

# Summarizing Noisy Documents

**Hongyan Jing**

IBM T.J. Watson Research Center  
Yorktown Heights, NY  
hjing@us.ibm.com

**Daniel Lopresti**

19 Elm Street  
Hopewell, NJ  
dpl@dlopresti.com

**Chilin Shih**

150 McMane Avenue  
Berkeley Heights, NJ  
cls@prosodies.org

## Abstract

*We investigate the problem of summarizing text documents that contain errors as a result of optical character recognition. Each stage in the process is tested, the error effects analyzed, and possible solutions suggested. Our experimental results show that current approaches, which are developed to deal with clean text, suffer significant degradation even with slight increases in the noise level of a document. We conclude by proposing possible ways of improving the performance of noisy document summarization.*

## 1 Introduction

Summarization aims to provide a user with the most important information gleaned from a document (or collection of related documents) [16]. A good summary can help the reader grasp key subject matter without requiring study of the entire document. This is especially useful nowadays as information-overload becomes a serious issue.

Much attention is currently being directed towards the problem of summarization [6, 25]. However, the focus to date has typically been on clean, well-formatted documents, i.e., documents that contain relatively few spelling and grammatical errors, such as news articles or published technical material. In this paper, we present a pilot study of noisy document summarization, motivated primarily by the impact of various kinds of physical degradation that pages may endure before they are scanned and processed using optical character recognition (OCR) software.

Understandably, summarizing documents that contain many errors is an extremely difficult task. In our study, we focus on analyzing how the quality of summaries is affected by the level of noise in the input document, and how each stage in summarization is impacted by the noise. Based on our analysis, we suggest possible ways of improving the performance of automatic summarization systems for noisy documents. We hope to use what we have learned from

this initial investigation to shed light on the directions future work should take.

What we ascertain from studying the problem of noisy document summarization can be useful in a number of other applications as well. Noisy documents constitute a significant percentage of documents we encounter in everyday life. The output from OCR and speech recognition (ASR) systems typically contain various degrees of errors, and even purely electronic media, such as email, are not error-free. To summarize such documents, we need to develop techniques to deal with noise, in addition to working on the core algorithms. Whether we can successfully handle noise will greatly influence the final quality of summaries of such documents.

A number of researchers have begun studying problems relating to information extraction from noisy sources. To date, this work has focused predominately on errors that arise during speech recognition, and on problems somewhat different from summarization. For example, Gotoh and Renals propose a finite state modeling approach to extract sentence boundary information from text and audio sources, using both n-gram and pause duration information [8]. They found that precision and recall of over 70% could be achieved by combining both kinds of features. Palmer and Ostendorf describe an approach for improving named entity extraction by explicitly modeling speech recognition errors through the use of statistics annotated with confidence scores [20]. Hori and Furui summarize broadcast news speech by extracting words from automatic transcripts using a word significance measure, a confidence score, linguistic likelihood, and a word concatenation probability [11].

There has been much less work, however, in the case of noise induced by optical character recognition. Early papers by Taghva, et al. show that moderate error rates have little impact on the effectiveness of traditional information retrieval measures [23, 24], but this conclusion does not seem to

apply to the task of summarization. Miller, et al. study the performance of named entity extraction under a variety of scenarios involving both ASR and OCR output [18], although speech is their primary interest. They found that by training their system on both clean and noisy input material, performance degraded linearly as a function of word error rates. They also note in their paper: “To our knowledge, no other information extraction technology has been applied to OCR material” (pg. 322).

An intriguing alternative to text-based summarization is Chen and Bloomberg’s approach to creating summaries without the need for optical character recognition [4]. Instead, they extract indicative summary sentences using purely image-based techniques and common document layout conventions. While this is effective when the final summary is to be viewed on-screen by the user, the issue of optical character recognition must ultimately be faced in most applications of interest (e.g., keyword-driven information retrieval).

