

Learning-Based Single-Document Summarization with Compression and Anaphoricity Constraints

Greg Durrett

Computer Science Division
UC Berkeley

gdurrett@cs.berkeley.edu

Taylor Berg-Kirkpatrick

School of Computer Science
Carnegie Mellon University

tberg@cs.cmu.edu

Dan Klein

Computer Science Division
UC Berkeley

klein@cs.berkeley.edu

Abstract

We present a discriminative model for single-document summarization that integrally combines compression and anaphoricity constraints. Our model selects textual units to include in the summary based on a rich set of sparse features whose weights are learned on a large corpus. We allow for the deletion of content within a sentence when that deletion is licensed by compression rules; in our framework, these are implemented as dependencies between subsentential units of text. Anaphoricity constraints then improve cross-sentence coherence by guaranteeing that, for each pronoun included in the summary, the pronoun’s antecedent is included as well or the pronoun is rewritten as a full mention. When trained end-to-end, our final system¹ outperforms prior work on both ROUGE as well as on human judgments of linguistic quality.

1 Introduction

While multi-document summarization is well-studied in the NLP literature (Carbonell and Goldstein, 1998; Gillick and Favre, 2009; Lin and Bilmes, 2011; Nenkova and McKeown, 2011), single-document summarization (McKeown et al., 1995; Marcu, 1998; Mani, 2001; Hira0 et al., 2013) has received less attention in recent years and is generally viewed as more difficult. Content selection is tricky without redundancy across multiple input documents as a guide and simple positional information is often hard to beat (Penn and Zhu, 2008). In this work, we tackle the single-document problem by training an expressive summarization model on a large nat-

urally occurring corpus—the New York Times Annotated Corpus (Sandhaus, 2008) which contains around 100,000 news articles with abstractive summaries—learning to select important content with lexical features. This corpus has been explored in related contexts (Dunietz and Gillick, 2014; Hong and Nenkova, 2014), but to our knowledge it has not been directly used for single-document summarization.

To increase the expressive capacity of our model we allow more aggressive compression of individual sentences by combining two different formalisms—one syntactic and the other discursive. Additionally, we incorporate a model of anaphora resolution and give our system the ability rewrite pronominal mentions, further increasing expressivity. In order to guide the model, we incorporate (1) constraints from coreference ensuring that critical pronoun references are clear in the final summary and (2) constraints from syntactic and discourse parsers ensuring that sentence realizations are well-formed. Despite the complexity of these additional constraints, we demonstrate an efficient inference procedure using an ILP-based approach. By training our full system end-to-end on a large-scale dataset, we are able to learn a high-capacity structured model of the summarization process, contrasting with past approaches to the single-document task which have typically been heuristic in nature (Daumé and Marcu, 2002; Hira0 et al., 2013).

We focus our evaluation on the New York Times Annotated corpus (Sandhaus, 2008). According to ROUGE, our system outperforms a document prefix baseline, a bigram coverage baseline adapted from a strong multi-document system (Gillick and Favre, 2009), and a discourse-informed method from prior work (Yoshida et al., 2014). Imposing discursive and referential constraints improves human judgments of linguistic clarity and referential structure—outperforming the method of

¹Available at <http://nlp.cs.berkeley.edu>

$$\begin{aligned}
& \max_{\mathbf{x}^{\text{UNIT}}, \mathbf{x}^{\text{REF}}} \left[\sum_i \left[x_i^{\text{UNIT}} (\mathbf{w}^\top \mathbf{f}(u_i)) \right] + \sum_{(i,j)} \left[x_{ij}^{\text{REF}} (\mathbf{w}^\top \mathbf{f}(r_{ij})) \right] \right] \\
& \text{subject to} \\
& \text{Grammaticality Constraints (Section 2.1)} \quad \forall i, k \quad x_i^{\text{UNIT}} \leq x_k^{\text{UNIT}} \text{ if } u_i \text{ requires } u_k \\
& \text{Length Constraint} \quad \sum_i x_i^{\text{UNIT}} |u_i| + \underbrace{\sum_{(i,j)} x_{ij}^{\text{REF}} (|r_{ij}| - 1)}_{\text{Length adjustment for explicit mention}} \leq k \\
& \text{Anaphora Constraints (Section 2.2)} \\
& \forall j \quad x_{ij}^{\text{REF}} = 1 \text{ iff no prior included textual unit mentions the entity that } r_{ij} \text{ refers to} \\
& \forall i, k \quad x_i^{\text{UNIT}} \leq x_k^{\text{UNIT}} \text{ if } u_i \text{ requires } u_k \text{ on the basis of pronoun anaphora}
\end{aligned}$$

Figure 1: ILP formulation of our single-document summarization model. The basic model extracts a set of textual units with binary variables \mathbf{x}^{UNIT} subject to a length constraint. These textual units \mathbf{u} are scored with weights \mathbf{w} and features \mathbf{f} . Next, we add constraints derived from both syntactic parses and Rhetorical Structure Theory (RST) to enforce grammaticality. Finally, we add anaphora constraints derived from coreference in order to improve summary coherence. We introduce additional binary variables \mathbf{x}^{REF} that control whether each pronoun is replaced with its antecedent using a candidate replacement r_{ij} . These are also scored in the objective and are incorporated into the length constraint.

Yoshida et al. (2014) and approaching the clarity of a sentence-extractive baseline—and still achieves substantially higher ROUGE score than either method. These results indicate that our model has the expressive capacity to extract important content, but is sufficiently constrained to ensure fluency is not sacrificed as a result.

