,

Tutorial on Abstractive Text Summarization

Advaith Siddharthan

NLG Summer School, Aberdeen, 22 July 2015

Tasks in text summarization

- Extractive Summarization (previous tutorial)
  - Sentence Selection, etc
- Abstractive Summarization
  - Mimicing what human summarizers do
  - Sentence Compression and Fusion
  - Regenerating Referring Expressions
- Template Based Summarization
  - Perform information extraction, then use NLG Templates

Cut and Paste in Professional Summarization

- Humans also reuse the input text to produce summaries
- But they don't just extract sentences, they do a lot of cut and paste
  - corpus analysis (Barzilay et al., 1999)
    - 300 summaries, 1,642 sentences
    - 81% sentences were constructed by cutting and pasting

Major Cut and Paste Operations

- Sentence Compression
  - ABA~~CDC~~DFDS~~GFG~~DA $\longrightarrow$ ABADFDSDA
  - Summarizing a sentence, e.g. for headline generation
  - Removes peripheral information from a sentence to shorten summary

Major Cut and Paste Operations

- Sentence Compression
  - ABA~~CDC~~DFDS~~GFG~~DA $\longrightarrow$ ABADFDSDA
  - Summarizing a sentence, e.g. for headline generation
  - Removes peripheral information from a sentence to shorten summary
- Sentence Fusion
  - ABA~~CDC~~DFDS~~GFG~~ + ~~CDC~~GFD~~GFGDA~~ $\longrightarrow$ ABAGFDDFDS
  - Merge information from multiple (similar) sentences.
  - Reduces redundancy in summary

Major Cut and Paste Operations

- Sentence Compression
  - ABA~~CDC~~DFDS~~GFG~~DA $\longrightarrow$ ABADFDSDA
  - Summarizing a sentence, e.g. for headline generation
  - Removes peripheral information from a sentence to shorten summary
- Sentence Fusion
  - ABA~~CDC~~DFDS~~GFG~~ + ~~CDC~~GFD~~GFGDA~~ $\longrightarrow$ ABAGFDDFDS
  - Merge information from multiple (similar) sentences.
  - Reduces redundancy in summary
- Syntactic Reorganization
  - ABADFGS $\longrightarrow$ DFGSABA
  - Often done to make the summary coherent (preserve focus, etc)

Major Cut and Paste Operations

- Sentence Compression
  - ABA~~CDC~~DFDS~~GFG~~DA $\longrightarrow$ ABADFDSDA
  - Summarizing a sentence, e.g. for headline generation
  - Removes peripheral information from a sentence to shorten summary
- Sentence Fusion
  - ABA~~CDC~~DFDS~~GFG~~ + ~~CDC~~GFD~~GFGDA~~ $\longrightarrow$ ABAGFDDFDS
  - Merge information from multiple (similar) sentences.
  - Reduces redundancy in summary
- Syntactic Reorganization
  - ABADFGS $\longrightarrow$ DFGSABA
  - Often done to make the summary coherent (preserve focus, etc)
- Lexical Paraphrase
  - ABACDFGDSFD $\longrightarrow$ ABAGHYGDSFD
  - Use simpler words that are easier to understand in the new context.

Sentence Compression

- A research topic in itself, too many approaches to discuss here in depth
- Typically viewed as producing a summary of a single sentence
  - Should be shorter
  - Should remain grammatical
  - Should keep the most important information

Sentence Compression

- (Grefenstette, 1998; Jing et al., 1998; Knight & Marcu, 2000; Riezler et al., 2003)...

  Former Democratic National Committee finance director Richard Sullivan faced more pointed questioning from Republicans during his second day on the witness stand in the Senate's fund-raising investigation.

- Richard Sullivan faced pointed questioning.
- Richard Sullivan faced pointed questioning from Republicans during day on stand in Senate fund-raising investigation.

Example: Reluctant Trimmer

- Developed by Nomoto (Angrosh et al., 2014) for Text Simplification (Siddharthan & Angrosh, 2014), rather than summarization.
  - Considers text as a whole and optimises global constraints for:
    - lexical density
    - ratio of difficult words
    - text length
  - Reluctant Trimmer is based on reluctant paraphrasing (Dras, 1999) "make as little change as possible to the text to satisfy a set of constraints"

