

Text Types in Hypermedia

Stephen W. Smoliar and James D. Baker

FX Palo Alto Laboratory, Inc.; 3400 Hillview Ave., Bldg. 4; Palo Alto, CA 94304

{smoliar,baker}@pal.xerox.com

Abstract

The discipline of narratology has long recognized the need to classify documents as instances of different text types. We have discovered that classification is as applicable to hypermedia as it is to any other document presentation. In this paper, following the work of Seymour Chatman, we shall consider three such text types: Description, Argument, and Narrative. The goal of a Description document is to describe some object or concept; this is usually achieved by describing component parts and then describing how those parts combine to constitute the entirety. An Argument document, on the other hand, is concerned with establishing some assertion or point of view; and it is based on supporting evidence, as well as possible refutations and justifications for defeating those refutations. Finally, a Narrative document recounts some sequence of events in time, addressing relationships, such as causality and contingency, among those events. We have analyzed these types through case studies that give an example of each as a hypermedia document: the preparation of a recipe (Description), a report of a group meeting (Argument), and recounting one of the versions of the Rashomon story (Narrative). We then argue that this classification provides an organizational framework that facilitates the construction of outlines that serve the writer in preparing the actual content of a document. Such outlines can also benefit the reader's understanding of the content that the writer intended to convey; and, if the writer does not make those outlines available explicitly to the reader, the reader can use knowledge of the document type to construct his own version of those outlines. Finally, we shall briefly review some early work in content-based indexing and search of multimedia documents; and we argue that knowledge of a document's type can be a significant asset in supporting indexing and retrieval that are actually based on content, rather than lower-level features of the media of the document.

1. Review of text types

The discipline of narratology has long recognized the need to classify documents as instances of different text types. We have discovered that classification is as

applicable to hypermedia as it is to any other document presentation. In this paper, following the work of Seymour Chatman [5], we shall consider three such text types: Description, Argument, and Narrative.

1.1. Description

Documents of the Description type are concerned with issues that, in a more formal framework, tend to be associated with knowledge representation [5]:

Descriptions render the properties of things—typically, though not necessarily, objects visible to or imaginable by the senses. They "portray," "depict," or "represent." Though the term can apply to abstractions and other nonvisible entities, Description is often considered the verbal analogue to painting or drawing.

Thus, descriptions are concerned with specifying entities and relationships among those entities, where the relationships that tend to be of greatest importance are instantiation (ISA) and decomposition (HAS-PART) [1]. Decomposition is particularly important because most entities are described as consisting of components, so the process of description involves first enumerating the components and then characterizing the relationships among them by which they combine to form the whole. This process may then be repeated for subcomponents of the components as is most appropriate for the description task at hand. An important characteristic of all relationships involved in description is that they are *time-independent*. Documents of the Description type are thus based on atemporal logical assertions.

1.2. Argument

Chatman [5] describes the Argument type as follows: "Argument is the text-type that relies on 'logic,' at least in the informal sense; it may employ not the strict 'demonstrative' logic of the syllogism but rather the softer one of the rhetorical enthymeme." Stephen Toulmin [15] has investigated the relationship between argument and logic in much greater detail, concluding that, beyond the usual conventions of syllogistic reasoning, argument is primarily concerned with establishing *claims*, where the

major sources of support for those claims consist in *data* and *warrants*. This approach has been applied to the systematic representation of the content of a complex article in *The New Republic* [3], and it has been subsequently adapted to the implementation of data structures that support the SEPIA authoring tool [13]. However, from Chatman's point of view, one of the most important qualities of the Argument type is that, as in the Description type, the relationship between a claim and any data and warrants that support it is time-independent. On the other hand the actual *presentation* of an argument usually depends on time-ordering. Thus, the presenter of the argument has to choose between beginning with a statement of a claim that is subsequently supported or beginning with data and warrants that gradually build up to the claim, delivered in a manner similar to the punch line of a joke. Both approaches can be equally effective, depending on the situation and the relationship between the person making the argument and those who have to be convinced.

1.3. Narrative

However, it is in documents of the Narrative type that time becomes a major issue [5]:

As has been clearly established in recent narratology, what makes Narrative unique among the text-types is its "chrono-logic," its doubly temporal logic. Narrative entails movement through time not only "externally" (the duration of the presentation of the novel, film, play) but also "internally" (the duration of the sequence of events that constitute the plot).

Thus, if the Argument type is concerned primarily with *assertions* and the Description type is concerned primarily with *entities*, the Narrative type is concerned primarily with *events*.

