

# Title: LdoD Visual - A Visual Reader for Fernando Pessoa's Book of Disquiet: An In-Out-In Metaphor<sup>1</sup>

Author: José Raposo (INESC-ID, IST, University of Lisbon)

## Abstract

There is an increasing use of information visualization techniques in Digital Humanities to support the analysis of literary works. On the other hand, the rise of the e-book has triggered a variety of hardware and software solutions to provide new reading environments and new material engagements with textual forms. In this paper, we propose a new metaphor, In-Out-In metaphor, to support the reading of literary works in a digital medium, which integrates information visualization techniques with plain text reading into the reading flow. This approach differs from the well-known dichotomy between close and distant reading because its emphasis is not on supporting the analysis of the texts being read, but on providing a smooth flow between the focus and digression, which naturally occurs during a reading experience.

In this context, we present a solution to read and explore an unfinished book that can be read in any order, *The Book of Disquiet* by Fernando Pessoa. The research and design decisions for this work are driven to provide an appealing, rich and interactive online reading experience through a web-based application, LdoD Visual<sup>2</sup>. The implemented features apply information visualization techniques with an emphasis on exploring the modular nature of the book. Since this is a fragmentary literary work, its reading can be fragmentary as well, but providing a smooth multiple flow reading experience. The developed application is an extension of the collaborative digital archive of the *The Book of Disquiet*, the LdoD Archive<sup>3</sup>.

## 1. Introduction

*The Book of Disquiet* (LdoD, from *Livro do Desassossego*, in Portuguese) by Fernando Pessoa is a fragmentary literary work that was left unedited and unpublished by the author. 47 years after Pessoa's death (1935), the book's first edition was published in 1982. Since then, several different editions of the book have been published, showcasing significant disparities of interpretation about how the book should be organized according to their respective editors.

The modularity of this work is, in fact, the contingent result of material and historical features that tend to reinforce each other: on one hand, each text is a self-contained unit that has no necessary narrative continuity with other texts; on the other hand, there is no defined ordering of the entire corpus of circa 600 pieces. Pessoa only defined a few textual clusters, and even those suffered changes over time. Given this contingent modular nature of the book, its fragments can be read without any particular order, offering the opportunity for a fragmentary reading of the *LdoD*, that is, a practice of reading that fosters divergent reading paths that can be built by the actual reader, instead of following a typical sequential reading that is bound to happen in a rigidly defined story sequence.

In this work we used the particularities of the *LdoD* structure and characteristics to explore, and smoothly integrate, two current research trends in Digital Humanities: information visualization techniques applied to texts, including the possibility of moving between close and distant reading scales (Jänicke et al., 2015), with the emergence of e-books as the new reading medium (Koolen et al. 2012). Therefore, we propose a new reading metaphor, the In-Out-In metaphor, which integrates information visualization techniques with the plain reading of the text, to support the focus-digress flow that structures the embodied activity of the reader. By *focus-digress flow* we mean that during the focus stage the reader is in the story, while the digress stage corresponds to moments of self-consciousness as a result of the reading.

Therefore, we have designed and implemented a web-based application to read the *LdoD* whose main objective is to give the reader a sense of choice and an active role while reading and browsing through the book's different fragments, in what should be a pleasurable and immersive user experience for both humanities researchers or just interested readers.

To achieve this goal, we integrate a set of meta-information about the *LdoD* with information visualization techniques. Information visualization is a broad field that can be applied to a large domain of different areas, providing insight based on input data with different perspectives, summarizing relationships, identifying and comparing patterns, trends, outliers, finding correlations, among other tasks. By integrating the existing meta-information with these information visualization techniques we developed a tool, the *LdoD Visual*, that supports the exploration and pleasurable reading of the *LdoD*.

The implemented solution provides new contributions for the reading and exploration of the *LdoD*, in particular, and to the concept of reading flow in Digital Humanities, in general, and it is built on top of the collaborative digital archive for the *LdoD* - the *LdoD Archive* (Portela and Rito Silva, 2014; Rito Silva and Portela, 2015).

## 2. Related Work

In this section we present some background concepts regarding the *LdoD Archive*, Visualization in Digital Humanities, and Classification in Digital Humanities, in which context this work was developed.

### 2.1. The *LdoD Archive*

The *LdoD Archive* provides every visitor the possibility of reading the *LdoD*

<sup>1</sup> Submitted to Digital Humanities Quarterly, with co-authors, António Rito Silva (INESC-ID/UL) and Manuel Portela (CLP/UC)

<sup>2</sup> <https://ldod.uc.pt/ldod-visual>

<sup>3</sup> <https://ldod.uc.pt/>

through four different available expert editions - the first edition, edited by Jacinto do Prado Coelho, which was published only in 1982, and other 3 major critical versions that have been published since then by Teresa Sobral Cunha (1990-91), Richard Zenith (1998) and Jerónimo Pizarro (2010). Each edition is composed of a different number of fragments and fragments are ordered differently in each edition. The user can read these fragments while comparing its respective position across these different expert editions, and they can also have access to other information that differs according to the edition, such as fragment date and heteronym attribution.

In the *LdoD Archive*, users can also create their own virtual editions of the *LdoD* by selecting fragments from each of the 4 expert editions, sorting, annotating and tagging them. This virtual edition creation can be produced in a collaborative manner, in which case several users are involved, and users can play the roles of editors and managers of a virtual edition. The tagging of a fragment, which can apply to the fragment as a whole, or to just a segment of the fragment, is done by means of a set of categories that is associated with the virtual edition. This set of categories is designated as the virtual edition's taxonomy.

Each fragment has different properties, such as its title, text, date and the edition and categories to which it belongs. There are many more features and possibilities that can be explored in the archive, such as additional reading tools, original documents, comparisons between textual transcriptions, fragment search, multiplayer classification games, and the writing of variations based on the existing fragments.

