LAMP-TR-062
CAR-TR-959
CS-TR-4212

February 2001

The architecture of TRUEVIZ: A groundTRUth/metadata Editing and VIsualiZing toolkit

Chang Ha Lee and Tapas Kanungo

The architecture of TRUEVIZ: A groundTRUth/metadata Editing and VIsualiZing toolkit

Chang Ha Lee and Tapas Kanungo

Language and Media Processing Laboratory
Center for Automation Research
University of Maryland
College Park, MD 20742
{chlee,kanungo}@cfar.umd.edu

Abstract

Tools for visualizing and creating groundtruth and metadata are crucial for document image analysis research. In this paper we describe TrueViz [LK00, KLCB01], which is a tool for visualizing and editing groundtruth/metadata. We first describe the groundtruthing task and the requirements for any interactive groundtruthing tool. Next we describe the system design of TrueViz and discuss how a user can use it to create groundtruth. TrueViz is implemented in the Java programming language and works on various platforms including Windows and Unix. TrueViz reads and stores groundtruth/metadata in XML format, and reads a corresponding image stored in TIFF image file format. Multilingual text editing, display, and search modules based on the Unicode representation for text are also provided. This software is being made available free of charge to researchers.

This research was funded in part by the Department of Defense under Contract MDA0949-6C-1250, Lockheed Martin under Contract 9802167270, the Defense Advanced Research Projects Agency under Contract N660010028910, and the National Science Foundation under Grant IIS9987944.

1 Introduction

In the document image analysis (DIA) research area, the term 'groundtruth' refers to various attributes associated with the text on the image — bounding box coordinates of words, lines, characters; font type; character size; direction of text; etc. Groundtruth data is crucial for document image analysis because it is impossible to train and test Optical Character Recognition (OCR) algorithms without it. Since groundtruth is created manually in most cases, tools for annotating and visualizing groundtruth are very important. In fact, at the MLOCR99 international workshop [mlo99] the consensus in the corpus working group was that our community needs i) a protocol for groundtruthing documents, ii) an XML-based groundtruth representation format, iii) a public-domain multilingual/multiplatform visualization and data-entry tool, and iv) a consortium for managing and distributing datasets.

In this paper we address two of the four issues raised by the working group: i) We describe an XML-based groundtruth representation format, and ii) we describe TrueViz, which is a public domain[1] annotation tool that we have developed at the University of Maryland.

This paper is organized as follows. In Section 2 we describe various existing annotation tools used in document image analysis and in related areas such as speech recognition, linguistics, and information retrieval. The desirable features of a document image groundtruthing tool are described in Section 3. In Section 4 we discuss design and implementation issues related to editing, visualization, and search. The XML data format for groundtruth is discussed in Section 5, where we also provide representative samples of XML files. The multilingual data entry, visualization, and search features of TrueViz are quite unique and are discussed in Section 6. Finally, in Section 7 we list the things that we hope the international DIA community will add to the public domain system.

2 Previous Work

There are many annotation and visualization tools in various domains. In this section we describe a few annotation tools commonly used in document image analysis, speech recognition, linguistics, information retrieval, video analysis, geographic systems, and statistics. In Table 1 we provide a comparison of these tools.

2.1 Document Image Visualization Tools

Visualization tools for displaying or editing a document image and groundtruth metadata have been developed for evaluating algorithms, creating document groundtruth, or browsing documents.

Pink Panther [YV98] is an environment for creating segmentation groundtruth files and for page segmentation benchmarking. Page segmentation is the process of decomposing a document page image into structural and logical units, such as images, paragraphs, headlines, tables, etc. The performance of a page segmentation algorithm is evaluated

---

[1] TrueViz is available at `http://www.cfar.umd.edu/~kanungo/software/software.html`

Table 1: Comparison of Visualization Tools

Name	Platform	Data Format	Domain
PinkPanther	Unix/X Windows System	ASCII	Document Image Groundtruth
Illuminator	Unix/X Windows System	DAFS	Document Image Groundtruth
Oulu Database Browser	Multi-Platform/Java	ASCII	Document Image Groundtruth
TrueViz	Multi-Platform/Java	XML Format	Document Image Groundtruth
Transcriber	Unix/Windows NT	XML Format	Speech Annotation
ATLAS	Unix/Windows NT	XML Format	Linguistic Annotation
Alembic Workbench	Unix system	SGML/PTF Format	Linguistic/Named Entities Annotation
ViPER	Multi-Platform/Java	ASCII	Video Sequence Groundtruth
XGobi	Unix/X Windows System	S Data Format/ASCII	Statistical Data
S-PLUS	Windows 95/98	Customized Data	Statistical Data
CLASP	Unix/Macintosh	Commonly Used Formats	Statistical Data
Mondrian	Multi-Platform/Java	ASCII/Databases	Categorical/ Geographical Data
PolyPaint+	SunOS/Solaris	netCDF	Geographical Data
Spotfire	MS Windows	Database/Spreadsheet/ ASCII	Decision Making by Data Analysis
Slicer Dicer	MS Windows	Binary/ASCII/ Commonly Used Formats	Medical/Scientific Data Defined on Grids

by running the algorithm on a set of document images, and comparing the output for each document to corresponding groundtruth metadata. Pink Panther consists of two parts: Grounds-Keeper and Cluzo. Grounds-Keeper is a tool for creating groundtruth metadata. It visualizes a document image and the corresponding metadata, and also allows users to zone the document image and specify the information for each zone. Groundtruth metadata created by Grounds-Keeper is stored in an ASCII file format. Cluzo is a benchmarking tool for collecting the locations, types and severities of segmentation errors on a page as well as information on segmentation performance. Pink Panther is implemented on the Unix and X Windows platforms and is written in C. While Grounds-Keeper allows the user to enter segmentation groundtruth, entering text groundtruth is not possible.

Illuminator [Fru95] is an editor developed by RAF Technology, Inc. for building document understanding test and training sets, for correction of OCR (Optical Character Recognition) errors, and for reverse-encoding the essential information and attributes of a document. Illuminator visualizes or edits a document image and its entities, which are specific regions of the image and the associated metadata. It is configured to handle text in major European languages and Japanese. Illuminator uses the DAFS (Document Attribute Format Specification) file format [Fru95] to store the document image and metadata. DAFS provides a format for breaking down a document into entities which have hierarchical structure, and for defining entity boundaries and attributes. Illuminator is implemented on the Unix and X Windows platforms and is written in C.