For the work we present in this paper, we performed a small pilot study in which we selected a set of documents and created noisy versions of them. These were generated both by OCR’ing real pages and by using a filter we have developed that injects various levels of noise into an original source document. The clean and noisy documents were then piped through a summarization system. We tested different modules that are often included in such systems, including sentence boundary detection, part-of-speech tagging, syntactic parsing, extraction, and editing of extracted sentences. The experimental results show that these modules suffer significant degradation as the noise level in the document increases. We discuss the errors made at each stage and how they affect the quality of final summaries.

In Section 2, we describe our experiment, including the data creation process and various tests we performed. In Section 3, we analyze the results of the experiment and correlate the quality of summaries with noise levels in the input document and the errors made at different stages of the summarization process. We then discuss some of the challenges in summarizing noisy documents and suggest possible methods for improving the performance of noisy document summarization. We conclude with a proposal for future work.

## 2 The Experiment

### 2.1 Data Creation

We selected a small set of four documents to study in our experiment. Three of four documents were from the data collection used in the Text REtrieval Conferences (TREC) [10] and one was from a Telecom-

munications corpus we collected ourselves [13]. All were professionally written news articles, each containing from 200 to 800 words (the shortest document was 9 sentences and the longest was 38 sentences).

For each document, we created 10 noisy versions. The first five corresponded to real pages that had been printed, possibly subjected to a degradation, scanned at 300 dpi using a UMAX Astra 1200S scanner, and then OCR’ed with Caere OmniPage Limited Edition. These included:

**clean** The page as printed.

**fax** A faxed version of the page.

**dark** An excessively dark (but legible) photocopy.

**light** An excessively light (but legible) photocopy.

**skew** The clean page skewed on the scanner glass.

Note that because the faxed and photocopied documents were processed by running them through automatic page feeders, these pages can also exhibit noticeable skew. The remaining five sample documents in each case were electronic copies of the original that had had synthetic noise (single-character deletions, insertions, and substitutions) randomly injected at predetermined rates: 5%, 10%, 15%, 20%, and 25%.

A summary was created for each document by human experts. For the three documents from the TREC corpus, the summaries were generated by taking a majority opinion. Each document was given to five people who were asked to select 20% of the original sentences as the summary. Sentences selected by three or more of the five human subjects were included in the summary of the document. (These summaries were created for our prior experiments studying summarization evaluation methodologies [14].) For the document from the Telecommunications corpus, an abstract of the document was provided by a staff writer from the news service. These human-created summaries were useful in evaluating the quality of the automatic summaries.

### 2.2 Summarization Pipeline Stages

We are interested in testing how each stage of a summarization system is affected by noise, and how this in turn affects the quality of the summaries. Many summarization approaches exist, and it would be difficult to study the effects of noise on all of them. However, the following pipeline is common to many summarization systems:

- **Step 1: Tokenization.** The main task here is to break the text into sentences. Tokens in the input text are also identified.

- Step 2: Preprocessing. This typically involves part-of-speech tagging and syntactic parsing. This step is optional; some systems do not perform tagging and parsing at all. Topic segmentation is deployed by some summarization systems, but not many.
- Step 3: Extraction. This is the main step in summarization, in which the automatic summarizer selects key sentences (sometimes paragraphs or phrases) to include in the summary. Many different approaches for sentence extraction have been proposed and various types of information are used to find summary sentences, including but not limited to: frequency, lexical cohesion, sentence position, cue phrases, discourse structures, and overlapping information in multiple documents.
- Step 4: Editing. Some systems post-edit the extracted sentences to make them more coherent and concise.

For each stage in the pipeline, we selected one or two systems that perform the task and tested their performance on both clean and noisy documents.

- For tokenization, we tested two tokenizers: one is a rule-based system that decides sentence boundaries based on heuristic rules encoded in the program, and the other one is a trainable tokenizer that uses a decision tree approach for detecting sentence boundaries and has been trained on a large amount of data.
- For part-of-speech tagging and syntactic parsing, we tested the English Slot Grammar (ESG) parser [17]. The outputs from both tokenizers were tested on ESG.
- For extraction, we used a program that relies on lexical cohesion, frequency, sentence positions, and cue phrases to identify key sentences [13]. The length parameter of the summaries was set to 20% of the number of sentences in the original document. The output from the rule-based tokenizer was used in this step.
- In the last step, we tested a cut-and-paste system that edits extracted sentences by simulating the revision operations often performed by professional abstractors [13]. The outputs from the three previous steps were used by the cut-and-paste system.