## 2.2. e-Books

The emergence of e-Books has significantly changed the landscape of reading. Several analyses have been done on the hardware and software features for reading resulting from the use of e-Books (Koolen et al. 2012) that highlight the differences between linear sequential reading and multicursal interactive reading, which is particularly relevant in the context of academic research or education. This tension pervades all solutions, where there is a tradeoff between the smooth flow of reading and the possibility of analysis and inquiry about what is being read.

### 2.2.1. Reading in the *LdoD Archive*

Prior to this work, the *LdoD Archive* provided two reading user interfaces: one that emphasizes comparison<sup>4</sup> and another that emphasizes the sequence of reading<sup>5</sup>.

In the comparison user interface, it is possible to have a detailed highlighting of the differences between the authorial sources and the editorial transcriptions, and also across the editorial transcriptions themselves. It is also possible to annotate particular texts (or textual passages) with comments and tags.

The differences between transcriptions among the different expert editions, and authorial sources, are highlighted by collocating the fragments side by side, and

using colors to identify which parts differ. This user interface also allows the comparison of meta-information assigned by the expert editors, like the given title, date of publication, heteronym attribution, page number, which can differ according to the edition. Additionally, this interface also allows users to compare the transcription with the original document sources, by presenting the documents' facsimiles, and providing a diplomatic transcription, which includes deletions, insertions and alterations. On the other hand, and since the *LdoD Archive* also supports definition of virtual editions, it is possible, through the comparison user interface, to compare the differences between the fragments, in terms of the assigned tags and annotations done. It is through the comparison interface that authenticated users can annotate and tag the fragments.

Due to the focus of the comparison user interface on the analysis aspects, another user interface was defined, in order to provide a smother reading experience, although still emphasizing the comparison between the different sequences of reading the *LdoD*. This interface has a look and feel closer to the printed typography of the period the *LdoD* was written, and intends to provide reading sequences according to each one of the four expert editions. Besides, it provides a recommendation interface that suggests the next fragment to read, given the fragment being read and a set of criteria, like, for instance, text similarity between fragments.

## 2.3. Visualization in Digital Humanities

Jänicke et al. (2015) address visual text analysis in Digital Humanities and categorize each of these techniques as distant reading techniques, which aim to generate an abstract view of textual content, and as close reading techniques, which rely mainly on annotations and the use of different colors and underlining styles to lead to deep comprehension and thorough interpretation of textual passages. They analyze how close and distant reading techniques can be combined to provide a bridge between these two perspectives so that the user can switch between close and distant reading. This is mainly achieved by side-by-side visualizations of the text being read and its meta-information, which make it possible to switch from a close reading of a particular fragment to some kind of overview associated with some of its metadata. On the other hand, it is possible to access a particular text from the distant reading view where the text is represented because it has properties that are analysed in that distant reading mode. This integration between close and distant reading is similar to our goal but our emphasis is on providing a reading experience, while the integration of close and distant reading is concerned with the analysis and understanding of the texts, such that close reading is frequently related with the annotation and generation of meta-information while the user is reading a text.

### 2.3.1. Visualization Techniques for *LdoD*

Given the specific characteristics of the *LdoD* we can analyze which visualization technique can be applied in its representation. These techniques belong to different areas and types of tasks, such as reading guidance, user navigation, global view, and exploratory analysis.

<sup>4</sup> As an example, see

[https://ldod.uc.pt/fragments/fragment/Fr080/inter/Fr080\\_WIT\\_ED\\_CRIT\\_Z](https://ldod.uc.pt/fragments/fragment/Fr080/inter/Fr080_WIT_ED_CRIT_Z)

<sup>5</sup> As an example, see

[https://ldod.uc.pt/reading/fragment/Fr080/inter/Fr080\\_WIT\\_ED\\_CRIT\\_Z](https://ldod.uc.pt/reading/fragment/Fr080/inter/Fr080_WIT_ED_CRIT_Z)

Reading guidance features consist of giving important information for the user regarding a specific fragment or even a collection or category of fragments, where the focus is on the fragment that is currently being read. User navigation features provide the user with an interface to navigate between different parts of the *LdoD*, in which the focus is on the relationship between the fragment being currently read and the remaining fragments of the book. Global view features have the objective of giving the user a summarized presentation or zoomed-out view of content from the *LdoD*, thus directing the focus to the edition that is being read. Exploratory analysis features present the user with information that is more focused on exploring relationships in the *LdoD* without necessarily converging to the task of reading.

Word clouds (Viégas et al, 2009; Mohammad, S. M. 2012; Heimerl et al, 2014) can be used to provide global views and reading guidance. Word clouds can be associated with the taxonomy categories of a virtual edition. After being presented with a word cloud of the available categories, and clicking on one of them, the user would be redirected to a list of fragments that corresponds to the selected category.

Also by exploring the categories associated with the virtual editions, it is possible to provide a global view of the edition, using colors to identify the fragments associated with a particular category. This feature would provide a global view of the fragments holding the category and the navigation to another fragment that hold the same category. This visualization technique, is, for instance, used in the Diggers diaries web site (Vilaplana and Pérez-Montoro, 2017).

Besides the categories criteria, other criteria are supported by the *LdoD Archive*, such as text and heteronym, and they can be used to define a distance value between fragments. Therefore, when reading a given fragment, the user could browse through similar fragments according to the selected criteria, and this relation can be visually represented using a network graph to show how close certain fragments are in relation to others. This would provide reading guidance. Alexander and Gleicher (2016) make use of network graphs to visualize text similarity between different books by William Shakespeare.

By highlighting some of the fragments' words according to their relevance in the context of the entire book it is possible to provide a guidance while reading a fragment. The weight of this highlighting could be made resorting to a metric like Term Frequency – Inverse Document Frequency (TF-IDF). Text Skimming (Brath and Banissi, 2014, 2015) is a technique used in knowledge maps and to support reading and that can be applied in this context.

There are some differences between the experience of reading a book on a browser or other portable virtual book readers and the actual physical book. In an attempt to close this gap, it would be interesting to provide the user with a graphical representation that makes the awareness of where the actual reading fragment is placed in the current edition. This type of feature belongs to the global view area.

