



**HAL**  
open science

## Document Sub-structure in Neural Machine Translation

Radina Dobрева, Jie Zhou, Rachel Bawden

► **To cite this version:**

Radina Dobрева, Jie Zhou, Rachel Bawden. Document Sub-structure in Neural Machine Translation. 12th Language Resources and Evaluation Conference, 2020, Marseille, France. pp.3657-3667. hal-02900568

**HAL Id: hal-02900568**

**<https://hal.science/hal-02900568>**

Submitted on 16 Jul 2020

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Document Sub-structure in Neural Machine Translation

Radina Dobрева<sup>1</sup>, Jie Zhou<sup>2</sup>, Rachel Bawden<sup>1</sup>

<sup>1</sup>School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh, United Kingdom

<sup>2</sup>Alibaba Group, 969 West Wen Yi Road, Hangzhou, China

rdobрева@ed.ac.uk, zj236040@alibaba-inc.com, rbawden@ed.ac.uk

## Abstract

Current approaches to machine translation (MT) either translate sentences in isolation, disregarding the context they appear in, or model context at the level of the full document, without a notion of any internal structure the document may have. In this work we consider the fact that documents are rarely homogeneous blocks of text, but rather consist of parts covering different topics. Some documents, such as biographies and encyclopedia entries, have highly predictable, regular structures in which sections are characterised by different topics. We draw inspiration from Louis and Webber (2014) who use this information to improve statistical MT and transfer their proposal into the framework of neural MT. We compare two different methods of including information about the topic of the section within which each sentence is found: one using side constraints and the other using a cache-based model. We create and release the data on which we run our experiments – parallel corpora for three language pairs (Chinese-English, French-English, Bulgarian-English) from Wikipedia biographies, which we extract automatically, preserving the boundaries of sections within the articles.

**Keywords:** machine translation, document structure, corpus creation, context, Wikipedia, parallel corpus

## 1 Introduction

Considerable progress has been made in machine translation (MT) thanks to the use of neural MT (NMT). While most NMT systems translate at the level of individual sentences, following similar practices in statistical MT (SMT), there has been significant interest in recent years in using document context to improve translation (Hardmeier, 2014; Bawden, 2018; Wang, 2019). However the intermediate level of the internal structure of documents, particularly for documents with regular sub-structure, could also provide useful information to improve MT.

Documents are rarely without internal structure, and certain document types (e.g. biographies, scientific articles and encyclopedia entries) are characterised by a more well-defined and regular structure than others. In such documents, the structure is explicitly defined by sections associated with headings dealing with different aspects of the main subject. These sections, many of which can be found across multiple documents, are often associated with specific vocabularies or even grammatical patterns. Figure 1 illustrates the classic structure of a Wikipedia article, with section headings (e.g. “early life”, “career”, “personal life”, etc.) that are likely to be found in other articles concerning high-profile people. This regularity in document sub-structure could be beneficial to the MT of such documents by providing additional information about the type of vocabulary used in different sections.

Exploiting document sub-structure in MT has previously been studied by Louis and Webber (2014) for SMT. They use topic models to integrate section information into a cache-based system and see improvements when applying their method to the translation of Wikipedia biographies for French→English. However MT techniques have since changed considerably, and the quality of MT has improved with the introduction of NMT. The change in paradigm provides us with new methods of integrating external information.

In this article, we draw inspiration from (Louis and Webber, 2014) to explore the effect of using predictable docu-

## Keanu Reeves

From Wikipedia, the free encyclopedia

*"Keanu" redirects here. For other uses, see Keanu (disambiguation).*

**Keanu Charles Reeves** (/kiˈɔːnuː/ *kee-AH-noo*<sup>[2][3][4]</sup> born September 2, 1964) is a Canadian<sup>[a]</sup> actor and musician. He gained fame for his starring roles in several blockbuster films, including comedies from the *Bill and Ted* franchise (1989–2020); action thrillers *Point Break* (1991), *Speed* (1994), the *John Wick* franchise (2014–present); psychological thriller *The Devil's Advocate* (1997); supernatural thriller *Constantine* (2005); and science fiction/action series *The Matrix* (1999–2003). He has also appeared in dramatic films such as *Dangerous Liaisons* (1988), *My Own Private Idaho* (1991), and *Little Buddha* (1993), as well as the romantic horror *Bram Stoker's Dracula* (1992).

