



Multi-candidate reduction: Sentence compression as a tool for document summarization tasks

David Zajic^{a,*}, Bonnie J. Dorr^a, Jimmy Lin^a, Richard Schwartz^b

^a University of Maryland, College Park, MD 20742, United States

^b BBN Technologies, 9861 Broken Land Parkway, Columbia, MD 21046, United States

Received 18 July 2006; received in revised form 3 January 2007; accepted 8 January 2007

Abstract

This article examines the application of two single-document sentence compression techniques to the problem of multi-document summarization—a “parse-and-trim” approach and a statistical noisy-channel approach. We introduce the multi-candidate reduction (MCR) framework for multi-document summarization, in which many compressed candidates are generated for each source sentence. These candidates are then selected for inclusion in the final summary based on a combination of static and dynamic features. Evaluations reveal that sentence compression is a valuable component of a larger multi-document summarization framework.

© 2007 Published by Elsevier Ltd.

PACS: 07.05.Mh; 89.20.Ff; 43.71.Sy

Keywords: Headline generation; Summarization; Parse-and-trim; Hidden Markov model

1. Introduction

This article presents an extension of two different single-document sentence compression methods to the problem of multi-document summarization. The first, a “parse-and-trim” approach, has been implemented in a system called Trimmer and its extended version called Topiary. The second, an HMM-based approach, has been implemented in a system called HMM Hedge.

These systems share the basic premise that a textual summary can be constructed by selecting a subset of words, in order, from the original text, with morphological variation allowed for some word classes. Trimmer selects subsequences of words using a linguistically motivated algorithm to trim syntactic constituents from sentences until a desired length has been reached. In the context of very short single-document summarization, or headline generation, this approach is applied to the lead sentence (the first non-trivial sentence) in order to

* Corresponding author. Tel.: +1 410 992 1444.

E-mail addresses: dmzajic@cs.umd.edu (D. Zajic), bonnie@umiacs.umd.edu (B.J. Dorr), jimmylin@umd.edu (J. Lin), schwartz@bbn.com (R. Schwartz).

produce a short headline tens of characters long. HMM Hedge (Hidden Markov Model HEadline GEnerator), on the other hand, is a statistical headline generation system that finds the subsequence of words most likely to be a headline for a given story. This approach is similar to techniques used in statistical machine translation in that the observed story is treated as the garbled version of an unseen headline transmitted through a noisy channel. In their original configurations, both Trimmer and HMM Hedge create sentence compressions that mimic the style of Headlines, a form of compressed English associated with newspaper headlines (Mårdh, 1980).

The goal of our work is to transition from a single-document summarization task with sentence-level character-based length constraint to a multi-document summarization task with summary-level word-based length constraint. In headline generation, summaries should fill as much of the available space as possible, without going over the limit. Headlines which fall below the length limit miss the opportunity to include additional information. Headlines which exceed the length limit must be truncated, resulting in ungrammatical sentences and loss of information. When the length constraint is applied to a collection of sentences instead of a single sentence, concerns about individual sentence length become less important. A short sentence does not result in wasted space, and it is acceptable to include a lengthy sentence as long as it includes important, non-redundant information.

To transition from single-document to multi-document summarization, we examined different combinations of possible sentence compressions to construct the best summary. A key innovation of this work—tested for the first time in 2005 at the Document Understanding Conference (DUC-2005) (Zajic, Dorr, Schwartz, Monz, & Lin, 2005b) and Multilingual Summarization Evaluation (MSE-2005) (Zajic, Dorr, Lin, & Schwartz, 2005a)—is the use of multiple compressed candidates for sentences in the source documents. Sentence compression is used for different purposes in single- and multi-document summarization tasks. In the first case, the goal is to fit long sentences into the available space while preserving important information. In multi-document summarization, sentence compression is used to generate multiple candidates that capture relevant, non-redundant information. Sentence selection algorithms can then be applied to determine which compressed candidates provide the best combination of topic coverage and brevity.

We introduce a framework for extractive multi-document summarization called multi-candidate reduction (MCR), whose architecture is shown in Fig. 1. Sentence compression techniques are employed in the intermediate stage of the processing pipeline to generate multiple compressed variants of source sentences. A sentence selector then builds the final summary by choosing among these candidates, based on features propagated from the sentence compression method, features of the candidates themselves, and features of the present summary state. Under this framework, our previously developed systems, Trimmer and HMM Hedge, are employed to generate the compressed candidates. The MCR architecture also includes a filtering module, which chooses the source sentences from which compressed candidates are generated. This work does not examine the filtering process in detail, as we only employ very simple approaches, e.g., retain first n sentences from each document.

Sentence compression supports extractive multi-document summarization by reducing the length of summary candidates while preserving their relevant content, thus allowing space for the inclusion of additional material. However, given the interaction of relevance and redundancy in a multi-sentence summary, it is

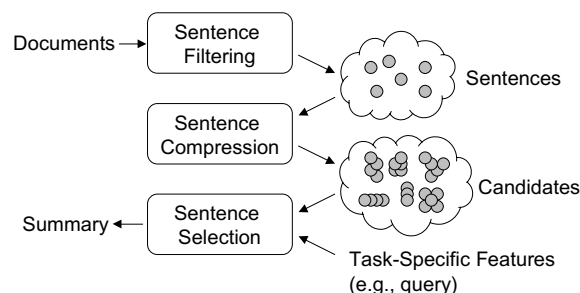


Fig. 1. The multi-candidate reduction (MCR) framework.

unlikely that a single algorithm or scoring metric can provide the “best” compression of a sentence. This is the motivation for MCR, which provides several alternative candidates for each source sentence. Subsequent processes can then select among many alternatives when constructing the final summary.

This article is organized as follows: The next section relates our approach to other existing summarization systems. In Section 3, we describe Trimmer, a variant of Trimmer called Topiary, and HMM Hedge. These systems implement two fundamentally different approaches to sentence compression. Section 4 describes how these tools can be applied to a multi-document summarization task. Section 5 describes evaluations of our framework. Our systems have also been applied to several different types of texts; some of these applications will be briefly covered in Section 6. Finally, we propose future work in the area before concluding.

2. Related work

We have developed a “parse-and-trim” approach to sentence compression, which is implemented in a system called Trimmer (Dorr, Zajic, & Schwartz, 2003b). The system generates a headline for a news story by compressing the *lead* (or main) topic sentence according to a linguistically motivated algorithm that operates on syntactic structures. Syntactic approaches to compression have been used in single-document summarization systems such as Cut-and-Paste (Jing & McKeown, 2000) and also in multi-document summarization systems such as SC (Blair-Goldensohn et al., 2004) and CLASSY (Conroy, Schlesinger, & Stewart, 2005; Conroy, Schlesinger, O’Leary, & Goldstein, 2006). The SC system pre-processes input to remove appositives and relative clauses. CLASSY uses an HMM sentence selection approach combined with a conservative sentence compression method based on shallow parsing to detect lexical cues to trigger phrase eliminations. The approach taken by Trimmer is most similar to that of Jing and McKeown (2000), in which algorithms for removing syntactic constituents from sentences are informed by analysis of human summarizers. Trimmer differs from these systems in that multiple compressed candidates of each sentence are generated. The potential of multiple alternative compressions has also been explored by Vanderwende, Suzuki, and Brockett (2006).

Summaries may also contain lists of words or short phrases that denote important topics or concepts in the document. In particular, extractive topic summaries consist of keywords or key phrases that occur in the document. We have built Topiary, an extension to Trimmer, that constructs headlines combining compressed sentences with topic terms. This approach is similar to the work of Euler et al. (2002), except that Euler uses topic lists to guide sentence selection and compression toward a query-specific summary, whereas Topiary uses topics to augment the concept coverage of a generic summary. Our system was demonstrated to be the top-performing single-document summarization system in the DUC-2004 evaluation (Zajic, Dorr, & Schwartz, 2004).

In contrast to Topiary, which combines sentence compression with topic terms, others have constructed summaries directly from topic terms. For example, Bergler, Witte, Khalife, Li, and Rudzicz (2003) choose noun phrases that represent the most important entities as determined by noun phrase coreference chains. Wang et al. (2005) propose a baseline system that constructs headlines from topic descriptors identified using term frequency counts; this system was reported to outperform LexTrim, their Topiary-style system. Zhou and Hovy (2003) construct fluent summaries from a topic list by finding phrase clusters early in the text that contain important topic words found throughout the text.

The task of assigning topic terms to documents is related to text categorization, in which documents are assigned to pre-defined categories. The categories can be labeled with topic terms, so that the decision to put a document in a category is equivalent to assigning that category’s label to the document. Assigning topic terms to documents by categorization permits the assignment of terms that do not occur in the document. Lewis (1999) describes a probabilistic feature-based method for assigning Reuters topics to news stories. Along these lines, OnTopicTM (Schwartz, Imai, Jubala, Nguyen, & Makhoul, 1997) uses an HMM to assign the topics to a document.

In addition to Trimmer, we have implemented HMM Hedge (Hidden Markov Model HEaDline Generator), a statistical headline generation system that finds the most likely headline for a given story. This approach is similar to techniques used in statistical machine translation in that the observed story is treated as the garbled version of an unseen headline transmitted through a noisy channel. The noisy-channel approach has been used for a wide range of Natural Language Processing (NLP) applications including speech recognition (Bahl, Jelinek, & Mercer, 1983); machine translation (Brown et al., 1990); spelling correction (Mays,

119 Damerau, & Mercer, 1990); language identification (Dunning, 1994); and part-of-speech tagging (Cutting,
120 Pedersen, & Sibun, 1992). Of those who have taken a noisy-channel approach to sentence compression for
121 text summarization Banko, Mittal, and Witbrock (2000), Knight and Marcu (2000, 2002), and Turner and
122 Charniak (2005) have used the technique to find the most probable short string that generated the observed
123 full sentence.