#### 2.4. Text Classification

The visualization techniques are built on top of the information to be presented plus some meta-information. In the context of the *LdoD Archive*, the information is the fragment transcription and its position in an edition, and the meta-information is the date, heteronym and categories assigned to a fragment. Additionally, some meta-information is generated from the fragment through the TF-IDF associated with it, by considering each fragment as a document and the words in the fragment its set of terms.

##### 2.4.1. Distance Measures in the *LdoD Archive*

The *LdoD Archive* implements a set of measures to calculate the distance between two fragments belonging to the same edition according to the 4 types of meta-information: date, heteronym, categories, and TF-IDF.

The distance between two fragments is calculated using the cosine similarity (Singhal, A., 2001), according to which a vector is defined for each of the fragments to be compared.

Associated with the heteronym, a vector of two cells is defined, one for each heteronym, Vicente Guedes and Bernardo Soares, in which the value 1.0 is assigned to the cell of the heteronym if the fragment is assigned by the editors to it.

For the date a vector is defined with the number of cells equal to  $1934-1913+1$ , which corresponds to the period when the fragments were written. For the cell corresponding to the year when a fragment was written, the value 1.0 is assigned and for the other cells a decay of  $\text{MAX}(0.0, 1.0 - N \times 0.1)$  is applied where  $N$  is the distance between the year associated with the cell and the closest cell with value 1.0.

For the text the set of the first TF-IDF terms associated with each one of the fragments that we intend to compare is calculated. Note that this set is between 100, in which the two fragments have the same first TF-IDF 100 terms, and 200, in which all terms are different. Then, for each fragment, a vector whose size equals the number of terms in the set is defined, and to the cell associated with a term the value 1.0 is assigned if the term belongs to the first 100 TF-IDF terms of the fragment, and 0.0 otherwise.

To calculate the distance between two fragments using the set of categories assigned to them a vector whose size is equal to the number of categories associated with the virtual edition taxonomy is defined. Then, for each fragment, its respective vector is filled with 1.0 for the cells corresponding to the categories the fragment has, and 0.0 for the other cells.

Note that the archive also uses a topic modelling algorithm, through the Mallet software<sup>6</sup>, which allows the virtual editions to be automatically classified using the set of generated topics.

### 3. The In-Out-In Metaphor

---

<sup>6</sup> <http://mallet.cs.umass.edu/>

Reading a book is a task in which the reader constantly switches between two states of focus. The reader is either completely engaged in the reading task or briefly disengages out of the act for various reasons. These moments of disengagement can be triggered, for example, by recalling what happened in previous chapters, picturing a description of a character or scenario, remembering a similar personal experience in relation to what has been read, linking a new event to what has happened in the past or even trying to make a prediction out of it, among other possibilities. Part of our solution model revolves around this phenomenon, as it is intended to materialize and direct this constant in and out that happens while emerging out of and submerging into the source text. It is important to highlight that, in comparison to a typical book where we have a rigidly defined story sequence, the *LdoD* is a fragmentary literary work. This means that these moments of disengagement can turn into a chance for the user to reposition herself in another part of the book. This nuance leads to a reading path and global mental picture of the book that can greatly vary from reader to reader, opening up space for providing the user with an immersive experience in which she can have a sense of choice and an active role while reading the *LdoD*.

#### 4. The LdoD Visual

The solution's interface is composed of a group of different main menus - edition selection, reading, current user activity, new user activity, and reading history menus. Before describing each menu in detail, it is important to explain the concepts that will be used throughout the rest of this solution proposal.

##### 4.1. Core Concepts

An *information visualization technique*, among other descriptions in previous sections, is a visual representation or metaphor that can provide insight and different perspectives about almost any type of input data, as we saw above from the multiple examples given in the related work applied to the context of *the LdoD*.

*Semantic criteria* is the set of attributes that can be encoded in the information visualization techniques in the context of *the LdoD*, namely word frequency, word relevance (TF-IDF), chronological order, heteronyms, fragment categories, and taxonomies. Text similarity is another semantic criterion that can be derived from word frequency and word relevance.

A *user activity* is an activity that can be selected by the user while reading a fragment from the *LdoD*, consisting of making use of an *information visualization technique* to encode different *semantic criteria*. An example of a user activity is "Find similar fragments by text similarity" to find fragments that are similar to the fragment that is currently being accessed, using a network graph as the information visualization technique that encodes text similarity - the semantic criterion that is chosen by the user while selecting the new user activity.

To achieve an intuitive and easy-to-use solution, independently of the user activity's information visualization technique and encoded semantic criteria, the user selects a user activity by clicking on a button with a high-level description of what the user activity was designed to achieve, e.g. "Find similar fragments by text similarity".

A *reading flow* is a sequence of accesses, done by a single user, such that the user progresses in the sequence of fragments through *in-in*, *in-out-in* and *in-out-out-in* subflows. An *in-in* subflow corresponds to a situation when the user is accessing a fragment and accesses another fragment, for instance, by switching to the previous page. This switch is done according to the *semantic criteria* the user chose for accessing the fragment, defined by the current user activity. For instance, if the semantic criteria correspond to the set of fragments of a given category, the reader will access another fragment belonging to the same category. An *in-out-in* subflow corresponds to a situation in which the user is accessing a fragment, moves to the current user activity menu (to be described in the next subsections) in which the reader has a complete view of the fragments according to the same semantic criteria used to access the fragment being read. For instance, if we consider the previous example, the complete view may show all fragments that have the same category in the context of all fragments. An *in-out-out-in* subflow corresponds to a situation when the reader stops reading, according to the current user activity, and chooses a new user activity in which she continues to read.