Reluctant Trimmer - Architecture

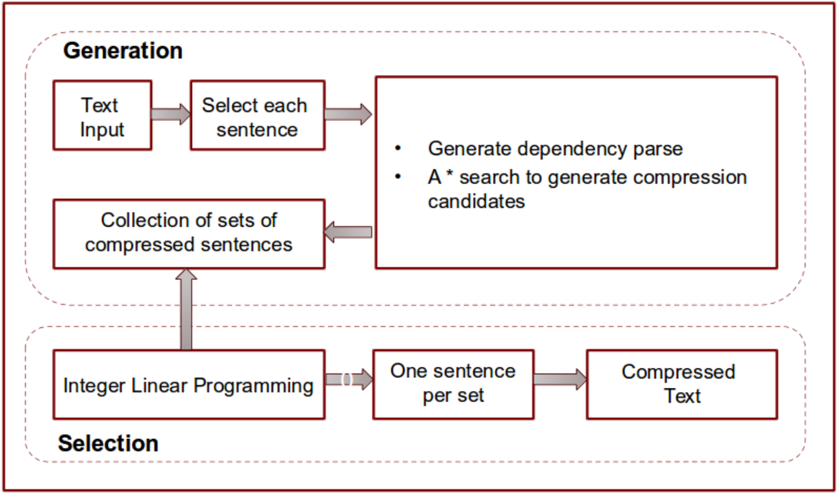

Reluctant Trimmer - Graphical View

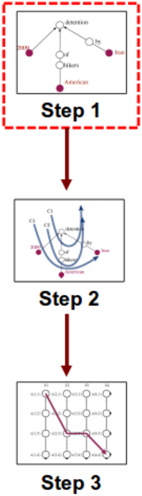

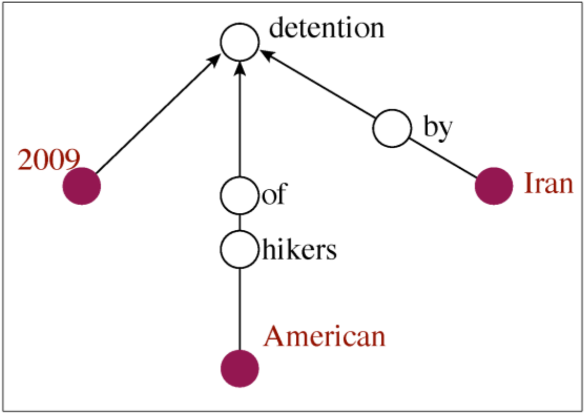

Step 1 - Dependency structure for "2009 detention of American hikers by Iran"

Reluctant Trimmer - Graphical View

Step 1

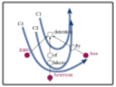

Step 2

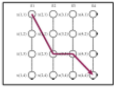

Step 3

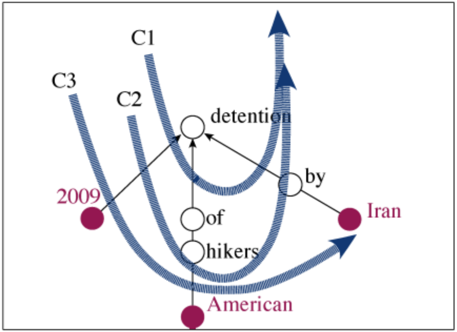

Step 2 – Trimming Dependency Tree

Sentences Generated:

C1: detention
C2: detention of hikers
C3: detention of hikers by Iran

Reluctant Trimmer - Graphical View

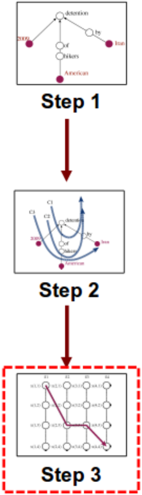

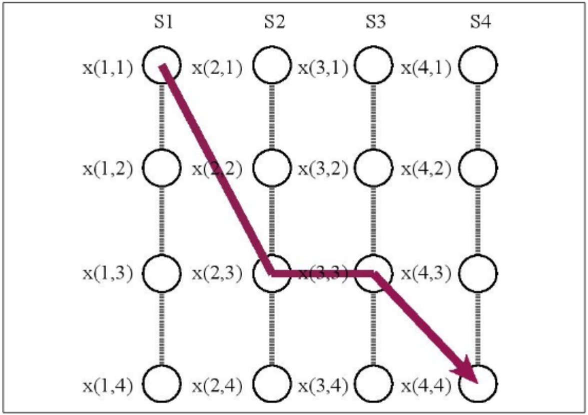

Step 3 – Decoding with ILP

Reluctant Trimmer

- Decoded using ILP
  - Constraints can be specified at the level of a text, not an individual sentence.
    - lexical density
    - ratio of difficult words
    - text length
  - While developed for text simplification, it can be adapted to summarisation tasks by changing the constraints, for example to take into account
    - some notion of topic

Sentence Fusion

1. IDF Spokeswoman did not confirm this, but said the Palestinians fired an antitank missile at a bulldozer.
2. The clash erupted when Palestinian militants fired machine guns and antitank missiles at a bulldozer that was building an embankment in the area to better protect Israeli forces.
3. The army expressed regret at the loss of innocent lives but a senior commander said troops had shot in self-defense after being fired at while using bulldozers to build a new embankment at an army base in the area.