124 Like other summarization systems based on the noisy-channel model, HMM Hedge treats the observed
125 data (the story) as the result of unobserved data (headlines) that have been distorted by transmission through
126 a noisy channel. The effect of the noisy channel is to add story words between the headline words. Our
127 approach differs from Knight and Marcu (2000), Banko et al. (2000), and Turner and Charniak (2005), in that
128 HMM Hedge does not require a corpus of paired stories and summaries. HMM Hedge uses distinct language
129 models of news stories and headlines, but does not require explicit pairings of stories and summaries.

130 To transition our single-sentence summarization techniques to the problem of multi-document summariza-
131 tion, we must consider how to select candidates for inclusion in the final summary. A common approach is to
132 rank candidate sentences according to a set of features and iteratively build the summary, appropriately re-
133 ranking the candidates at each step to avoid redundancy. MEAD (Radev et al., 2004) scores source sentences
134 according to a linear combination of features including centroid, position and first-sentence overlap. These
135 scores are then refined to consider cross-sentence dependencies, such as redundancy, chronological order
136 and source preferences. MCR differs in that multiple variants of the same source sentences are available as
137 candidate for inclusion in the final summary.

138 Minimization of redundancy is an important element of a multi-document summarization system. Carbo-
139 nell and Goldstein (1998) propose a technique called Maximal Marginal Relevance (MMR) for ranking doc-
140 uments returned by an information retrieval system so that the front of the ranked list will contain diversity as
141 well as high relevance. Goldstein, Mittal, Carbonell, and Kantrowitz (2000) extend MMR to multi-document
142 summarization. MCR borrows the ranking approach of MMR, but uses a different set of features. Like
143 MEAD, these approaches use feature weights that were optimized to maximize an automatic metric on train-
144 ing data.

145 Several researchers have shown the importance of summarization in domains other than written news
146 (Clarke & Lapata, 2006; Muresan, Tzoukermann, & Klavans, 2001). Within the MCR framework, we discuss
147 the portability of Trimmer and HMM Hedge to a variety of different texts: written news, broadcast news tran-
148 scriptions, email threads, and text in foreign language.

149 3. Single-sentence compression

150 Our general approach to the generation of a summary from a single document is to produce a *headline* by
151 selecting words in order from the text of the story. Consider the following excerpt from a news story and cor-
152 responding headline:

- 153 (1) (i) After months of debate following the September 11 terrorist hijackings, the Transportation Depart-
154 ment has decided that airline **pilots will not be allowed to have guns in the cockpits**.
155 (ii) Pilots not allowed to have guns in cockpits.

156 The bold words in (1i) form a fluent and accurate headline, as shown in (1ii).

157 This basic approach has been realized in two ways. The first, Trimmer, uses a linguistically motivated algo-
158 rithm to remove grammatical constituents from the lead sentence until a length threshold is met. Topiary is a
159 variant of Trimmer that combines fluent text from a compressed sentence with topic terms to produce head-
160 lines. The second, HMM Hedge, employs a noisy-channel model to find the most likely headline for a given
161 story. The remainder of this section will present Trimmer, Topiary and HMM Hedge in more detail.

163 3.1. Trimmer

164 Our first approach to sentence compression involves iteratively removing grammatical constituents from
165 the parse tree of a sentence using linguistically motivated rules until a length threshold has been met. When

applied to the lead sentence, or first non-trivial sentence of a story, our algorithm generates a very short summary, or headline. This idea is implemented in our Trimmer system, which can leverage the output of any constituency parser that uses the Penn Treebank conventions. At present we use Charniak's parser (Charniak, 2000).

The insights that form the basis and justification for the Trimmer rules come from our previous study, which compared the relative prevalence of certain constructions in human-written summaries and lead sentences in stories. This study used 218 human-written summaries of 73 documents from the TIPSTER corpus (Harman & Liberman, 1993) dated January 1, 1989. The 218 summaries and the lead sentences of the 73 stories were parsed using the BBN SIFT parser (Miller, Ramshaw, Fox, & Weischedel, 2000). The parser produced 957 noun phrases (NP nodes in the parse trees) and 315 clauses (S nodes in the parse trees) for the 218 summaries. For the 73 lead sentences, the parser produced 817 noun phrases and 316 clauses.

At each level (sentence, clause, and noun phrase), different types of linguistic phenomena were counted.

- At the sentence level, the numbers of preposed adjuncts, conjoined clauses, and conjoined verb phrases were counted. Children of the root S node that occur to the left of the first NP are considered to be preposed adjuncts. The bracketed phase in “[According to police] the crime rate has gone down” is a prototypical example of a preposed adjunct.
- At the clause level, temporal expressions, trailing SBAR nodes, and trailing PP nodes were counted. Trailing constituents are those not designated as an argument of a verb phrase.
- At both the sentence and clause levels, conjoined S nodes and conjoined VP nodes were counted.
- At the NP level, determiners and relative clauses were counted.

The counts and prevalence of the phenomena in the human-generated headlines and lead sentences are shown in Table 1. The results of this analysis illuminated the opportunities for trimming constituents, and guided the development of our Trimmer rules, detailed below.

3.1.1. Trimmer algorithm

Trimmer applies syntactic compression rules to a parse tree according the following algorithm:

- (1) Remove temporal expressions;
- (2) Select Root S node;
- (3) Remove preposed adjuncts;
- (4) Remove some determiners;
- (5) Remove conjunctions;
- (6) Remove modal verbs;
- (7) Remove complementizer *that*;
- (8) Apply the XP over XP rule;
- (9) Remove PPs that do not contain named entities;

Table 1
Counts and prevalence of phenomena found in summaries and lead sentences

Level	Phenomenon	Summary		Lead sentence	
Sentence	Preposed adjuncts	0/218	0%	2/73	2.7%
	Conjoined S	1/218	0.5%	3/73	4%
	Conjoined VP	7/218	3%	20/73	27%
Clause	Temporal expression	5/315	1.5%	77/316	24%
	Trailing PP	165/315	52%	184/316	58%
	Trailing SBAR	24/315	8%	49/316	16%
Noun phrase	Relative clause	3/957	0.3%	29/817	3.5%
	Determiner	31/957	3%	205/817	25%

- 201 (10) Remove all PPs under SBARs;
- 202 (11) Remove SBARs;
- 203 (12) Backtrack to state before Step 9;
- 204 (13) Remove SBARs;
- 205 (14) Remove PPs that do not contain named entities;
- 206 (15) Remove all PPs.

207

208 Steps 1 and 4 of the algorithm remove low-content units from the parse tree.

209 Temporal expressions—although certainly not content-free—are not usually vital for summarizing the con-
 210 tent of an article. Since the goal is to provide an informative headline, the identification and elimination of
 211 temporal expressions (Step 1) allows other more important details to remain in the length-constrained head-
 212 line. The use of BBN's *IdentiFinder*TM (Bikel, Schwartz, & Weischedel, 1999) for removal of temporal expres-
 213 sions is described in Section 3.1.2.

214 The determiner rule (Step 4) removes leaf nodes that are assigned the part-of-speech tag DT and have the
 215 surface form *the*, *a* or *an*. The intuition for this rule is that the information carried by articles is expendable in
 216 summaries, even though this makes the summaries ungrammatical for general English. Omitting articles is one
 217 of the most salient features of newspaper headlines. Sentences (2) and (3), taken from the New York Times
 218 website on September 27, 2006, illustrate this phenomenon. The italicized articles did not occur in the actual
 219 newspaper headlines.

220 (2) *The* Gotti Case Ends With *a* Mistrial for *the* Third Time in a Year.

221 (3) *A* Texas Case Involving Marital Counseling Is *the* Latest to Test *the* Line Between Church and State.

222

223 Step 2 identifies nodes in a parse tree of a sentence that could serve as the root of a compression for the
 224 sentence. Such nodes will be referred to as Root S nodes. A node in a tree is a Root S node if it is labeled
 225 S in the parse tree and has children labeled NP and VP, in that order. The human-generated headlines we stud-
 226 ied always conform to this rule. It has been adopted as a constraint in the Trimmer algorithm that the lowest
 227 leftmost Root S node is taken to be the root node of the headline. An example of this rule application is shown
 228 in (4). The boldfaced material in the parse is retained and the italicized material is eliminated.

229 (4) (i) Input: Rebels agreed to talks with government officials, international observers said Tuesday.

230 (ii) Parse: [S [S [NP **Rebels**][VP **agreed to talks with government officials**]], *international observers said*
 231 *Tuesday*].

232 (iii) Output: Rebels agreed to talks with government officials.

233

234 When the parser produces a correct parse tree, this rule selects a valid starting point for compression. How-
 235 ever, the parser sometimes produces incorrect output, as in the cases below (from DUC-2003):

236 (5) (i) Parse: [S[SBAR **What started as a local controversy**][VP **has evolved into an international scandal**]].

237 (ii) Parse: [NP[NP **Bangladesh**][CC **and**][NP[NP **India**][VP **signed a water sharing accord**]].

238

239 In (5i), an S exists, but it does not conform to the requirements of the Root S rule because it does not have
 240 as children an NP followed by a VP. The problem is resolved by selecting the lowest leftmost S, ignoring the
 241 constraints on the children. In (5ii), no S is present in the parse. This problem is resolved by selecting the root
 242 of the entire parse tree as the root of the headline. These parsing errors occur infrequently—only 6% of the
 243 sentences in the DUC-2003 evaluation data exhibit these problems, based on parses generated by the BBN
 244 SIFT parser.

245 The motivation for removing preposed adjuncts (Step 3) is that all of the human-generated headlines omit
 246 what we refer to as the *preamble* of the sentence. Preposed adjuncts are constituents that precede the first NP
 247 (the subject) under the Root S chosen in Step 2; the preamble of a sentence consists of its preposed adjuncts.
 248 The impact of preposed adjunct removal can be seen in example (6).