Past work has explored various kinds of structure for summarization. Some work has focused on improving content selection using discourse structure (Louis et al., 2010; Hirao et al., 2013), topical structure (Barzilay and Lee, 2004), or related techniques (Mithun and Kosseim, 2011). Other work has used structure primarily to reorder summaries and ensure coherence (Barzilay et al., 2001; Barzilay and Lapata, 2008; Louis and Nenkova, 2012; Christensen et al., 2013) or to represent content for sentence fusion or abstraction (Thadani and McKeown, 2013; Pighin et al., 2014). Similar to these approaches, we appeal to structures from upstream NLP tasks (syntactic parsing, RST parsing, and coreference) to restrict our model’s capacity to generate. However, we go further by optimizing for ROUGE subject to these constraints with end-to-end learning.

2 Model

Our model is shown in Figure 1. Broadly, our ILP takes a set of textual units $\mathbf{u} = (u_1, \dots, u_n)$ from a document and finds the highest-scoring extractive summary by optimizing over variables

$\mathbf{x}^{\text{UNIT}} = x_1^{\text{UNIT}}, \dots, x_n^{\text{UNIT}}$, which are binary indicators of whether each unit is included. Textual units are contiguous parts of sentences that serve as the fundamental units of extraction in our model. For a sentence-extractive model, these would be entire sentences, but for our compressive models we will have more fine-grained units, as shown in Figure 2 and described in Section 2.1. Textual units are scored according to features \mathbf{f} and model parameters \mathbf{w} learned on training data. Finally, the extraction process is subject to a length constraint of k words. This approach is similar in spirit to ILP formulations of multi-document summarization systems, though in those systems content is typically modeled in terms of bigrams (Gillick and Favre, 2009; Berg-Kirkpatrick et al., 2011; Hong and Nenkova, 2014; Li et al., 2015). For our model, type-level n -gram scoring only arises when we compute our loss function in max-margin training (see Section 3).

In Section 2.1, we discuss grammaticality constraints, which take the form of introducing dependencies between textual units, as shown in Figure 2. If one textual unit *requires* another, it cannot be included unless its prerequisite is. We will show that different sets of requirements can capture both syntactic and discourse-based compression schemes.

Furthermore, we introduce anaphora constraints (Section 2.2) via a new set of variables that capture the process of rewriting pronouns to make them

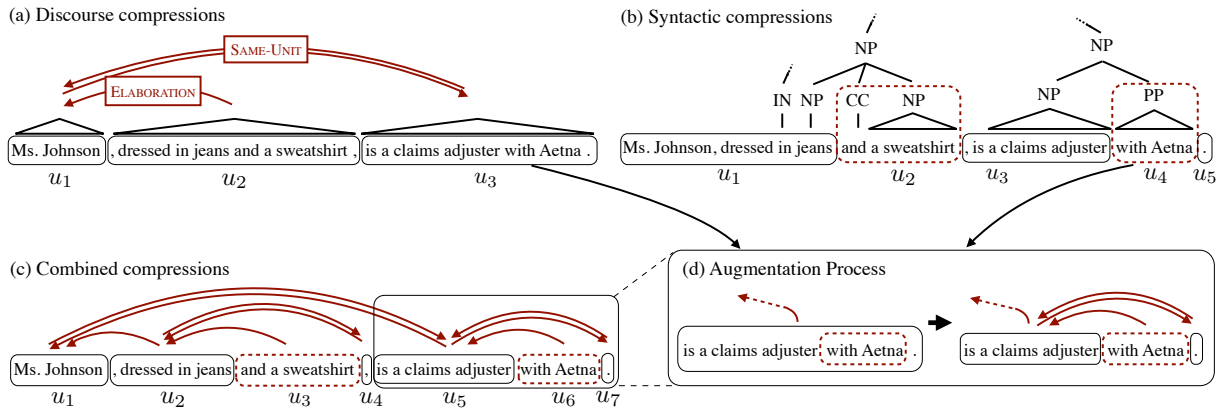


Figure 2: Compression constraints on an example sentence. (a) RST-based compression structure like that in Hirao et al. (2013), where we can delete the ELABORATION clause. (b) Two syntactic compression options from Berg-Kirkpatrick et al. (2011), namely deletion of a coordinate and deletion of a PP modifier. (c) Textual units and requirement relations (arrows) after merging all of the available compressions. (d) Process of augmenting a textual unit with syntactic compressions.

explicit mentions. That is, $x_{ij}^{\text{REF}} = 1$ if we should rewrite the j th pronoun in the i th unit with its antecedent. These pronoun rewrites are scored in the objective and introduced into the length constraint to make sure they do not cause our summary to be too long. Finally, constraints on these variables control when they are used and also require the model to include antecedents of pronouns when the model is not confident enough to rewrite them.

2.1 Grammaticality Constraints

Following work on isolated sentence compression (McDonald, 2006; Clarke and Lapata, 2008) and compressive summarization (Lin, 2003; Martins and Smith, 2009; Berg-Kirkpatrick et al., 2011; Woodsend and Lapata, 2012; Almeida and Martins, 2013), we wish to be able to compress sentences so we can pack more information into a summary. During training, our model learns how to take advantage of available compression options and select content to match human generated summaries as closely possible.² We explore two ways of deriving units for compression: the RST-based compressions of Hirao et al. (2013) and the syntactic compressions of Berg-Kirkpatrick et al. (2011).