Overall, the reading flow corresponds to the implementation of the *in-out* metaphor, where the level of disruption is minimal in *in-in* subflows, and maximal in *in-out-out-in* subflows. Note that to reduce the disruption of the reading flow, all changes can be done using shortcuts, pressing a single key, and most menu elements fade-out in the reading interface when the reader does not use the mouse so that the screen only contains the text of the fragment.

##### 4.2. User Interface Structure: Menus and Navigation

In the context of an edition, the *LdoD Visual* has three main types of elements: selection menus, reading view, and global views.

###### 4.2.1. User Activity Selection Menu

Once a context of an edition and one of its fragments is chosen, the main selection menu allows users to choose between the available user activities, given the context, and access the corresponding global view. The user activities are ordered by semantic criteria - edition order, chronological order, text similarity, categories/taxonomy, and heteronym.

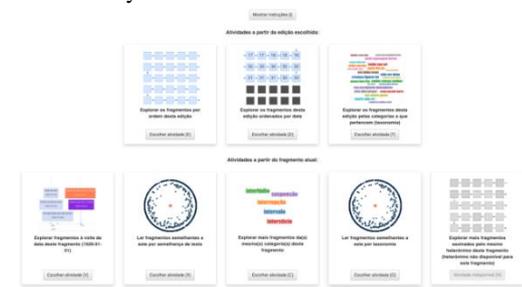


Figure 1 - User activities that can be chosen in the context of a specific fragment, for an edition that does not assign heteronyms to the fragments (the last user activity is not available).

As shown in Figure 1, each user activity is represented by a card with a title, a button with the corresponding keyboard shortcut between straight parentheses that will

take the user to the activity and an image that represents the information visualization technique for that user activity, which can also be clicked to access the user activity.

The image and button will be greyed out if the activity is unavailable - for example, if the currently selected edition has no taxonomy, user activities that use this semantic criteria category will not be available, displaying a message explaining why it is not available in the title and its selection button is also unavailable.

The same semantic criteria can be used with different information visualization techniques. In Figure 1, we can observe that the top 2<sup>nd</sup> and bottom 1<sup>st</sup> cards use different information visualization techniques for chronological order, respectively, custom square map and timeline; the top 3<sup>rd</sup>, bottom 3<sup>rd</sup> and 4<sup>th</sup> cards use the semantic criteria associated with the taxonomies, where top and bottom 3<sup>rd</sup> use the same information visualization technique, word cloud, the former in the context of the selected edition, and the latter in the context of the selected fragment; while the bottom 4<sup>th</sup> uses a different visualization technique, network graph.

#### 4.2.2. Reading View

In this view, the focus is on reading the text of a selected fragment, as represented in Figure 2, where the fragment title and text are presented with a reading progress bar. Since the text will be read on a screen, the used text font is a sans serif font. The fragment title font and text font are the same ones used in one of the LdoD Archive reading sections.



Figure 2 - Plain reading view of a fragment, in which it is possible to navigate to the next and previous fragments using the keyboard arrow keys and in which there is a progress reading bar at the bottom.

Besides the progress reading bar, the reading view features some elements of established browser book readers, see Figure 3. There are previous/next fragment arrow buttons that retreat/advance through the selected edition fragments depending on the currently selected edition. It also allows users to navigate between fragments according to the user selected category or heteronym, using the yellow arrows.

There are some detail features, such as smooth animation between fragment change and smooth page scroll back to the top for the same matter. This smooth page scroll can also be triggered by clicking on its matching button or pressing "T" on the keyboard.



Figure 3 - Reading view of a fragment, in which the menu is visible, the most relevant words are highlighted, and it is possible to navigate to the next/previous fragment in the edition, using the black arrows, and to the next/previous fragment according to the selected user category or heteronym, using the yellow arrows.

When the user stays inactive for more than 3 seconds, everything in this menu disappears with the exception of the reading progress bar and the fragment's title and text, as shown in Figure 2. Inactivity is timed when the user is not interacting with any button, progress bar or previous/next fragment arrow buttons. It is then possible to trigger everything to appear again by doing a mouse hover on any of these elements of the menu. When the user scrolls down the page and the top buttons overlap the text, there is also a transparency trigger. These features are activated as an effort to maximize the user focus on reading the text and are another extra step towards achieving our metaphor implementation.

In this menu, there is a user activity "under the hood" in the sense that it resembles a user activity that does not break the reading flow: by clicking on the matching button ("Highlight the most relevant words"), left button on top (shortcut [R]), the 4 most relevant words are highlighted in blue, as presented in Figure 3. This number was decided after extensive experimentation and fine tuning, trying to find a balance between actual relevant content highlighting and not highlighting too many words to the point it becomes irrelevant to do so. This option makes use of an information visualization technique - text skimming - to encode a semantic criterion - word relevance (more specifically, TF-IDF) -, even though it is not accounted as a user activity in the sense that it can be applied independently of what is the current user's activity.

The reader can also see its own history of reading by pressing the button on the right in the group of three buttons on top (shortcut [H]), as presented in Figure 4.

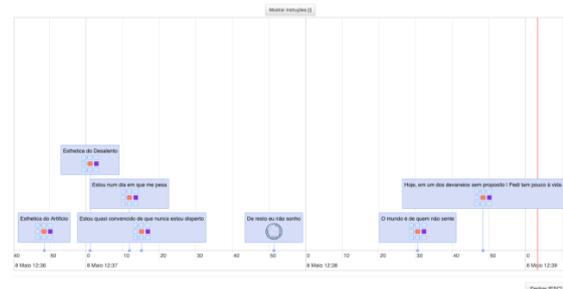


Figure 4 - Reading history view of an edition in which the reader as accessed one fragment (De resto eu não sonho) in the context of a different type of user activity.

In this view, the reader is presented with an interactive timeline that displays the fragments that have been read for the currently selected edition. Each history entry is a fragment, displaying its title and a thumbnail that represents the information visualization technique of the user activity through which that fragment was reached. Besides, having a similar interaction to the user activity that makes use of the timeline in the global view, the user can navigate back to a given fragment that was previously read by clicking on its entry.

