corpora. *Visualization and Computer Graphics, IEEE Transactions On*, 16(6), 1172–1181.
- Chen, C. (2010). Information visualization. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4), 387–403.
- Collins, C., & Carpendale, S. (2007). VisLink: Revealing relationships amongst visualizations. *Visualization and Computer Graphics, IEEE Transactions On*, 13(6), 1192–1199.
- Collins, C. M. (2010). *Interactive Visualizations of natural language*. University of Toronto, Toronto.
- Collins, C., Penn, G., & Carpendale, S. (2009). Bubble sets: Revealing set relations with isocontours over existing visualizations. *Visualization and Computer Graphics, IEEE Transactions On*, 15(6), 1009–1016.
- Culy, C., & Lyding, V. (2011). Corpus clouds-facilitating text analysis by means of visualizations. In *Human Language Technology. Challenges for Computer Science and Linguistics* (pp. 351–360). Springer.
- Feng, V. W., & Hirst, G. (2014). A Linear-Time Bottom-Up Discourse Parser with Constraints and Post-Editing. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1, 511–521.
- Heer, J., Viégas, F. B., & Wattenberg, M. (2007). Voyagers and voyeurs: supporting asynchronous collaborative information visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 1029–1038).
- Hernault, H., Prendinger, H., & Ishizuka, M. (2010). HILDA: a discourse parser using support vector machine classification. *Dialogue & Discourse*, 1(3).
- Hunston, S., & Francis, G. (2000). *Pattern grammar: a corpus-driven approach to the lexical*

grammar of English. Amsterdam; Philadelphia: John Benjamins Pub. Co.

Isenberg, P., Elmqvist, N., Scholtz, J., Cernea, D., Ma, K.-L., & Hagen, H. (2011). Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization, 10*(4), 310–326.

Keim, D. A., Mansmann, F., Schneidewind, J., Thomas, J., & Ziegler, H. (2008). Visual analytics: Scope and challenges. In S. J. Simoff, M. H. Böhlen, & A. Mazeika (Eds.), *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Vol. 4404, pp. 76–90). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-71080-6_6

Koh, K., Lee, B., Kim, B., & Seo, J. (2010). Maniwordle: Providing flexible control over wordle. *Visualization and Computer Graphics, IEEE Transactions On, 16*(6), 1190–1197.

Leydesdorff, L., & Welbers, K. (2011). The semantic mapping of words and co-words in contexts. *Journal of Informetrics, 5*(3), 469–475.

Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text, 8*(3), 243–281.

Mazza, R. (2008). *Introduction to Information Visualisation*. Guildford, Surrey: Springer London.

Nilsson, J., & Nivre, J. (2008). MaltEval: an Evaluation and Visualization Tool for Dependency Parsing. In *LREC*.

O'Donnell, M. (2000). RSTTool 2.4: a markup tool for Rhetorical Structure Theory. In *Proceedings of the first international conference on Natural language generation-Volume 14* (pp. 253–256).

O'Donnell, M. (2008). Demonstration of the UAM CorpusTool for text and image annotation. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Demo Session* (pp. 13–16).

Rockwell, G., Sinclair, S. G., Ruecker, S., & Organisciak, P. (2010). Ubiquitous text analysis. *Paj: The Journal of the Initiative for Digital Humanities, Media, and Culture, 2*(1).

Rohrdantz, C., Mayer, T., Butt, M., Plank, F., & Keim, D. A. (2010). Comparative visual analysis of cross-linguistic features. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)* (pp. 1–7).

Stasko, J., Görg, C., & Liu, Z. (2008). Jigsaw: supporting investigative analysis through interactive visualization. *Information Visualization, 7*(2), 118–132.

Teich, E. (1999). *Systemic functional grammar in natural language generation linguistic description and computational representation*. London; New York: Cassell.

Viegas, F. B., Wattenberg, M., Van Ham, F., Kriss, J., & McKeon, M. (2007). Manyeyes: a site for visualization at internet scale. *Visualization and Computer Graphics, IEEE*

Transactions On, 13(6), 1121–1128.