RST compressions Figure 2a shows how to derive compressions from Rhetorical Structure Theory (Mann and Thompson, 1988; Carlson et al., 2001). We show a sentence broken into elemen-

tary discourse units (EDUs) with RST relations between them. Units marked as SAME-UNIT must both be kept or both be deleted, but other nodes in the tree structure can be deleted as long as we do not delete the parent of an included node. For example, we can delete the ELABORATION clause, but we can delete neither the first nor last EDU. Arrows depict the constraints this gives rise to in the ILP (see Figure 1): u_2 requires u_1 , and u_1 and u_3 mutually require each other. This is a more constrained form of compression than was used in past work (Hirao et al., 2013), but we find that it improves human judgments of fluency (Section 4.3).

Syntactic compressions Figure 2b shows two examples of compressions arising from syntactic patterns (Berg-Kirkpatrick et al., 2011): deletion of the second part of a coordinated NP and deletion of a PP modifier to an NP. These patterns were curated to leave sentences as grammatical after being compressed, though perhaps with damaged semantic content.

Combined compressions Figure 2c shows the textual units and requirement relations yielded by combining these two types of compression. On this example, the two schemes capture orthogonal compressions, and more generally we find that they stack to give better results for our final system (see Section 4.3). To actually synthesize textual units and the constraints between them, we start from the set of RST textual units and introduce syntactic compressions as new children when they don't cross existing brackets; because syntactic compressions are typically narrower in scope, they are usually completely contained in EDUs.

²The features in our model are actually rich enough to learn a sophisticated compression model, but the data we have (abstractive summaries) does not directly provide examples of correct compressions; past work has gotten around this with multi-task learning (Almeida and Martins, 2013), but we simply treat grammaticality as a constraint from up-stream models.

Figure 2d shows an example of this process: the possible deletion of *with Aetna* is grafted onto the textual unit and appropriate requirement relations are introduced. The net effect is that the textual unit is wholly included, partially included (*with Aetna* removed), or not at all.

Formally, we define an RST tree as $T_{\text{rst}} = (S_{\text{rst}}, \pi_{\text{rst}})$ where S_{rst} is a set of EDU spans (i, j) and $\pi : S \rightarrow 2^S$ is a mapping from each EDU span to EDU spans it depends on. Syntactic compressions can be expressed in a similar way with trees T_{syn} . These compressions are typically smaller-scale than EDU-based compressions, so we use the following modification scheme. Denote by $T_{\text{syn}(kl)}$ a nontrivial (supports some compression) subtree of T_{syn} that is completely contained in an EDU (i, j) . We build the following combined compression tree, which we refer to as the *augmentation* of T_{rst} with $T_{\text{syn}(kl)}$:

$$T_{\text{comb}} = (S \cup S_{\text{syn}(kl)} \cup \{(i, k), (l, j)\}, \pi_{\text{rst}} \cup \pi_{\text{syn}(kl)} \cup \{(i, k) \rightarrow (l, j), (l, j) \rightarrow (i, k), (k, l) \rightarrow (i, k)\})$$


That is, we maintain the existing tree structure except for the EDU (i, j) , which is broken into three parts: the outer two depend on each other (*is a claims adjuster* and *.* from Figure 2d) and the inner one depends on the others and preserves the tree structure from T_{syn} . We augment T_{rst} with all maximal subtrees of T_{syn} , i.e. all trees that are not contained in other trees that are used in the augmentation process.

This is broadly similar to the combined compression scheme in Kikuchi et al. (2014) but we use a different set of constraints that more strictly enforce grammaticality.³

2.2 Anaphora Constraints

What kind of cross-sentential coherence do we need to ensure for the kinds of summaries our system produces? Many notions of coherence are useful, including centering theory (Grosz et al., 1995) and lexical cohesion (Nishikawa et al., 2014), but one of the most pressing phenomena to deal with is pronoun anaphora (Clarke and Lapata, 2010). Cases of pronouns being “orphaned” during extraction (their antecedents are deleted) are

³We also differ from past work in that we do not use cross-sentential RST constraints (Hirao et al., 2013; Yoshida et al., 2014). We experimented with these and found no improvement from using them, possibly because we have a feature-based model rather than a heuristic content selection procedure, and possibly because automatic discourse parsers are less good at recovering cross-sentence relations.



Replacement (2.2.1): If $\max(p_1, p_2, p_3) > \alpha$:

Allow pronoun replacement with the predicted antecedent and add the following constraint:

$$x_{2,1}^{\text{REF}} = 1 \text{ iff } x_1^{\text{UNIT}} = 0 \text{ and } x_2^{\text{UNIT}} = 1$$

Replace the first pronoun in the second textual unit

Otherwise (i.e. if no replacement is possible):

Antecedent inclusion (2.2.2): If $p_1 + p_2 > \beta$

Add the following constraint: $x_2^{\text{UNIT}} \leq x_1^{\text{UNIT}}$

Figure 3: Modifications to the ILP to capture pronoun coherence. *It*, which refers to *Kellogg*, has several possible antecedents from the standpoint of an automatic coreference system (Durrett and Klein, 2014). If the coreference system is confident about its selection (above a threshold α on the posterior probability), we allow for the model to explicitly replace the pronoun if its antecedent would be deleted (Section 2.2.1). Otherwise, we merely constrain one or more probable antecedents to be included (Section 2.2.2); even if the coreference system is incorrect, a human can often correctly interpret the pronoun with this additional context.

relatively common: they occur in roughly 60% of examples produced by our summarizer when no anaphora constraints are enforced. This kind of error is particularly concerning for summary interpretation and impedes the ability of summaries to convey information effectively (Grice, 1975). Our solution is to explicitly impose constraints on the model based on pronoun anaphora resolution.⁴

Figure 3 shows an example of a problem case. If we extract only the second textual unit shown, the pronoun *it* will lose its antecedent, which in this case is *Kellogg*. We explore two types of constraints for dealing with this: rewriting the pronoun explicitly, or constraining the summary to include the pronoun’s antecedent.