While it is important to be aware of this morphology, Chatman [5] emphasizes that a document is rarely of a single type. One type may be the dominant form, but other types may be *embedded* within that form. Thus, in the following Section we shall see an example of a Description document that presents a recipe for cooking quail; and there are important temporal elements that cannot be ignored in preparing that recipe (or just about any other recipe, for that matter). Because cooking is a process, one cannot describe a recipe without accounting for how certain events are sequenced; and this need to account for events means that this particular document of the Description type must include elements that are of the Narrative type. Most important, however, is that no one of these types dominates the others. Thus, one cannot plan to use SEPIA as a general-purpose hypermedia authoring system because one cannot accept the assumption, implicit in SEPIA's approach to data structures, that all documents can be reduced to the

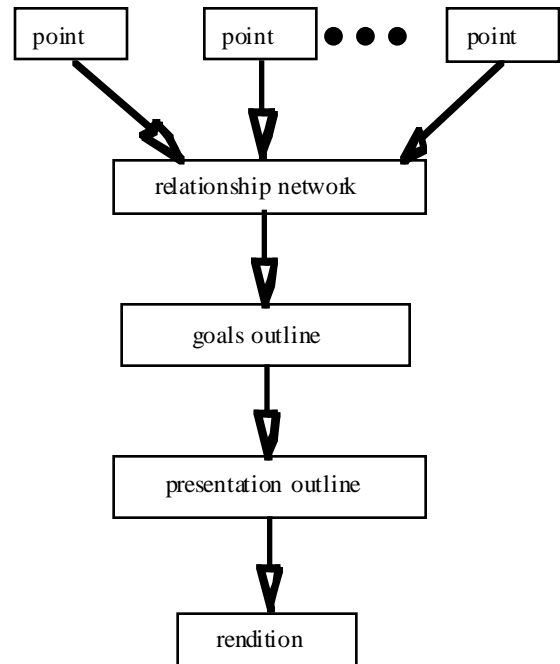


Figure 1: Model of reading and writing.

Argument type. Rather, the models and technology that support authoring must give equal priority to all three types of documents.

2. Examples of text types in hypermedia

We have analyzed all of these document types through the construction of case studies that give an example of each as a hypermedia document. We shall now review one case study for each type: the preparation of a recipe, as cited above (Description), a report of a group meeting (Argument), and recounting one of the versions of the story that forms the basis for the film *Rashomon* [9] (Narrative). Each of these examples will now be discussed in turn.

2.1. Description: Preparing a recipe

The basis for this case study was an excerpt from the television series *World Class Cuisine*¹ aired on the Discovery Channel. This excerpt presented a French chef's recipe for cooking quail. The production technique was not that of a conventional "how to" cooking program. The camera essentially played the role of a discreet guest visiting the kitchen of one of the great chefs of the world, and the chef ignores its presence entirely. All explanation was provided in the form of narration on an audio track that was added after all the video had been collected and

¹ Used with permission of the producers, Arna Vodenos Productions.

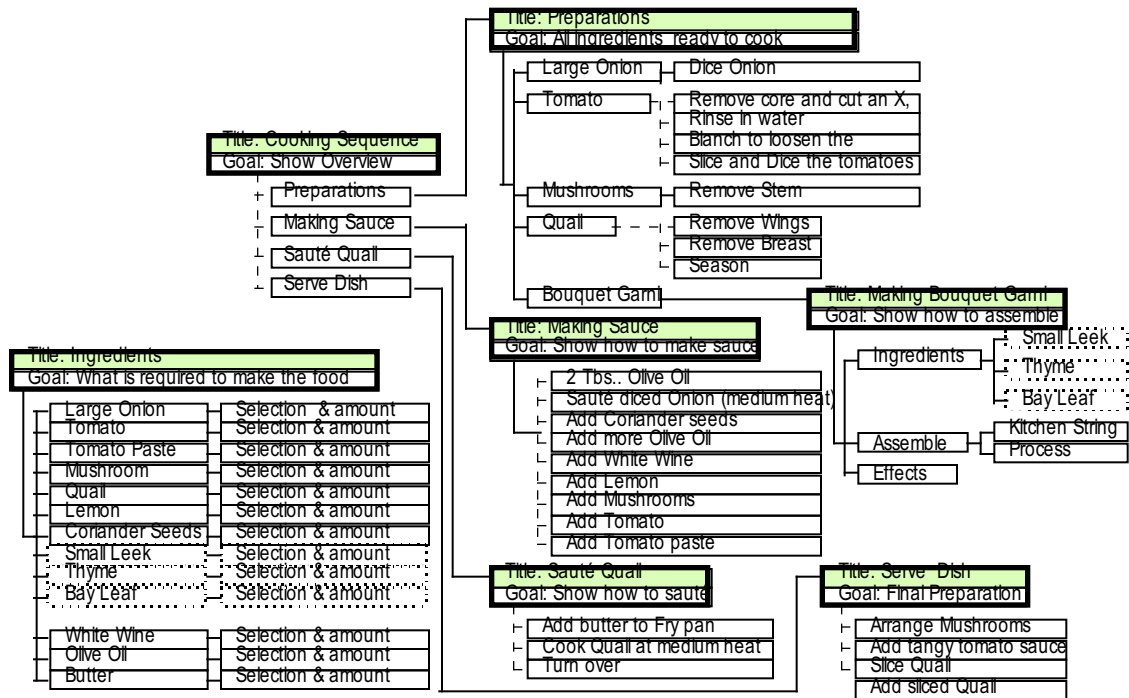


Figure 2: The goals outline of the cooking video.

edited. We felt this was a useful exercise in communication, since cooking videos are designed to convey instructional information to the audiences in relatively short time and, while the intention of the chef was definitely *not* instructional, a viewer could "read" the resulting video for instructional purposes.