The buttons on top right allow the user to see the instructions for this menu and select another edition.

The two first buttons in the group of three buttons allow the user to access the global view associated with the currently selected user activity (shortcut [A]) and to navigate to the new user activity selection menu (shortcut [N]).

#### 4.2.3 Global Views

The global views correspond to the *out* phase of the metaphor and are a combination of semantic criteria with an information visualization technique. In the following subsections we show how the global views are presented in the *LdoD Visual* in the context of the various information visualization techniques.

##### 4.2.3.1. Network Graph

This information visualization technique consists of a network graph with hidden edges, as shown in Figure 5. Each node is a circle that represents a fragment from the currently selected edition. If the user clicks on a circle, she will be taken back to the *Reading Menu* with the corresponding fragment of that circle open to be read.

When accessing the view, the central orange circle represents the fragment in whose context the user activity was started, and the purple circle represents the fragment being currently read. This strategy was adopted so that the reading occurring during a user activity is contextualized by the initially read fragment. The blue circles represent the other fragments from the edition, and their relative distance to the central circle expresses the encoded semantic criterion.

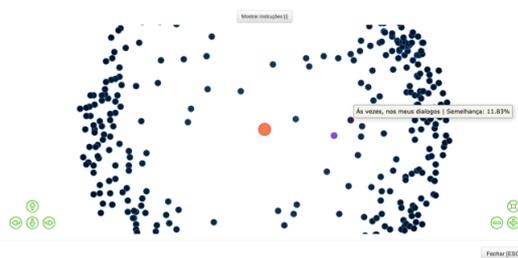


Figure 5 – Global View using a Network Graph for the semantic criteria of text similarity. The closest circle will always be in the same position, be it 5% or 100% similar, to allow have a sparser representation, note that, for instance, when considering text similarity, the most similar fragments are less than 20% similar. Color encodes the similarity in percentage. There is a blue color range - a 100% similar fragment will have its corresponding circle colored in light blue, while a 0% similar fragment will have its corresponding circle colored in dark blue.

The user can put the mouse cursor over the circle so that a label is displayed, showcasing the title and the value of similarity in percentage. Besides being able to drag and zoom the network graph using the mouse pointer and mouse wheel, the user can also use the green navigation buttons to drag, re-size, zoom in and out of the network graph.

This information visualization technique is used to encode similarity by text and taxonomy, which in the case of Figure 5 is by text.

##### 4.2.3.2. Word Cloud

Word clouds are used to encode semantic criteria associated with the categories of a certain fragment or taxonomy (group of categories) of the selected edition, as shown in Figure 6. It displays each category using different colors and the font size varies according to the number of fragments that belong to that category.



Figure 6 - Global view using a Word Cloud for the categories of the taxonomy of the selected edition.

The Word Cloud, and its semantic criteria, can be applied in two cases: (1) the set of categories of the taxonomy of an edition, which corresponds to the case presented in Figure 6; (2) The categories of the current fragment. In both cases, by selecting one of the categories, the user is redirected to a custom squares map that highlights the fragments of the select category, as described in the next subsection.

Heimerl et al. (2014) developed a system that uses word clouds as its central visualization method for interactive text analysis. Its results showed that even though word clouds are aesthetically pleasing with its words disposed in different angles, users do not find them very functional and easy to use. Thus, results showed that participants preferred sequential layouts, where words are placed horizontally without any kind of inclination or angles for aesthetic purposes. For this reason, we decided to follow the same layout in the *LdoD Visual*.

##### 4.2.3.3. Custom Squares Map

This information visualization technique consists of squares linked by arrows, as shown in Figure 7. It is similar to the network graph if we picture the squares as nodes and the arrows as edges. Likewise, each square represents a fragment from the currently selected edition.

This group of squares linked by arrows offers a global view of the currently selected edition, and suggests a certain order derived from the arrows. This order depends

on the encoded criterion. Each square's color and highlighting also depends on the encoded criterion.

Similarly to network graphs, there is a special color highlighting for each square. The orange square corresponds to the fragment where the user activity was initiated, and the purple square to the fragment currently being read.

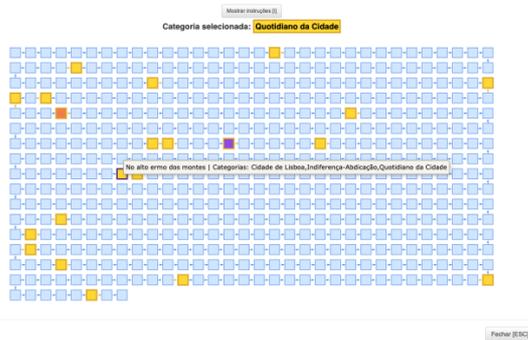


Figure 7 - Global View using a Custom Square Map to show the fragments associated with the category “Quotidiano da Cidade” in the context of the edition fragments.

This information visualization technique is used to encode the following semantic criteria:

- **Edition order.** Fragment squares are sorted by order of the currently selected edition. If the user places the mouse cursor over a square, it will display a label with the fragment title and its position in the edition, almost as if it was a page number. This global view allows users to explore the order of the fragments in the edition.
- **Chronological order.** Fragment squares are sorted by the date in which they were created or published. Each square displays a two-digit number that represents the year from that date - for example, a fragment from 1927 will have its corresponding square with the number 27. Fragments without date have their corresponding squares greyed out. If the user places the mouse cursor over a square, a label will be displayed containing the fragment title and its date. This global view allows users to explore the edition ordered by date.
- **Categories/taxonomy.** Fragments are displayed exactly as if the semantic criterion was the edition's order, with the addition of highlighting in yellow the squares of the fragments from the selected category. If the user places the mouse cursor over a square, a label will be displayed with the fragment title and the categories to which it belongs, as shown in Figure 7. As previously described, smaller yellow previous/next fragment arrow buttons will appear under the black previous/next buttons in the reading menu.
- **Heteronym.** It has the same behaviour as the categories/taxonomy semantic criteria, but it highlights in yellow the squares of the fragments that were signed by a certain heteronym. The yellow arrow buttons are used to move exclusively to the previous/next fragment belonging to the selected

heteronym. The global view allows users to explore the fragments assigned to a particular heteronym.