Contents [hide]	
1	Early life
2	Career
2.1	Early career: 1980–1986
2.2	Breakthrough: 1986–1994
2.3	Rise of prominence in film: 1994–1999
2.4	Hollywood stardom and <i>The Matrix</i> trilogy: 1999–2009
2.5	Eclectic filmmaking and <i>John Wick</i> : 2009–present
2.6	Future projects
3	Personal life
3.1	Family and views
3.2	Legal incidents
3.3	Philanthropy and business
4	Filmography
5	Notes
6	References
7	Further reading
8	External links

Figure 1: An extract of a Wikipedia biography article, in which the document structure is summarised.

ment structure to improve NMT. As in their work, we focus on the domain of Wikipedia biographies, which are segmented into sections covering different aspects of the life of the person they are describing – their childhood, career, personal life, old age, etc. We automatically create datasets of parallel Wikipedia biographies in three language pairs - French-English (Fr-En), Bulgarian-English (Bg-En) and Chinese-English (Zh-En) - preserving document and document sub-structure information. We also collect monolingual Wikipedia biography corpora in the four languages.<sup>1</sup> To test the usefulness of exploiting document substructure for NMT, we conduct experiments to compare methods us-

<sup>1</sup>Available at <https://github.com/radidd/Doc-substructure-NMT> under a CC-BY-SA 3.0 licence.

ing topic information at the level of document sections. As in (Louis and Webber, 2014), we use topic modelling to model article sections. We compare two methods of integrating this information into NMT. The first uses side constraints (Sennrich et al., 2016a) and involves prepending topic information to source sentences. The second, adapted from (Kuang et al., 2018), uses caches containing relevant vocabulary. It is similar to the approach used by Louis and Webber (2014) but applied to NMT. We test these methods on three language directions (Fr→En, Zh→En, Bg→En). Our main contributions can be summarised as follows:

- The automatic creation of parallel and monolingual datasets of Wikipedia biographies with document sub-structure information for Fr-En, Bg-En and Zh-En
- Experiments comparing two methods of exploiting document sub-structure in NMT, applied to three language directions: Fr→En, Zh→En and Bg→En.

The remainder of the paper is organised as follows. Section 2 presents relevant work on modelling document structure and background on the approaches we compare. Section 3 describes the collection and processing of our datasets. In Section 4 we present two methods of integrating document structure information. Section 5 gives the experimental setup and results, and Section 6 presents analysis of these results. Finally, in Section 7 we provide a conclusion and ideas for further research.

## 2 Related Work

The exploitation of document sub-structure has previously been studied in the context of SMT by Louis and Webber (2014). The authors focus on the domain of Wikipedia biographies, presenting domain-adapted MT models. As they discuss, the biography domain is interesting because it consists of documents with a regular structure. Most biographies consist of sections that discuss topics such as their early life, career (which can fall into many categories based on the person), personal life and later life. To demonstrate the usefulness of document structure for MT, they compare models using document-wide context with models using context at the level of individual sections. They use two caches to pass information about a sentence’s topic and its preceding context to the MT model. In their structured model they load the topic cache with words that are specific to the given section, as opposed to the whole document. They also clear and reload the caches at section boundaries instead of at document boundaries. They show that the structured model has an advantage over the model that treats documents as a whole, with topic information found to be particularly useful.

While current work in NMT has not made use of document structure, there have been many efforts to supply NMT models with document-level context to improve translation (Voita et al., 2018; Maruf et al., 2019; Zhang et al., 2018; Miculicich et al., 2018). One way to integrate structural information into NMT is to adapt approaches to document-level NMT to consider section boundaries.

Of particular relevance is the work of Kuang et al. (2018), who supply contextual information using caches. Their

model involves two caches, a topic cache and a dynamic cache, which contain words that are important for the document-level context. The topic cache consists of words related to the document’s topic, while the dynamic cache is updated to contain the translations of previous sentences within the same document and the current sentence up to the current time step. Using both the topic and the dynamic cache provides the model with document-wide information ensuring consistency, as well as information from the preceding context which allows for better coherence.

The topic information used in (Louis and Webber, 2014) and (Kuang et al., 2018) is supplied using topic models that identify the most important words for any given topic. The method they use is Latent Dirichlet Allocation (LDA) (Blei et al., 2003). It is important to note that topic modelling can be done on entire documents as in (Kuang et al., 2018) or on sections within documents (as in this paper). In both cases, the text is treated as a “bag of words” and the important information comes from the frequency of the words in the segments. That is, topic modelling done on sections within documents will learn topics that are relevant to the text in the sections. Furthermore, while topic models treat documents as a mixture of topics with different probabilities, for downstream tasks it can be useful to only consider one or several of the most probable topics. For instance, Louis and Webber (2014) only consider the topic with the highest probability for each section of a document.