All of the summaries produced in this experiment were generic, single-document summaries (i.e., the summary was about the main topic conveyed

in a document, rather than some specific information that is relevant to particular interests defined by a user). Multiple document summaries are more complex, and we did not study them in this experiment. Neither did we study translangual or query-based summarization. However, we are very interested in studying translangual, multi-document, or query-based summarization of noisy documents in the future.

### 3 Results and Analysis

In this section, we present results at each stage of summarization, analyzing the errors made and their effects on the quality of summaries.

#### 3.1 OCR performance

We begin by examining the overall performance of the OCR process. Using standard edit distance techniques [7], we can compare the output of OCR to the ground-truth to classify and quantify the errors that have arisen. We then compute, on a per-character and per-word basis, a figure for average precision (percentage of characters or words recognized that are correct) and recall (percentage of characters or words in the input document that are correctly recognized). As indicated in Table 1, OCR performance varies widely depending on the type of degradation. Precision values are generally higher than recall because, in certain cases, the OCR system failed to produce output for a portion of the page in question. Since we are particularly interested in punctuation due to its importance in delimiting sentence boundaries, we tabulate a separate set of precision and recall values for such characters. Note that these are uniformly lower than the other values in the table. Recall, in particular, is a serious issue; many punctuation marks are missed in the OCR output.

#### 3.2 Sentence boundary errors

Since most summarization systems rely on sentence extraction, it is important to identify sentence boundaries correctly. For clean text, sentence boundary detection is not a big problem; the reported accuracy is usually above 95% [19, 21, 22]. However, since such systems typically depend on punctuation, capitalization, and words immediately preceding and following punctuation to make judgments about potential sentence boundaries, detecting sentence boundaries in noisy documents is a challenge due to the unreliability of such features. Punctuation errors arise frequently in the OCR'ing of degraded page images, as we have just noted.

We tested two tokenizers: one is a rule-based system that relies on heuristics encoded in the program, and the other is a decision tree system that has been

Table 1: OCR performance relative to ground-truth (average precision and recall).

	Per-Character				Per-Word	
	All Symbols		Punctuation		Prec.	Recall
	Prec.	Recall	Prec.	Recall		
OCR.clean	0.990	0.882	0.869	0.506	0.963	0.874
OCR.light	0.897	0.829	0.556	0.668	0.731	0.679
OCR.dark	0.934	0.739	0.607	0.539	0.776	0.608
OCR.fax	0.969	0.939	0.781	0.561	0.888	0.879
OCR.skew	0.991	0.879	0.961	0.496	0.963	0.869

trained on a large amount of data. We are interested in how well these systems perform on noisy documents and the kinds of errors they make.

The experimental results show that for the clean text, the two systems perform almost equally well. We manually checked the results for the four documents and found that both tokenizers made very few errors. There should be 90 sentence boundaries in total. The decision tree tokenizer correctly identified 88 of the sentence boundaries and missed two. The rule-based tokenizer correctly identified 89 of the boundaries and missed one. Neither system made any false positive errors (i.e., they did not break sentences at non-sentence boundaries).

For the noisy documents, however, both tokenizers made significant numbers of errors. The types of errors they made, moreover, were quite different. While the rule-based system made many false negative errors, the decision tree system made many false positive errors. Therefore, the rule-based system identified far fewer sentence boundaries than the truth, while the decision tree system identified far more than the truth.

Table 2 shows the number of sentences identified by each tokenizer for different versions of the documents. As we can see from the table, the noisier the documents, the more errors the tokenizers made. This relationship was demonstrated clearly by the results for the documents with synthetic noise. As the noise rate increases, the number of boundaries identified by the decision tree tokenizer gradually increases, and the number of boundaries identified by the rule-based tokenizer gradually decreases. Both numbers diverge from truth, but they err in opposite directions.