#### 4.2.3.4. Timeline

This information visualization technique consists of an interactive timeline, in which the user is presented with the timeline centered on the time window around the currently selected fragment, as shown in Figure 8.

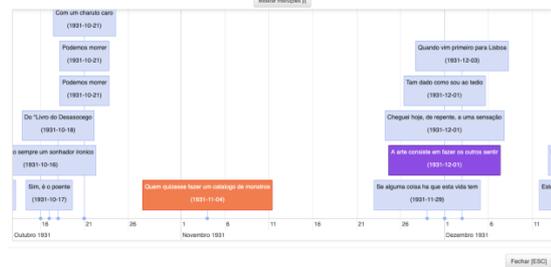


Figure 8 - Global View using a Timeline to represent the chronological order assigned to the fragments.

There is only one global view that uses this information visualization technique to encode the chronological order as the semantic criterion, in order to explore the fragments around the date of the current fragment. Each timeline entry represents a fragment from the currently selected edition. Being consistent with the color scheme for other visualizations, an entry will be orange if it represents the fragment that was initially selected in the context of the user activity, and purple if it is the fragment currently being read.

The user can navigate to a certain fragment by clicking on its entry.

## 5. Results and Discussion

The evaluation of the *LdoD Visual* was done by two types of tests: usability tests, which assess the system's usability, independently of the its particular purpose; and, utility tests, which assess the quality of the system, considering its purpose as a reading tool for the *LdoD*.

### 5.1. Usability Tests

To assess how usable, efficient and satisfying the interaction with the *LdoD Visual* is, we have performed tests with 11 volunteer users who have no expert knowledge on the *LdoD Archive* and on the Book of Disquiet (*LdoD*). The results of this type of test are objective and quantitative.

Each usability test consisted of 4 stages:

- Introduce the user to the *LdoD Visual's* concepts and briefly demonstrate the system in about 5 minutes;
- Let the user explore the system freely for about 5 minutes;
- Ask the user to execute 10 tasks;
- Ask the user to fill and submit a SUS (system usability scale) questionnaire.

In the third stage, for each task, we counted the number of errors and time taken to complete the task. The tasks assess the usability and intuition of the user activities' interaction with the information visualization techniques in order to explore and navigate through the fragments from the selected editions. For instance, to select any fragment from an edition, select the user activity that

makes use of the squares map encoding the edition's order, choose the first fragment, highlight the most relevant words and say out loud what are the most relevant words of the selected fragment, which was, actually, task 9.

These tasks were always executed in the same order for each user, exactly from the same menus.

User	Task #1		Task #2		Task #3		Task #4		Task #5		Task #6		Task #7		Task #8		Task #9		Task #10	
	Time	Errors	Time	Errors																
1	6,1	0	5,89	0	6	0	3,07	0	12,15	0	4,86	0	15,08	0	14,81	0	14,05	0	25	0
2	9,44	1	19,55	0	30	0	24	1	36	0	28	0	50	1	15	0	9,7	0	45,77	0
3	43,77	0	10,1	0	17,5	0	10,22	0	40,7	0	14,61	0	15,78	0	40,79	0	42	1	37	0
4	2,07	0	3,52	0	15,73	1	2,66	0	21,51	0	12,96	0	12,87	0	33,58	0	9,55	0	24,76	2
5	8,79	0	9,14	0	14,13	0	6,02	0	26,65	0	16,4	0	27,28	0	26,16	0	17,8	0	36,76	0
6	8,24	0	10,22	0	13,71	0	7,35	0	17,28	0	11,15	0	12,37	0	24,78	0	11,54	0	30,39	0
7	6,68	0	5,39	0	7,77	0	4,98	0	18,57	0	7,9	0	18,18	0	17,92	0	8,62	0	12,55	0
8	2,81	0	2,43	0	3,7	0	2,38	0	7,06	0	4,6	0	7,03	0	10,38	0	5,44	0	12,53	0
9	4,71	0	5,96	0	7,57	0	5,66	0	20,77	0	9,9	0	20,37	0	17,39	0	4,86	0	22,23	0
10	6,03	0	3,88	0	8,64	0	5,95	0	17,57	0	10,97	0	12,13	0	14,33	0	11,05	0	27,35	0
11	3,76	0	3,53	0	4,33	0	7	0	9,1	0	8,43	0	11,17	0	15,79	0	9,07	0	18,81	0
Average	9,31	0,09	7,24	0,00	11,73	0,09	7,21	0,09	20,67	0,00	11,80	0,00	18,39	0,09	20,99	0,00	13,06	0,09	26,65	0,18

Table 1 - Time to perform each task in seconds and respective number of errors.

Table 1 presents the results from which we conclude that, on average, every task was completed without exceeding the expected time and with basically no errors, with the exception of the last task, which was more difficult. Task 10 was conceived to confirm if users understood why a user activity was not available for certain fragments and what they should do in order to select other fragments that supported that previously unavailable user activity. This was the most complex task, which explains why it took the most time out of all tasks, even though most users ended up figuring out what they had to do.

For the fourth stage, the users filled and submitted a SUS (system usability scale) questionnaire. This is a quick questionnaire that is commonly used to figure out if any aspect was not considered by the tests. This questionnaire involved 10 statements for the user to answer how much he would agree with each one of them in a scale from 1 (completely disagree) to 5 (completely agree):

1. I think I would like to use the *LdoD Visual* regularly.
2. I think the *LdoD Visual* was unnecessarily complex.
3. I think the *LdoD Visual* was easy to use.
4. I think I would need help from a person with technical knowledge to use the *LdoD Visual*.
5. I think that all of the *LdoD Visual* features were well integrated.
6. I think the *LdoD Visual* shows a lot of inconsistency.
7. I think people will learn how to use the *LdoD Visual* easily.
8. I think the *LdoD Visual* was very confusing to use.
9. I felt confident while using the *LdoD Visual*.
10. I needed to learn a lot of new things before I could use *LdoD Visual*.