Whitelaw, C., & Argamon, S. (2004). Systemic functional features in stylistic text classification. *Ms., Sydney Language Technology Research Group, University of Sydney, Sydney, Australia.*

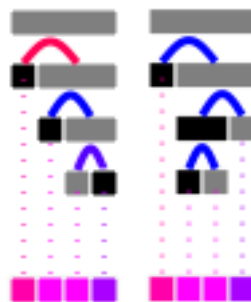
Whitelaw, C., Garg, N., & Argamon, S. (2005). Using appraisal groups for sentiment analysis. In *Proceedings of the 14th ACM international conference on Information and knowledge management* (pp. 625–631).

Yan, H. (2014). Automatic labelling of transitivity functional roles. *Journal of World Languages*, 1(2), 157–170. <https://doi.org/10.1080/21698252.2014.937563>

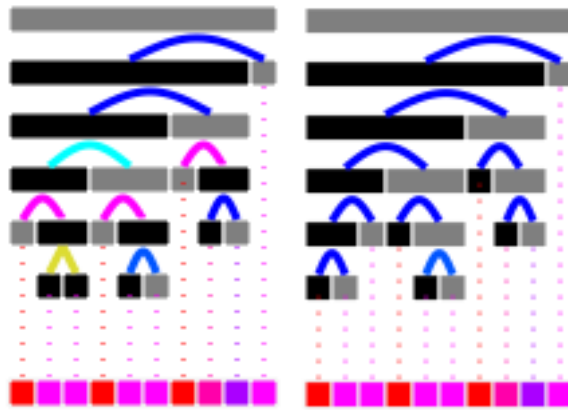
Yan, H., & Webster, J. (2013). A Corpus-based Approach to Linguistic Function. In *Proceedings of the 27th Pacific Asia Conference on Language, Information, and Computation (PACLIC 27)* (pp. 215–221). Retrieved from <http://www.aclweb.org/anthology/Y13-1020>

Zhao, J., Chevalier, F., Collins, C., & Balakrishnan, R. (2012). Facilitating Discourse Analysis with Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), 2639–2648. <https://doi.org/10.1109/TVCG.2012.226>

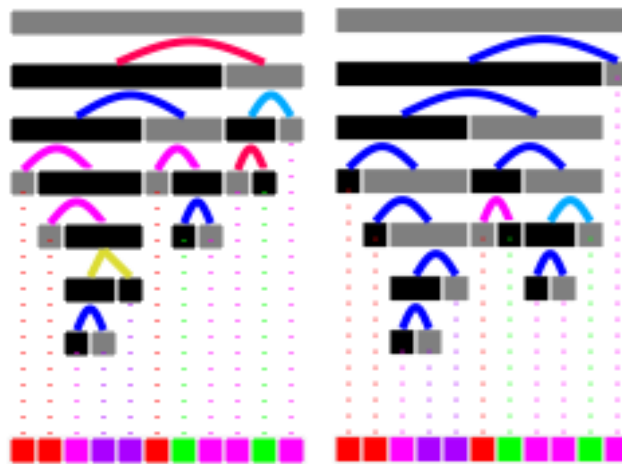
Appendix - Comparative Visualizations of Manually and Automatically Generated Functional-semantic Structures



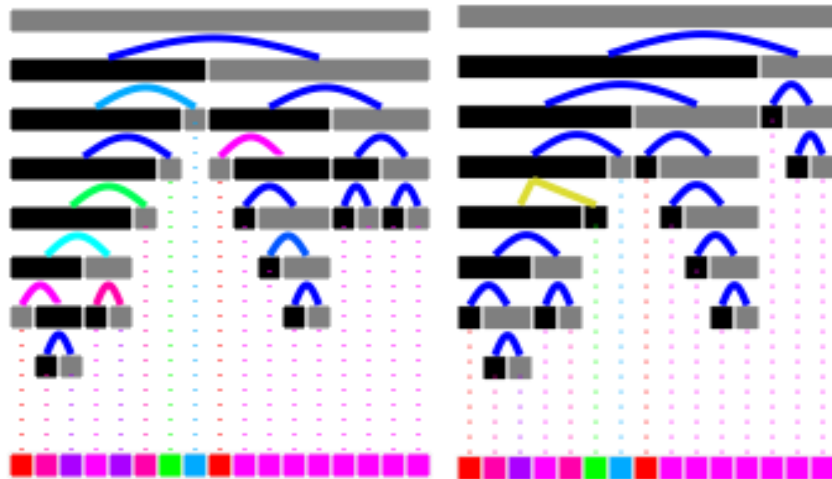
Visualizations of file *wsj_0644* (4 clauses, 57 tokens)