- (6) (i) Input: According to a now finalized blueprint described by US officials and other sources, the Bush administration plans to take complete, unilateral control of a post-Saddam Hussein Iraq.
- (ii) Parse: [S [PP According to a now finalized blueprint described by US officials and other sources], [Det the] **Bush administration plans to take complete, unilateral control of [Det a] post-Saddam Hussein Iraq**].
- (iii) Output: Bush administration plans to take complete unilateral control of post-Saddam Hussein Iraq.

The remaining steps of the algorithm remove linguistically peripheral material through successive deletions of constituents until the sentence is shorter than a length threshold. Each stage of the algorithm corresponds to the application of one of the rules. Trimmer first finds the pool of nodes in the parse to which a rule can be applied. The rule is then iteratively applied to the deepest, rightmost remaining node in the pool until the length threshold is reached or the pool is exhausted. After a rule has been applied at all possible nodes in the parse tree, the algorithm moves to the next step.

In the case of a conjunction with two children (Step 5), one of the children will be removed. If the conjunction is *and*, the first child is removed. If the conjunction is *but*, the second child is removed. This rule is illustrated by the following examples, where the struck out text is trimmed.

- (7) When Sotheby's sold a Harry S Truman signature that turned out to be a reproduction, the prestigious auction house apologized *and bought it back*.
- (8) President Clinton *expressed sympathy after a car-bomb explosion in a Jerusalem market wounded 24 people but* said the attack should not derail the recent land-for-security deal between Israel and the Palestinians.

The modal verb rule (Step 6) applies to verb phrases in which the head is a modal verb, and the head of the child verb phrase is a form of *have* or *be*. In such cases, the modal and auxiliary verbs are removed. Sentences (9) and (10) show examples of this rule application. Note that although in Sentence (10), the omission of trimmed material changes the meaning, given a tight space constraint, the loss of the modality is preferable to the loss of other content information.

- (9) People's palms and fingerprints *may be* used to diagnose schizophrenia.
- (10) Agents *may have* fired potentially flammable tear gas canisters.

The complementizer rule (Step 7) removes the word *that* when it occurs as a complementizer. Sentence (11) shows an example in which two complementizers can be removed.

- (11) Hoffman stressed *that* study is only preliminary and can't prove *that* treatment useful.

The XP-over-XP rule (Step 8) is a linguistic generalization that allows a single rule to cover two different phenomena. XP in the name of the rule is a variable that can take two values: NP and VP. In constructions of the form [XP [XP ...] ...], the other children of the higher XP are removed. Note that the child XP must be the first child of the parent XP. When XP = NP the rule removes relative clauses (as in Sentence (12)) and appositives (as in Sentence (13)).

- (12) Schizophrenia patients *whose medication couldn't stop the imaginary voices in their heads* gained some relief.
- (13) A team led by Dr. Linda Brzustowicz, *assistant professor of neuroscience at Rutgers University's Center for Molecular and Behavioral Neuroscience in Newark*, studied DNA of dozens of members of 22 families.

The rules that remove prepositional phrases and subordinate clauses (Steps 9–15) are sometimes prone to removing important content. Thus, these rules are applied last, only when there are no other types of rules to apply. Moreover, these rules are applied with a backoff option to avoid over-trimming the parse tree. First, the

PP rule is applied (Steps 9 and 10),¹ followed by the SBAR rule (Step 11). If the desired sentence length has not been reached, the system reverts to the parse tree as it was before any prepositional phrases were removed (Step 12) and applies the SBAR rule (Step 13). If the desired length still has not been reached, the PP rule is applied again (Steps 14 and 15).

The intuition behind this ordering is that, when removing constituents from a parse tree, it is preferable to remove smaller fragments before larger ones, and prepositional phrases tend to be smaller than subordinate clauses. Thus, Trimmer first attempts to achieve the desired length by removing smaller constituents (PPs), but if this cannot be accomplished, the system restores the smaller constituents, removes a larger constituent, and then resumes the deletion of the smaller constituents. To reduce the risk of removing prepositional phrases that contain important information, BBN's *IdentiFinder* is used to distinguish PPs containing temporal expressions and named entities, as described next.

3.1.2. Use of BBN's *IdentiFinder* in Trimmer

BBN's *IdentiFinder* is used both for the elimination of temporal expressions and for conservative deletion of PPs containing named entities. The elimination of temporal expressions (Step 1) is a two-step process: (a) use *IdentiFinder* to mark temporal expressions; and (b) remove [PP ... [NP [X] ...] ...] and [NP [X]] where X is tagged as part of a temporal expression. The following examples illustrate the application of temporal expression removal rule:

- (14) (i) Input: The State Department on Friday lifted the ban it had imposed on foreign fliers.
 (ii) Parse: [S [NP [Det The] **State Department**] [PP [IN on] [NP [NNP Friday]]] [VP lifted [Det the] **ban it had imposed on foreign fliers**]].
 (iii) Output: State Department lifted ban it had imposed on foreign fliers.
 (15) (i) Input: An international relief agency announced Wednesday that it is withdrawing from North Korea.
 (ii) Parse: [S [NP [Det An] **international relief agency**] [VP announced [NP [NNP Wednesday]]] **that it is withdrawing from North Korea**]].
 (iii) Output: International relief agency announced that it is withdrawing from North Korea.

IdentiFinder is also used to ensure that prepositional phrases containing named entities are not removed during the first round of PP removal (Step 9). However, prepositional phrases containing named entities that are descendants of SBARs are removed before the parent SBAR is removed, since we should remove a smaller constituent before removing a larger constituent that subsumes it. Sentence (16) shows an example of a SBAR subsuming two PPs, one of which contains a named entity.

- (16) The commercial fishing restrictions in Washington will not be lifted [SBAR unless the salmon population increases [PP to a sustainable number] [PP in the Columbia River]].

If the PP rule were not sensitive to named entities, the PP *in the Columbia River* would be the first prepositional phrase to be removed, because it is the lowest rightmost PP in the parse. However, this PP provides an important piece of information: the location of the salmon population. The rule in Step 9 will skip the last prepositional phrase and remove the penultimate PP *to a sustainable number*.

This concludes an overview of the Trimmer rules and our syntactic sentence compression algorithm. Given a length limit, the system will produce a single compressed version of the target sentence. Multiple compressions can be generated by setting the length limit to be very small and storing the state of the sentence after each rule application as a compressed variant. In Section 4, we will describe how multiple compressed candidates generated by Trimmer are used as a component of a multi-document summarization system.

¹ The reason for breaking PP removal into two stages is discussed in Section 3.1.2.

340 3.2. Topiary

341 We have used the Trimmer approach to compression in another variant of single-sentence summarization
 342 called Topiary. This system combines Trimmer with a topic discovery approach (described next) to produce a
 343 fluent summary along with additional context.

344 The Trimmer algorithm is constrained to build a headline from a single sentence. However, it is often the
 345 case that no single sentence contains all the important information in a story. Relevant information can be
 346 spread over multiple sentences, linked by anaphora or ellipsis. In addition, the choice of lead sentence may
 347 not be ideal and our trimming rules are imperfect.

348 On the other hand, approaches that construct headlines from lists of topic terms (Lewis, 1999; Schwartz
 349 et al., 1997) also have limitations. For example, Unsupervised Topic Discovery (UTD)—described below—
 350 rarely generates any topic terms that are verbs. Thus, topic lists are good at indicating the general subject
 351 but rarely give any direct indication of what events took place. Intuitively, we need both fluent text to tell what
 352 happened and topic terms to provide context.

353 3.2.1. Topic term generation: UTD and OnTopic

354 OnTopic (Schwartz et al., 1997) uses an HMM to assign topics to a document; topic models are derived
 355 from an annotated corpus. However, it is often difficult to acquire such data, especially for a new genre or
 356 language. UTD (Sista, Schwartz, Leek, & Makhoul, 2002) was developed to overcome this limitation: it takes
 357 as input a large unannotated corpus and automatically creates a set of topic models with meaningful names.

358 The UTD algorithm has several stages. First, it analyzes the corpus to find multi-word sequences that can
 359 be treated as single tokens. It does this using two methods. One method is a minimum description length cri-
 360 terion, which detects phrases that occur frequently relative to the individual words. The second method uses
 361 BBN's IdentiFinder to detect multi-word names. These names are added to the text as additional tokens. They
 362 are also likely to be chosen as potential topic names. In the second stage of UTD, we find those terms (both
 363 single-word and multi-word) with high *tf.idf*. Only those topic names that occur as high-content terms in at
 364 least four different documents are kept. The third stage trains topic models corresponding to these topic terms.
 365 The modified Expectation Maximization procedure of BBN's OnTopic system is used to determine which
 366 words in the documents often signify these topic names. This produces topic models. Fourth, these topic mod-
 367 els are used to find the most likely topics for each document, which is equivalent to assigning the name of the
 368 topic model to the document as a topic term. This often assigns topics to documents where the topic name
 369 does not occur in the document text.

370 We found, in various experiments (Sista et al., 2002), that the topic names derived by this procedure were
 371 usually meaningful and that the topic assignment was about as good as when the topics were derived from a
 372 corpus that was annotated by people. We have also used this procedure on different languages and shown the
 373 same behavior. Since UTD is unsupervised, it can run equally well on a new language, as long as the docu-
 374 ments can be divided into strings that approximate words.

375 The topic list in (17) was generated by UTD and OnTopic for a story about the FBI investigation of the
 376 1998 bombing of the US embassy in Nairobi.

377 (17) BIN LADEN, EMBASSY, BOMBING, POLICE OFFICIALS, PRISON, HOUSE, FIRE, KABILA

378
 379 Topiary uses UTD to generate topic terms for the collection of documents to be summarized, and OnTopic
 380 to assign the topic terms to the documents. The next section will describe how topic terms and sentence com-
 381 pressions are combined to form Topiary summaries.