(Barzilay & McKeown, 2005; Marsi & Krahmer, 2005; Filippova & Strube, 2008; Thadani & McKeown, 2013)

Graph Intersection

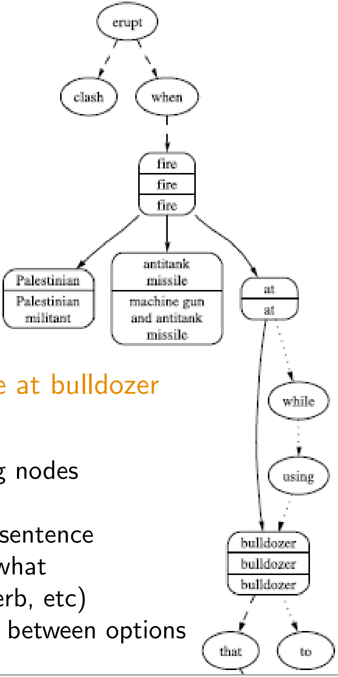

Palestian militants fired antitank missile at bulldozer

- (Barzilay & McKeown, 2005)
  - Merge Sentences by aligning nodes
  - Identify Intersection
  - Linearise graph to contruct sentence
  - Some hand coded rules on what cannot be cut (subject of verb, etc)
  - Use language model to pick between options

Extensions to this approach

- Marsi & Krahmer (2005) allow union as well as intersection
  1. Posttraumatic stress disorder (PTSD) is a psychological disorder which is classified as an anxiety disorder in the DSM-IV.
  2. Posttraumatic stress disorder (abbrev. PTSD) is a psychological disorder caused bya mental trauma (also called psychotrauma) that can develop after exposure to a terrifying event.

Intersection: Posttraumatic stress disorder (PTSD) is a psychological disorder.

Union: Posttraumatic stress disorder (PTSD) is a psychological disorder, which is classified as an anxiety disorder in the DSM-IV, caused by a mental trauma (also called psychotrauma) that can develop after exposure to a terrifying event.

Extensions to this approach

- (Filippova & Strube, 2008)
  - Include topic model for deciding which nodes to keep
  - Encode semantic constraints for union through coordination: Coordinated concepts have to be related, but not synonyms or hyponyms, etc.
- (Thadani & McKeown, 2013)
  - Supervised approach based on corpus of fused sentences

Computational Approaches to Summarization

- Bottom-Up
  What is in these texts? Give me the gist.
  - User needs: anything that is important
  - System needs: generic importance metrics
  - Techniques: Extractive summarization, sentence compression and fusion, etc.

Computational Approaches to Summarization

- Bottom-Up
  What is in these texts? Give me the gist.
  - User needs: anything that is important
  - System needs: generic importance metrics
  - Techniques: Extractive summarization, sentence compression and fusion, etc.
- Top-Down
  I know what I want – Find it for me.
  - User needs: only certain types of information
  - System needs: particular criteria of interest, used to focus search
  - Techniques: Information Extraction and Template-based generation

Top-Down Summaries

- Information Extraction (IE)
- Create Template for a particular type of story
  - Fields and values
  - Instantiate Fields from documents
  - Use Natural Language Generation to generate sentences from Template

IE Summarisation Strategy

- Instantiate Template by finding evidence – Pattern matching on text
  - Thousands of people are feared dead following a powerful earthquake that hit Afghanistan today. The quake registered 6.9 on the Richter scale.

Template for Natural Disasters

Disaster Type: earthquake

- location: Afghanistan
- magnitude: 6.9
- epicenter: a remote part of the country
- Damage:
  - human-effect:
    - number: Thousands of people
    - outcome: dead
    - confidence: medium
    - confidence-marker: feared
  - physical-effect:
    - object: entire villages
    - outcome: damaged
    - confidence: medium
    - confidence-marker: reports say

RIPTIDES (White et al., 2001)