The MediaTeam Oulu Document Database [SK98] is a collection of scanned documents with corresponding groundtruth for the physical and logical structure of the documents. It was developed by the University of Oulu MediaTeam. The document database browser is a visualization tool for exploring the contents of the database. The browser is written in the Java programming language and allows visualization of document images and corresponding metadata simultaneously. The browser can explore the database and select particular documents for visualization. The browser also provides a window to list attributes of the document. Document images which were originally stored in TIFF image format are stored in JPEG image format and metadata is stored in an ASCII file format.

Pink Panther and Illuminator work only on the Unix platform. Because there are many tools that are executable only on the Windows platform, this is a limitation. The Oulu document database browser is written in the Java programming language, and can be run on various platforms. However, the Oulu document database supports JPEG image format only, while TIFF is the most popular image format for document images. Furthermore, the file representation of the groundtruth is non-standard. In fact, all the above tools store document metadata in their own file formats. To provide data compatibility, a standard file format, or a file format to which other file formats can be easily converted, is needed.

A prototype system for visualizing and editing groundtruth is currently being built at the University of Fribourg, Switzerland [HRI00]. This system allows users to edit the hierarchical structure of the document. However, the system does not provide a compatible OCR evaluation package to visualize OCR segmentation results.

2.2 Other Visualization Tools

We surveyed visualization tools in other data domains to find out the best way to provide multi-platform and data compatibility. In this section we summarize features of visualization tools in various domains such as statistical, categorical, geographical, and medical data as well as linguistic data and speech signals.

Transcriber [BGWL00, GBBW00, BGWL98] is a tool for segmenting, labeling and transcribing speech signals. It supports most common audio formats and stores the transcription in XML format. It was developed in the Tcl/Tk and C programming languages, and works on Unix and Windows NT platforms.

ATLAS [BDH$^+$00] is an architecture and tool for linguistic analysis based on a formal model for annotating linguistic artifacts. It uses an XML-based ATLAS Interchange Format (AIF) for storing annotated corpora, and was developed in the C++, Perl, Tcl/Tk and Java programming languages.

Alembic Workbench [DAH$^+$97] is a new set of integrated tools that uses a mixed-initiative approach to bootstrapping the manual tagging process with the goal of reducing the overhead associated with corpus development. The Alembic Workbench is developed using the Tcl/Tk, Perl, C and Lisp programming languages, and works on the Unix platform. Alembic uses the SGML and PTF (Parallel Tag File) formats for source text and annotations.

ViPER (Video Processing Evaluation Resource) [DM00] consists of three main components: ViPER-GT, ViPER-PE, and ViPER-Viz. ViPER-GT contains modules for configuring and producing groundtruth information which describes a video sequence. The ViPER-PE module provides performance evaluation capabilities for comparing computed results with appropriate groundtruth information. ViPER-Viz enables a user to visualize groundtruth, analysis results, performance evaluation results, or an entire video clip. ViPER was developed in the Java programming language, and groundtruth and results are stored in ASCII file format.

XGobi [SCB98, SHB91, SCB92] is an X Window application for interactively exploring statistical data. Its current functionalities include brushing, identification, and editing of connected lines, as well as rotation and the grand tour, with several interactive projection pursuit indices. Several functions can be linked so that actions in one window are promptly reflected in another.

S-PLUS [VR99] is a desktop data analysis tool that provides data analysis and visualization capabilities to identify trends in data. It allows data import and export from spreadsheets such as Excel, as well as from a wide range of relational and other data sources.

The Common Lisp Analytical Statistics Package (CLASP) [AWC$^+$95] is a tool for visualizing and statistically analyzing data. CLASP provides an interactive environment for data manipulation and statistical analysis and a variety of descriptive and hypothesis-testing statistics. It includes many features that facilitate exploratory data analysis.

Mondrian [Uni] is a data-visualization system written in Java. Its main emphasis is on visualization techniques for categorical data and geographical aata. Mondrian provides various plots such as mosaic plots, maps, barcharts, and parallel coordinates, which are fully linked and allow various interrogations.

PolyPaint+ [Nat] is an interactive scientific visualization tool that displays complex structures within three-dimensional data fields. It provides color shaded-surface display, as well as simple volumetric rendering in either index or true color. PolyPaint+ routines first compute the polygon set that describes a desired surface within the 3D data volume, and these polygons are then rendered as continuously shaded surfaces. Objects rendered volumetrically may be viewed along with shaded surfaces. Additional data sets can be overlaid on shaded surfaces by color coding the data according to a specified color map.

Spotfire [AS94] is a decision analysis workspace that uses the connectivity of the Web to provide a workspace in which to access large amounts of complex data from wherever it resides, to visually explore and analyze the data, and to share results.

Slicer Dicer [PIX] provides tools for analysis, interpretation and documentation of complex data defined in three or more dimensions. It helps in exploring the data visually by "slicing and dicing" to create arbitrary orthogonal and oblique slices, rectilinear blocks and cutouts, isosurfaces, and projected volumes. It also provides animation sequences featuring continuous rotation, moving slices, blocks, parametric variation (time animation), oblique slice rotation, and varying transparency.

A more detailed review and taxonomy of visualization tools can be found in an article by Shneiderman [Shn96], and a good general reference for user interfaces is Shneiderman's book [Shn98].

3 Desired GUI Functionalities

Since TrueViz will be used by different researchers for different tasks, we first summarize the functionalities that are desired of such a tool. The simplest task that the tool could be used for is to visualize and input multilingual text. Next, it could be used to mark regions of a scanned document image as text or graphics, and assign labels to regions. A researcher wanting to look at the results obtained by a DIA system might want to search for all the incorrectly recognized characters and then zoom into the image at those locations. A researcher interested in extracting the logical structure of a document might want to label the reading order of the text areas, or the hierarchy of the text regions corresponding to sections and subsections.

After studying the various tasks for which a user might want to use the to-be-designed tool, we formulated the following set of requirements for the graphical user interface:

Entities: Users should be able to visualize and edit zone-, line-, word-, and character-level geometric groundtruth. Furthermore, they should be able to establish their own entity structure. For each entity, they should be able to define attributes (e.g. bounding boxes) and specify their values.

Scale: Users should be able to zoom in and out of the image and overlaid groundtruth so that they can study the image and OCR error results at the page, paragraph, line, word, or character level.

Color: It should be possible to display entities that have different attributes in different colors. For example, image zones could be shown in one color and table or text zones in another. Thus if a DIA system incorrectly recognizes a table zone as an image zone, the error would be easily identifiable from the color coding.