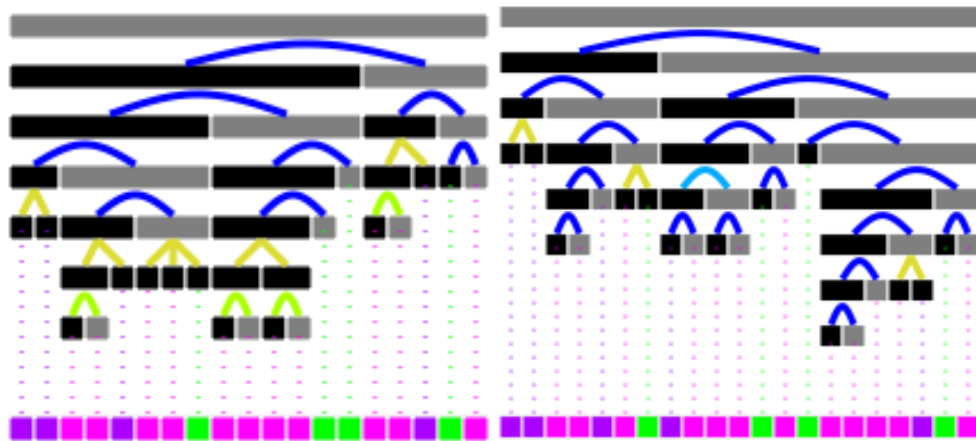
Visualizations of file *wsj_0684* (10 clauses, 76 tokens)



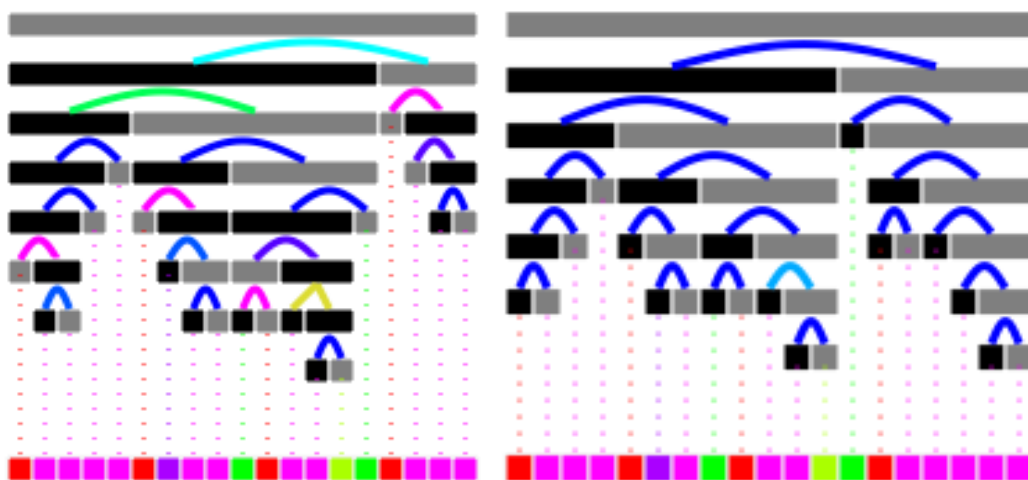
Visualizations of file *wsj_1354* (11 clauses, 123 tokens)