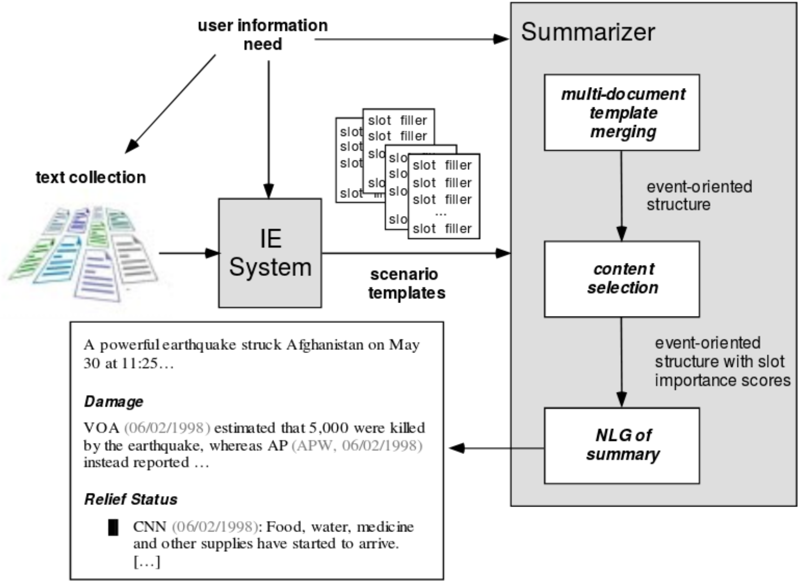

Problems with Template approach

- Templates are domain dependent
  - Manual effort in creating a template
  - Manual effort in designing a system that can generate sentences from a template
  - Cannot create a template for every possible news story this way
- Recent work attempts to learn such templates
  - Template Bank from historical texts (Schilder et al., 2013)

Templates, Generation and Reference

- Error correction for Multilingual Summarization
  - Extractive approaches are limited in how they can address noisy input (output of machine transation)
    - Replace sentences with similar ones from extraneous English Documents (Evans et al., 2004)
    - Improves Readability
    - Exact Matches hard to find, so can change meaning/emphasis
- Siddharthan & McKeown (2005); Siddharthan & Evans (2005):

  Apply a template approach to clean up referring expressions

References to People

- Distribution on premodifying Words
  - In initial references to people in DUC human summaries (monolingual task 2001-2004) Siddharthan et al. (2004)

71%	Role: *Prime Minister* or *Physicist* ‘ Time: *former* or *designate*
22%	Country, State, Location or Organization

- Our task is to:
  1. Collect all references to the person in different translations of each document in the set
  2. Identify above attributes, filtering any noise
  3. Generate a reference

Automatic semantic tagging

- organization, location, person name
  - BBN's IdentiFinder
- country, state
  - CIA factsheet: includes adjectival forms
    eg. *United Kingdom/U.K./British/Briton*
- role
  - WordNet hyponyms of *person*
  - 2371 entries including multiword expressions
    eg. *chancellor of the exchequer, brother in law* etc.
  - Sequences of roles are conflated
- temporal modifier
  - Also from WordNet, eg. *former, designate*

Example of Analysis

...<NP> <ROLE> representative </ROLE> of <COUNTRY> Iraq </COUNTRY> of the <ORG> United Nations </ORG> <PERSON> Nizar Hamdoon </PERSON> </NP> that <NP> thousands of people </NP> killed or wounded in <NP> the <TIME> next </TIME> few days four of the aerial bombardment of <COUNTRY> Iraq </COUNTRY> </NP>...

$$\begin{bmatrix} \mathit{name} & \text{Nizar Hamdoon} \\ \mathit{role} & \text{representative} \\ \mathit{country} & \text{Iraq } (arg1) \\ \mathit{organization} & \text{United Nations } (arg2) \end{bmatrix}$$

Identifying redundancy

- Coreference by comparing AVMs

$$\begin{bmatrix} \textit{name} & \text{Nizar Hamdoon}(2) \\ \textit{role} & \text{representative}(2) \\ \textit{country} & \text{Iraq}(2)\ (arg1) \\ \textit{organization} & \text{United Nations}(2)\ (arg2) \end{bmatrix}$$

- Numbers in brackets represent the counts of this value across all references
- The *arg* values now represent the most frequent ordering in the input references

Another Example

$$\begin{bmatrix} \textit{name} & \text{Zeroual(24), Liamine Zeroual(20)} \\ \textit{role} & \text{president(23), leader(2)} \\ \textit{country} & \text{Algeria(18)} \ (arg1) \\ \textit{organization} & \text{Renovation Party(2)} \ (arg1), \\ & \text{AFP(1)} \ (arg1) \\ \textit{time} & \text{former(1)} \end{bmatrix}$$

- Common issues:
  - Multiple roles and affiliations
  - Noise due to Errors from Tokenization, chunking, NE tools etc.