The objective of the case study was to restructure the content of this video as a hypermedia document. The resulting document could then be said to constitute a "reading" of the content of the original video. Our approach to reading the video is based on the model illustrated in **Figure 1**. As has been discussed elsewhere, this model is based on several different, but related, approaches to model the process of writing [10]; but it is also consistent with the model of strategic discourse processing developed by Teun van Dijk and Walter Kintsch [6].

In our model one collects individual units of content, which we call *points*, not only when one is preparing to write but also when one is trying to understand what one is reading; and the resulting collection is the product of what van Dijk and Kintsch call *propositional strategies*. As points are accumulated, they are organized into a *relationship network*, following what van Dijk and Kintsch call *local coherence strategies*. The resulting organization is important, since the points are not particularly manageable as an unstructured aggregate; and it can then assist either the writer or the reader in coming

to terms with the goals of the document—what the document is trying to communicate. This understanding of goals arises from what van Dijk and Kintsch call *macrostrategies*; and, in our model, it results in a *goals outline* that provides a hierarchical representation of the goal/subgoal structure. How this structure is actually realized in a document, however, is based on what van Dijk and Kintsch call *schematic strategies*; and our model has postulated a *presentation outline* to represent the product of those strategies. This need not have the neat hierarchical structure of the goals outline. Rather, it constitutes the skeleton of the hypermedia presentation and must be structured according to the writer's presentational objectives; so the *rendition* of the document, itself, becomes a relatively straightforward matter of "fleshing out" this skeleton according to what van Dijk and Kintsch call *production strategies*.

The goals outline for both the video and hypermedia versions is depicted in **Figure 2** as a forest of two goal trees. Each leaf of this forest constitutes a single point, and the connectivity of the graph was used as a basis for defining links. (Note that the leaves in dotted boxes are assigned to two different goals.) This representation also illustrates, using dashed lines, temporal orderings that serve as constraints on the presentation of particular points. These lines indicate that their associated points must be treated as events within an embedded Narrative structure.

We have discovered that presentation outlines may be structured according to three organizing principles, which we call *logical*, *spatial*, and *temporal* [12]. Thus, a temporal outline of the video is a time table that gives the time codes for the video segments corresponding to the leaves of the goals outline, while the temporal outline of our hypermedia version is a "Road Map" (illustrated in **Figure 3**), which summarizes the temporal constraints of the cooking process and is "hot" with links to the appropriate points. **Figure 4** is an example of a spatial outline, a single frame from the source video encompassing the entire kitchen, where both ingredients and utensils are "hot" and link the reader to respective details concerned with both required ingredients and steps of preparation. An example of a logical outline is **Figure**

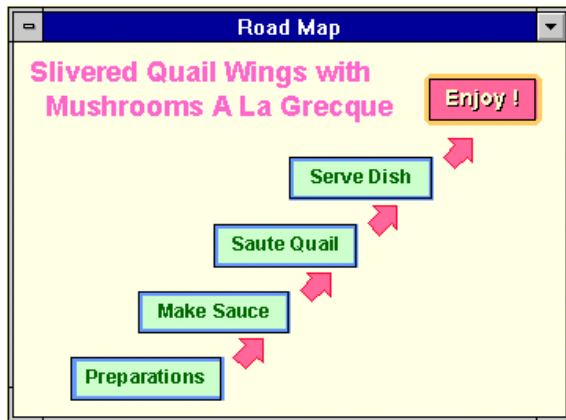


Figure 3: "Road Map" for the cooking process.



Figure 4: The kitchen as "view" of the cooking video.

Ingredients List

Quail with Mushrooms A La Grecque

- 1 Quail
- 1 tbs. Coarse Salt
- 4 tbsps. Butter

— SAUCE —

- 1/3 Large Onion
- 4 tbsns. Extra Virgin Olive Oil
- 2 Ripe Tomatos
- 2 tbsns. Tomato Paste
- 1/2 lb. Button Mushrooms
- 1/2 Lemon
- 1 tbsn. Coriander Seeds
- 1 1/2 cups French Wine

— Bouquet Garni —

- 1 Bay Leaf
- Few Sprigs Thyme
- 1 Small Leek

Figure 5: List of ingredients.

Menu

Quail with Mushrooms A La Grecque

- Ingredient List
- Kitchen
- Process Road-Map
- Time Table

Figure 6: Hypermedia user interface.

5, which translates that portion of the goal tree that summarizes the necessary ingredients. Finally, because it is important that every document have a well-defined beginning, another logical outline, illustrated in **Figure 6**, integrates the content of the other presentation outlines into a single "first page" for the reader.

2.2. Argument: Reporting on a group meeting

As another example of source material, we have been making videotape recordings of a wide variety of meetings conducted at our laboratory. Our long-term goal is to explore the use of the Coral confederation of tools that support the real-time capture of and subsequent access to

meetings [8]. However, in preparation for this activity, we have begun to videotape all meetings that take place in our large conference room.