2.2.1 Pronoun Replacement

One way of dealing with these pronoun reference issues is to explicitly replace the pronoun with what it refers to. This replacement allows us to maintain maximal extraction flexibility, since we

⁴We focus on pronoun coreference because it is the most pressing manifestation of this problem and because existing coreference systems perform well on pronouns compared to harder instances of coreference (Durrett and Klein, 2013).

can make an isolated textual unit meaningful even if it contains a pronoun. Figure 3 shows how this process works. We run the Berkeley Entity Resolution System (Durrett and Klein, 2014) and compute posteriors over possible links for the pronoun. If the coreference system is sufficiently confident in its prediction (i.e. $\max_i p_i > \alpha$ for a specified threshold $\alpha > \frac{1}{2}$), we allow ourselves to replace the pronoun with the first mention of the entity corresponding to the pronoun’s most likely antecedent. In Figure 3, if the system correctly determines that *Kellogg* is the correct antecedent with high probability, we enable the first replacement shown there, which is used if u_2 is included the summary without u_1 .⁵

As shown in the ILP in Figure 1, we instantiate corresponding pronoun replacement variables \mathbf{x}^{REF} where $x_{ij}^{\text{REF}} = 1$ implies that the j th pronoun in the i th sentence should be replaced in the summary. We use a candidate pronoun replacement if and only if the pronoun’s corresponding (predicted) entity hasn’t been mentioned previously in the summary.⁶ Because we are generally replacing pronouns with longer mentions, we also need to modify the length constraint to take this into account. Finally, we incorporate features on pronoun replacements in the objective, which helps the model learn to prefer pronoun replacements that help it to more closely match the human summaries.

2.2.2 Pronoun Antecedent Constraints

Explicitly replacing pronouns is risky: if the coreference system makes an incorrect prediction, the intended meaning of the summary may be damaged. Fortunately, the coreference model’s posterior probabilities have been shown to be well-calibrated (Nguyen and O’Connor, 2015), meaning that cases where it is likely to make errors are signaled by flatter posterior distributions. In this case, we enable a more conservative set of constraints that include additional content in the summary to make the pronoun reference clear without explicitly replacing it. This is done by requiring the inclusion of any textual unit which contains

⁵If the proposed replacement is a proper mention, we replace the pronoun just with the subset of the mention that constitutes a named entity (rather than the whole noun phrase). We control for possessive pronouns by deleting or adding ‘s as appropriate.

⁶Such a previous mention may be a pronoun; however, note that that pronoun would then be targeted for replacement unless its antecedent were included somehow.

possible pronoun references whose posteriors sum to at least a threshold parameter β . Figure 3 shows that this constraint can force the inclusion of u_1 to provide additional context. Although this could still lead to unclear pronouns if text is stitched together in an ambiguous or even misleading way, in practice we observe that the textual units we force to be added almost always occur very recently before the pronoun, giving enough additional context for a human reader to figure out the pronoun’s antecedent unambiguously.

2.3 Features

The features in our model (see Figure 1) consist of a set of surface indicators capturing mostly lexical and configurational information. Their primary role is to identify important document content. The first three types of features fire over textual units, the last over pronoun replacements.

Lexical These include indicator features on non-stopwords in the textual unit that appear at least five times in the training set and analogous POS features. We also use lexical features on the first, last, preceding, and following words for each textual unit. Finally, we conjoin each of these features with an indicator of bucketed position in the document (the index of the sentence containing the textual unit).

Structural These features include various conjunctions of the position of the textual unit in the document, its length, the length of its corresponding sentence, the index of the paragraph it occurs in, and whether it starts a new paragraph (all values are bucketed).

Centrality These features capture rough information about the centrality of content: they consist of bucketed word counts conjoined with bucketed sentence index in the document. We also fire features on the number of times of each entity mentioned in the sentence is mentioned in the rest of the document (according to a coreference system), the number of entities mentioned in the sentence, and surface properties of mentions including type and length

Pronoun replacement These target properties of the pronoun replacement such as its length, its sentence distance from the current mention, its type (nominal or proper), and the identity of the pronoun being replaced.

3 Learning

We learn weights \mathbf{w} for our model by training on a large corpus of documents \mathbf{u} paired with reference summaries \mathbf{y} . We formulate our learning problem as a standard instance of structured SVM (see Smith (2011) for an introduction). Because we want to optimize explicitly for ROUGE-1,⁷ we define a ROUGE-based loss function that accommodates the nature of our supervision, which is in terms of abstractive summaries \mathbf{y} that in general cannot be produced by our model. Specifically, we take:

$$\ell(\mathbf{x}^{\text{NGRAM}}, \mathbf{y}) = \max_{\mathbf{x}^*} \text{ROUGE-1}(\mathbf{x}^*, \mathbf{y}) - \text{ROUGE-1}(\mathbf{x}^{\text{NGRAM}}, \mathbf{y})$$

i.e. the gap between the hypothesis’s ROUGE score and the oracle ROUGE score achievable under the model (including constraints). Here $\mathbf{x}^{\text{NGRAM}}$ are indicator variables that track, for each n -gram type in the reference summary, whether that n -gram is present in the system summary. These are the sufficient statistics for computing ROUGE.