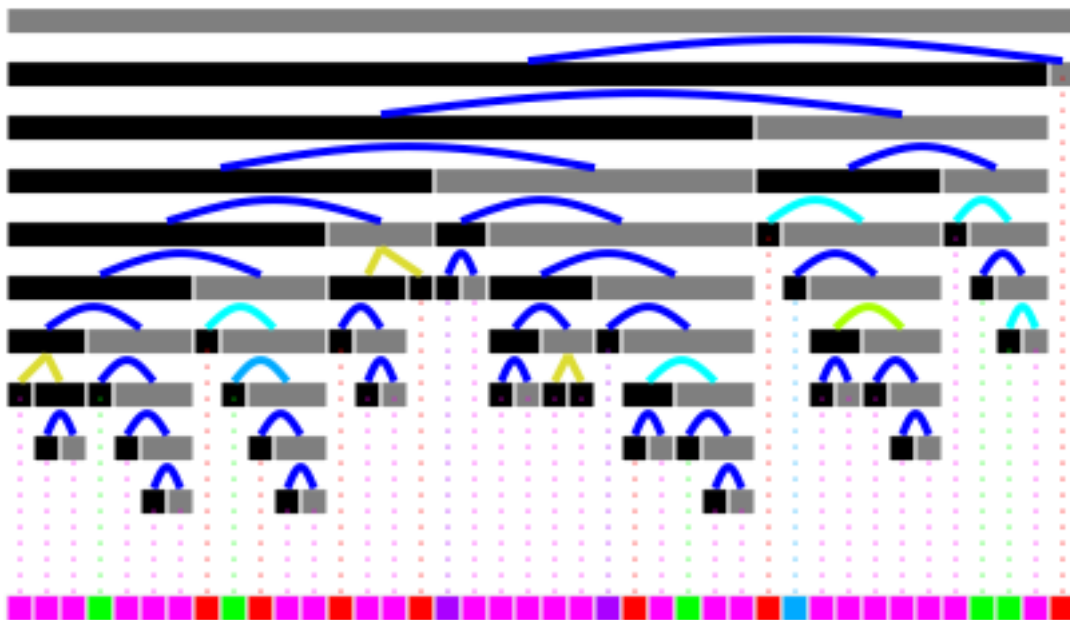
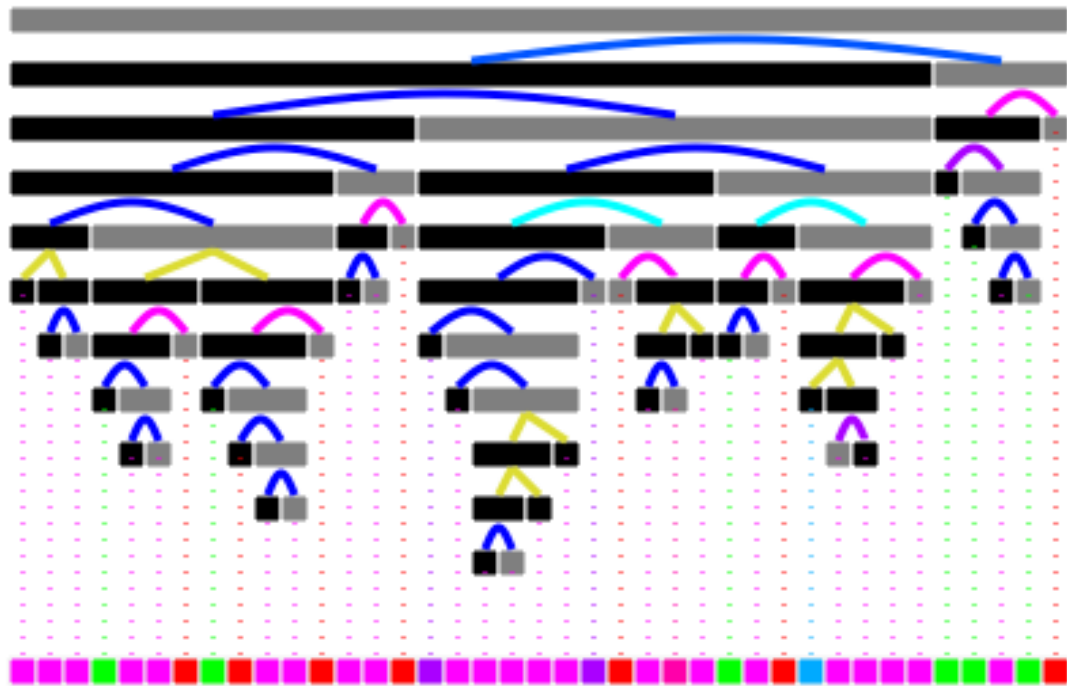
Visualizations of file *wsj_0667* (17 clauses, 207 tokens)