The two tokenizers behaved less consistently on the OCR’ed documents. For OCR.light, OCR.dark, and OCR.fax, the decision tree tokenizer produced more sentence boundaries than the rule-based tokenizer. But for OCR.clean and OCR.skew, the decision tree tokenizer produced fewer sentence boundaries. This may be related to the noise level in the document. OCR.clean and OCR.skew contain

fewer errors than the other noisy versions (recall Table 1). According to our computations, 97% of the words that occurred in OCR.clean or OCR.skew also appeared in the original document, while other OCR’ed documents have a much lower word overlap, as shown in Table 4. This seems to indicate that the decision tree tokenizer tends to identify fewer sentence boundaries than the rule-based tokenizer for clean text or documents with very low levels of noise, but more sentence boundaries when the documents have a relatively high level of noise.

Errors made at this stage are extremely detrimental, since they will propagate to all of the other modules in a summarization system. When a sentence boundary is incorrectly marked, the part-of-speech tagging and the syntactic parsing are likely to fail. Sentence extraction may become problematic; for example, one of the documents in our test set contains 24 sentences, but for one of its noisy versions (OCR.dark), the rule-based tokenizer missed most sentence boundaries and divided the document into only three sentences, making extraction at the sentence level difficult at best.

Since sentence boundary detection is important to summarization, the development of robust techniques that can handle noisy documents is worthwhile. We will return to this point in Section 4.

### 3.3 Parsing errors

Some summarization systems use a part-of-speech tagger or a syntactic parser in their preprocessing steps. To study the errors made at this stage, we piped the results from both tokenizers to the ESG parser, which requires as input divided sentences and returns a parse tree for each input sentence. The parse tree also includes a part-of-speech tag for each word in the sentence.

We computed the percentage of sentences that ESG failed to return a complete parse tree, and used that value as one way of measuring the performance of the parser on the noisy documents. As we can see from Table 3, a significant percentage of noisy sentences were not parsed. Even for the documents

Table 2: Sentence boundary detection results: total number of sentences detected and average words per sentence for two tokenizers. The ground-truth is represented by *Original*.

	Tokenizer 1 (Decision tree)		Tokenizer 2 (Rule-based)	
	Sentences	Avg. words/sent.	Sentences	Avg. words/sent.
Original	88	23	89	22
Snoise.05	95	20	70	27
Snoise.10	97	20	69	28
Snoise.15	105	19	65	30
Snoise.20	109	17	60	31
Snoise.25	121	15	51	35
OCR.clean	77	23	82	21
OCR.light	119	15	64	28
OCR.dark	70	21	46	33
OCR.fax	78	26	75	27
OCR.skew	77	23	82	21

with synthetic noise at a 5% rate, around 60% of the sentences cannot be handled by the parser. For the sentences that were handled, the returned parse trees may not be correct. For example, the sentence “*Internet sites found that almost 90 percent collected personal information from youngsters*” was transformed to “*uInternet sites fo6ndha alQmostK0 percent coll / 9ed pe?*” after adding synthetic noise at a 25% rate. For this noisy sentence, the parser returned a complete parse tree that marked the word “*sites*” as the main verb of the sentence, and tagged all the other words in the sentence as nouns.<sup>1</sup> Although a complete parse tree is returned in this case, it is incorrect. This may explain the phenomenon that the parser returned a higher percentage of complete parse trees for documents with synthetic noise at the 25% rate than for documents with lower levels of noise.

The above results indicate that syntactic parsers may be very vulnerable to noise in a document. Even low levels of noise tend to lead to a significant drop in performance. For documents with high levels of noise, it may be better not to rely on syntactic parsing at all since it will likely fail on a large portion of the text, and even when results are returned, they will be unreliable.

### 3.4 Extract quality versus noise level

In the next step, we studied how the sentence extraction module in a summarization system is affected by noise in the input document. For this, we used a sentence extraction system we had developed previously [13]. The sentence extractor relies on lexical links between words, word frequency, cue phrases, and sentence positions to identify key sentences. We

<sup>1</sup>One reason might be that the tagger is likely to tag unknown words as nouns, since most out-of-vocabulary words are nouns.

Table 3: Percentage of sentences with incomplete parse trees from the ESG parser. Sentence boundaries were first detected using Tokenizer 1 and Tokenizer 2, and divided sentences were given to ESG as input.