One form of meeting that we feel is in particular need of documentation is what we call our "roundtable discussions." Each of these discussions has a leader who prepares a set of slides on a specific topic; but the basic "rule of the game" for these meetings is that open discussion may commence with the very first slide. As a result issues frequently arise that were not anticipated by the leader, may induce significant (and often beneficial) digressions, and may oblige the leader to reorder or eliminate some of the slides prepared for the presentation. It is because these discussions are so free-wheeling that documenting them is especially valuable. In addition we have discovered a need to report the proceedings of many of these meetings to several of our colleagues working at other laboratories in Japan. A simple video record is not terribly effective for either of these needs, so we have been

exploring the hypermedia document as a more viable alternative.

The context of this project is very similar to the cooking exercise just discussed. In both cases we had no production control over the content of the video source material. We could only work with what the camera and microphone recorded without any luxury of repeating any of the events for the sake of setting up better shots or having the speakers articulate more clearly. Indeed, because these meetings are so spontaneous, the video source material is far less polished than what we used for the cooking exercise. However, again, the exercise was one of restructuring material originally in video form into a hypermedia document; and we shall now examine how the result both resembles and differs from the result of our first experiment.

In the case of the cooking exercise, it was a straightforward matter to begin with a goals outline, because that outline could be based on general knowledge

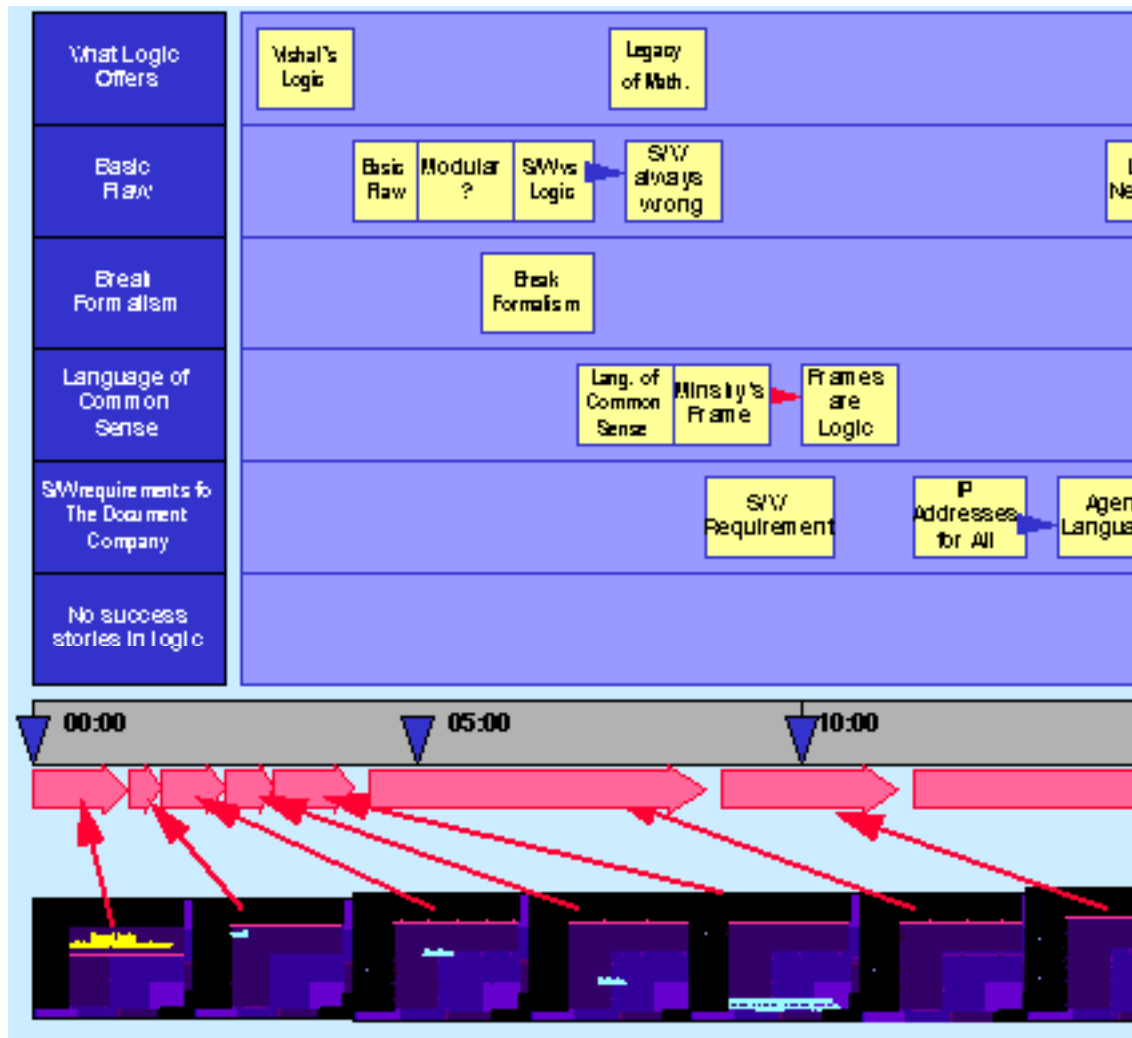


Figure 7: Beginning of time table for the roundtable discussion.