Visualizations of file *wsj_0607* (19 clauses, 156 tokens)



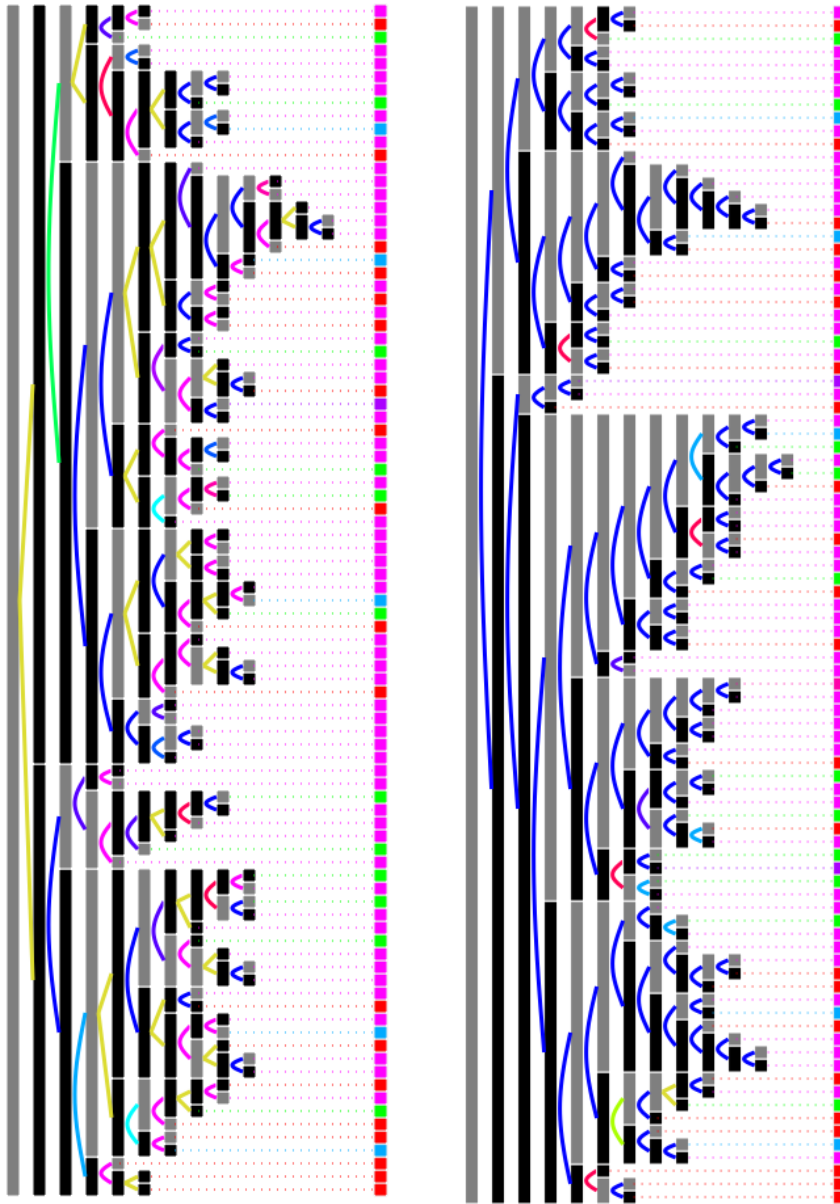
Visualizations of file *wsj_0654* (19 clauses, 208 tokens)