The score of the SUS questionnaire is obtained by transforming the 1 to 5 scale in a 0 to 4 scale, and converting the values to be consistent, e.g., the odd-number questions express a positive opinion and even-

number questions express a negative opinion. Table 2 presents the results.

User	Question										Sum	SUS Final Score
	1	2	3	4	5	6	7	8	9	10		
1	3	1	3	4	4	1	4	1	3	1	29	72,5
2	5	2	3	3	5	2	4	2	4	1	31	77,5
3	5	1	5	3	5	2	5	1	4	2	35	87,5
4	5	1	4	1	5	1	4	1	5	1	38	95
5	4	1	5	1	5	1	5	1	5	1	39	97,5
6	5	1	5	1	4	1	4	1	4	1	37	92,5
7	5	2	4	2	5	1	4	2	5	1	35	87,5
8	5	1	5	1	5	1	4	1	5	1	39	97,5
9	5	1	5	3	5	1	5	1	4	3	35	87,5
10	5	2	5	1	5	1	4	2	4	1	36	90
11	4	1	4	1	4	1	4	1	4	1	35	87,5
Average	3,6	3,7	3,4	3,1	3,7	3,8	3,3	3,7	3,3	3,7	35,36	88,41

Table 2 - SUS (system usability scale) questionnaire results

Analyzing the SUS score, we can observe that the lowest scoring answer is number 4, *I think I would need help from a person with technical knowledge to use LdoD Visual*. This makes sense as the *LdoD Visual* is relatively complex to users that have no experience with the *LdoD* and the *LdoD Archive*, which is the case of the participants. Regarding the final score, when a system scores a SUS score of 80.3 or more, it is considered to be in the top 10% of scores and the testing users are more likely to recommend it to others. This means that if we are considering a SUS scale, the *LdoD Visual* is a system with good usability, having a SUS score of 88.41.

Besides these 10 SUS questions, we also added questions in order to know about the users. In terms of the characterization of the usability test participants: 9.1% where younger than 18, 36.4% between 18-24, 45.5% between 25-34, and 9.1% between 45-54; and 9.1% have secondary education, 45.5% have a bachelor degree, and 45.5% a master degree. Most users had already used other software to read books in the past and most users did not read the *LdoD*. It is also possible to observe that most users

strongly agreed that they would have a good reading experience using the *LdoD Visual*.

### 5.2. Utility Tests

The purpose of the utility tests is to obtain information about the utility and quality of the experience while using the system. These tests are important in the sense they allow us to assess subjective and qualitative aspects of the *LdoD Visual*, which are difficult to be measured quantitatively.

The tests were done with 3 users that have expert knowledge about *LdoD* and the *LdoD Archive*. The user of the first case study is a postdoctoral researcher with a PhD thesis on *The Book of Disquiet*. This user has also worked, between 2012 and 2015, with the XML encoding of the fragments for the *LdoD Archive*. The user of the second case study is a PhD student from the PhD Program of MATLIT (Materialities of Literature), which is a doctoral program that addresses the material and technological mediations of literary practices. One of its research fields are digital humanities. This user has a BSc and MSc in graphic design. The user of the third case study is another PhD student from the PhD Program of MATLIT, who also has a BSc and MSc in graphic design. This user is interested in the aspects of usability of the *LdoD Archive* since her PhD project involves creative practices through the situated use of the archive. She is one of *LdoD Archive* most regular users and has been organizing a plan of activities to teach *LdoD Archive*'s users how to use it and fully explore its various functionalities, including the creation of virtual editions.

The testing protocol for this type of testing is more flexible in comparison to the usability tests' protocol. The users were contacted by email, introduced to the *LdoD Visual* metaphor, and encouraged to explore the system before we had our actual meeting to do the case study.

During the case study the users were encouraged to use the *LdoD Visual* while "thinking out loud". We asked them several questions regarding the real utility, potential and reading experience that they thought the *LdoD Visual* provides. We also asked them about the differences they felt as more important, in comparison to other tools and electronic readers available, as well as comparing the reading experience of the *LdoD Visual* to the other reading features available in the *LdoD Archive*.

Considering the feedback received from these case studies, we can conclude that the reception of the *LdoD Visual* is positive, both in terms of concept, execution and utility.

The users liked the approach on how to read a book. They thought that it is a good match for the fragmentary idea of the *LdoD*. The attempt to output a visual representation of the editions and fragments brings a new perspective on how to read and explore the *LdoD* and use the *LdoD Archive*, being a relevant work for the area of digital humanities.

The second user also described her experience with user activities in a way that suggested that the *LdoD Visual*'s metaphor is successfully implemented. She stated: "The real reading customization strength comes from the activities around the currently selected fragment. If the

reader stops and tries, for example, to continue her reading path with an user activity that is related with reading similar fragments by text to the currently selected fragment, there is already a reflexive question that is specific to the reader and the navigation method centered around the selected fragment becomes the medium for the user to be able to read around an idea that is interesting to her. In *LdoD Visual*, the reader is able to navigate through the fragments according to what is desired, having a visual reference of the type of navigation and exploration that is being performed in *The Book of Disquiet*."

The users also suggested that the *LdoD Visual* is relatively complex and that there should exist more ways of welcoming non-expert users who are not aware of the problem and structure of the *LdoD* and the *LdoD Archive*, such as "selling the concept" both as a ludic and research tool, and showcasing the different uses and possibilities in videos scattered around the application.

The users also acknowledged the potentiality of the *LdoD Visual* concepts such that they could be developed and expanded to explore other alternatives of interaction with the *LdoD*.

We could observe that each user spent time with different parts and user activities of the *LdoD Visual*. Our analysis on this is that the *LdoD Visual* seems to provide different affordances for different users, and that the objective of customizing the exploration and reading experience of the *The Book of Disquiet* was met.