Logical information: The visualization tool should allow users to visualize and edit the logical reading order of text zones, and also to specify the hierarchy of the text zones. For example, it should be possible to visually specify that a subsection is contained in a section.

Multilingual Visualization: Since DIA systems are being developed for various languages and scripts, users should be able to visualize groundtruth text in these languages and scripts. The use of a standard encoding such as Unicode is highly desirable.

Multilingual Data Entry: While regular English text can be entered by regular keyboards, keyboard mappings that allow other languages and scripts to be entered should also be available.

XML-based Representation: The XML markup language would be ideal for representing page layout groundtruth since it is the current industry standard and various parsers, syntax checkers and editors are publicly available for it.

Converters: Converters to convert standard datasets such as the University of Washington dataset (in DAFS format) into the XML representation would help bootstrap research by providing seed datasets.

Search: Users should be able to search for strings in the groundtruth and find the locations where they appear in the image. The search module should work in any language and users should be able to specify edit distances for approximate searching, which is essential when searching for strings in noisy OCR text.

Evaluation: The tool should have a built-in OCR evaluation module or should be compatible with one, so that users are able to visualize OCR evaluation results easily.

Multiplatform: Since researchers and data entry persons work on various platforms such as UNIX, PC and Mac, the tool should be platform-independent so that users need not spend time learning how to use it on a platform that they are not familiar with.

Public Domain: In order for the community to take full advantage of it, the tool should be freely available.

4 Design and Implementation

4.1 Overview

The TrueViz display is vertically split into two panels (see Figure 1). The left panel is an image panel for displaying a document image and corresponding geometric metadata, and the right panel is a tree view for displaying textual metadata structure.

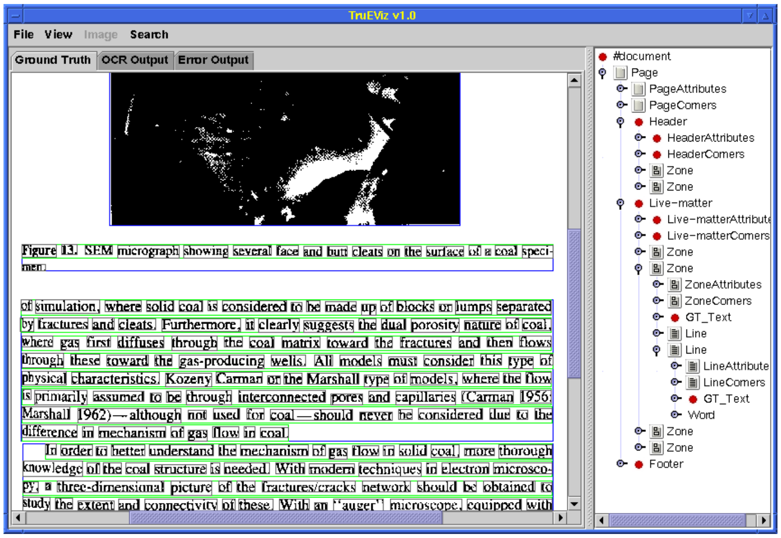

Figure 13. SEM micrograph showing several face and butt cleats on the surface of a coal specimen.

of simulation, where solid coal is considered to be made up of blocks or lumps separated by fractures and cleats. Furthermore, it clearly suggests the dual porosity nature of coal, where gas first diffuses through the coal matrix toward the fractures and then flows through these toward the gas-producing wells. All models must consider this type of physical characteristics. Kozeny Carman or the Marshall type of models, where the flow is primarily assumed to be through interconnected pores and capillaries (Carman 1956; Marshall 1962)—although not used for coal—should never be considered due to the difference in mechanism of gas flow in coal.

In order to better understand the mechanism of gas flow in solid coal, more thorough knowledge of the coal structure is needed. With modern techniques in electron microscopy, a three-dimensional picture of the fractures/cracks network should be obtained to study the extent and connectivity of these. With an "auger" microscope, equipped with

Figure 1: TrueViz consists of an image panel (left) and a tree view (right).

The image panel displays a document image and overlays geometric metadata on the image. Currently, three kinds of geometric metadata can be visualized: Bounding boxes, logical relationships, and an Infopanel. The bounding box of an entity is visualized as a polygon whose color represents the type of the entity. "Logical relationship" refers to logical reading order, and is visualized using an arrow from one entity to the next. The Infopanel is a small window for displaying a few important attributes of the entity. The image and metadata visualization can be scaled to various resolutions.

The tree view displays the XML-based groundtruth metadata in a tree structure of expandable and collapsible nodes. The attribute values can be edited in the tree nodes and the groundtruth text can be edited in the separate multilingual text editor.

4.1.1 Metadata Visualization

Entities can be classified into four categories: Zones, Lines, Words and Characters. Entities are hierarchical in nature, so a Zone is contained within a Page, a Line is contained within a Zone, a Word is contained within a Line, and a Character is contained within a Word. Because of the hierarchical nature of the entities, it is necessary to change views in order to view specific portions of the structure. There are five views: Image Only, Page, Zone, Line, Word and Character. The Image Only view shows only the image without any groundtruth visualization. The Page view shows metadata for all entities, from the highest level to the lowest level. This view is not editable or selectable. The Zone view shows only Zone metadata. A Zone's data can be accessed by clicking on the Zone. This causes the Zone to be active (selected) and highlighted, and the Infopanel to pop up. The Infopanel is a small window for displaying important metadata for the active entity (see Figure 8). The corresponding node in the tree view will also be selected. Similarly, the Line view shows all Line metadata (see Figure 2 (a)), the Word view shows all Word metadata (see Figure 2 (b)), and the Character view shows all Character metadata. As