We train the model via stochastic subgradient descent on the primal form of the structured SVM objective (Ratliff et al., 2007; Kummerfeld et al., 2015). In order to compute the subgradient for a given training example, we need to find the most violated constraint on the given instance through a loss-augmented decode, which for a linear model takes the form $\arg \max_{\mathbf{x}} \mathbf{w}^\top \mathbf{f}(\mathbf{x}) + \ell(\mathbf{x}, \mathbf{y})$. To do this decode at training time in the context of our model, we use an extended version of our ILP in Figure 1 that is augmented to explicitly track type-level n -grams:

$$\max_{\mathbf{x}^{\text{UNIT}}, \mathbf{x}^{\text{REF}}, \mathbf{x}^{\text{NGRAM}}} \left[\sum_i \left[x_i^{\text{UNIT}} (\mathbf{w}^\top \mathbf{f}(u_i)) \right] + \sum_{(i,j)} \left[x_{ij}^{\text{REF}} (\mathbf{w}^\top \mathbf{f}(r_{ij})) \right] - \ell(\mathbf{x}^{\text{NGRAM}}, \mathbf{y}) \right]$$

subject to all constraints from Figure 1, and

$$x_i^{\text{NGRAM}} = 1 \text{ iff an included textual unit or replacement contains the } i\text{th reference } n\text{-gram}$$

These kinds of variables and constraints are common in multi-document summarization systems

⁷We found that optimizing for ROUGE-1 actually resulted in a model with better performance on both ROUGE-1 and ROUGE-2. We hypothesize that this is because framing our optimization in terms of ROUGE-2 would lead to a less nuanced set of constraints: bigram matches are relatively rare when the reference is a short, abstractive summary, so a loss function based on ROUGE-2 will express a flatter preference structure among possible outputs.

that score bigrams (Gillick and Favre, 2009 *inter alia*). Note that since ROUGE is only computed over non-stopword n -grams and pronoun replacements only replace pronouns, pronoun replacement can never remove an n -gram that would otherwise be included.

For all experiments, we optimize our objective using AdaGrad (Duchi et al., 2011) with ℓ_1 regularization ($\lambda = 10^{-8}$, chosen by grid search), with a step size of 0.1 and a minibatch size of 1. We train for 10 iterations on the training data, at which point held-out model performance no longer improves. Finally, we set the anaphora thresholds $\alpha = 0.8$ and $\beta = 0.6$ (see Section 2.2). The values of these and other hyperparameters were determined on a held-out development set from our New York Times training data. All ILPs are solved using GLPK version 4.55.

4 Experiments

We primarily evaluate our model on a roughly 3000-document evaluation set from the New York Times Annotated Corpus (Sandhaus, 2008). We also investigate its performance on the RST Discourse Treebank (Carlson et al., 2001), but because this dataset is only 30 documents it provides much less robust estimates of performance.⁸ Throughout this section, when we decode a document, we set the word budget for our summarizer to be the same as the number of words in the corresponding reference summary, following previous work (Hirao et al., 2013; Yoshida et al., 2014).

4.1 Preprocessing

We preprocess all data using the Berkeley Parser (Petrov et al., 2006), specifically the GPU-accelerated version of the parser from Hall et al. (2014), and the Berkeley Entity Resolution System (Durrett and Klein, 2014). For RST discourse analysis, we segment text into EDUs using a semi-Markov CRF trained on the RST treebank with features on boundaries similar to those of Hernault et al. (2010), plus novel features on spans including span length and span identity for short spans.

To follow the conditions of Yoshida et al. (2014) as closely as possible, we also build a discourse parser in the style of Hirao et al. (2013), since their parser is not publicly available. Specifically,

⁸Tasks like DUC and TAC have focused on multi-document summarization since around 2003, hence the lack of more standard datasets for single-document summarization.

NYT50 article:

Federal officials reported yesterday that students in 4th, 8th and 12th grades had scored modestly higher on an American history test than five years earlier, although **more than half of high school seniors still showed poor command of basic facts** like the effect of the cotton gin on the slave economy or the causes of the Korean War. Federal officials said they considered the results encouraging because at each level tested, student performance had improved since the last time the exam was administered, in 2001. "In U.S. history there were higher scores in 2006 for all three grades," said Mark Schneider, commissioner of the National Center for Education Statistics, which administers the test, at a Boston news conference that the Education Department carried by Webcast. **The results were less encouraging on a national civics test, on which only fourth graders made any progress.** The best results in the history test were also in fourth grade, where 70 percent of students attained the basic level of achievement or better. **The test results in the two subjects are likely to be closely studied, because Congress is considering the renewal of President Bush's signature education law, the No Child Left Behind Act.** A number of studies have shown that because No Child Left Behind requires states...

Summary:

National Center for Education Statistics reports students in 4th, 8th and 12th grades scored modestly higher on American history test than five years earlier. Says more **than half of high school seniors still show poor command of basic facts.** Only 4th graders made any progress in civics test. New exam results are another ingredient in debate over renewing Pres Bush's signature No Child Left Behind Act.

Filtered article:

Long before President Bush's proposal to rethink Social Security became part of the national conversation, Westchester County came up with its own dialogue to bring issues of aging to the forefront. Before the White House Conference on Aging scheduled in October, **the county's Office for the Aging a year ago started Speak-Up**, which stands for Student Participants Embrace Aging Issues of Key Concern, to reach students in the county's 13 colleges and universities. Through a variety of events **to bring together the elderly and college students**, organizers said they hoped to have by this spring a series of recommendations that could be given to Washington...

Summary:

Article on Speak-Up, program begun by Westchester County Office for the Aging to bring together elderly and college students.