	Tokenizer 1	Tokenizer 2
Original	10%	5%
Snoise.05	59%	58%
Snoise.10	69%	71%
Snoise.15	66%	81%
Snoise.20	64%	66%
Snoise.25	58%	76%
OCR.clean	2%	3%
OCR.light	46%	53%
OCR.dark	37%	43%
OCR.fax	37%	30%
OCR.skew	5%	6%

set the summary length parameter as 20% of the number of sentences in the original document. This sentence extraction system does not use results from part-of-speech tagging or syntactic parsing, only the output from the rule-based tokenizer.

Evaluation of noisy document summaries is an interesting problem. Intrinsic evaluation (i.e., asking human subjects to judge the quality of summaries) can be used, but this appears much more complex than intrinsic evaluation for clean documents. When the noise rate in a document is high, even when a summarization system extracts the right sentences, a human subject may still rank the quality of the summary as very low due to the noise. Extrinsic evaluation (i.e., using the summaries to perform certain tasks and measuring how much the summaries help in performing the tasks) is also difficult since the

noise level of extracted sentences can significantly affect the result.

We employed three measures that have been used in the Document Understanding Conference [6] for assessing the quality of generated summaries: unigram overlap between the automatic summary and the human-created summary, bigram overlap, and the simple cosine. These results are shown in Table 4. The unigram overlap is computed as the number of unique words occurring both in the extract and the ideal summary for the document, divided by the total number of unique words in the extract. Bigram overlap is computed similarly, replacing words with bigrams. The simple cosine is computed as the cosine of two document vectors, the weight of each element in the vector being  $1/\sqrt{N}$ , where  $N$  is the total number of elements in the vector.

Not surprisingly, summaries of noisier documents generally have a lower overlap with human-created summaries. However, this can be caused by either the noise in the document or poor performance of the sentence extraction system. To separate these effects and measure the performance of sentence extraction alone, we also computed the unigram overlap, bigram overlap, and cosine between each noisy document and its corresponding original text. These numbers are included in Table 4 in parentheses; they are an indication of the average noise level in a document. For instance, the table shows that 97% of words that occurred in OCR.clean documents also appeared in the original text, while only 62% of words that occurred in OCR.light appeared in the original. This confirms that OCR.clean is less noisy than OCR.light.

### 3.5 Abstract generation for noisy documents

To generate more concise and coherent summaries, a summarization system may edit extracted sentences. To study how this step in summarization is affected by noise, we tested a cut-and-paste system that edits extracted sentences by simulating revision operations often used by human abstractors, including the operations of removing phrases from an extracted sentence, and combining a reduced sentence with other sentences. This cut-and-paste stage relies on the results from sentence extraction in the previous step, the output from ESG, and a co-reference resolution algorithm.

For the clean text, the cut-and-paste system performed sentence reduction on 59% of the sentences that were extracted in the sentence extraction step, and sentence combination on 17% of the extracted sentences. For the noisy text, however, the system applied very few revision operations to the extracted (noisy) sentences. Since the cut-and-paste system

relies on the output from ESG and co-reference resolution, which failed on most of the noisy text, it is not surprising that it did not perform well under these circumstances. Editing sentences requires a deeper understanding of the document and, as the last step in the summarization pipeline, relies on results from all of the previous steps. Hence, it is affected most severely by noise in the input document.

## 4 Challenges in Noisy Document Summarization

In the previous section, we have presented and analyzed errors at each stage of summarization when applied to noisy documents. The results show that the methods we tested at every step are fragile, susceptible to failures and errors even with slight increases in the noise level of a document. Clearly, much work needs to be done to achieve acceptable performance in noisy document summarization. We need to develop summarization algorithms that do not suffer significant degradation when used on noisy documents. We also need to develop the robust natural language processing techniques that are required by summarization. For example, sentence boundary detection systems that can reliably identify sentence breaks in noisy documents are clearly important. One way to achieve this might be to retrain an existing system on noisy documents so that it will be more tolerant of noise. However, this is only applicable if the noise level is low. Significant work is needed to develop robust methods that can handle documents with high noise levels.