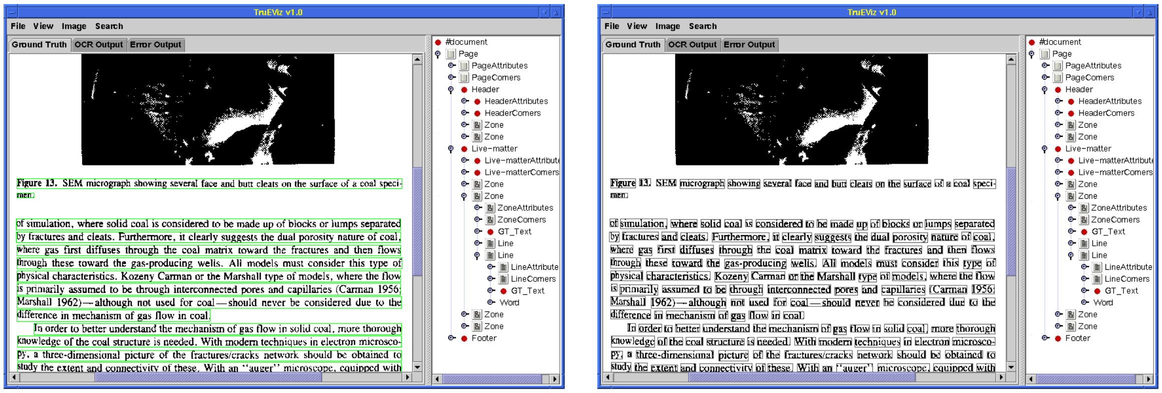

Figure 13. SEM micrograph showing several face and butt cleats on the surface of a coal specimen.

of simulation, where solid coal is considered to be made up of blocks or lumps separated by fractures and cleats. Furthermore, it clearly suggests the dual porosity nature of coal, where gas first diffuses through the coal matrix toward the fractures and then flows through these toward the gas-producing wells. All models must consider this type of physical characteristics. Kozeny Carman or the Marshall type of models, where the flow is primarily assumed to be through interconnected pores and capillaries (Carman 1956; Marshall 1962)—although not used for coal—should never be considered due to the difference in mechanism of gas flow in coal.

In order to better understand the mechanism of gas flow in solid coal, more thorough knowledge of the coal structure is needed. With modern techniques in electron microscopy, a three-dimensional picture of the fractures/cracks network should be obtained to study the extent and connectivity of these. With an ''auger'' microscope, equipped with

(a) Line view displays all Line entities. (b) Word view displays all Word entities.

Figure 2: Hierarchical display.

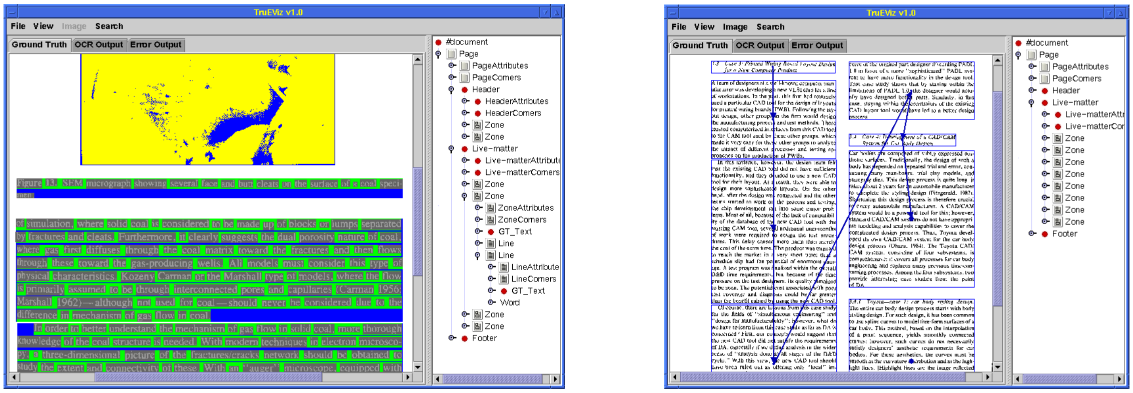

(a) Fill Bounding Boxes. (b) Logical Relations.

Figure 3: View options

in the Zone view, metadata can be selected, and the Infopanel for the active entity is popped up.

There are two options for views: 'Fill Bounding Boxes' and 'Logical Relations'. If the 'Fill Bounding Boxes' option is checked, all entities are painted in colors corresponding to their types (see Figure 3 (a)). Otherwise, entities are displayed using polygonal outlines whose colors also represent their types. This option is useful when the document is displayed at a large scale, because a user can see the type of an entity from its color even if the bounding box is too large to fit on the screen. If the 'Show Logical Relations' option is selected, the logical reading order relations are visualized using arrows from each entity to the next logical entity (see Figure 3 (b)).

4.1.2 Metadata Editing

Groundtruth metadata can be edited in two ways: graphical editing and text editing. All metadata can be edited within the attribute value node in the tree view. Because the

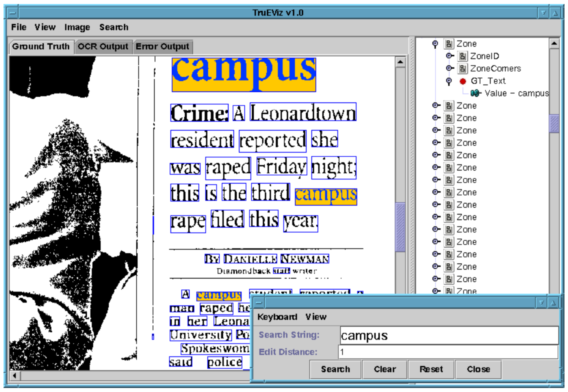

Search String: campus

Edit Distance: 1

Figure 4: Search string "campus" and edit distance 1 are specified in the search window (in the lower right corner), and the matching entities are highlighted.

groundtruth text may contain multilingual text, it is edited in the separate multilingual text editor. The metadata visualized in the image panel can also be edited graphically. It is very difficult to correct bounding boxes of entities by editing their coordinates. Therefore, TrueViz enables users to change the coordinates of bounding boxes graphically. In addition to the bounding boxes, the logical relationships can be changed graphically. The image panel can also be used used to create and delete entities.

4.2 Search

TrueViz provides a multilingual approximate search functionality. A search string and edit distance can be specified in the search window. TrueViz provides multilingual input for a search string. The edit distance is the minimum number of substitutions, insertions and deletions required to transform one string into another. The maximum edit distance allowed during the search can be specified [Gus97]. After the search is finished, all entities containing the search string within the specified edit distance are highlighted (see Figure 4).