We can also conclude that the *LdoD Visual* is a tool that mainly offers two different possibilities: the possibility of navigating through the *LdoD* with a focus on the pleasurable assisted user-driven reading experience, and the possibility of navigating with a humanities researcher mindset, finding patterns, correlations and textual data insights. Both of these reading practices are possible resorting to the same engine, the user activities' information visualization techniques, even though the first type of reading practice is the main focus of the *LdoD Visual*.

## 6. Conclusion

*The Book of Disquiet* by Fernando Pessoa is an unfinished book that can be read in any order. The modular nature of this fragmentary literary work brings endless possibilities on how it can be traversed.

The goal of this work was to design and implement a web-based application to read *The Book of Disquiet*, giving the reader an active role while reading and browsing through the book's different fragments, defining user activities that make use of information visualization techniques to encode different semantic criteria.

In order to achieve this, we researched on the use of information visualization techniques applied to textual data and literature, and on the concept and tools for e-books. Afterwards, we defined the *them-out-in* metaphor that was used as the guiding point for whether or not a certain architecture or feature choice met the objectives and served the overall purpose of this work.

According to this metaphor, we defined the user activity concept, which establishes a perspective on how to read

the book. A user activity is a combination of an information visualization technique with a semantic criterion, where the latter is used to classify the fragments and the former to visualize them according to the classification. The solution architecture considers three types of user interfaces: a user interface in which the user activity is selected; a global view in which the fragments are visualized according to the chosen user activity; and a reading view to support the reading and navigation between fragments, according to the chosen semantic criteria.

The final system was tested for usability and utility, and it got positive feedback from both non-expert and expert users of the *LdoD* and the *LdoD Archive*. The usability tests have shown that it is possible to complete the defined tasks in the expected time and without errors, and the utility tests that the proposed metaphor meets the experts' vision of the book and how to explore its multiple paths of reading.

As a web-based application for exploring the process of reading, the *LdoD Visual* contains several features that correspond to what the *LdoD Archive* describes as "the simulation principle", that is, the ability to provide interactors with a reflexive engagement with the textual environment. Through a recursive process of going into the text and coming out of the text, readers are able to explore the *Book of Disquiet* as a particular kind of reading experience and they are also able to see how their acts of engaging with the text become registered in the visualizations. The implemented visualization techniques thus bring together the modularization of the text and the modularization of the reading of the text.

The *LdoD Visual* is a tool that has the potential to grow and become an improved version of what it already accomplishes. One major improvement would be the possibility of saving the reading history state to a user account, preferably in the *LdoD Archive*. This would be a logical extension since the process of reading a book can take days, weeks or months. Another improvement would be developing a feature to generate new virtual editions according to specific reading paths. This would explore the *LdoD Archive* goal of fostering the construction of new editions of the *LdoD*, virtual editions. In its reflexive graphical representation of reading activities, this visual reader contributes to the experimental textual rationale of the *LdoD Archive*.

In order to keep track of the project's evolution, the source code is publicly available in a GitHub repository<sup>7</sup>.

## References

- Alexander, E, and Gleicher, M. (2016). Task-driven comparison of topic models. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):320–329.
- Brath. R., and Banissi, E. (2014). Using font attributes in knowledge maps and information retrieval. *Proceedings of the First Workshop on Knowledge Maps and Information Retrieval (KMIR 2014)* co-located with International Conference on Digital Libraries (DL 2014), pp. 23–30.
- Brath. R., and Banissi, E. (2015). Using text in visualizations for micro/macro readings. *4th Workshop on Visual Text Analytics*. ACM, 2015.
- Jänicke, S.; Franzini, G.; Cheema, M.F.; Scheuermann, G. On close and distant reading in digital humanities: A survey and future challenges. In *Proceedings of the Eurographics Conference on Visualization (EuroVis) 2015*, Cagliari, Italy, 25–29 May 2015.
- Heimerl F., Lohmann, S., Lange, S., and Ertl, T. 2014. Word cloud explorer: Text analytics based on word clouds. *47th Hawaii International Conference on System Sciences*, Jan 2014, pp. 1833–1842.
- Koolen, C., Garnett, A., Siemens, R. 2012. Electronic Environments for Reading: An Annotated Bibliography of Pertinent Hardware and Software (2011). *Scholarly and Research Communication*: 3(4).
- Mohammad, S. M. 2012. From once upon a time to happily ever after: Tracking emotions in mail and books. *Decision Support Systems*, 53(4):730 – 741
- Portela, Manuel and Rito Silva, António. A model for a virtual LdoD. 2014. *Digital Scholarship in the Humanities* September 1, 2015 30(3):354. <http://dx.doi.org/10.1093/llc/fqu004>
- Rito Silva, António and Portela, Manuel, « TEI4LdoD: Textual Encoding and Social Editing in Web 2.0 Environments », *Journal of the Text Encoding Initiative* [Online], Issue 8 | December 2014 - December 2015, Online since 09 April 2015, connection on 15 November 2018. URL : <http://journals.openedition.org/jtei/1171> ; DOI : 10.4000/jtei.1171
- Singhal, A. (2001). Modern Information Retrieval: A Brief Overview. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* 24 (4): 35–43.
- Viégas, F.B., Wattenberg, M., and Feinberg, J. 2009. Participatory visualization with wordle. *IEEE Transactions on Visualization & Computer Graphics*, 15:1137–1144. [doi.ieeecomputersociety.org/10.1109/TVCG.2009.171](https://doi.org/10.1109/TVCG.2009.171)
- Vilaplana, J. N., and Pérez-Montoro, M. (2017). Diggersdiaries: Using text analysis to support exploration and reading in a large document collection. *EuroVis 2017 - Posters*, A. P. Puig and T. Isenberg, Eds. The Eurographics Association.

---

<sup>7</sup><https://github.com/socialsoftware/edition>