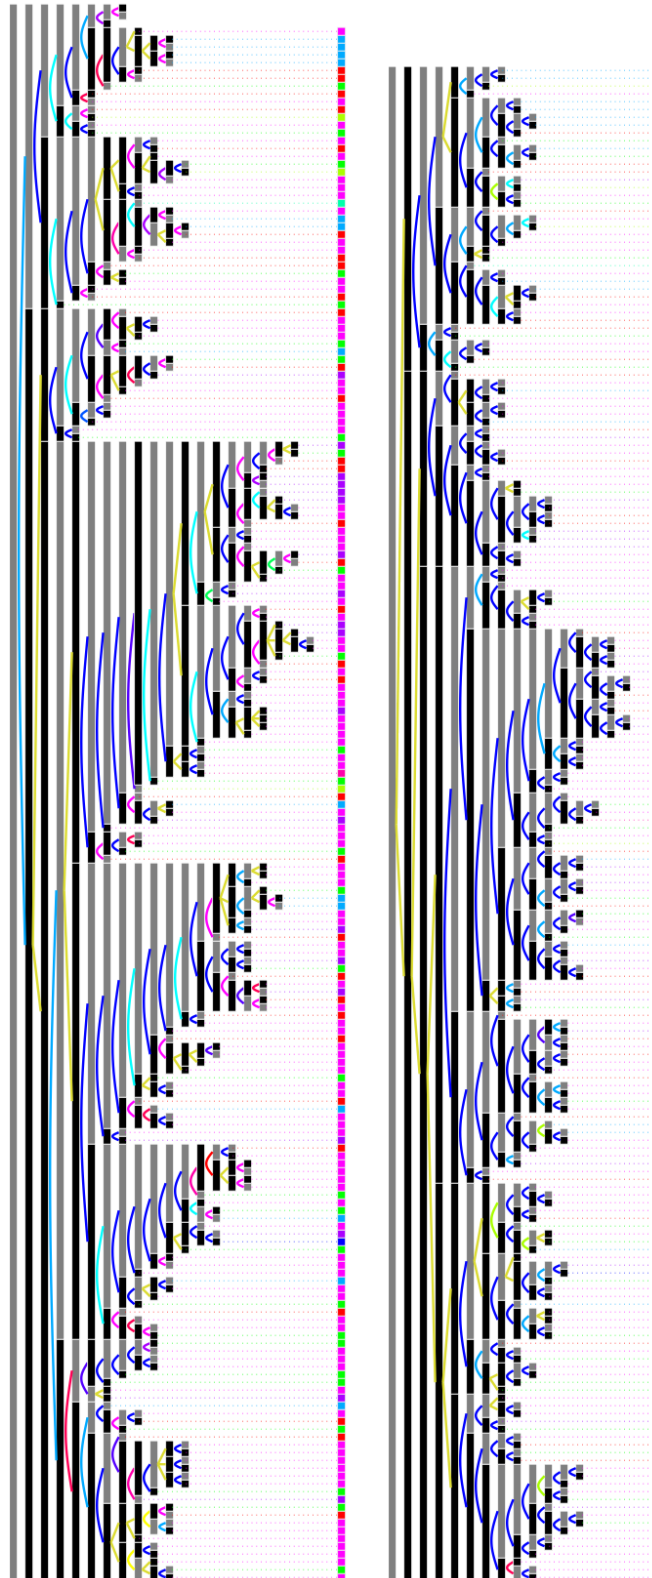
Visualizations of file *wsj_1365* (39 clauses, 335 tokens)



Visualizations of file *wsj_1148* (43 clauses, 426 tokens)



Visualizations of file *wsj_1189* (91 clauses, 731 tokens)



Visualizations of file *wsj_1376* (202 clauses, 1628 tokens)

Copyright Disclaimer

Copyright for this article is retained by the author(s), with first publication rights granted to the journal.

This is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).