of how to describe a recipe without any specific regard to the actual video content. However, even with the benefit of copies of slides prepared by the speaker, figuring out what the roundtable discussion was actually about was far more challenging. Consequently, instead of beginning with the goals outline, we began by trying to segment the video according to the points that were being made by all the participants. We then proceeded to define clusters of these points according to topical groupings. We anticipated that such a grouping could be hierarchical; but our actual experiences yielded a "flat" list of topics. We then decided to represent a time table for the video as a two-dimensional graph, plotting time on the horizontal axis and partitioning the vertical axis with labels for the clusters. The horizontal axis is further supplemented with arrows, each of which represents the period during which a specific slide, whose image is provided in iconic form, was displayed. The beginning portion of the result is illustrated in **Figure 7**.

As we reviewed this material, the following goals outline, based on arguing about the value of logic as a technology, emerged as our "reading" of the recorded discussion:

- I. What constitutes the technology of logic?
- II. Past successes
- III. Hypothesized successes
 - A. Can logic be used to express common sense?
 - 1. Refutations
 - 2. Rebuttals
 - B. Can logic fulfill software technology needs?

- 1. Refutations
- 2. Rebuttals

What is most important about this particular outline is the value of logic is presented in the form of hypotheses for which both supporting and refuting evidence has been offered. In the course of the discussion, as it was conducted, no firm conclusions were reached. For both hypotheses both sides of the questions were reviewed, but no solid agreement was established to confirm or deny either of them.

This makes the discussion rather different from the traditional form of an argument. The whole purpose of an argument, as discussed by Toulmin [15], is to establish a claim. If refutations are introduced, then either those refutations need to be refuted by a rebuttal argument or the claim needs to be qualified to account for them. A document of the Argument type does not usually allow for the conclusions to be inconclusive, so to speak. We thus introduced a modification to Toulmin's structure, mapping the content of the discussion with the diagram illustrated in **Figure 8**. We describe the focus of the discussion as a *position*, rather than a *claim*, because the significance of the discussion lies in how that position is both supported and opposed.

While the cooking exercise was based on providing the reader with several different views of the content, the roundtable discussion was reported primarily in terms of the view in **Figure 8**. Each of the nodes of this map can then be unfolded to lead the reader to the specific points made during the discussion. These can be presented in

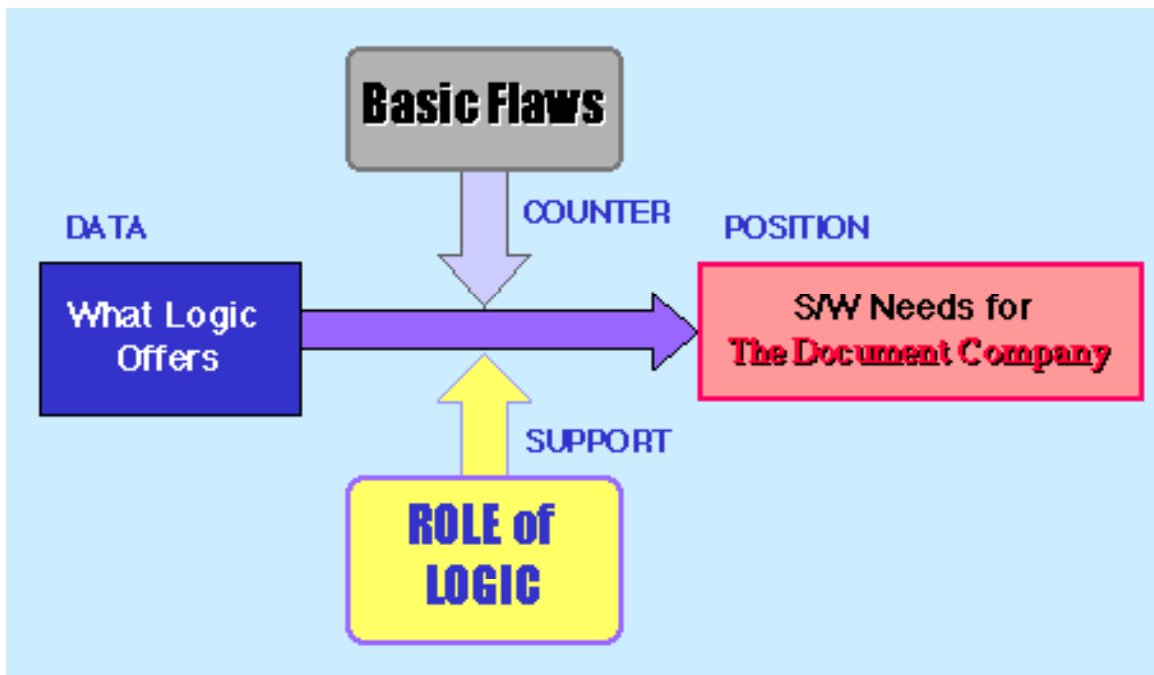


Figure 8: Argument structure for roundtable discussion.

several forms: text summary, video excerpt, and supporting slide image. In addition the reader interested in the actual video content can trace any excerpt back through the time table illustrated in **Figure 7** as a means of establishing the context within which that point was presented. Similarly, the document has a spatial outline, organized according to the same principles as **Figure 4**. In this case the image is a "salient still" [14], a composite image of the entire conference room showing all the participants in the discussion; and each participant serves as an index to his contributions to the discussion. Since both the time table and this image of all participants index the entire video source, the reader may also use them to draw his own conclusions from the discussion. Consequently, if necessary, that reader can construct a new document based on an alternative approach to arguing the source material. The inconclusive nature of this content can thus serve as an encouragement for all readers to develop and document their own conclusions.