Removing Noise

1. Select the most frequent name with more than one word (this is the most likely full name).
2. Select the most frequent role.
3. Prune the AVM of values that occur with a frequency below an empirically determined threshold.

$$\begin{bmatrix} \mathit{name} & \text{Zeroual(24), Liamine Zeroual(20)} \\ \mathit{role} & \text{president(23), leader(2)} \\ \mathit{country} & \text{Algeria(18)}\ (arg1) \\ \mathit{organization} & \text{Renovation Party(2)}\ (arg1), \text{AFP(1)}\ (arg1) \\ \mathit{time} & \text{former(1)} \end{bmatrix}$$

Generating references

- Input Semantics:

$$\begin{bmatrix} \mathit{name} & \text{Nizar Hamdoon} \\ \mathit{role} & \text{representative} \\ \mathit{country} & \text{Iraq } (arg1) \\ \mathit{organization} & \text{United Nations } (arg2) \end{bmatrix}$$

$$\begin{bmatrix} \mathit{name} & \text{Liamine Zeroual} \\ \mathit{role} & \text{president} \\ \mathit{country} & \text{Algeria } (arg1) \end{bmatrix}$$

- To Generate,
    - Need knowledge of syntax
    - Determined by syntactic frames of role

Acquiring frames

- Acquire Frames for each role from semantic analysis of the Reuters News corpus

  ROLE=ambassador

(prob=.35)	COUNTRY ambassador PERSON
(.18)	ambassador PERSON
(.12)	COUNTRY ORG ambassador PERSON
(.12)	COUNTRY ambassador to COUNTRY PERSON
(.06)	ORG ambassador PERSON
(.06)	COUNTRY ambassador to LOCATION PERSON
(.06)	COUNTRY ambassador to ORG PERSON
(.03)	COUNTRY ambassador in LOCATION PERSON
(.03)	ambassador to COUNTRY PERSON

- Frames provide us with the required syntactic information
  - Word Order, Preposition Choice
- Use most probable frame that matches

Example

the representative of Iraq in the United Nations Nizar Hamdoon

`+1`

`representative of Iraq of the United Nations Nizar HAMDOON`

$\downarrow$

$$\begin{bmatrix} \textit{name} & \text{Nizar Hamdoon} \\ \textit{role} & \text{representative} \\ \textit{country} & \text{Iraq } (arg1) \\ \textit{organization} & \text{United Nations } (arg2) \end{bmatrix}$$

$\downarrow$

`Iraqi United Nations representative Nizar Hamdoon`

Automatic Evaluation

- Compared with Model References:
  - First References to same person in Human translation
- Data: DUC 2004 multilingual task
  - 24 sets
  - 6 used for development
  - 18 used for evaluation
- Baselines
  - Base1: most frequent initial reference to the person
  - Base2: randomly selected initial reference to the person

Results

1-GRAMS	P_{av}	R_{av}	F_{av}
Generated	0.847*@	0.786	0.799*@
Base1	0.753*	0.805	0.746*
Base2	0.681	0.767	0.688
2-GRAMS	P_{av}	R_{av}	F_{av}
Generated	0.684*@	0.591	0.615*
Base1	0.598*	0.612	0.562*
Base2	0.492	0.550	0.475
3-GRAMS	P_{av}	R_{av}	F_{av}
Generated	0.514*@	0.417	0.443*
Base1	0.424*	0.432	0.393*
Base2	0.338	0.359	0.315

@ Significantly better than Base1
* Significantly better than Base2
(unpaired t-test at 95% confidence)

Redundancy and Error Correction

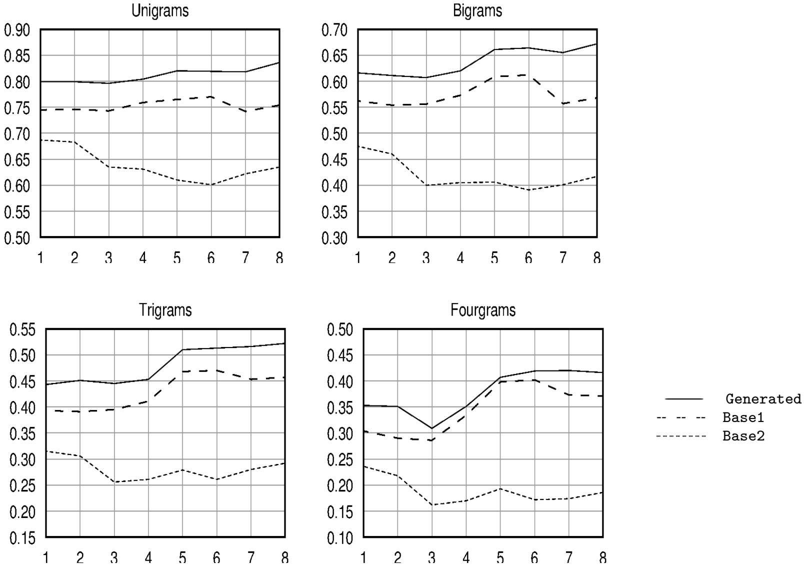

Back to Monolingual Multi-Doc Summarization

Nenkova et al. (2005); Siddharthan et al. (2011)