5 The Data Format

5.1 Overview

Groundtruth metadata is stored in XML file format [Har99, BPSM98, WHA+99, McL00] (see Figure 5), and document images are stored in TIFF image file format. The tree view reflects the XML data file, and an internal data structure is created to visualize the groundtruth metadata. The internal data structure consists of Region of Interest (ROI) nodes. A ROI is a generic term used to describe any area of the image that the user deems of interest. The internal data forms a directed acyclic graph with ROIs as nodes and hierarchical or logical links as edges.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE Page SYSTEM "Trueviz.dtd">
<Page>
        <PageID Value="P000"> </PageID>
        <PageType Value="Journal"> </PageType>
        <PageNumber Value="1"> </PageNumber>
        <PageColumns Value="1"> </PageColumns>
        <Font Size="9-12" Spacing="Undefined" Style="Normal" Type="Serif"> </Font>
        <Zone>
                <ZoneID Value="Z000"/>
                <ZoneNext Value="Z001"/>
                <CharacterOrientation Type="String" Value="up-right"/>
                <DominantFontSize Type="String" Value="9-12"/>
                <DominantFontSpacing Type="String" Value="proportional"/>
                <DominantFontStyle Type="String" Value="plain"/>
                <DominantFontType Type="String" Value="serif"/>
                <Language Type="String" Value="English"/>
                <TextAlignment Type="String" Value="justified"/>
                <TextReadingDirection Type="String" Value="left-right"/>
                <ZoneCorners>
                        <Vertex x="1281" y="3136"></Vertex>
                        <Vertex x="1296" y="3136"></Vertex>
                        <Vertex x="1296" y="3169"></Vertex>
                        <Vertex x="1281" y="3169"></Vertex>
                </ZoneCorners>
                <GT_Text Value="a"></GT_Text>
                <Line>
                        <LineID Value="Z000L000"/>
                        <LineCorners>
                                <Vertex x="1281" y="3136"></Vertex>
                                <Vertex x="1296" y="3136"></Vertex>
                                <Vertex x="1296" y="3169"></Vertex>
                                <Vertex x="1281" y="3169"></Vertex>
                        </LineCorners>
                        <GT_Text Value="a"></GT_Text>
                        <Word>
                                <WordID Value="Z000L000W000"/>
                                <WordCorners>
                                        <Vertex x="1281" y="3136"></Vertex>
                                        <Vertex x="1296" y="3136"></Vertex>
                                        <Vertex x="1296" y="3169"></Vertex>
                                        <Vertex x="1281" y="3169"></Vertex>
                                </WordCorners>
                                <GT_Text Value="a"></GT_Text>
                                <Character>
                                        <CharacterID Value="Z000L000W000C000"/>
                                        <CharacterCorners>
                                                <Vertex x="1281" y="3136"></Vertex>
                                                <Vertex x="1296" y="3136"></Vertex>
                                                <Vertex x="1296" y="3169"></Vertex>
                                                <Vertex x="1281" y="3169"></Vertex>
                                        </CharacterCorners>
                                        <GT_Text Value="a"></GT_Text>
                                </Character>
                        </Word>
                </Line>
        </Zone>
        <Zone>
                <ZoneID Value="Z001"/>
                <ZoneNext Value=""/>
                <CharacterOrientation Type="String" Value="up-right"/>
                <DominantFontSize Type="String" Value="9-12"/>
                <DominantFontSpacing Type="String" Value="proportional"/>
                <DominantFontStyle Type="String" Value="italic"/>
                <DominantFontType Type="String" Value="serif"/>
                <Language Type="String" Value="English"/>
                <TextAlignment Type="String" Value="justified"/>
                <TextReadingDirection Type="String" Value="left-right"/>
                <ZoneCorners>
                        <Vertex x="2281" y="3136"></Vertex>
                        <Vertex x="2296" y="3136"></Vertex>
                        <Vertex x="2296" y="3169"></Vertex>
                        <Vertex x="2281" y="3169"></Vertex>
                </ZoneCorners>
                <GT_Text Value="b"></GT_Text>
        </Zone>
</Page>
```

Figure 5: An example XML file.

5.2 XML Data Format

The groundtruth data is organized in a hierarchical structure. The highest-level and therefore most inclusive entity is the Document. A Document is, in its simplest form, a collection of individual units, known as Pages, which are related to or support a specific topic or purpose (e.g. a report or manual). A Page is the next level down in the hierarchy and represents individual units of a Document. Each Page has an associated image that represents the original hard copy. A Page contains one or more Zones. A Zone is usually a rectangular area definable by its horizontal and vertical coordinates within a page. The purpose of a Zone is to identify a key area of the page such as title, heading, graphic, page number, etc. Each Zone may contain one or more Lines. A Line is an individual line of text. A Line can be broken down into one or more Words, each of which may contain one or more Characters. Each tag in the XML file represents an entity or attribute. An entity name can be any alphanumeric word, but the only entities that can be graphically edited in TrueViz are Zone, Line, Word and Character.

An entity's attributes can be listed under the entity's tag in the XML file. While any attribute name can be listed, some built-in attributes are crucial for the visualization of groundtruth data.

ID: ID is the identification of the entity. The attribute name for ID is combined with the entity name. For example, the ID of a Zone entity is represented as ZoneID, and similarly we use LineID for Line, WordID for Word, and CharacterID for Character.

Corners: Corners represent the bounding box of the entity. The upper left, upper right, lower right, and lower left vertices are listed inside a Corners tag in order. Like the ID, the attribute name is combined with the entity name.

Next: Next stores the ID of the logically following entity. As with the ID, the attribute name is combined with the entity name.

GT_Text: GT_Text stores the groundtruth text of the entity.

The following example shows a simple entity.