2.3. Narrative: Recounting a story

When we turn to examples of narrative, the exercise of restructuring a given document as hypermedia becomes more problematic. Because the elements of descriptions and arguments are not dominated by temporal relationships, one may develop data structures that concentrate on all other relationships among those elements, without worrying about how time should be either represented conceptually or presented to the reader. On the other hand since, as was observed in Section 1.3, there are actually *two* senses of time at work in a narrative, in designing a hypermedia document we must address which, if either, of them requires explicit representation.

We decided to investigate the film *Rashomon* for several reasons. One was that we could draw upon not only the film itself but also the literary sources that were adapted by its director, Akira Kurosawa. Furthermore, the film was segmented into individual tellings of a story in a very straightforward manner. Indeed, the basic theme of the film revolves around the inconsistencies among these stories, all of which purport to be describing the same events. Thus, we had the opportunity to examine a single story in isolation prior to addressing the relationships among these conflicting accounts. The scope of that single story was then comparable to the scope of our other case studies.

The first of the tellings that serves as a relatively complete account of the events being examined is told by the bandit who perpetrated most of those events. This version may be summarized as follows [9]:

He was asleep under a tree when the man and his wife went past; the wind blew her veil and he saw her and decided he wanted her. He tricked the husband into following him, tied him up, went back, got the wife, raped her in

front of the husband, and then turned to go, when she stopped him, saying that her honor demanded that the men now fight. In the resulting duel the bandit killed the husband and the woman ran away.

This corresponds to a single sequence from the Kurosawa film. (Since we currently do not have rights to reproduce any material from this film, we cannot, in this document, illustrate our observations with selected frames from the source material.) Within that sequence, there is a linear flow of time, which means that we can prepare a time table similar to the one we prepared for the cooking exercise. Furthermore, because the events are presented in the order in which they occurred, we do not have to worry about constructing one time table for the telling of the story and another for the events themselves. What is missing, however, is any means of representing the relationships among those events. Also, there is more to the story than just the events. There are also what Chatman [4] calls the *existents*, a single term intended to capture the relevant features concerned with both character and setting.

Thus, in some respects we may treat this particular narrative in a manner very much like our treatment of the cooking exercise, perhaps to the point of providing a variety of views into the content (still of just this one telling) similar to the views summarized in **Figure 6**. From the point of view of a logical outline, instead of an "ingredient list," one provides a summary of the characters and a representation of their connections to the different events in which they are involved. Similarly, one can provide a spatial view through the setting, just as the kitchen is the setting of the cooking presentation. Because the setting for this particular narrative is too broad to be summarized in a single camera shot, one may not necessarily be able to construct a user interface similar to **Figure 4**. On the other hand a "salient stills" [14] technique may again be used to construct a composite image of the entire setting based on appropriate frames selected from Kurosawa's original source material. Both of these perspectives may then be set against the temporal outline of the actual sequencing of events, which, in turn, may be used as an index into both the Kurosawa film and the text by Ryunosuke Akutagawa on which the film was based [9].

Unlike the previous case studies, this exercise is still in progress; but it is already raising some interesting questions. If we can, indeed, treat the content of this particular telling of the story in a manner so similar to our Description document for a cooking process, might that mean that Narrative documents are actually a special case of Description documents? As was observed in Section 1.3, this is not the case because of what Chatman [5] calls the "chrono-logic" of Narrative. Time only figures in describing a cooking process when it comes to sequencing the preparation events; but the role of time in Narrative can be far more subtle, as Chatman illustrates with many

examples. However, this subtlety does not become a significant issue in the bandit's account of the *Rashomon* story, because the order of events in his telling corresponds straightforwardly with the order of events as they occurred. In this respect, then, a single story from *Rashomon* does not really raise issues of how a Narrative document is significantly different from a Description document; so, as it has proceeded thus far, our final case study has not yet really begun to explore the nature of the Narrative document. Consequently, further work on the project will address broadening the scope to the entirety of the film, where "chrono-logic" becomes one of the major elements of the content, rather than continuing to focus at the level of single tellings of the story. At the same time we anticipate that this exercise will take much longer than our other case studies, since, as we have already seen in the case of setting, being able to generate the most appropriate material for presentation has proven to be very time consuming.