382 3.2.2. Topiary algorithm

383 As each Trimmer rule is applied to a sentence, the resulting state of the sentence is stored as a compressed
 384 variant of the source sentence. Topiary selects from the variants the longest one such that there is room to
 385 prepend the highest scoring non-redundant topic term. Suppose the highest scoring topic term is “terrorism”
 386 and the length threshold is 75 characters. To make room for the topic “terrorism”, the length threshold is low-
 387 ered by 10 characters: 9 characters for the topic and 1 character as a separator. Thus, Topiary seeks the longest

388 trimmed variant under 65 characters that does not contain the word “terrorism”. If there is no such candidate,
 389 i.e., all the trimmed variants contain the word terrorism, Topiary would consider the second highest scoring
 390 topic word, “bomb”. Topiary would seek the longest trimmed variant under 70 characters that does not con-
 391 tain the word “bomb”. After Topiary has selected a trimmed variant and prepended a topic to it, it checks to
 392 see how much unused space remains under the threshold. Additional topic words are added between the first
 393 topic word and the compressed sentence until all space is exhausted.

394 This process results in a headline that contains one or more main topics about the story and a short sen-
 395 tence that says what happened concerning them. The combination is often more concise than a fully fluent
 396 sentence and compensates for the fact that the information content from the topic and the compressed sen-
 397 tence do not occur together in any single sentence from the source text.

398 As examples, sentences (18) and (19) are the outputs of Trimmer and Topiary, respectively, for the same
 399 story in which UTD selected the topic terms in (17).

400 (18) FBI agents this week began questioning relatives of the victims.

401 (19) BIN LADEN, EMBASSY, BOMBING: FBI agents this week began questioning relatives.

402

403 By combining topics and parse-and-trim compression, Topiary achieved the highest score on the single-doc-
 404 ument summarization task (i.e., headline generation task) in DUC-2004 (Zajic et al., 2004).

405 3.3. HMM Hedge

406 Our second approach to sentence compression, implemented in HMM Hedge, is to treat the observed data
 407 (the story) as the result of unobserved data (headlines) that have been distorted by transmission through a
 408 noisy channel. The effect of the noisy channel is to add story words between the headline words. The model
 409 is biased by parameters to make the resulting headlines more like Headlines, the observed language of news-
 410 paper headlines created by copy editors.

411 Formally, we consider a story S to be a sequence of N words. We want to find a headline H , a subsequence
 412 of words from S , that maximizes the likelihood that H generated the story S , or

$$414 \quad \operatorname{argmax}_H P(H|S)$$

415 It is difficult to directly estimate $P(H|S)$, but this probability can be expressed in terms of other probabilities
 416 that are easier to compute, using Bayes' rule:

$$418 \quad P(H|S) = \frac{P(S|H)P(H)}{P(S)}$$

419 Since the goal is to maximize this expression over H , and $P(S)$ is constant with respect to H , the denominator
 420 of the above expression can be omitted. Thus we wish to find:

$$422 \quad \operatorname{argmax}_H P(S|H)P(H)$$

423 Let H be a headline consisting of words h_1, h_2, \dots, h_n . Let the special symbols *start* and *end* represent the begin-
 424 ning and end of a headline, respectively. $P(H)$ can be estimated using a bigram model of Headlines:

$$426 \quad P(H) = P(h_1|start)P(h_2|h_1) \dots P(h_n|end)$$

427 The bigram probabilities of the words in the headline language were computed from a corpus of English head-
 428 lines of 242,918 AP newswire stories from the TIPSTER corpus. The headlines contain 2,848,194 words from
 429 a vocabulary of 88,627 distinct words.

430 Given a story S and a headline H , the action of the noisy channel is to form S by adding non-headline
 431 words to H . Let G be the non-headline words added by the channel to the headline: g_1, g_2, \dots, g_m . For the
 432 moment, we assume that the headline words are transmitted through the channel with probability 1. We esti-
 433 mate $P(S|H)$, the probability that the channel added non-headline words G to headline H to form story S .
 434 This is accomplished using a unigram model of newspaper stories, which we will refer to as the general lan-
 435 guage, in contrast to the headline language. Let $P_{\text{gl}}(g)$ be the probability of non-headline word g in the general

language, and $P_{\text{ch}}(h) = 1$ be the probability that headline word h is transmitted through the channel as story word h .

$$P(S|H) = P_{\text{gl}}(g_1)P_{\text{gl}}(g_2) \dots P_{\text{gl}}(g_m)P_{\text{ch}}(h_1)P_{\text{ch}}(h_2) \dots P_{\text{ch}}(h_n) = P_{\text{gl}}(g_1)P_{\text{gl}}(g_2) \dots P_{\text{gl}}(g_m)$$

The unigram probabilities of the words in the general language were computed from 242,918 English AP news stories in the TIPSTER corpus. The stories contain 135,150,288 words from a vocabulary of 428,633 distinct words.

The process by which the noisy channel generates a story from a headline can be represented by a Hidden Markov Model (HMM) (Baum, 1972). An HMM is a weighted finite-state automaton in which each state probabilistically emits a string. The simplest HMM for generating headlines consists of two states: an H state which emits words that occur in the headline, and a G state which emits all the other words in the story.

Since we use a bigram model of headlines, each state which emits headline words must “remember” the previously emitted headline word. If we did not constrain headline words to actually occur in the story, we would need an H state for each word in the headline vocabulary. However, because headline words are chosen from the story words, it is sufficient to have an H state for each story word. For any story, the HMM consists of a start state S , end state E , an H state for each word in the story, a corresponding G state for each H state, and a state G_{start} that emits words which occur before the first headline word in the story. An H state can emit only the word it represents. The corresponding G state remembers which word was emitted by its H state and can emit any word in the story language. A headline corresponds to a path through the HMM from S to E that emits all the words in the story in the correct order. In practice the HMM is constructed with states for only the first N words of the story, where N is a constant (60), or N is the number of words in the first sentence.²

In example (1i), given earlier, the H states will emit the words in bold (pilots, not, allowed, to, have, guns, in, cockpits), and the G states will emit all the other words. The HMM will transition between the H and G states as needed to generate the words of the story. In the current example, the model will have states $Start, G_{\text{start}}, End$ and 28 H states with 28 corresponding G states.³ The headline given in example (1ii) corresponds to the following sequence of states: $Start, G_{\text{start}}$ 17 times, $H_{\text{pilots}}, G_{\text{pilots}}, H_{\text{not}}, G_{\text{not}}, H_{\text{allowed}}, H_{\text{to}}, H_{\text{have}}, H_{\text{guns}}, H_{\text{in}}, G_{\text{in}}, H_{\text{cockpits}}, End$. This path is not the only one that could generate the story in (1i). Other possibilities are

(20) (i) Transportation Department decided airline pilots not to have guns.

(ii) Months of the terrorist has to have cockpits.

Although (20i) and (20ii) are possible headlines for (1i), the conditional probability of (20ii) given (1i) will be lower than the conditional probability of (20i) given (1i).

The Viterbi algorithm (Viterbi, 1967) is used to select the most likely headline for a given story. We use length constraints to find the most likely headlines consisting of W words, where W ranges from 5 to 15. Multiple backpointers are used so that we can find the n most likely headlines at each length.

HMM Hedge is enhanced by three additional decoding parameters to help the system choose outputs that best mimic actual headlines: a position bias, a clump bias and a gap bias. The incorporation of these penalties changes the score produced by the decoder from a probability to a relative desirability score. The three parameters were motivated by analysis of system output and their values were set by trial and error. A logical extension to this work would be to learn the best setting of these penalties, e.g., through Expectation Maximization.

The position bias favors headlines which include words near the front of the story. This reflects our observations of human-constructed headlines, in which headline words tend to appear near the front of the story. The initial position bias p is a positive number less than one. The story word in the n th position is assigned a

² Limiting consideration of headline words to the early part of the story is justified in Dorr, Zajic, and Schwartz (2003a) where it was shown that more than half of the headline words are chosen from the first sentence of the story. Other methods for selecting the window of story words are possible and will be explored in future research.

³ The subscript of a G state indicates the H state it is associated with, not the story word it emits. In the example, G_{pilots} emits story word *will*, G_{not} emits story word *be*, and G_{in} emits story word *the*.

position bias of $\log(p'')$. When an H state emits a story word, the position bias is added to the desirability score. Thus, words near the front of the story carry less of a position bias than words farther along. Note that this generalization often does not hold in the case of human interest and sports stories, which may start with a hook to get the reader's attention, rather than a topic sentence.

We also observed that human-constructed headlines tend to contain contiguous blocks of story words. Example (1ii), given earlier, illustrates this with the string “allowed to have guns”. The string bias is used to favor “clumpiness”, i.e., the tendency to generate headlines composed of clumps of contiguous story words. The log of the clump bias is added to the desirability score with each transition from an H state to its associated G state. With high clump biases, the system will favor headlines consisting of fewer but larger clumps of contiguous story words.

The gap bias is used to disfavor headline “gappiness”, i.e., large gaps of non-headline words in the story between clumps of headline words. Although humans are capable of constructing fluent headlines by selecting widely spaced words, we observed that HMM Hedge was more likely to combine unrelated material by doing this. At each transition from a G state to an H state, corresponding to the end of a sequence of non-headline words in the story, a gap bias is applied that increases with the size of the gap between the current headline and the last headline word to be emitted. This can also be seen as a penalty for spending too much time in one G state. With high gap biases, the system will favor headlines with few large gaps.

One characteristic difference between newspaper headline text and newspaper story text is that headlines tend to be in present tense while story sentences tend to be in the past tense. Past tense verbs occur more rarely in the headline language than in the general language. HMM Hedge mimics this aspect of Headlines by allowing morphological variation between headline verbs and the corresponding story verbs. Morphological variation for verbs is achieved by creating an H state for each available variant of a story verb. These H states still emit the story verb but they are labeled with the variant. HMM Hedge can generate a headline in which *proposes* is the unobserved headline word that emits the observed story word *proposed*, even though *proposes* does not occur in the story.