- Task definition
  - In the Document Understanding Conference context:
    - Input : Cluster of 10 news reports on same event(s)
    - Output: 100 Word (or 665 byte) Summary
  - Data compression of around 50:1

Scope for post-editing extractive summaries

- News Reports
  - Av. Sentence Length: 21.4 words
- Human Summaries
  - Av. Sentence Length: 17.4 words
- Machine Summaries
  - Av. Sentence Length: 28.8 words

Data source: Document Understanding Conference (DUC) 2001–2004

Sentence Compression

- Grefenstette (1998), Knight & Marcu (2000), Riezler et al. (2003)...

Former Democratic National Committee finance director Richard Sullivan faced more pointed questioning from Republicans during his second day on the witness stand in the Senate's fund-raising investigation.

- Richard Sullivan faced pointed questioning.
- Richard Sullivan faced pointed questioning from Republicans during day on stand in Senate fund-raising investigation.

But...

- Lin (2003) showed that statistical sentence-shortening approaches like Knight & Marcu (2000) do *not* improve content selection in summaries.
- Shortening approaches appear to remove the wrong words from a summary...
- Q: What are the right words to remove?

Syntactic Simplification

PAL, which has been unable to make payments on dlrs 2.1 billion in debt, was devastated by a pilots' strike in June and by the region's currency crisis, which reduced passenger numbers and inflated costs.

- PAL has been unable to make payments on dlrs 2.1 billion in debt
- PAL was devastated by a pilots' strike in June and by the region's currency crisis.
- The crisis reduced passenger numbers and inflated costs.

Does Syntactic Simplification help?

The Summary Genre

- News Reports
  - One appositive or relative clause every 3.9 sentences
- Human Summaries
  - One appositive or relative clause every 8.9 sentences
- Machine Summaries
  - One appositive or relative clause every 3.6 sentences

Data source: Document Understanding Conference (DUC) 2001–2004

Results (Siddharthan et al., 2004)

- Removing Parentheticals improves content selection

PAL, which has been unable to make payments on dlrs 2.1 billion in debt, was devastated by a pilots' strike in June and by the region's currency crisis, which reduced passenger numbers and inflated costs.

Shorter Sentences ⟶ Tighter Clusters:

1. PAL was devastated by a pilots' strike in June and by the region's currency crisis.
2. In June, PAL was embroiled in a crippling three-week pilots' strike.
3. The majority of PAL's pilots staged a devastating strike in June.
4. In June, PAL was embroiled in a crippling three-week pilots' strike.

Description and Content

- Do machine summarizers err on side of too much description?
- Removing Relative Clauses and Apposition from Input:
  - Siddharthan et al. (2004) and Conroy & Schlesinger (2004) report significant improvement.
- Removing Parentheticals improves content selection - Possibly at expense of Coherence
- Referring expressions require a formal treatment
  - inclusion of parentheticals just one aspect...

Referring Expressions in Summaries

A Machine Summary

Turkey has been trying to form a new government since a coalition government led by Yilmaz collapsed last month over allegations that he rigged the sale of a bank. Ecevit refused even to consult with Kutan during his efforts to form a government. Demirel consulted Turkey's party leaders immediately after Ecevit gave up.

- Familiarity?
- Minimal Description?

Multi-Doc Summarization

- In 100 words
  - Important events need to be summarized
  - Protagonists need to be described
- There is therefore a tradeoff
  - Too little description $\longrightarrow$ Incoherence
  - Too much description $\longrightarrow$ Compromised content
- What is the ideal level of description?
  - How much reference shortenning can we get away with?

Information Status

- Inferring Information Status for Referring Expression Generation
  a. Federal Reserve Chairman Alan Greenspan suggested that the Senate make the tax-cut permanent.
  b. Greenspan suggested that the Senate make the tax-cut permanent.
  c. The Federal Reserve Chairman suggested that the Senate make the tax-cut permanent.
  - Discourse new / Discourse old
  - Hearer new / Hearer old
  - Major / Minor
- In 100 word summaries, you don't want to waste space describing entities that are hearer old
- Or naming minor characters

Can information status be learnt?