## 3 Datasets

We follow Louis and Webber (2014) in using Wikipedia biographies to test whether section information can be useful for NMT. We agree that data from Wikipedia is particularly illustrative of regularly structured documents as it contains separate sections with section headings. It is reasonable to assume that some sections within different biographies might share similar vocabulary, while other sections in the same documents may be very different. For instance, a biography of a composer and one of a basketball player may both contain information about their early years, but the rest of the documents may be different. Therefore, providing more information about the recurring topic may be useful, as opposed to providing only document-level information which may be too general.

We use a similar method to Louis and Webber (2014) for data creation. However, we extend the process to cover three language pairs (Fr-En, Zh-En, Bg-En), and considerably more data,<sup>2</sup> which is necessary as NMT models require more data to be trained effectively.

Data was collected and processed using the following steps, which are described in more detail below:

1. Extraction of biographies from Wikipedia dumps and separation into parallel and monolingual data.
2. Text extraction, sentence-splitting but with preservation of document and section structure.
3. Sentence-level alignment of parallel data.

---

<sup>2</sup>Louis and Webber (2014) use a dataset of 1000 monolingual articles for each language and parallel development and test sets of 15 and 30 articles respectively.

	Zh-En		Fr-En		Bg-En	
	training	validation	training	validation	training	validation
Total #documents	693	75	23,648	274	1,899	149
Total #sentences	26,505	2,020	179,270	2,002	29,348	2,004
Average #sentences/section (source)	7.54	6.14	3.33	3.08	5.35	5.09
Average #sentences/section (target)	7.22	5.87	3.29	3.10	4.14	4.00
Average #sections/document (source)	5.07	4.38	2.28	2.37	2.89	2.64
Average #sections/document (target)	5.29	4.59	2.30	2.35	3.74	3.36

Table 1: Parallel training and validation data statistics. Source refers to one of {Zh, Bg, Fr} and target refers to En.

Lang. pair	Zh-En		Fr-En		Bg-En	
	Original lang.	En	Zh	En	Fr	En
Total #docs	22	30	78	82	154	25
Total #sents	1147	875	1130	1198	2539	273
Ave. #sents/sec	9.89	6.58	4.56	4.83	5.37	3.41
Ave. #secs/doc	5.27	4.43	3.18	3.02	3.06	3.2

Table 2: Parallel test set statistics. We split the set in two, depending on which language was the original text’s language (as opposed to the translation).

	Zh	Fr	Bg	En
Ave. #sents/sec	7.05	5.55	5.99	7.39
Ave. #secs/doc	3.29	2.82	3.06	3.97
Total #docs	68,433	167,484	56,275	99,106
Total #sents	1,586,194	2,619,842	1,029,626	2,904,641

Table 3: Monolingual corpus statistics.

- Cleaning of aligned parallel articles to avoid very long sentences or mismatches between source and target length using Moses scripts (Koehn et al., 2007).
- Division of data into training, validation and test sets

**Extraction of biographies** Relevant Wikipedia articles were obtained by filtering Wikimedia dumps using their metadata. We extracted only biographies using the category of the articles with keywords related to people (e.g. “person”, “writer”, “politician”) and using the presence of a “Biography” section as indicators. Parallel biography data was obtained by selecting articles that were indicated in the metadata as translations (where either language was the original). Monolingual data contains all biography data for a given language, but does not include the parallel test sets.

**Text extraction and sentence splitting** We extracted the text from the articles, splitting it into sentences and preserving information about the document and section each sentence belongs to, as well as sentence order in the text.

**Alignment of parallel data** Sentence-level alignment is not trivial, especially since Wikipedia is open for anyone to edit and articles can change significantly after being translated. This is illustrated in Figure 2, which shows some examples of non-exact sentence alignment due to non-exact translation or post-editions. Some degree of such noise is therefore expected in the data. We attempt to reduce it as much as possible by filtering based on automatic alignment scores. For Fr-En, we used Hunalign (Varga et al., 2005), while for Zh-En and Bg-En we used previously trained MT

models to translate the source text into English and align the sentences based on similarity for Zh-En and BLEU score for Bg-En using Bleualign (Sennrich and Volk, 2011). Future versions of the corpus can explore how to further reduce the level of noise left in the data.

**Dataset partitions** Parallel data is divided into training, validation and test sets. For each language pair, we created two test sets, based on the original language the articles were written in (e.g. for Bg-En, one test set contains articles originally written in English and translated into Bulgarian, and the other contains articles originally written in Bulgarian and translated into English).<sup>3</sup> Statistics for the training and validation sets are given in Table 1 and for the test sets in Table 2. Monolingual data statistics are given in Table 3. Parallel training data sizes range from 29,348 sentences for Bg-En to 179,270 sentences for Fr-En. More monolingual data is available, from 1,029,626 sentences for Bulgarian to 2,904,641 sentences for English.