- (21) (i) A group has proposed awarding \$1 million in every general election to one randomly chosen voter.
 (ii) Group proposes awarding \$1 million to randomly chosen voter.

Finally, we constrain each headline to contain at least one verb. That is to say, we ignore headlines which do not contain at least one verb, no matter how desirable the decoding is.

Although we have described an application of HMM Hedge to blocks of story words without reference to sentence boundaries, it is also possible to use HMM as a single sentence compressor by limiting the block to the words in a single sentence.

Also, as we will see shortly, multiple alternative compressions of a sentence may be generated with HMM Hedge. The Viterbi algorithm is capable of discovering n -best compressions of a window of story words and can be constrained to consider only paths that include a specific number of H states, corresponding to compressions that contain a specific number of words. We use HMM Hedge to generate 55 compressions for each sentence by computing the five best headlines at each length, from 5 to 15 words. In Section 4 we will describe how HMM Hedge is used as a component of a multi-document summarization system.

4. Multi-document summarization

The sentence compression tools we developed for single-document summarization have been incorporated into our multi-candidate reduction framework for multi-document summarization. MCR produces a textual summary from a collection of relevant documents in three steps. First, sentences are selected from the source documents for compression. The most important information occurs near the front of the stories, so we select the first five sentences of each document for compression. Second, multiple compressed versions of each sentence are produced using Trimmer or HMM Hedge to create a pool of candidates for inclusion in the summary. Finally, a sentence selector constructs the summary by iteratively choosing from the pool of candidates based on a linear combination of features until the summary reaches a desired length.

At present, weights for the features are determined by manually optimizing on a set of training data to maximize the ROUGE-2 recall score (Lin & Hovy, 2003), using ROUGE version 1.5.5. A typical summarization task might call for the system to generate a 250-word summary from a couple of dozen news stories. These summaries may be query-focused, in the sense that the summaries should be responsive to a particular information need, or they may be generic, in that a broad overview of the documents is desired.

Our sentence selector adopts certain aspects of Maximal Marginal Relevance (MMR) (Carbonell & Goldstein, 1998), an approach that attempts to balance relevance and anti-redundancy. In MCR's selection module, the highest scoring sentence from the pool of eligible candidates is chosen for inclusion in the summary. Features that contribute to a candidate's score can be divided into two types: dynamic and static. When a candidate is chosen, all other compressed variants of that sentence are eliminated. After a candidate is added to the summary, the dynamic features are re-computed, and the candidates are re-ranked. Candidates are added to the summary until the desired length is achieved. The ordering of candidates in the summary is the same as the order in which they were selected for inclusion. The final sentence of the summary is truncated if it causes the summary to exceed the length limit.

4.1. Static features

Static features are calculated before sentence selection begins, and do not change during the process of summary construction:

- Position. The zero-based position of the sentence in the document.
- Sentence relevance. The relevance score of the sentence to the query.
- Document relevance. The relevance score of the document to the query.
- Sentence centrality. The centrality score of the sentence to the topic cluster.
- Document centrality. The centrality score of the document to the topic cluster.
- Scores from compression modules:
 - Trim rule application counts. For Trimmer-based MCR, the number of trimmer rule instances applied to produce the candidate.
 - Negative log desirability. For HMM-based MCR, the relative desirability score of the candidate.

We use the uniform retrieval architecture (URA), University of Maryland's software infrastructure for information retrieval tasks, to compute relevance and centrality scores for each compressed candidate. There are four such scores: the relevance score between a compressed sentence and the query, the relevance score between the document containing the compressed sentence and the query, the centrality score between a compressed sentence and the topic cluster, and the centrality score between the document containing the compressed sentence and the topic cluster. We define the topic cluster to be the entire collection of documents relevant to this particular summarization task. Centrality is a concept that quantifies how similar a piece of text is to all other texts that discuss the same general topic. We assume that sentences having higher term overlap with the query and sources more "central" to the topic cluster are preferred for inclusion in the final summary.

The relevance score between a compressed sentence and the query is an *idf*-weighted count of overlapping terms (number of terms shared by the two text segments). Inverse document frequency (*idf*), a commonly used measure in the information retrieval literature, roughly captures term salience. The *idf* of a term t is defined by $\log(N/c_t)$, where N is the total number of documents in a particular corpus, and c_t is the number of documents containing term t ; these statistics were calculated from one year's worth of LA Times articles. Weighting term overlap by inverse document frequency captures the intuition that matching certain terms is more important than matching others.

Lucene, a freely available off-the-shelf information retrieval system, is used to compute the three other scores. The relevance score between the document containing the compressed sentence and the query is computed using Lucene's built-in similarity function. The centrality score between the compressed sentence and the topic cluster is the mean of the similarity between the sentence and each document comprising the cluster

(once again, as computed by Lucene's built-in similarity function). The document-cluster centrality score is also computed in much the same way, by taking the mean of the similarity of the particular document with every other document in the cluster. In order to obtain an accurate distribution of term frequencies to facilitate the similarity calculation, we indexed all relevant documents (i.e., the topic cluster) along with a comparable corpus (one year of the LA Times)—this additional text essentially serves as a background model for non-relevant documents.

Some features are derived from the sentence compression modules used to generate candidates. For Trimmer, the rule application count feature of a candidate is the number of rules that were applied to a source sentence to produce the candidate. The rules are not presumed to be equally effective, so the rule application counts are broken down by rule type. For HMM Hedge, we use the relative desirability score calculated by the decoder, expressed as a negative log.

The features discussed in this section are assigned to the candidates before summary generation begins, and are fixed throughout the process of summary sentence selection. The next section discusses how candidate features are assigned new values as summary generation proceeds.

4.2. Dynamic features

Dynamic features change during the process of sentence selection to reflect changes in the state of the summary as sentences are added.⁴ The dynamic features are

- Redundancy. A measure of how similar the sentence is to the current summary.
- Sentence-from-doc. The number of sentences already selected from the sentence's document.

The intuition behind our redundancy measure is that candidates containing words that occur much more frequently in the current state of the summary than they do in general English are redundant to the summary. We imagine that sentences in the summary are generated by the underlying word distribution of the summary rather than the distribution of words in the general language. If a sentence appears to have been generated by the summary rather than by the general language, we take it to be redundant to the summary. Suppose we have a summary about earthquakes. The presence in a candidate of words like *earthquake*, *seismic*, and *Richter Scale*, which have a high likelihood in the summary, will make us think that the candidate is redundant to the summary.

To estimate the extent to which a candidate is more likely to have been generated by a summary than by the general language, we consider the probabilities of the words in the candidate. We estimate that the probability that a word w occurs in a candidate generated by the summary is

$$P(w) = \lambda P(w|D) + (1 - \lambda)P(w|C)$$

where D is the summary, C is the general language corpus,⁵ λ is a parameter estimating the probability that the word was generated by the summary and $(1 - \lambda)$ is the probability that the word was generated by the general language. We have set $\lambda = 0.3$, as a general estimate of the portion of words in a text that are specific to the text's topic. We estimate the probabilities by counting the words⁶ in the current summary and the general language corpus:

$$P(w|D) = \frac{\text{count of } w \text{ in } D}{\text{size of } D}$$

$$P(w|C) = \frac{\text{count of } w \text{ in } C}{\text{size of } C}$$

⁴ At present the dynamic features are properties of the candidates, calculated with respect to the current summary state. There are no features directly relating to the amount of space left in the summary, so there is no mechanism that would affect the distribution of compressed candidates over the iterations of the sentence selector. This issue will be addressed as future work in Section 7.

⁵ The documents in the set being summarized are used to estimate the general language model.

⁶ Actually, preprocessing for redundancy includes stopword removal and applying the Porter Stemmer (Porter, 1980).

We take the probability of a sentence to be the product of the probabilities of its words, so we calculate the probability that a sentence was generated by the summary, i.e. our redundancy metric, as

$$\text{Redundancy}(S) = \prod_{s \in S} \lambda P(s|D) + (1 - \lambda)P(s|C)$$

For ease of computation, we actually use log probabilities:

$$\sum_{s \in S} \log(\lambda P(s|D) + (1 - \lambda)P(s|C))$$

Redundancy is a dynamic feature because the word distribution of the current summary changes with every iteration of the sentence selector.

4.3. Examples of system output

We applied our MCR framework to test data from the DUC-2006 evaluation (Dang & Harman, 2006). Given a topic description and a set of 25 documents related to the topic (drawn from AP newswire, the New York Times, and the Xinhua News Agency English Service), the system's task was to create a 250-word summary that addressed the information need expressed in the topic. One of the topic descriptions is shown in Fig. 2. The 25 documents in the document set have an average size of 1170 words, so a 250-word summary represents a compression ratio of 0.86%.

Figs. 3–5 show examples of MCR output using Trimmer compression, HMM Hedge compression, or no compression. For readability, we use \circ as a sentence delimiter; this is not part of the actual system output. The sentences compressed by Trimmer mimic Headlines by omitting determiners and auxiliary verbs. For example, the first sentence in Fig. 3 is a compression of the following source sentence:

Seeking to get a more accurate count of the country's American Indian population, the Census Bureau is turning to tribal leaders and residents on reservations to help overcome long-standing feelings of wariness or anger toward the federal government.

Three determiners and a form of *be* have been removed from the source sentence in the compression that appears in the summary. The removal of this material make the sentence appear more like a headline.

In comparison with Trimmer compressions, HMM compressions are generally less readable and more likely to be misleading. Consider the final sentence in Fig. 4.

(22) Main purpose of reservation to pay American Indians by poverty proposals.

This is a compression of the following source sentence:

(23) But the main purpose of the visit—the first to a reservation by a president since Franklin Roosevelt—was simply to pay attention to American Indians, who are so raked by grinding poverty that Clinton's own advisers suggested he come up with special proposals geared specifically to the Indians' plight.

Because HMM Hedge uses a bigram model of Headlines, it is unable to capture sentence-level grammaticality. The same limitation makes it difficult to prevent misleading or incorrect compressions.