```
<Zone>
  <ZoneID Value="Z001"/>
  <ZoneNext Value="Z002"/>
  <GT_Text Value="Hello, world">
  <ZoneCorners>
    <Vertex x="10" y="10"/>
    <Vertex x="100" y="10"/>
    <Vertex x="100" y="30"/>
    <Vertex x="10" y="30"/>
  </ZoneCorners>
<Zone>
```

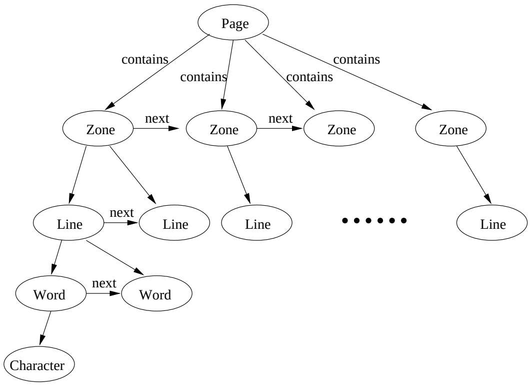

Figure 6: Entity structure.

5.3 Internal Data Structure

The groundtruth metadata is stored in XML file format, which is essentially a tree. The entities, on the other hand, form a directed acyclic graph structure. Each entity contains child entities and has a next logical entity. The graph representing the entity structure can be expressed by equation (1) (see Figure 6):

$$G = (V, E) \quad \text{where} \quad V = \{Zone, Word, Line, Character\},\ E = \{contains, next\} \tag{1}$$

Because of the difference between the entity structure and the XML structure, TrueViz has an internal data structure that is a little different from the XML structure. The internal data structure consists of Region of Interest (ROI) nodes, and the ROIs form a directed acyclic graph as described in equation (1). A next logical entity is stored as an attribute of an entity in the XML file, and is converted into a link from a ROI to the next ROI in the internal data structure.

For parsing XML files and converting XML structures into internal structures, Java APIs (Application Program Interfaces) were used. Two kinds of Java APIs can be used for XML parsing: SAX (Simple API for XML) and DOM (Document Object Model) [McL00]. SAX is an event-based framework for parsing XML data. It reads through the XML document, breaks down the data into usable parts, and defines the events that occur at each step of the process. DOM provides a data representation of an XML document as a tree, which can be traversed and manipulated. A DOM parser was used in TrueViz because TrueViz has an internal data structure that needs to be kept in memory.

5.4 Flexible Entity Structure

Various entity hierarchies are used [AS99, KA99, dli99], depending on the type of document. Users may want to build document metadata using their own structures. TrueViz provides flexible entity structures so that users can build their own entity structures and DTD (Document Type Definition) files for defining and verifying the structures of their

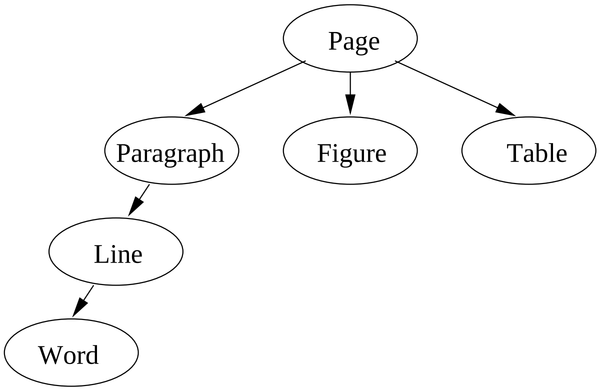

Figure 7: User-defined entity structure.

XML files. The entity structure is extracted from the XML file, and the DTD file can be used to verify that the XML file conforms to the corresponding entity structure. The DTD file can be created and edited using any existing public domain editor.

If an element has an "Entity" attribute and its value is "True," the element is recognized as an entity when the XML file is parsed. An entity structure for an XML file is automatically built by TrueViz from the recognized entities and their level information. The following is an example XML file with a user-defined entity structure, and Figure 7 is the entity structure extracted from the XML file. If there are no elements with the attribute "Entity," the default entity structure (see Section 5.2) is used.