Figure 4: Examples of an article kept in the NYT50 dataset (top) and an article removed because the summary is too short. The top summary has a rich structure to it, corresponding to various parts of the document (bolded) and including some text that is essentially a direct extraction.




Figure 5: Counts on a 1000-document sample of how frequently both a document prefix baseline and a ROUGE oracle summary contain sentences at various indices in the document. There is a long tail of useful sentences later in the document, as seen by the fact that the oracle sentence counts drop off relatively slowly. Smart selection of content therefore has room to improve over taking a prefix of the document.

we use the first-order projective parsing model of McDonald et al. (2005) and features from Soricut and Marcu (2003), Hernault et al. (2010), and Joty et al. (2013). When using the same head annotation scheme as Yoshida et al. (2014), we outperform their discourse dependency parser on unlabeled dependency accuracy, getting 56% as opposed to 53%.

4.2 New York Times Corpus

We now provide some details about the New York Times Annotated corpus. This dataset contains

110,540 articles with abstractive summaries; we split these into 100,834 training and 9706 test examples, based on date of publication (test is all articles published on January 1, 2007 or later). Examples of two documents from this dataset are shown in Figure 4. The bottom example demonstrates that some summaries are extremely short and formulaic (especially those for obituaries and editorials). To counter this, we filter the raw dataset by removing all documents with summaries that are shorter than 50 words. One benefit of filtering is that the length distribution of our resulting dataset is more in line with standard summarization evaluations like DUC; it also ensures a sufficient number of tokens in the budget to produce nontrivial summaries. The filtered test set, which we call NYT50, includes 3,452 test examples out of the original 9,706.

Interestingly, this dataset is one where the classic document prefix baseline can be substantially outperformed, unlike in some other summarization settings (Penn and Zhu, 2008). We show this fact explicitly in Section 4.3, but Figure 5 provides additional analysis in this regard. We compute oracle ROUGE-1 sentence-extractive summaries on a 1000-document subset of the training set and look at where the extracted sentences lie in the document. While they certainly skew earlier in the document, they do not all fall within the doc-

	R-1 ↑	R-2 ↑	CG ↑	UP ↓
Baselines				
First sentences	28.6	17.3	8.21	0.28
First k words	35.7	21.6	—	—
Bigram Frequency	25.1	9.8	—	—
Past work				
Tree Knapsack	34.7	19.6	7.20	0.42
This work				
Sentence extraction	38.8	23.5	7.93	0.32
EDU extraction	41.9	25.3	6.38	0.65
Full	42.2	25.9	*†7.52	*0.36
Ablations from Full				
No Anaphoricity	42.5	26.3	7.46	0.44
No Syntactic Compr	41.1	25.0	—	—
No Discourse Compr	40.5	24.7	—	—

Table 1: Results on the NYT50 test set (documents with summaries of at least 50 tokens) from the New York Times Annotated Corpus (Sandhaus, 2008). We report ROUGE-1 (R-1), ROUGE-2 (R-2), clarity/grammaticality (CG), and number of unclear pronouns (UP) (lower is better). On content selection, our system substantially outperforms all baselines, our implementation of the tree knapsack system (Yoshida et al., 2014), and learned extractive systems with less compression, even an EDU-extractive system that sacrifices grammaticality. On clarity metrics, our final system performs nearly as well as sentence-extractive systems. The symbols * and † indicate statistically significant gains compared to No Anaphoricity and Tree Knapsack (respectively) with $p < 0.05$ according to a bootstrap resampling test. We also see that removing either syntactic or EDU-based compressions decreases ROUGE.

ument prefix summary. One reason for this is that many of the articles are longer-form pieces that begin with a relatively content-free lede of several sentences, which should be identifiable with lexicosyntactic indicators as are used in our discriminative model.

4.3 New York Times Results

We evaluate our system along two axes: first, on content selection, using ROUGE⁹ (Lin and Hovy, 2003), and second, on clarity of language and referential structure, using annotators from Amazon Mechanical Turk. We follow the method of Gillick and Liu (2010) for this evaluation and ask Turkers to rate a summary on how grammatical it is using a 10-point Likert scale. Furthermore, we ask how many unclear pronouns references there were in the text. The Turkers do not see the original document or the reference summary, and rate each summary in isolation. Gillick and Liu (2010) showed that for linguistic quality judgments (as opposed to content judgments), Turkers reproduced the ranking of systems according to expert judgments.

To speed up preprocessing and training time

⁹We use the ROUGE 1.5.5 script with the following command line arguments: `-n 2 -x -m -s`. All given results are macro-averaged recall values over the test set.

on this corpus, we further restrict our training set to only contain documents with fewer than 100 EDUs. All told, the final system takes roughly 20 hours to make 10 passes through the subsampled training data (22,000 documents) on a single core of an Amazon EC2 r3.4xlarge instance.

Table 1 shows the results on the NYT50 corpus. We compare several variants of our system and baselines. For baselines, we use two variants of first k : one which must stop on a sentence boundary (which gives better linguistic quality) and one which always consumes k tokens (which gives better ROUGE). We also use a heuristic sentence-extractive baseline that maximizes the document counts (term frequency) of bigrams covered by the summary, similar in spirit to the multi-document method of Gillick and Favre (2009).¹⁰ We also compare to our implementation of the Tree Knapsack method of Yoshida et al. (2014), which matches their results very closely on the RST Discourse Treebank when discourse trees are controlled for. Finally, we compare several variants of our system: purely extractive systems operating over sentences and EDUs respectively, our full system, and ablations removing either the anaphoricity component or parts of the compression module.