In the remainder of this section, we discuss several issues in noisy document summarization, identifying the problems and proposing possible solutions. We regard this as a first step towards a more comprehensive study on the topic of noisy document summarization.

### 4.1 Choosing an appropriate granularity

It is important to choose an appropriate unit level to represent the summaries. For clean text, sentence extraction is a feasible goal since we can reliably identify sentence boundaries. For documents with very low levels of noise, sentence extraction is still possible since we can probably improve our programs to handle such documents. However, for documents with relatively high noise rates, we believe it is better to forgo sentence extraction and instead favor extraction of keywords or noun phrases, or generation of headline-style summaries. In our experiment, when the synthetic noise rate reached 10% (which is representative of what can happen when real-world documents are degraded), it was already

Table 4: Unigram overlap, bigram overlap, and simple cosine between extracts and human-created summaries (the numbers in parentheses are the corresponding values between the documents and the original text).

	Unigram overlap	Bigram overlap	Cosine
Original	0.85 (1.00)	0.75 (1.00)	0.51 (1.00)
Snoise.05	0.55 (0.61)	0.38 (0.50)	0.34 (0.65)
Snoise.10	0.41 (0.41)	0.22 (0.27)	0.25 (0.47)
Snoise.15	0.25 (0.26)	0.10 (0.13)	0.20 (0.31)
Snoise.20	0.17 (0.19)	0.04 (0.07)	0.14 (0.23)
Snoise.25	0.18 (0.14)	0.04 (0.04)	0.09 (0.16)
OCR.clean	0.86 (0.97)	0.78 (0.96)	0.50 (0.93)
OCR.light	0.62 (0.63)	0.47 (0.55)	0.36 (0.65)
OCR.dark	0.81 (0.70)	0.73 (0.65)	0.38 (0.66)
OCR.fax	0.77 (0.84)	0.67 (0.79)	0.48 (0.86)
OCR.skew	0.84 (0.97)	0.74 (0.96)	0.48 (0.93)

difficult for a human to recover the information intended to be conveyed from the noisy documents.

Keywords, noun phrases, or headline-style summaries are informative indications of the main topic of a document. For documents with high noise rates, extracting keywords or noun phrases is a more realistic and attainable goal than sentence extraction. Still, it may be desirable to correct the noise in the extracted keywords or phrases. There has been past work on correcting spelling mistakes and errors in OCR output; these techniques would be useful in noisy document summarization.

To choose an appropriate granularity for summary presentation, we need to have an assessment of the noise level in the document. In subsection 4.3, we discuss ways to measure this quantity.

## 4.2 Using other information sources

In addition to text, target documents contain other types of useful information that could be employed in creating summaries. As noted previously, Chen and Bloomberg’s image-based summarization technique avoids many of the problems we have been discussing by exploiting document layout features. A possible approach to summarizing noisy documents, then, might be to use their method to create an image summary and then apply OCR afterwards to the resulting page. We note, though, that it seems unlikely this would lead to an improvement of the overall OCR results, a problem which may almost certainly must be faced at some point in the process.

## 4.3 Assessing error rates without ground-truth

The quality of summarization is directly tied to the level of noise in a document. Summarization results are not seriously impacted in the presence of mi-

nor errors, but as errors increase, the summary may range from being difficult to read to incomprehensible.

In this context, it would be useful to develop methods for assessing document noise levels without having access to the ground-truth. Such measurements could be incorporated into summarization algorithms for the purpose of avoiding problematic regions, thereby improving the overall readability of the summary. Past work on attempting to quantify document image quality for predicting OCR accuracy [2, 3, 9] addresses a related problem, but one which exhibits some significant differences.

Intuitively, OCR may create errors that cause the output text to deviate from “normal” text. Therefore, one way of evaluating OCR output, in the absence of the original ground-truth, is to compare its features against features obtained from a large corpus of correct text. Letter trigrams [5] are commonly used to correct spelling and OCR errors [1, 15, 26], and can be applied to evaluate OCR output.