For example, the third sentence from the end of Fig. 4 seems to say that a court legalized gambling on Indian reservations:

Title: Native American Reservation System—pros and cons
Narrative Description: Discuss conditions on American Indian reservations or among Native American communities. Include the benefits and drawbacks of the reservation system. Include legal privileges and problems.

Fig. 2. Topic D0601A from the DUC-2006 multi-document summarization task.

Seeking to get more accurate count of country's American Indian population, Census Bureau turning to tribal leaders and residents on reservations to help overcome long-standing feelings. ◦ American Indian reservations would get infusion. ◦ Smith and thousands seeking help for substance abuse at American Indian Community House, largest of handful of Native American cultural institutions. ◦ Clinton going to Pine Ridge Reservation for visit with Oglala Sioux nation and to participate in conference on Native American homeownership and economic development. ◦ Said Glen Revere, nutritionist with Indian Health Services on 2.8 million-acre Wind River Reservation, about 100 miles east of Jackson, Wyo. "Then we came up with idea for this community garden, and it been bigger than we ever expected." ◦ Road leading into Shinnecock Indian reservation is not welcoming one But main purpose of visit – first to reservation by president since Franklin Roosevelt – was simply to pay attention to American Indians, who raked by grinding poverty Clinton's own advisers suggested he come up with special proposals geared specifically to Indians' plight. ◦ "This highlights what going on out there, since beginning of reservation system," said Sidney Harring, professor at City University of New York School of Law and expert on Indian crime and criminal law. ◦ American Indians are victims. ◦ President Clinton turned attention to arguably poorest, most forgotten ◦ U.S. citizens: American Indians. ◦ When American Indians began embracing gambling, Hualapai tribe moved quickly to open casino. ◦ members of Northern Arapaho Tribe on Wind River Reservation started seven-acre community garden with donated land, seeds and

Fig. 3. MCR summary for DUC-2006 Topic D0601A, using Trimmer for sentence compression.

David Rocchio deputy legal counsel to Vermont Gov. Howard Dean who has been involved in discussions on Indian gambling through the National Governors Association said that the concern that governors have is not with the benefit casinos bring to tribes ◦ Native Americans living on reservations that maintain 50 percent or more unemployment are exempt from the national five year family limit on welfare benefits ◦ Smith and thousands like her are seeking help for their substance abuse at the American Indian Community House the largest of a handful of Native American cultural institutions in the New York area ◦ Juvenile crime is one strand in the web of social problems facing urban and reservation Indian communities the report said ◦ Soldierwolf's family represents the problems that plague many of the 1.3 million American Indians who live on reservations of whom 49 percent are unemployed ◦ Powless said the Onondaga people want to work with the community outside the reservation to improve the economy of the region perhaps creating tourism destinations that might include Indian culture or setting up a free trade zone at unused manufacturing sites ◦ As Indian communities across the nation struggle with short funds and a long list of problems they are watching the Navajo Nation's legal battle with the federal government ◦ recognize Indians not only Native Americans as Americans ◦ go on reservation system Harring Indian ◦ Supreme Court allows legalized gambling Indian reservations ◦ American Indian reservations tribal colleges rise faster than ◦ main purpose of reservation to pay American Indians by poverty proposals

Fig. 4. MCR summary for DUC-2006 Topic D0601A, using HMM Hedge for sentence compression.

655 (24) Supreme Court allows legalized gambling Indian reservations.

656

657 However, it is a compression of the following source sentence:

658 (25) Only Monday, the California Supreme Court overturned a ballot measure that would have allowed
659 expansion of legalized gambling on Indian reservations.

Seeking to get a more accurate count of the country's American Indian population, the Census Bureau is turning to tribal leaders and residents on reservations to help overcome long-standing feelings of wariness or anger toward the federal government. ◦ American Indian reservations would get an infusion of \$1.2 billion in federal money for education, health care and law enforcement under President Clinton's proposed 2001 budget ◦ Smith and thousands like her are seeking help for their substance abuse at the American Indian Community House, the largest of a handful of Native American cultural institutions in the New York area. ◦ Clinton was going to the Pine Ridge Reservation for a visit with the Oglala Sioux nation and to participate in a conference on Native American homeownership and economic development. ◦ said Glen Revere, a nutritionist with the Indian Health Services on the 2.8 million-acre Wind River Reservation, about 100 miles east of Jackson, Wyo. "Then we came up with the idea for this community garden, and it's been bigger than we ever expected in so many ways." ◦ The road leading into the Shinnecock Indian reservation is not a welcoming one ◦ But the main purpose of the visit – the first to a reservation by a president since Franklin Roosevelt – was simply to pay attention to American Indians, who are so raked by grinding poverty that Clinton's own advisers suggested he come up with special proposals geared specifically to the Indians' plight. ◦ "This highlights what has been going on out there for 130 years,

Fig. 5. MCR summary for DUC-2006 Topic D0601A, with no sentence compression.

660

661 Nevertheless, we can see from the examples that sentence compression allows a summary to include more
662 material from other sources. This increases the topic coverage of system output.

663 5. System evaluations

664 We tested four single-document summarization systems on the DUC-2003 Task 1 test set:

- 665 • HMM Hedge using the first sentence of each document (HMM Sentence).
- 666 • HMM Hedge using the first 60 words of each document (HMM 60 block).
- 667 • Trimmer.
- 668 • Topiary.

669

670 Task 1 from DUC-2003 was to construct generic 75-byte summaries for 624 documents drawn from AP
671 Newswire and the New York Times. The average size of the documents was 3997 bytes, so a 75-byte summary
672 represents a compression ratio of 1.9%.

673 An automatic summarization evaluation tool, ROUGE (Lin & Hovy, 2003), was used to evaluate the results.
674 The system parameters were optimized by hand to maximize the ROUGE-1 recall on a comparable training cor-
675 pus, 500 AP Newswire and New York Times articles from the DUC-2004 single-document short summary test
676 data.

677 Results show that HMM Hedge 60 scored significantly lower than most other systems, and that Topiary
678 scored higher than all other systems for all R1 measures. In addition, HMM Hedge Sentence scored signifi-
679 cantly higher than Trimmer for the R1 measures (Table 2).

680 We also evaluated Trimmer and HMM Hedge as components in our multi-candidate reduction framework,
681 along with a baseline that uses the same sentence selector but does not use sentence compression. All three
682 systems considered the first five sentences of each document, and used the sentence selection algorithm pre-
683 sented in Section 4. The feature weights were manually optimized to maximize ROUGE-2 recall on a comparable
684 training corpus, 1593 Financial Times and Los Angeles Times articles grouped into 50 topics from the DUC-
685 2005 query-focused multi-document summarization test data. The systems were used to generate query-
686 focused, 250-word summaries using the DUC-2006 test data, described in Section 4.3.

Table 2

ROUGE scores and 95% confidence intervals for 624 documents from DUC-2003 test set

	HMM sentence	HMM 60 block	Trimmer	Topiary
R1 Recall	0.23552 (0.23014–0.24082)	0.21381 (0.20912–0.21827)	0.21014 (0.20436–0.21594)	0.25143 (0.24632–0.25663)
R1 Precision	0.21896 (0.21386–0.22384)	0.18882 (0.18444–0.19301)	0.20183 (0.19627–0.20722)	0.23038 (0.22567–0.23522)
R1 F	0.22496 (0.21983–0.22978)	0.19966 (0.19505–0.20391)	0.20179 (0.19612–0.20718)	0.23848 (0.23373–0.24328)
R2 Recall	0.06838 (0.06546–0.07155)	0.06133 (0.05848–0.06414)	0.06337 (0.06030–0.06677)	0.06637 (0.06345–0.06958)
R2 Precision	0.06287 (0.06017–0.06576)	0.05351 (0.05097–0.05588)	0.06230 (0.05887–0.06617)	0.06024 (0.05747–0.06326)
R2 F	0.06488 (0.06209–0.06785)	0.05686 (0.05420–0.05942)	0.06079 (0.05788–0.06401)	0.06252 (0.05976–0.06561)

Table 3

ROUGE scores and 95% confidence intervals for 50 DUC-2006 test topics, comparing three MCR variants

	Trimmer	HMM Hedge	No compression
R1 Recall	0.29391 (0.28560–0.30247)	0.27311 (0.26554–0.28008)	0.27576 (0.26772–0.28430)
R2 Recall	0.06718 (0.06332–0.07111)	0.06251 (0.05873–0.06620)	0.06126 (0.05767–0.06519)

Table 4

Official DUC-2006 automatic metrics for our MCR submission (System 32)

	ROUGE-2	ROUGE-SU4	BE-HM
MCR score	0.0805	0.1360	0.0413
Higher	1	1	0
Not different	23	24	27
Range	0.0678–0.0899	0.1238–0.1475	0.0318–0.0508
Lower	11	10	8

The systems were evaluated using ROUGE, configured to omit stopwords from the calculation.⁷ Results are shown in Table 3. MCR using Trimmer compressions scored significantly higher than MCR using HMM Hedge compressions and the baseline for ROUGE-1, but there was not a significant difference among the three systems for ROUGE-2.

Finally, the University of Maryland and BBN submitted a version of MCR to the official DUC-2006 evaluation. This version used Trimmer as the source of sentence compressions. Results show that use of sentence compression hurt the system on human evaluation of grammaticality. This is not surprising, since Trimmer aims to produce compressions that are grammatical in Headlines, rather than standard English. Our MCR run scored significantly lower than 23 systems on NIST's human evaluation of grammaticality. However, the system did not score significantly lower than any other system on NIST's human evaluation of content responsiveness. A second NIST evaluation of content responsiveness asked evaluators to take readability into consideration. In this evaluation, MCR scored significantly lower than only two systems. The evaluators recognized that Trimmer compressions are not grammatical in standard English; yet, the content coverage was not significantly different from the best automatic systems, and only two systems were found to be significantly more readable.