The Experiment

- Assumptions
  - Writers of news reports have some idea of who the intended readership is familiar with
  - This is reflected in how they describe people in the story
  - Information status can be learnt
- Methodology
  - Label data with Information Status (this is the clever bit)
  - Perform lexical and syntactic analysis of references in news reports
  - Learn information status using features derived from above

Acquiring Labeled Data

- 120 document sets (10 news reports each) and manual summaries from DUC 2001–2004
- In manual summaries:
  - Hearer Old/New
    - Marked entities as hearer old if first mention was title+last name or only name.
    - Marked the rest as hearer new
  - Major/Minor Character
    - Marked entities as major if mentioned by name in at least one summary
    - Marked as minor if not mentioned by name in any summary
- 118 examples of hearer-old, 140 of hearer-new.
- 258 examples of major characters, 3926 of minor.

Syntactic Analysis

[IR] Nobel laureate Andrei D. Sakharov ; [CO] Sakharov ; [CO] Sakharov ; [CO] Sakharov ; [CO] Sakharov ; [PR] his ; [CO] Sakharov ; [PR] his ; [CO] Sakharov ; [RC] who acted as an unofficial Kremlin envoy to the troubled Transcaucasian region last month ; [PR] he ; [PR] He ; [CO] Sakharov ;

[IR] Andrei Sakharov ; [AP] , 68 , a Nobel Peace Prize winner and a human rights activist , ; [CO] Sakharov ; [IS] a physicist ; [PR] his ; [CO] Sakharov ;

Information collected for *Andrei Sakharov* from two news report.

IR = initial reference	CO = subsequent noun co-reference
PR = pronoun reference	AP = apposition
RC = relative clause	IS = Copula

Lexical Analysis

- Unigram and Bigram models of Premodifiers
  - Obtained from 2 months worth of news articles from the web
  - Independent of DUC data - from Newsblaster logs
- Formed list of 20 most frequent premodifying unigrams and bigrams
- Intuition:
  - Presidents more likely to be hearer old than judges...
  - Americans more likely to be hearer old than Turks...

Classification Results

- Major or Minor?

Algorithm	Accuracy
WEKA (J48)	0.96
Majority class prediction	0.94

- Familiarity (Hearer old or new?)

Algorithm	Accuracy
SVM (SMO Algorithm)	0.76
Majority class prediction (always hearer new)	0.54

The Generation Task

- Two aspects to deciding initial references:
  - What (if any) premodifiers to use
  - What (if any) postmodifiers to use
- Analysis of Premodifiers in DUC Human summaries
  - 72% words were:
    - Role or Title (eg.Prime Minister, Physicist or Dr)
    - Or reference modifying adjectives such as former that have to be included with the role.
  - DUC summarisers tended to follow journalistic convention and incude these words for everyone.
  - But for greater compression, the role or title can be omitted for hearer-old persons; eg. Margaret Thatcher instead of Former Prime Minister Margaret Thatcher.

The Generation Algorithm

1. IF Minor Character THEN:
    1. EXCLUDE name from reference and only INCLUDE role, temporal modification and affiliation
2. ELSE IF Major Character:
    1. INCLUDE name
    2. INCLUDE role and any temporal modifier, to follow journalistic conventions
    3. IF Hearer-old THEN:
        1. EXCLUDE other modifiers including affiliation
        2. EXCLUDE any post-modification such as apposition or relative clauses
    4. ELSE IF Hearer-new THEN:
        1. IF the person's affiliation has already been mentioned AND is the most salient organization in the discourse at the point where the reference needs to be generated THEN EXCLUDE affiliation ELSE

Predictive Power

- Successfully modelled variations in the initial references used by different human summarizers for the same document set

1. Brazilian President Fernando Henrique Cardoso was re-elected in the...
   [*hearer new* and Brazil not in context]

2. Brazil's economic woes dominated the political scene as President Cardoso...
   [*hearer new* and Brazil most salient country in context]

Predictive Power

- Successfully models variations in the initial reference to the same person across summaries of different document sets

1. It appeared that Iraq's President Saddam Hussein was determined to solve his countries financial problems and territorial ambitions...
[*hearer new* for this document set and Iraq not in context]

2. ...A United States aircraft battle group moved into the Arabian Sea. Saddam Hussein warned the Iraqi populace that United States might attack...
[*hearer old* for this document set]

Predictive Power

- For predicted hearer-old people, there was no postmodification in *any* gold standard summary.