We computed trigram tables (including symbols and punctuation marks) for 10 days of AP news articles and evaluated the documents used in our experiment. As expected, OCR errors create rare or previously unseen trigrams that lead to higher trigram scores in noisy documents. As indicated in Table 5, the ground-truth (original) documents have the lowest average trigram score. These scores provide a relative ranking that reflects the controlled noise levels (Snoise.05 through Snoise.25), as well as certain of the real OCR data (OCR.clean, OCR.dark, and OCR.light).

Different texts have very different baseline trigram scores. The ranges of scores for clean and noisy text overlap. This is because some documents contain more instances of frequent words than others (such as “the”), which bring down the average scores. This

Table 5: Average trigram scores.

	Trigram score
Original	2.30
Snoise.05	2.75
Snoise.10	3.13
Snoise.15	3.50
Snoise.20	3.81
Snoise.25	4.14
OCR.clean	2.60
OCR.light	3.11
OCR.dark	2.98
OCR.fax	2.55
OCR.skew	2.40

issue makes it impractical to use trigram scores in isolation to judge OCR output.

It may be possible to identify some problems if we scan larger units and incorporate contextual information. For example, a window of three characters is too small to judge whether the symbol @ is used properly: *a@b* seems to be a potential OCR error, but is acceptable when it appears in an email address such as *lsa@bbb.com*. Increasing the unit size will create sparse data problems, however, which is already an issue for trigrams.

In the future, we plan to experiment with improved methods for identifying problematic regions in OCR text, including using language models and incorporating grammatical patterns. Many linguistic properties can be identified when letter sequences are encoded in broad classes. For example, long consonant strings are rare in English text, while long number strings are legal. These properties can be captured when characters are mapped into carefully selected classes such as symbols, numbers, upper- and lower-case letters, consonants, and vowels. Such mappings effectively reduce complexity, allowing us to sample longer strings to scan for abnormal patterns without running into severe sparse data problems.

Our intention is to establish a robust index that measures whether a given section of text is “summarizable.” This problem is related to the general question of assessing OCR output without ground-truth, but we shift the scope of the problem to ask whether the text is summarizable, rather than how many errors it may contain.

We also note that documents often contain logical components that go beyond basic text. Pages may include photographs and figures, program code, lists, indices, etc. Tables, for example, can be detected, parsed, and reformulated so that it becomes possible to describe their overall structure and even allow

users to query them [12]. Developing appropriate ways of summarizing such material is another topic of interest.

## 5 Conclusions and Future Work

In this paper, we have discussed some of the challenges in summarizing noisy documents. In particular, we broke down the summarization process into four steps: sentence boundary detection, preprocessing (part-of-speech tagging and syntactic parsing), extraction, and editing. We tested each step on noisy documents and analyzed the errors that arose. We also studied how the quality of summarization is affected by the noise level and the errors made at each stage of processing.

To improve the performance of noisy document summarization, we suggest extracting keywords or phrases rather than full sentences, especially when summarizing documents with high levels of noise. We also propose using other sources of information, such as document layout cues, in combination with text when summarizing noisy documents. In certain cases, it will be important to be able to assess the noise level in a document; we have begun exploring this question as well. Our plans for the future include developing robust techniques to address the issues we have outlined in this paper.

Lastly, we regard presentation and user interaction as a crucial component in real-world summarization systems. Given that noisy documents, and hence their summaries, may contain errors, it is important to find the best ways of displaying such information so that the user may proceed with confidence, knowing that the summary is truly representative of the document(s) in question.

## References

- [1] R. Angell, G. Freund, and P. Willet. Automatic spelling correction using a trigram similarity measure. *Information Processing and Management*, 19(4):255–261, 1983.
- [2] L. R. Blando, J. Kanai, and T. A. Nartker. Prediction of OCR accuracy using simple image features. In *Proceedings of the Third International Conference on Document Analysis and Recognition*, pages 319–322, Montréal, Canada, August 1995.
- [3] M. Cannon, J. Hochberg, and P. Kelly. Quality assessment and restoration of typewritten document images. Technical Report LA-UR 99-1233, Los Alamos National Laboratory, 1999.
- [4] F. R. Chen and D. S. Bloomberg. Summarization of imaged documents without OCR.