```
<Page>
  <Paragraph Entity="True">
    <Line Entity="True">
      <Word Entity="True">
        <Character Entity="True">
        </Character>
      </Word>
    </Line>
  </Paragraph>
  <Figure Entity="True">
  </Figure>
  <Table Entity="True">
  </Table>
</Page>
```

6 Multilingual Features

6.1 Multilingual Text Data

Java programs running on JDK1.1 or JDK1.2 can display any Unicode [Con97] character which can be rendered with a host font. TrueViz displays multilingual text using Java Unicode facilities (see Figure 8). TrueViz can read Unicode characters from the XML

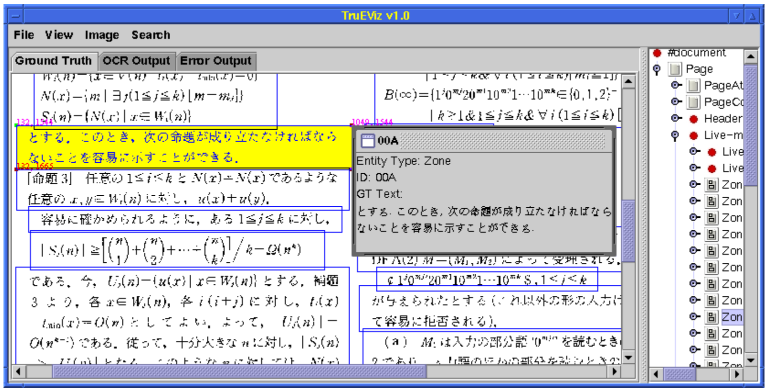

00A

Entity Type: Zone

ID: 00A

GT Text:

とする. このとき, 次の命題が成り立たなければならないことを容易に示すことができる.

Figure 8: Infopanel and multilingual display.

file, and saves the XML file in the Unicode UTF8 format [Con97]. However Java does not provide a multilingual input method. We therefore developed such a method, which is described in Section 6.2.

6.2 Multilingual Input System

TrueViz provides a multilingual input system. Some languages like Chinese, Japanese, or Korean use more characters than can be input by a regular keyboard. To handle such languages, a sequence of several characters needs to be typed to construct a single character. While this composition process is going on, the input system accepts the sequence of characters, and produces composed text and committed text. The composed text is the intermediate text which is being processed to produce the intended text. The final text is called committed text (see Figure 10). Input capabilities for various languages can be easily added using this common interface. In addition to the default English language input, Russian input is also currently implemented. The input system can be used anywhere multilingual text input is needed (see Figure 9). For example, TrueViz supports multilingual text input in the search window for multilingual search. For people who are not familiar with the keyboard mapping, TrueViz provides a keyboard mapping display. In addition to keyboard input, TrueViz provides Unicode character input using a code table, so that any Unicode character can be selected and inserted into a text.

6.3 Adding New Input Capabilities

Input capabilities for various languages can be easily added using the common interface *did.gui.DIDInputMethod*. A new input capability can be added by implementing the following member functions of the interface.

public DIDKeyBDisplay getKeyBDisplay(): The function for getting the keyboard mapping display.

public String getComposingText(): The function for getting the current composed text.

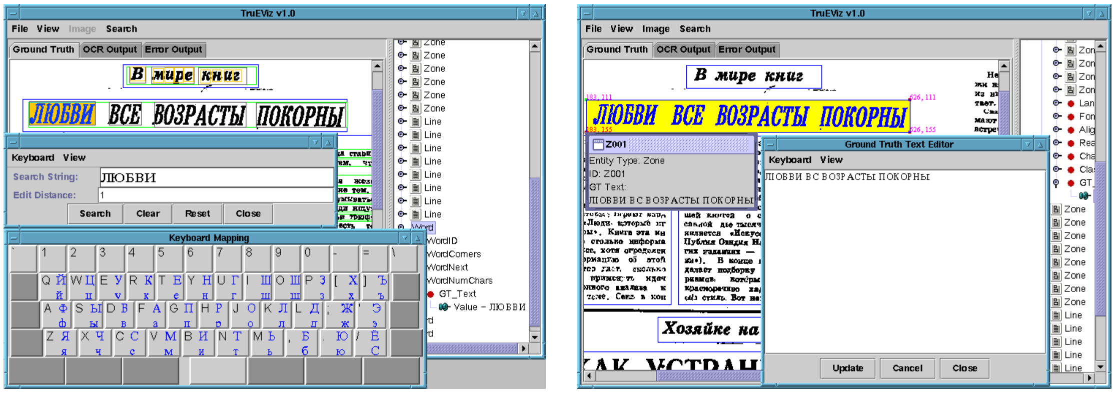

Search String: ЛЮБВИ

Edit Distance: 1

Value – ЛЮБВИ

Z001

Entity Type: Zone

ID: Z001

GT Text:

ЛЮБВИ ВСЕ ВОЗРАСТЫ ПОКОРНЫ

(a) (b)

Figure 9: Multilingual input. (a) Russian input in search. For users who are not familiar with the keyboard mapping, a keyboard mapping display window is provided. (b) Russian input in groundtruth editor.

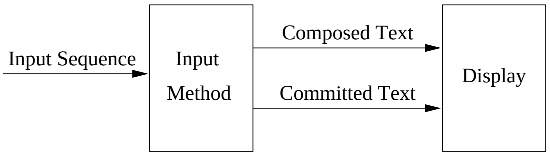

Figure 10: Input system.

public String getCommittedText(): The function for getting the current committed text.

public void keyTyped(char ch): The function for sending a typed character to the input.

public void showKeyboard(): The function for showing keyboard mappings.

7 Future Directions

TrueViz provides basic OCR groundtruthing functionalities. We hope that researchers in the international community will volunteer to add other features.

Currently TrueViz provides only English and Russian input. Other languages such as Korean, Japanese, and Chinese would be useful in multilingual OCR. A public-domain Java package with keyboards for various scripts/languages [LLW98a, LLW98b] that could be incorporated into TrueViz would be of great benefit to researchers.

Since tables are not trees, existing XML validation programs cannot verify table groundtruth data. Thus convenient ways for representing, annotating, and validating tables is needed.

Converters for DAFS to XML and XML to DAFS are currently implemented. This makes the XML representation compatible with the public domain performance evaluation toolkit PSET [MK00c, MK00b, MK01, MK00a], and allows researchers to visualize segmentation evaluation results using TrueViz. Converters from SGT format (produced by the Pink Panther groundtruthing tool), XDOC (the Xerox representation for groundtruth), and the Caere representation would be helpful.

Document images contain huge amounts of data, and XML files can require more disk space than binary-formatted files. If compressed XML files could be saved and read, the file size of the XML files would not be a concern.

Zooming is a very integral part of any document image groundtruth visualization tool. A more "zoom-centric" design using the zoomable user interface package Jazz [BMG00] could be explored.

TrueViz was tested by several members of our research group. A more thorough quantitative user evaluation using questionnaires would be desirable [CDN88].

Since the OCR community currently does not have annotation standards similar to the Corpus Encoding Standard [Exp], it would be beneficial to start a working group to build such a standard and also ensure that TrueViz is compatible with this new encoding standard.

Acknowledgments

The authors would like to thank the participants of MLOCR99 for valuable discussions; Jeff Czorapinski and Ivan Bella for their help in the initial phases of this project; Song Mao for discussion and user testing; Ben Bederson for comments on the user interface; Thomas Baby for suggesting the method of handling flexible entity structures; and Azriel Rosenfeld for editorial comments.

This research was funded in part by the Department of Defense under Contract MDA0949-6C-1250, Lockheed Martin under Contract 9802167270, the Defense Advanced Research Projects Agency under Contract N660010028910, and the National Science Foundation under Grant IIS9987944.

References

[AS94] C. Ahlberg and B. Shneiderman. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Proceedings of the ACM CHI94 Conference*, pages 313–317, Boston, MA, April 1994. `http://www.spotfire.com/products/spotfire_net.asp`.

[AS99] R. B. Allen and J. Schalow. Metadata and data structures for the historical newspaper digital library. In *Proceedings of the Eighth International Conference on Information Knowledge Management*, pages 147–153, Kansas City, MO, November 1999.

[AWC+95] S. D. Anderson, D. L. Westbrook, A. Carlson, D. M. Hart, and P. R. Cohen. *Common Lisp Analytical Statistics Package: User Manual.* University of Massachusetts, 1995. `http://eksl-www.cs.umass.edu/clasp.html`.

[BDH+00] S. Bird, D. Day, J. Garofolo J. Henderson, C. Laptun, and M. Liberman. ATLAS: A flexible and extensible architecture for linguistic annotation. In *Proceedings of the Second International Language Resources and Evaluation Conference*, pages 1699–1706, Athens, Greece, May 2000.