In terms of content selection, we see that all of the systems that incorporate end-to-end learning (under “This work”) substantially outperform our various heuristic baselines. Our full system using the full compression scheme is substantially better on ROUGE than ablations where the syntactic or discourse compressions are removed. These improvements reflect the fact that more compression options give the system more flexibility to include key content words. Removing the anaphora resolution constraints actually causes ROUGE to increase slightly (as a result of granting the model flexibility), but has a negative impact on the linguistic quality metrics.

On our linguistic quality metrics, it is no surprise that the sentence prefix baseline performs the best. Our sentence-extractive system also does well on these metrics. Compared to the EDU-extractive system with no constraints, our constrained compression method improves substantially on both linguistic quality and reduces the

¹⁰Other heuristic multi-document approaches could be compared to, e.g. He et al. (2012), but a simple term frequency method suffices to illustrate how these approaches can underperform in the single-document setting.

	ROUGE-1	ROUGE-2
First k words	23.5	8.3
Tree Knapsack	25.1	8.7
Full	26.3	8.0

Table 2: Results for RST Discourse Treebank (Carlson et al., 2001). Differences between our system and the Tree Knapsack system of Yoshida et al. (2014) are not statistically significant, reflecting the high variance in this small (20 document) test set.

number of unclear pronouns, and adding the pronoun anaphora constraints gives further improvement. Our final system is approaches the sentence-extractive baseline, particularly on unclear pronouns, and achieves substantially higher ROUGE score.

4.4 RST Treebank

We also evaluate on the RST Discourse Treebank, of which 30 documents have abstractive summaries. Following Hirao et al. (2013), we use the gold EDU segmentation from the RST corpus but automatic RST trees. We break this into a 10-document development set and a 20-document test set. Table 2 shows the results on the RST corpus. Our system is roughly comparable to Tree Knapsack here, and we note that none of the differences in the table are statistically significant. We also observed significant variation between multiple runs on this corpus, with scores changing by 1-2 ROUGE points for slightly different system variants.¹¹

5 Conclusion

We presented a single-document summarization system trained end-to-end on a large corpus. We integrate a compression model that enforces grammaticality as well as pronoun anaphoricity constraints that enforce coherence. Our system improves substantially over baseline systems on ROUGE while still maintaining good linguistic quality.

Our system and models are publicly available at <http://nlp.cs.berkeley.edu>

¹¹The system of Yoshida et al. (2014) is unavailable, so we use a reimplementation. Our results differ from theirs due to having slightly different discourse trees, which cause large changes in metrics due to high variance on the test set.

Acknowledgments

This work was partially supported by NSF Grant CNS-1237265 and a Google Faculty Research Award. Thanks to Tsutomu Hirao for providing assistance with our reimplementation of the Tree Knapsack model, and thanks the anonymous reviewers for their helpful comments.

References

- Miguel Almeida and Andre Martins. 2013. Fast and Robust Compressive Summarization with Dual Decomposition and Multi-Task Learning. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Regina Barzilay and Mirella Lapata. 2008. Modeling Local Coherence: An Entity-based Approach. *Computational Linguistics*, 34(1):1–34, March.
- Regina Barzilay and Lillian Lee. 2004. Catching the Drift: Probabilistic Content Models, with Applications to Generation and Summarization. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Regina Barzilay, Noemie Elhadad, and Kathleen R. McKeown. 2001. Sentence Ordering in Multidocument Summarization. In *Proceedings of the International Conference on Human Language Technology Research*.
- Taylor Berg-Kirkpatrick, Dan Gillick, and Dan Klein. 2011. Jointly Learning to Extract and Compress. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Jaime Carbonell and Jade Goldstein. 1998. The Use of MMR, Diversity-based Reranking for Reordering Documents and Producing Summaries. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurowski. 2001. Building a Discourse-tagged Corpus in the Framework of Rhetorical Structure Theory. In *Proceedings of the Second SIGDIAL Workshop on Discourse and Dialogue*.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2013. Towards Coherent Multi-Document Summarization. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- James Clarke and Mirella Lapata. 2008. Global Inference for Sentence Compression an Integer Linear Programming Approach. *Journal of Artificial Intelligence Research*, 31(1):399–429, March.