- Computer Vision and Image Understanding*, 70(3):307–320, 1998.
- [5] K. Church and W. Gale. Probability scoring for spelling correction. *Statistics and Computing*, 1:93–103, 1991.
- [6] Document Understanding Conference (DUC): Workshop on Text Summarization, 2002. <http://tides.nist.gov/>.
- [7] J. Esakov, D. P. Lopresti, and J. S. Sandberg. Classification and distribution of optical character recognition errors. In *Proceedings of Document Recognition I (IS&T/SPIE Electronic Imaging)*, volume 2181, pages 204–216, San Jose, CA, February 1994.
- [8] Y. Gotoh and S. Renals. Sentence boundary detection in broadcast speech transcripts. In *Proceedings of ISCA Tutorial and Research Workshop ASR-2000*, Paris, France, 2000.
- [9] V. Govindaraju and S. N. Srihari. Assessment of image quality to predict readability of documents. In *Proceedings of Document Recognition III (IS&T/SPIE Electronic Imaging)*, volume 2660, pages 333–342, San Jose, CA, January 1996.
- [10] D. Harman and M. Liberman. *TIPSTER Complete*. Linguistic Data Consortium, University of Pennsylvania, 1993. LDC catalog number: LDC93T3A. ISBN: 1-58563-020-9.
- [11] C. Hori and S. Furui. Advances in automatic speech summarization. In *Proceedings of the 7th European Conference on Speech Communication and Technology*, pages 1771–1774, Aalborg, Denmark, 2001.
- [12] J. Hu, R. Kashi, D. Lopresti, and G. Wilfong. A system for understanding and reformulating tables. In *Proceedings of the Fourth IAPR International Workshop on Document Analysis Systems*, pages 361–372, Rio de Janeiro, Brazil, December 2000.
- [13] H. Jing. *Cut-and-paste Text Summarization*. PhD thesis, Department of Computer Science, Columbia University, New York, NY, 2001.
- [14] H. Jing, R. Barzilay, K. McKeown, and M. Elhadad. Summarization evaluation methods: experiments and analysis. In *Working Notes of AAAI Symposium on Intelligent Summarization*, Stanford University, CA, March 1998.
- [15] K. Kuckich. Techniques for automatically correcting words in text. *ACM Computing Surveys*, 24(4):377–439, 1992.
- [16] I. Mani. *Automatic Summarization*. John Benjamins Publishing Company, Amsterdam/Philadelphia, 2001.
- [17] M. McCord. *English Slot Grammar*. IBM, 1990.
- [18] D. Miller, S. Boisen, R. Schwartz, R. Stone, and R. Weischedel. Named entity extraction from noisy input: Speech and OCR. In *Proceedings of the 6th Applied Natural Language Processing Conference*, pages 316–324, Seattle, WA, 2000.
- [19] D. Palmer and M. Hearst. Adaptive multilingual sentence boundary disambiguation. *Computational Linguistics*, 23(2):241–267, June 1997.
- [20] D. D. Palmer and M. Ostendorf. Improving information extraction by modeling errors in speech recognizer output. In J. Allan, editor, *Proceedings of the First International Conference on Human Language Technology Research*, 2001.
- [21] J. C. Reyner and A. Ratnaparkhi. A maximum entropy approach to identifying sentence boundaries. In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, Washington D.C., 1997.
- [22] M. Riley. Some applications of tree-based modelling to speech and language. In *Proceedings of the DARPA Speech and Natural Language Workshop*, pages 339–352, Cape Cod, MA, 1989.
- [23] K. Taghva, J. Borsack, and A. Condit. Effects of OCR errors on ranking and feedback using the vector space model. *Information Processing and Management*, 32(3):317–327, 1996.
- [24] K. Taghva, J. Borsack, and A. Condit. Evaluation of model-based retrieval effectiveness with OCR text. *ACM Transactions on Information Systems*, 14:64–93, January 1996.
- [25] Translingual Information Detection, Extraction and Summarization (TIDES). <http://www.darpa.mil/iao/tides.htm>.
- [26] E. Zamora, J. Pollock, and A. Zamora. The use of trigram analysis for spelling error detection. *Information Processing and Management*, 17(6):305–316, 1981.