[BGWL98] C. Barras, E. Geoffrois, Z. Wu, and M. Liberman. Transcriber: A free tool for segmenting, labeling and transcribing speech. In *Proceedings of the First International Conference on Language Resources and Evaluation*, pages 1373–1376, Granada, Spain, May 1998.

[BGWL00] C. Barras, E. Geoffrois, Z. Wu, and M. Liberman. Transcriber: development and use of a tool for assisting speech corpora production. *Speech Communication special issue on Speech Annotation and Corpus Tools*, 33(1-2), January 2000.

[BMG00] B. Bederson, J. Meyer, and L. Good. Jazz: An extensible zoomable user interface graphics toolkit in Java. Technical Report CS-TR-4137, UMIACS-TR-2000-30, University of Maryland, College Park, MD, May 2000. `http://www.cs.umd.edu/hcil/jazz/`.

[BPSM98] T. Bray, J. Paoli, and C. M. Sperberg-McQueen. *Extensible Markup Language (XML).* W3C, 1998. `http://www.w3.org/TR/REC-xml`.

[CDN88] J. P. Chin, V. A. Diehl, and K. L. Norman. Development of an instrument measuring user satisfaction of the human-computer interface. In *Proceedings of SIGCHI '88*, pages 213–218, New York, NY, October 1988. `http://www.lap.umd.edu/QUIS/index.html`.

[Con97] The Unicode Consortium. *The Unicode Standard, Version 2.0.* Addison Wesley Developers Press, 1997.

[DAH$^+$97] D. Day, J. Aberdeen, L. Hirschman, R. Kozierok, P. Robinson, and M. Vilain. Mixed-initiative development of language processing systems. In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, Washington, DC, March-April 1997.

[dli99] *Proceedings of the IAPR Workshop on Document Layout Interpretation and its Applications*, Bangalore, India, September 1999.

[DM00] D. Doermann and D. Mihalcik. Tools and techniques for video performance evaluation. In *Proceedings of the 15th International Conference on Pattern Recognition*, pages 167–170, Barcelona, Spain, September 2000. `http://documents.cfar.umd.edu/LAMP/Media/Projects/ViPER/`.

[Exp] Expert Advisory Group on Language Engineering Standards. *Corpus Encoding Standard - Document CES 1, Version 1.5.* `http://www.cs.vassar.edu/CES/`.

[Fru95] T. Fruchterman. DAFS: A standard for document and image understanding. In *Proceedings of the Symposium on Document Image Understanding Technology*, pages 94–100, Bowie, MD, October 1995.

[GBBW00] E. Geoffrois, C. Barras, S. Bird, and Z. Wu. Transcribing with annotation graphs. In *Proceedings of the Second International Conference on Language Resources and Evaluation*, pages 1517–1521, Athens, Greece, May-June 2000.

[Gus97] D. Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology.* Cambridge University Press, 1997.

[Har99] E. R. Harold. *XML Bible.* IDG Books, Foster City, CA, 1999.

[HRI00] O. Hitz, L. Robadey, and R. Ingold. An architecture for editing document recognition results using XML technology. In *Proceedings of the Fourth IAPR International Workshop on Document Analysis Systems*, pages 385–396, Rio de Janeiro, Brazil, December 2000.

[KA99] T. Kanungo and R. B. Allen. Full-text access to historical newspapers. Technical Report CS-TR-4014, Laboratory for Language and Media Processing, University of Maryland, College Park, MD, April 1999.

[KLCB01] T. Kanungo, C. H. Lee, J. Czorapinski, and I. Bella. TRUEVIZ: a groundtruth/metadata editing and visualizing toolkit for OCR. In *Proceedings of the SPIE Conference on Document Recognition and Retrieval*, pages 1–12, San Jose, CA, January 2001.

[LK00] C. H. Lee and T. Kanungo. *TRUEVIZ User's Manual*, August 2000.

[LLW98a] K. Y. Leong, H. Liu, and O. P. Wu. Java input method engine. In *Proceedings of the Seventh International World Wide Web Conference*, Brisbane, Australia, April 1998. `http://www7.scu.edu.au/programme/fullpapers/1915/com1915.htm`.

[LLW98b] K. Y. Leong, H. Liu, and O. P. Wu. Web internationalization and Java keyboard input methods. In *Proceedings of INET 98*, pages 21–24, Geneva, Switzerland, July 1998.

[McL00] B. McLaughlin. *Java and XML*. O'Reilly, Sebastopol, CA, 2000.

[MK00a] S. Mao and T. Kanungo. Empirical performance evaluation of page segmentation algorithms. In *Proceedings of the SPIE Conference on Document Recognition and Retrieval*, pages 303–314, January 2000.

[MK00b] S. Mao and T. Kanungo. PSET: A page segmentation evaluation toolkit. In *Fourth IAPR International Workshop on Document Analysis Systems*, pages 451–462, Rio de Janeiro, Brazil, December 2000.

[MK00c] S. Mao and T. Kanungo. Software architecture of PSET: A page segmentation evaluation toolkit. Technical Report CAR-TR-955, University of Maryland, College Park, MD, September 2000. `http://www.cfar.umd.edu/~kanungo/pubs/trpset.ps`. Software is available at `http://www.cfar.umd.edu/~kanungo/software/software.html`.

[MK01] S. Mao and T. Kanungo. Empirical performance evaluation methodology and its application to page segmentation algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(3):242–256, 2001. (To appear.).

[mlo99] Working group notes, international workshop on performance evaluation issues in multilingual OCR, September 1999. `http://www.cfar.umd.edu/~kanungo/workshop/reco.html`.

[Nat] National Center for Atmospheric Research/Mesoscale and Microscale Meteorology. *PolyPaint User Manual, Version 3.0.* `http://lasp.colorado.edu/polypaint/home.html`.

[PIX] PIXOTEC, LLC. *Slicer Dicer.* `http://www.slicerdicer.com/`.

[SCB92] D. F. Swayne, D. Cook, and A. Buja. XGobi: interactive dynamic graphics in the X window system with a link to S. In *Proceedings of the American Statistical Association Meetings*, 1992. `http://www.research.att.com/areas/stat/xgobi/`.

[SCB98] D. F. Swayne, D. Cook, and A. Buja. Xgobi: Interactive dynamic data visualization in the X window system. *Journal of Computational and Graphical Statistics*, 7, 1998. `http://www.research.att.com/areas/ stat/xgobi/`.

[SHB91] D. F. Swayne, N. Hubbell, and A. Buja. XGobi meets S: Integrating software for data analysis. In *Proceedings of the Symposium on the Interface*, pages 430–434, 1991. `http://www.research.att.com/areas/stat/xgobi/`.

[Shn96] B. Shneiderman. The eyes have it: A task by data type taxonomy of information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages*, pages 336–343, September 1996. `http://otal.umd.edu/Olive/`.

[Shn98] B. Shneiderman. *Designing the User Interface.* Addison Wesley, Reading, MA, 1998.

[SK98] J. Sauvola and H. Kauniskangas. *MediaTeam Oulu Document Database.* MediaTeam, University of Oulu, Finland, 1998. `http://www.mediateam.oulu.fi/ MTDB/`.

[Uni] University of Augsburg. *Mondrian.* `http://jetta.math.uni-augsburg.de/ Mondrian/`.

[VR99] W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S-Plus.* Springer, Berlin, 1999. `http:// www.splus.mathsoft.com/`.

[WHA+99] L. Wood, A. L. Hors, V. Apparao, L. Cable, M. Champion, J. Kesselman, P. L. Hegaret, T. Pixley, J. Robie, P. Sharpe, and C. Wilson. *Document Object Model (DOM).* W3C, 1999. `http://www.w3.org/TR/DOM-Level-2/`.

[YV98] B. A. Yanikoglu and L. Vincent. Pink Panther: A complete environment for ground-truthing and benchmarking document page segmentation. *Pattern Recognition*, 31:1191–1204, 1998.