Reference Accuracy

Generation Decision	Prediction Accuracy
Discourse-new references	
Include Name	.74 (rising to .92 when there is unanimity among human summarizers)
Include Role & temporal mods	.79
Include Affiliation	.79
Include Post-Modification	.72 (rising to 1.00 when there is unanimity among human summarizers)
Discourse-old references	
Include Only Surname	.70

Impact on Summaries

Preference of experimental participants for one summary-type over the other (140 comparisons):

	More informative	More coherent	More preferred
Extractive	46	22	37
Rewritten	23	79	69
No difference	71	39	34

Rewriting References:

- Shortened Summaries by 11 words on average
- Led to more coherent summaries (p¡0.01)
- Led to more preferred summaries (p¡0.01)
- Led to less informative summaries - but correlated with length of summary rho=0.8; $p<0.001$).

'

References

Angrosh, M., T. Nomoto, & A. Siddharthan. 2014. Lexico-syntactic text simplification and compression with typed dependencies. *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*.

Barzilay, R., & K. McKeown. 2005. Sentence fusion for multidocument news summarization. *Computational Linguistics* 31(3): 297–328.

Barzilay, R., K. R. McKeown, & M. Elhadad. 1999. Information fusion in the context of multi-document summarization. *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*.

Conroy, J. M., & J. D. Schlesinger. 2004. Left-Brain/Right-Brain Multi-Document Summarization. *4th Document Understanding Conference (DUC 2004) at HLT/NAACL 2004, Boston, MA*.

Dras, M. 1999. Tree adjoining grammar and the reluctant paraphrasing of text. Ph.D. thesis, Macquarie University NSW 2109 Australia.

Filippova, K., & M. Strube. 2008. Sentence fusion via dependency graph compression. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Grefenstette, G. 1998. Producing Intelligent Telegraphic Text Reduction to Provide an Audio Scanning Service for the Blind. *Intelligent Text Summarization, AAAI Spring Symposium Series*.

Jing, H., R. Barzilay, K. McKeown, & M. Elhadad. 1998. Summarization Evaluation Methods: Experiments and Analysis. *AAAI Symposium on Intelligent Summarization*.

Knight, K., & D. Marcu. 2000. Statistics-Based Summarization — Step One: Sentence Compression. *Proceeding of The American Association for Artificial Intelligence Conference (AAAI-2000)*.

Lin, C.-Y. 2003. Improving Summarization Performance by Sentence Compression - A Pilot Study. *In Proceedings of the Sixth International Workshop on Information Retrieval with Asian Languages (IRAL 2003)*.

Marsi, E., & E. Krahmer. 2005. Explorations in sentence fusion. *Proceedings of the European Workshop on Natural Language Generation*.

Nenkova, A., A. Siddharthan, & K. McKeown. 2005. Automatically learning cognitive status for multi-document summarization of newswire. *Proceedings of HLT/EMNLP 2005*.

,

Riezler, S., T. H. King, R. Crouch, & A. Zaenen. 2003. Statistical Sentence Condensation using Ambiguity Packing and Stochastic Disambiguation Methods for Lexical-Functional Grammar. *Proceedings of the Human Language Technology Conference and the 3rd Meeting of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL'03)*.

Schilder, F., B. Howald, & R. Kondadadi. 2013. Gennext: A consolidated domain adaptable nlg system. *Proceedings of the 14th European Workshop on Natural Language Generation*.

Siddharthan, A., A. Nenkova, & K. McKeown. 2004. Syntactic Simplification for Improving Content Selection in Multi-Document Summarization. *Proceedings of the 20th International Conference on Computational Linguistics (COLING 2004)*.

Siddharthan, A., & M. Angrosh. 2014. Hybrid Text Simplification using Synchronous Dependency Grammars with Hand-written and Automatically Harvested Rules. *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL'14)*.

Siddharthan, A., & D. Evans. 2005. Columbia University at MSE2005. *2005 Multilingual Summarization Evaluation Workshop, Ann Arbor, MI, June 29th 2005*.

Siddharthan, A., & K. McKeown. 2005. Improving Multilingual Summarization: Using Redundancy in the Input to Correct MT errors. *Proceedings of HLT/EMNLP 2005*.

Siddharthan, A., A. Nenkova, & K. McKeown. 2011. Information status distinctions and referring expressions: An empirical study of references to people in news summaries. *Computational Linguistics* 37(4): 811–842.

Thadani, K., & K. McKeown. 2013. Supervised sentence fusion with single-stage inference. *Proceedings of the Sixth International Joint Conference on Natural Language Processing*.

White, M., T. Korelsky, C. Cardie, V. Ng, D. Pierce, & K. Wagstaff. 2001. Multidocument summarization via information extraction. *Proceedings of the first international conference on Human language technology research*.