3. Benefits of text types

What do we gain from being able to classify text types? We shall now examine three approaches to answering this question. We first argue that this classification provides an organizational framework that facilitates the construction of outlines that serve the writer in preparing the actual content of a document. As we saw in Section 2.1, such outlines can also benefit the reader's understanding of the content that the writer intended to convey; and, if the writer does not make those outlines available explicitly to the reader, the reader can use knowledge of the document type to construct his own version of those outlines. Finally, we shall briefly review some initial research concerned with a new approach to content-based indexing and search of multimedia documents; and we argue that knowledge of a document's type can be a significant asset in supporting indexing and retrieval that are actually based on content, rather than lower-level features of the media of the document.

3.1. Organization for writing

One of the most important lessons we have learned from the case studies we have pursued is that outlining is as important in the authoring of hypermedia documents as it is in conventional writing. As a matter of fact, we would argue that outlining is even *more* important, because we feel our case studies have provided support for our model of reading and writing as illustrated in **Figure 1**, which postulates that proper authoring requires more than one outline. While the goals outline provides the general plan for the *content* of the document, it is still necessary to plan out how the hypermedia document, itself, will be organized; and this latter plan is represented by one or more presentation outlines. In the verbal text medium a table of contents tends to serve

reasonably well as a presentation outline, but when other media are involved the principles behind a table of contents cannot be generalized. This is because other media have to deal with spatial and temporal elements of presentation, while a logical presentation outline is sufficient for verbal text. Consequently, storyboards generally serve as temporal presentation outlines for films; and sketches or cartoons tend to be used as spatial presentation outlines for images. Thus, the choice of medium tends to influence which types of presentation outlines are required; but, inversely, the *content* of the material being presented, by virtue of being logical, spatial, and/or temporal in nature, will influence which media are most suitable for presentation. Furthermore, when hypermedia links are also involved, a presentation outline must accommodate not only the *appearance* of the document in the media being engaged but also its *affordances for interaction*. Thus, an outline for **Figure 4** must provide not only a sketch of the kitchen but also a specification of anchor points for links and some information regarding the destinations of those links.

We thus see outline construction as a critical task in hypermedia authoring, sufficiently important that many writers would benefit from assistance. We see classifying the text type of a document as a means of providing some "first cut" versions of the outlines required for writing them. For example the author of an Argument document may benefit from being provided with templates for the necessary outlines. The template for the goals outline would assist in articulating the position being argued, the data that form the basis for the argument, and the points that need to be addressed that both support and oppose the position. This would then correspond to a presentation outline that would allow the author to develop a user interface similar to the one illustrated in **Figure 8**. On the other hand the goals outline for a Description document would specify how an entity is assembled from components, which, in turn, may be decomposed into their own components, and so on to that level of granularity that satisfies the author. In this case the author may be presented with different templates for a presentation outline corresponding to the different approaches to presentation discussed in Section 2.1.

3.2. Organization for reading

According to the model illustrated in **Figure 1**, these outlines are as valuable to the reader as they are to the writer; and, while most readers can manage by constructing outlines in their heads, often the complexity of a text is such that the reader will benefit from actually *writing* an outline if none has been provided by the author. This is basically what the readers of *The New Republic* were doing in the exercise discussed in Section 1.2. However, what is important about this exercise is that, because the readers could establish *a priori* that the document was of the Argument type, they could draw upon the fact that an outline based on the structural

analysis of arguments proposed by Toulmin [15] would be particularly effective. In other words, in the absence of an outline, the reader can use knowledge of document type as a source of guidelines for how an outline may best be constructed.

3.3. Organization for search

This is a more speculative application, but it is also one that addresses a particularly critical need in contemporary multimedia technology [7]. While significant effort has gone into the development of "content-based" retrieval systems for multimedia documents, the techniques employed by these systems have little, if anything, to do with what may be legitimately called "document content" [11]. Instead, these systems are based on means of detecting and indexing *features* of different media types, such as keywords for text and distributions of colors and textures for images. To draw upon the terminology of semiology [2], these systems perform indexing and retrieval on the plane of *expression*, rather than the plane of *content*.

Classifying a document by type provides a possible escape from this semiological conundrum. As we have seen from their impact on outline structure, document types introduce structure on the plane of content. They provide a means of organizing those elements that actually pertain to document content with respect to critical content-related questions, such as, for example:

- What's the point?
- What are the contributing elements?
- How are those elements structured?

Once those elements have been identified and organized with respect to the content-related structure of the document, *then* they may be searched on the plane of expression.

This may be illustrated by an example in image retrieval. If a user is looking for a tomato with current image retrieval technology, most likely he can do no better than specify a circular region of a reddish hue, allowing for some level of variation in the specification of that red with respect to the dimensions of some color space [7]. Such a query will yield tomatoes, but only of a particular size, as well as red balls, stop lights, and sunsets when the sun is at a particular distance from the horizon.

However, if the search space is filtered on the plane of content before issuing a query on the plane of expression, the results are likely to be less frustrating. Thus, in this case one can begin by restricting search to Description documents; and one may further narrow the search to documents of the recipe genre, since that genre implies a descriptive structure which includes a list of ingredients—one of the best places to look for tomatoes. In other words both document type and genre provide a powerful means to define a content-based context for expression-based search; and search within a context is more likely to yield useful positive results.