NIST computed three “official” automatic evaluation metrics for the DUC-2006: ROUGE-2, ROUGE-SU4 and BE-HM. Table 4 shows the official scores of the submitted MCR system for these three metrics, along with numbers of systems that scored significantly higher, significantly lower, or were not significantly different from

⁷ This is a change in the ROUGE configuration from the official DUC-2006 evaluation. We note that the removal of non-essential stopwords (typical of Headlines) is an important component of Trimmer-based sentence compression. For internal system comparisons, we configure ROUGE in a way that will allow us to detect system differences relevant to our research focus. For reporting of official ROUGE results on submitted systems we use the community's accepted ROUGE configurations.

our MCR run. Also shown is the range of scores for the systems that were not significantly different from MCR. These results show that the performance of our MCR run was comparable to most other systems submitted to DUC-2006.

The evaluation in Table 3 suggest that Trimmer sentence compression is preferable to HMM Hedge sentence compression for generation of English summaries of collections of document in English. However, HMM Hedge may prove to have value with noisier data, as we discuss in the next section. Nevertheless, sentence compression appears to be a valuable component of our framework for multi-document summarization, thus validating the ideas behind multi-candidate reduction.

6. Applications to different types of texts

We have applied the MCR framework to summarizing different types of texts. In this section we briefly touch on genre-specific issues that are the subject of ongoing work. Trimmer, Topiary, and HMM Hedge were designed for summarization of written news. In this genre, the lead sentence is almost always the first non-trivial sentence of the document. More sophisticated methods for finding lead sentences did not outperform the baseline of simply selecting the first sentence for AP wire “hard” news stories. However, some types of articles, such as sports stories, opinion pieces, and movie reviews often do not have informative lead sentences, and will require additional work in finding the best sentence for compression.

MCR has also been applied to summarizing transcripts of broadcast news—another input form where lead sentences are often not informative. The conventions of broadcast news introduce categories of story-initial light content sentences, such as “I’m Dan Rather” or “We have an update on the story we’ve been following”. These present challenges for the filtering stage of our MCR framework.

Such texts are additionally complicated by a range of problems not encountered in written news: noise introduced by automatic speech recognizers or other faulty transcription, issues associated with sentence boundary detection and story boundary detection. If word error rate is high, parser failures can prevent Trimmer from producing useful output. In this context, HMM Hedge becomes more attractive, since our language models are more resilient to noisy input.

We have performed an initial evaluation of Trimmer, Topiary and a baseline consisting of the first 75 characters of a document, on the task of creating 75-character headlines for broadcast news transcriptions (Zajic, 2007). The corpus for this task consisted of 560 broadcast news stories from ABC, CNN, NBC, Public Radio International and Voice of America. We used ROUGE-1 recall to evaluate the summaries, and found that both systems scored higher than the baseline, and that Topiary scored higher than Trimmer. However there were no significant differences among the system.

Another application of our framework is the summarization of email threads—collections of emails that share a common topic or were written as responses to each other. This task can essentially be treated as a multi-document summarization problem, albeit email thread structure introduces some constraints with respect to the ordering of summary sentences. Noisy data is inherent in this problem, and the pre-processing to remove quoted text, attachments and headers is crucial. We have found that metadata, such as the name of the sender of each included extract help make email summaries easier to read.

We performed an initial evaluation of HMM Hedge and Trimmer as the source of sentence compressions for an email thread summarization system based on the MCR framework (Zajic, 2007). The corpus for this task consisted of 10 manually constructed email threads from the Enron Corpus (Klimt & Yang, 2004). We used ROUGE-1 and ROUGE-2 recall with jackknifing to compare the automatic systems and the human summarizers. We did not observe a significant difference between the two systems, but we found that the task of summarizing email threads was extremely difficult for the humans (one summarizer scored significantly worse than the automatic systems). This application of MCR to email thread summarization is an initial effort. The difficulty of the task for the humans suggests that the community needs to develop a clearer understanding of what makes a good email thread summary, and explore practical uses for them.

Finally, Trimmer and HMM Hedge have been applied to Hindi-English cross-language summarization. In this case, Trimmer was applied to the output of machine translation. We adapted HMM Hedge to cross-

lingual summarization by using the mechanism developed for morphological variation to represent translation probabilities from Hindi story words to English headline words. For more details, see [Dorr et al. \(2003a\)](#).

7. Future work

Future work on text summarization under the multi-candidate reduction framework will focus on the three main components of the architecture: sentence filtering, sentence compression, and candidate selection.

For single document summarization, the simple technique of selecting the first non-trivial sentence of a document for compression remains the best approach. However, for human interest stories or sports articles, this approach is less effective. In broadcast news transcripts, the first sentence often does not contain important information. Currently, filtering for multi-document summarization also relies on the assumption that important information tends to appear near the front of documents—the first five sentences of each document are retained to generate compressed candidates. An interesting area of future work is to explore other approaches to filtering, such as using query relevance and document centrality, to move beyond the baseline of selecting the first n sentences. For HMM Hedge, these methods can be used to determine the optimal blocks of text on which to apply the decoder.

Currently, Trimmer produces multiple compressions by applying rules in a fixed order; the state of the compressed sentence after each rule application becomes a candidate. A richer pool of candidates can be produced by modifying Trimmer rules to operate in order-independent combinations, rather than a fixed sequence. We believe that the sentence selector can produce better summaries if it has larger pools of candidates to choose from. Naturally, different sentence compressions are not the only techniques for enriching the candidate pool—other possibilities include merging sentences and resolving anaphora. Topiary will also be enhanced by using multiple combinations of compressions and topic terms in the context of headline generation.

We also plan to enrich the candidate selector by taking into account more features of the current summary state. Possibilities include sentence selector iteration count and remaining summary space, as well as feature weights that change during the progress of summary generation. These extensions will allow us to study the distribution of compressed and uncompressed sentences across sentence selector iterations. System output can potentially be improved by finer-grained control of this distribution. These features might also help avoid the current problem in which the final sentence is truncated due to length restrictions (e.g., by selecting a final sentence of more appropriate length).

Proper setting of parameters is another important area for future work. Systematic optimization of parameter values in HMM Hedge and the sentence selector could lead to significant improvements in output quality. A logical extension to this work would be to learn the best parameter settings, e.g., through Expectation Maximization.

At present, MCR focuses exclusively on summary content selection and does not take sentence ordering into consideration when constructing the summary. Naturally, high-quality summaries should read fluently in addition to having relevant content. Recent work in this area that can be applied to MCR includes [Conroy et al. \(2006\)](#), [Barzilay, Elhadad, and McKeown \(2002\)](#), [Okazaki, Matsuo, and Ishizuka \(2004\)](#), [Lapata \(2003\)](#), and [Dorr and Gaasterland \(this special issue\)](#). Within the MCR architecture, fluency considerations can be balanced with other important factors such as relevance and anti-redundancy through appropriate feature weighting.

8. Conclusion

This work presents multi-candidate reduction, a general architecture for multi-document summarization. The framework integrates successful single-document compression techniques that we have previously developed. MCR is motivated by the insight that multiple candidate compression of source sentences should be made available to subsequent processing modules, which may have access to more information for summary construction. This is implemented in a dynamic feature-based sentence selector that iteratively builds a summary from compressed variants. Evaluations show indeed that sentence compression plays an important role in multi-document summarization, and that our MCR framework is both flexible and extensible.

800 Acknowledgement

801 This work has been supported, in part, under the GALE program of the Defense Advanced Research Pro-
 802 jects Agency, Contract No. HR0011-06-2-0001, the TIDES program of the Defense Advanced Research Pro-
 803 jects Agency, BBNT Contract No. 9500006806, and the University of Maryland Joint Institute for Knowledge
 804 Discovery. Any opinions, findings, conclusions or recommendations expressed in this paper are those of the
 805 authors and do not necessarily reflect the views of DARPA. The first author would like to thank Naomi for
 806 proofreading, support and encouragement. The second author would like to thank Steve, Carissa, and Ryan
 807 for their energy enablement. The third author would like to thank Esther and Kiri for their kind support.