- James Clarke and Mirella Lapata. 2010. Discourse Constraints for Document Compression. *Computational Linguistics*, 36(3):411–441, September.
- Hal Daumé, III and Daniel Marcu. 2002. A Noisy-Channel Model for Document Compression. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, 12:2121–2159, July.
- Jesse Dunietz and Daniel Gillick. 2014. A New Entity Salience Task with Millions of Training Examples. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Greg Durrett and Dan Klein. 2013. Easy Victories and Uphill Battles in Coreference Resolution. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, October.
- Greg Durrett and Dan Klein. 2014. A Joint Model for Entity Analysis: Coreference, Typing, and Linking. In *Transactions of the Association for Computational Linguistics (TACL)*.
- Dan Gillick and Benoit Favre. 2009. A Scalable Global Model for Summarization. In *Proceedings of the Workshop on Integer Linear Programming for Natural Language Processing*.
- Dan Gillick and Yang Liu. 2010. Non-Expert Evaluation of Summarization Systems is Risky. In *Proceedings of the NAACL Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*.
- H.P. Grice. 1975. Logic and Conversation. *Syntax and Semantics 3: Speech Acts*, pages 41–58.
- Barbara J. Grosz, Scott Weinstein, and Aravind K. Joshi. 1995. Centering: A Framework for Modeling the Local Coherence of Discourse. *Computational Linguistics*, 21(2):203–225, June.
- David Hall, Taylor Berg-Kirkpatrick, John Canny, and Dan Klein. 2014. Sparser, Better, Faster GPU Parsing. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Zhanying He, Chun Chen, Jiajun Bu, Can Wang, Lijun Zhang, Deng Cai, and Xiaofei He. 2012. Document Summarization Based on Data Reconstruction. In *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*.
- Hugo Hernault, Helmut Prendinger, David A. Duverle, Mitsuru Ishizuka, and Tim Paek. 2010. HILDA: A discourse parser using support vector machine classification. *Dialogue and Discourse*, 1:1–33.
- Tsutomu Hirao, Yasuhisa Yoshida, Masaaki Nishino, Norihito Yasuda, and Masaaki Nagata. 2013. Single-Document Summarization as a Tree Knapsack Problem. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kai Hong and Ani Nenkova. 2014. Improving the Estimation of Word Importance for News Multi-Document Summarization. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Shafiq Joty, Giuseppe Carenini, Raymond Ng, and Yashar Mehdad. 2013. Combining Intra- and Multi-sentential Rhetorical Parsing for Document-level Discourse Analysis. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Yuta Kikuchi, Tsutomu Hirao, Hiroya Takamura, Manabu Okumura, and Masaaki Nagata. 2014. Single Document Summarization based on Nested Tree Structure. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Jonathan K. Kummerfeld, Taylor Berg-Kirkpatrick, and Dan Klein. 2015. An Empirical Analysis of Optimization for Max-Margin NLP. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Chen Li, Yang Liu, and Lin Zhao. 2015. Using External Resources and Joint Learning for Bigram Weighting in ILP-Based Multi-Document Summarization. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Hui Lin and Jeff Bilmes. 2011. A Class of Submodular Functions for Document Summarization. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Chin-Yew Lin and Eduard Hovy. 2003. Automatic Evaluation of Summaries Using N-gram Co-occurrence Statistics. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Chin-Yew Lin. 2003. Improving Summarization Performance by Sentence Compression: A Pilot Study. In *Proceedings of the International Workshop on Information Retrieval with Asian Languages*.
- Annie Louis and Ani Nenkova. 2012. A Coherence Model Based on Syntactic Patterns. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.
- Annie Louis, Aravind Joshi, and Ani Nenkova. 2010. Discourse Indicators for Content Selection in Summarization. In *Proceedings of the SIGDIAL 2010 Conference*.

- Inderjeet Mani. 2001. *Automatic Summarization*. John Benjamins Publishing.
- William C. Mann and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8(3):243–281.
- Daniel Marcu. 1998. Improving summarization through rhetorical parsing tuning. In *Proceedings of the Workshop on Very Large Corpora*.
- Andre Martins and Noah A. Smith. 2009. Summarization with a Joint Model for Sentence Extraction and Compression. In *Proceedings of the Workshop on Integer Linear Programming for Natural Language Processing*.
- Ryan McDonald, Koby Crammer, and Fernando Pereira. 2005. Online Large-margin Training of Dependency Parsers. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Ryan McDonald. 2006. Discriminative Sentence Compression With Soft Syntactic Evidence. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Kathleen McKeown, Jacques Robin, and Karen Kuchich. 1995. Generating Concise Natural Language Summaries. *Information Processing and Management*, 31(5):703–733, September.
- Shamima Mithun and Leila Kosseim. 2011. Discourse Structures to Reduce Discourse Incoherence in Blog Summarization. In *Proceedings of Recent Advances in Natural Language Processing*.
- Ani Nenkova and Kathleen McKeown. 2011. Automatic summarization. *Foundations and Trends in Information Retrieval*, 5(2?3):103–233.
- Khanh Nguyen and Brendan O’Connor. 2015. Posterior Calibration and Exploratory Analysis for Natural Language Processing Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Hitoshi Nishikawa, Kazuho Arita, Katsumi Tanaka, Tsutomu Hirao, Toshiro Makino, and Yoshihiro Matsuo. 2014. Learning to Generate Coherent Summary with Discriminative Hidden Semi-Markov Model. In *Proceedings of the International Conference on Computational Linguistics (COLING)*.
- Gerald Penn and Xiaodan Zhu. 2008. A Critical Re-assessment of Evaluation Baselines for Speech Summarization. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. Learning Accurate, Compact, and Interpretable Tree Annotation. In *Proceedings of the Conference on Computational Linguistics and the Association for Computational Linguistics (ACL-COLING)*.
- Daniele Pighin, Marco Cornolti, Enrique Alfonseca, and Katja Filippova. 2014. Modelling Events through Memory-based, Open-IE Patterns for Abstractive Summarization. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Nathan J. Ratliff, Andrew Bagnell, and Martin Zinkevich. 2007. (Online) Subgradient Methods for Structured Prediction. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.
- Evan Sandhaus. 2008. The New York Times Annotated Corpus. In *Linguistic Data Consortium*.
- Noah A. Smith. 2011. *Linguistic Structure Prediction*. Morgan & Claypool Publishers, 1st edition.
- Radu Soricut and Daniel Marcu. 2003. Sentence Level Discourse Parsing Using Syntactic and Lexical Information. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Kapil Thadani and Kathleen McKeown. 2013. Supervised Sentence Fusion with Single-Stage Inference. In *Proceedings of the International Joint Conference on Natural Language Processing (IJCNLP)*.
- Kristian Woodsend and Mirella Lapata. 2012. Multiple Aspect Summarization Using Integer Linear Programming. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.
- Yasuhisa Yoshida, Jun Suzuki, Tsutomu Hirao, and Masaaki Nagata. 2014. Dependency-based Discourse Parser for Single-Document Summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.