4. Conclusions

This last example finally allows us to relate issues of text type to those of genre. While both text type and genre are ultimately concerned with establishing frameworks that describe the nature of communication, genre tends to concentrate on communicative *actions*. Thus, for example, we have the following characterization [16]:

Genres (e.g., the memo, the proposal, and the meeting) are typified communicative actions characterized by similar substance and form and taken in response to recurrent situations.

Text type, on the other hand, addresses how those actions are ultimately rendered in some form of *text* (where the concept of text can accommodate a variety of media, including both the printed and the spoken word, images, sounds, video, and other media which have yet to be anticipated). The key difference appears to be that text type tends to concentrate on the *intent* of the agent doing the communicating; and, once that intention has been established, it may then be realized through one of several different generic forms. Thus, the intent to describe how to cook a quail is realized from the recipe genre; and that genre is one of several addressing descriptive intentions. (Description is actually a good example for considering the variety of media. For example, if we want to describe the structure of a copy machine, the most appropriate genre is that of the exploded diagram, which tends not to benefit from verbal text. Indeed, this genre is best realized through interactive three-dimensional graphics, although it still tends to be rendered more frequently on the printed page.) What is important about these examples is that both text type and genre impact our organizational thoughts about documents; and, furthermore, that organization is hierarchical, structured first by text type and then by genre.

5. Acknowledgements

This paper is the product of extensive (and enjoyable) hours of discussion and implementation within the Hypermedia Authoring project at the FX Palo Alto Laboratory (FXPAL). It would be very difficult to list all the participants who helped contribute to the ideas we have now documented. However, we feel it is important to thank the contributors to the specific case studies we have examined. The cooking exercise was primarily implemented by Takeshi Shimizu, with abundant assistance from Eiji Ishida. Much of the effort towards documenting our roundtable discussion was assumed by Takeshi Shimizu, John Boreczky, and Naomi Friedlander; and the initial work on *Rashomon* was conducted by Ben Arnette and Tina Schneider, both of

whom, along with Ms. Friedlander, spent their summer interning at our facility.² Last, but far from least, we wish to thank Seymour Chatman, of the University of California at Berkeley, for reviewing all of our case studies, hearing us out as we ventured into territory far more familiar to him than it was to us, and helping us out with many valuable suggestions and guideposts.

6. References

[1] A. Barr and E. A. Feigenbaum, editors, Representation of Knowledge, in *The Handbook of Artificial Intelligence*, Vol. I, Chapter III, William Kaufmann, Los Altos, 1981, pp. 141–222.

[2] R. Barthes, *Elements of Semiology*, A. Lavers and C. Smith, translators, Hill and Wang, New York, 1973.

[3] J. S. Brown, "Process Versus Product: A Perspective on Tools for Communal and Informal Electronic Learning", *Journal of Educational Computing Research*, Vol. 1, No. 2, 1985, pp. 179–201.

[4] S. Chatman, *Story and Discourse: Narrative Structure in Fiction and Film*, Cornell University Press, Ithaca, 1978.

[5] S. Chatman, *Coming to Terms: The Rhetoric of Narrative in Fiction and Film* Cornell University Press, Ithaca, 1990.

[6] T. A. van Dijk and W. Kintsch, *Strategies of Discourse Comprehension*, Academic Press, San Diego, 1983.

[7] B. Furht, S. W. Smoliar, and H.-J. Zhang, *Video and Image Processing in Multimedia Systems*, Kluwer Academic Publishers, Boston, 1995.

[8] S. Minneman *et al.*, A Confederation of Tools for Capturing and Accessing Collaborative Activity, Proceedings: ACM Multimedia '95 (November 1995), pp. 523–534.

[9] D. Richie, Rashomon, in *Rashomon*, D. Richie, editor, Rutgers University Press, New Brunswick, 1987, pp. 1–21.

[10] S. W. Smoliar and J. D. Baker, Hypermedia Authoring Needs Semiology, submitted to Multimedia Computing and Networking 1997 (February 1997).

[11] S. W. Smoliar *et al.*, Multimedia Search: An Authoring Perspective, Proceedings: First International Workshop on Image Databases and Multi Media Search (August 1996), pp. 1–8.

[12] S. W. Smoliar *et al.*, The Effort of Authoring Hypermedia Documents, submitted to Hypertext 97 (April 1997).

[13] N. A. Streitz, J. Hannemann, and M. Thüring, From Ideas and Arguments to Hyperdocuments: Travelling through Activity Spaces, Hypertext '89 Proceedings (November 1989), pp. 343–364.

[14] L. Teodosio and W. Bender, Salient Video Stills: Content and Context Preserved, Proceedings: ACM Multimedia '93 (August 1993), pp. 39–46.

[15] S. E. Toulmin, *The Uses of Argument*, Cambridge University Press, Cambridge, 1958.

[16] J. Yates and W. J. Orlikowski, "Genres of Organizational Communication: A Structural Approach to Studying Communication and Media", *Academy of Management Review*, Vol. 17, No. 2, 1992, pp. 299–326.

² Ms. Schneider was employed by Xerox PARC but was working with FXPAL researchers.