808 References

- 809 Bahl, L., Jelinek, F., & Mercer, R. (1983). A maximum likelihood approach to continuous speech recognition. *IEEE Transactions on*
 810 *Pattern Analysis and Machine Intelligence*, 5(2), 179–190.
- 811 Banko, M., Mittal, V., & Witbrock, M. (2000). Headline generation based on statistical translation. In *Proceedings of the 38th annual*
 812 *meeting of the association for computational linguistics (ACL 2000)*, Hong Kong (pp. 318–325).
- 813 Barzilay, R., Elhadad, N., & McKeown, K. (2002). Inferring strategies for sentence ordering in multidocument news summarization.
 814 *Journal of Artificial Intelligence Research*, 17, 35–55.
- 815 Baum, L. (1972). An inequality and associated maximization technique in statistical estimation of probabilistic functions of a Markov
 816 process. *Inequalities*, 3, 1–8.
- 817 Bergler, S., Witte, R., Khalife, M., Li, Z., & Rudzicz, F. (2003). Using knowledge-poor coreference resolution for text summarization. In
 818 *Proceedings of the HLT-NAACL 2003 Text summarization workshop and document understanding conference (DUC 2003)*, Edmonton,
 819 Alberta (pp. 85–92).
- 820 Bikel, D., Schwartz, R., & Weischedel, R. (1999). An algorithm that learns what's in a name. *Machine Learning*, 34(1/3), 211–231.
- 821 Blair-Goldensohn, S., Evans, D., Hatzivassiloglou, V., McKeown, K., Nenkova, A., & Passonneau, R. (2004). Columbia University at
 822 DUC 2004. In *Proceedings of the 2004 document understanding conference (DUC 2004) at HLT/NAACL 2004*, Boston, Massachusetts
 823 (pp. 23–30).
- 824 Brown, P., Cocke, J., Pietra, S., Pietra, V., Jelinek, F., Lafferty, J., et al. (1990). A statistical approach to machine translation.
 825 *Computational Linguistics*, 16(2), 79–85.
- 826 Carbonell, J., & Goldstein, J. (1998). The use of MMR, diversity-based reranking for reordering documents and producing summaries. In
 827 *Proceedings of the 21st annual international ACM SIGIR conference on research and development in information retrieval (SIGIR 1998)*,
 828 Melbourne, Australia (pp. 335–336).
- 829 Charniak, E. (2000). A maximum-entropy-inspired parser. In *Proceedings of the first meeting of the North American chapter of the*
 830 *association for computational linguistics (NAACL 2000)*, Seattle, Washington (pp. 132–139).
- 831 Clarke, J., & Lapata, M. (2006). Models for sentence compression: a comparison across domains, training requirements and evaluation
 832 measures. In *Proceedings of the 21st international conference on computational linguistics and 44th annual meeting of the association for*
 833 *computational linguistics (COLING/ACL 2006)*, Sydney, Australia (pp. 377–384).
- 834 Conroy, J., Schlesinger, J., & Stewart, J. G. (2005). CLASSY query-based multi-document summarization. In *Proceedings of the 2005*
 835 *document understanding conference (DUC-2005) at NLT/EMNLP 2005*, Vancouver, Canada.
- 836 Conroy, J., Schlesinger, J., O'Leary, D., & Goldstein, J. (2006) Back to basics: CLASSY 2006. In *Proceedings of the 2006 document*
 837 *understanding conference (DUC 2006) at HLT/NAACL 2006*, New York, NY.
- 838 Cutting, D., Pedersen, J., & Sibun, P. (1992). A practical part-of-speech tagger. In *Proceedings of the third conference on applied natural*
 839 *language processing*. Trento, Italy.
- 840 Dang, H., & Harman, D. (2006). In *Proceedings of the 2006 document understanding conference (DUC 2006) at HLT/NAACL 2006*.
- 841 Dorr, B. J., & Gaasterland, T. Exploiting aspectual features and connecting words for summarization-inspired temporal-relation
 842 extraction. *Information Processing and Management*, this special issue.
- 843 Dorr, B., Zajic, D., & Schwartz, R. (2003a). Cross-language headline generation for Hindi. *ACM Transactions on Asian Language*
 844 *Information Processing (TALIP)*, 2(3), 270–289.
- 845 Dorr, B. J., Zajic, D., & Schwartz, R. (2003b). Hedge trimmer: a parse-and-trim approach to headline generation. In *Proceedings of the*
 846 *HLT-NAACL 2003 text summarization workshop and document understanding conference (DUC 2003)*, Edmonton, Alberta (pp. 1–8).
- 847 Dunning, T. (1994). Statistical identification of language. Tech. Rep. MCCS 94-273, New Mexico State University.
- 848 Euler, T. (2002). Tailoring text using topic words: selection and compression. In *Proceedings of 13th international workshop on database*
 849 *and expert systems applications (DEXA 2002)*, Aix-en-Provence, France (pp. 215–222).
- 850 Goldstein, J., Mittal, V., Carbonell, J., & Kantrowitz, M. (2000). Multi-document summarization by sentence extraction. In *Proceedings of*
 851 *ANLP/NAACL 2000 workshop on automatic summarization* (pp. 40–48).
- 852 Harman, D., & Liberman, M. (1993). *TIPSTER complete*. Philadelphia: Linguistic Data Consortium (LDC).
- 853 Jing, H., & McKeown, K. (2000). Cut and paste based text summarization. In *Proceedings of the first meeting of the North American*
 854 *chapter of the association for computational linguistics (NAACL 2000)*, Seattle, Washington (pp. 178–185).
- 855 Klimt, B., & Yang, Y. (2004). Introducing the Enron corpus. In *Proceedings of the first conference on email and anti-spam (CEAS)*.
 856 Mountain View, CA.

- 22 D. Zajic et al. / *Information Processing and Management xxx (2007) xxx–xxx*
- 857 Knight, K., & Marcu, D. (2000). Statistics-based summarization—step one: sentence compression. In *Proceedings of the seventeenth*
858 *national conference on artificial intelligence (AAAI-2000)*, Austin, TX.
- 859 Knight, K., & Marcu, D. (2002). Summarization beyond sentence extraction: a probabilistic approach to sentence compression. *Artificial*
860 *Intelligence*, 139(1), 91–107.
- 861 Lapata, M. (2003). Probabilistic text structuring: experiments with sentence ordering. In *Proceedings of the 42nd annual meeting of the*
862 *association for computational linguistics (ACL 2004)*, Barcelona, Spain (pp. 545–552).
- 863 Lewis, D. D. (1999). An evaluation of phrasal and clustered representations on a text categorization task. In *Proceedings of the 15th annual*
864 *international ACM SIGIR conference on research and development in information retrieval (SIGIR 1992)*, Copenhagen, Denmark (pp.
865 37–50).
- 866 Lin, C. -Y., & Hovy, E. (2003). Automatic evaluation of summaries using n-gram co-occurrence statistics. In *Proceedings of the 2003*
867 *human language technology conference and the North American chapter of the association for computational linguistics annual meeting*
868 *(HLT/NAACL 2003)*, Edmonton, Alberta (pp. 71–78).
- 869 Mårdh, I. (1980). Headlines: On the grammar of english front page headlines, Malmo.
- 870 Mays, E., Damerau, F., & Mercer, R. (1990). Context-based spelling correction. In *Proceedings of IBM natural language ITL*, Paris,
871 France (pp. 517–522).
- 872 Miller, S., Ramshaw, L., Fox, H., & Weischedel, R. (2000). A novel use of statistical parsing to extract information from text. In
873 *Proceedings of the first meeting of the North American chapter of the association for computational linguistics (NAACL 2000)*, Seattle,
874 Washington (pp. 226–233).
- 875 Muresan, S., Tzoukermann, E., & Klavans, J. L. (2001). Combining linguistic and machine learning techniques for email. In *Proceedings*
876 *of the ACL/EACL 2001 workshop on computational natural language learning (ConLL)*, Toulouse, France (pp. 290–297).
- 877 Okazaki, N., Matsuo, Y., & Ishizuka, M. (2004). Improving chronological sentence ordering by precedence relation. In *Proceedings of the*
878 *20th international conference on computational linguistics (COLING 2004)*, Geneva, Switzerland (pp. 750–756).
- 879 Porter, M. (1980). An algorithm for suffix stripping. *Program*, 14(3), 130–137.
- 880 Radev, D., Allison, T., Blair-Goldensohn, S., Blitzer, J., Çelebi, A., Dimitrov, S., et al. (2004). MEAD—a platform for multidocument
881 multilingual text summarization. In *Proceedings of the 4th international conference on language resources and evaluation (LREC 2004)*,
882 Lisbon, Portugal.
- 883 Schwartz, R., Imai, T., Jubala, F., Nguyen, L., & Makhoul, J. (1997). A maximum likelihood model for topic classification of broadcast
884 news. In *Proceedings of the fifth European speech communication association conference on speech communication and technology*
885 *(Eurospeech-97)*, Rhodes, Greece.
- 886 Sista, S., Schwartz, R., Leek, T. R., & Makhoul, J. (2002). An algorithm for unsupervised topic discovery from broadcast news stories. In
887 *Proceedings of the 2002 human language technology conference (HLT)*, San Diego, California (pp. 99–103).
- 888 Turner, J., & Charniak, E. (2005). Supervised and unsupervised learning for sentence compression. In *Proceedings of the 43rd annual*
889 *meeting of the association for computational linguistics (ACL 2005)*, Ann Arbor, Michigan (pp. 290–297).
- 890 Vanderwende, L., Suzuki, H., & Brockett, C. (2006). Microsoft research at DUC2006: task-focused summarization with sentence
891 simplification and lexical expansion. In *Proceedings of the 2006 document understanding conference (DUC 2006) at HLT/NAACL 2006*.
892 New York, NY.
- 893 Viterbi, A. (1967). Error bounds for convolution codes and an asymptotically optimal decoding algorithm. *IEEE Transactions on*
894 *Information Theory*, 13, 260–269.
- 895 Wang, R., Stokes, N., Doran, W., Newman, E., Carthy, J., & Dunnion, J. (2005). Comparing topiary-style approaches to headline
896 generation. In *Lecture notes in computer science: Advances in information retrieval: 27th European conference on IR research (ECIR*
897 *2005)*. Vol. 3408. Berlin/Heidelberg: Springer, Santiago de Compostela, Spain.
- 898 Zajic, D. M. (2007). Multiple alternative sentence compressions (MASC) as a tool for automatic summarization tasks. PhD thesis,
899 University of Maryland, College Park.
- 900 Zajic, D., Dorr, B. J., & Schwartz, R. (2004). BBN/UMD at DUC-2004: topiary. In *Proceedings of the 2004 document understanding*
901 *conference (DUC 2004) at NLT/NAACL 2004*, Boston, MA (pp. 112–119).
- 902 Zajic, D. M., Dorr, B. J., Lin, J., & Schwartz, R. (2005a). UMD/BBN at MSE2005. In *Proceedings of the MSE2005 track of the ACL 2005*
903 *workshop on intrinsic and extrinsic evaluation measures for MT and/or summarization*, Ann Arbor, Michigan.
- 904 Zajic, D. M., Dorr, B. J., Schwartz, R., Monz, C., & Lin, J. (2005b). A sentence-trimming approach to multi-document summarization. In
905 *Proceedings of the 2005 document understanding conference (DUC-2005) at NLT/EMNLP 2005*, Vancouver, Canada (pp. 151–158).
- 906 Zhou, L., & Hovy, E. (2003). Headline summarization at ISI. In *Proceedings of the HLT-NAACL 2003 text summarization workshop and*
907 *document understanding conference (DUC 2003)*, Edmonton, Alberta (pp. 174–178).