## 4 Integrating Document Sub-structure

We compare two methods of integrating document structure information, both relying on the assumption that different sections of the text cover different topics. Here, we consider different sections to be those that are delimited by headings (whether these are section, subsection or subsection headings), which creates a flat hierarchy of delimited areas of text. In both cases, we follow Louis and Webber (2014) in using topic modelling to learn section topic representations, which we describe in Section 4.1. The methods differ in how this learnt topic information is integrated. The first, described in Section 4.2.1, uses side constraints (Sennrich et al., 2016a) to incorporate information about the section topic associated with each sentence. For the second method we use Kuang et al. (2018)’s cache-based neural model to provide the model with topic information and previous context within the boundaries of the current section. While our first method is simple to implement and train, the second provides more fine-grained information about each section within an article.

### 4.1 Modelling Section Topics

While Wikipedia section headings are a useful indication of the section boundaries, they are not necessarily optimal for determining the granularity of different topics. As can be

<sup>3</sup>It is interesting to separate out these two translation directions, as the translation direction can have an effect on ease of translation due to the “translationese” effect (Zhang and Toral, 2019; Graham et al., 2019)

Source	Source gloss	Reference
Ils se marient le 30 juin 1894 à Paris.	‘They married on 30 June 1894 in Paris.’	They married on 30 June 1894 in Paris <b>and had two daughters.</b>
Le 11 décembre 2013, <b>elle</b> est nommée déléguée générale du Québec à New York.	‘On 11 December 2013, <b>she</b> was appointed Quebec Delegate General in New York’	On 11 December 2013 <b>Poirier</b> was appointed Quebec Delegate General in New York.
毛贻昌选中了罗一秀。 1908年, 罗一秀与毛泽东举办了婚礼	‘ <b>Mao Yichang</b> selected Luo Yixiu.’ ‘In 1908, <b>Luo Yixiu and Mao Zedong</b> held a wedding.’	<b>He</b> selected Luo Yixiu <b>in either late 1907 or 1908.</b> The wedding took place in 1908.
Първият му гол в Серия А е на 16 септември 2012 г. срещу Парма.	‘His first goal in Serie A is on 16 September 2012 against Parma.’	On 16 September 2012, he scored his first goal in Serie A, <b>after entering as a substitute for Edinson Cavani in a 3-1 home win over Parma.</b>
Два пъти, <b>през 1962 и 1967 година</b> , получава званието “Герой на социалистическия труд”.	‘Twice, <b>in 1962 and 1967</b> , he received the title “Hero of socialist labour”.’	<b>Georgiev</b> was twice awarded the title Hero of Socialist Labour.

Figure 2: Examples of non-exact or paraphrased references in the biography data, most likely due to translation divergences and/or subsequent modification. Differences between source and target sentences are indicated in bold.

seen in Figure 1, some headings are too general, (e.g. “Career” can refer to any occupation and the contents of this section can vary immensely between a sports person and a politician), whereas other headings are too specific to particular people and do not allow for any useful generalisation (e.g. “Eclectic filmmaking and John Wick: 2009-present”). We therefore choose not to use the headings themselves to determine our topics, and instead use topic modelling to induce topics from article sections.

**Training topic models** As in (Louis and Webber, 2014), we use Latent Dirichlet Allocation (LDA) (Blei et al., 2003) to train topic models on our monolingual biography data, taking into account section boundaries. The resulting topic models are then used to obtain a distribution over topics for each section in the parallel data, and the topic of the section is selected based on the topic with the highest probability. Each topic is associated with its own distribution over the language’s vocabulary.

LDA places certain assumptions on the generation of topic probabilities for each document, as well as on the generation of words given a topic. These assumptions can be regulated with the two hyperparameters of LDA, alpha and beta, which control the distribution of topics in documents and of words in documents. Setting them to very small values encourages sparsity in the model (i.e. it encourages the model to assign higher probabilities to only a few topics for each document and only some words in each topic). This is useful for MT as the topic models need to be able to differentiate between sections of different topics and provide the most salient words for the particular topic. Another important hyperparameter of LDA is the number of topics to be produced by the model. It is chosen empirically and depends on the sections to be modelled.

**Aligning language-specific topics** Since topic models for the two languages in the parallel data are trained separately, there is no direct correspondence between topics on the two sides. In practice, when translating new documents, only the text and topic on the source side is available and yet in our proposed cached model (cf. Section 4.2.2), words from the target side topics are used. It is therefore necessary to be able to predict a target topic using the topic of the source

language. Topics in the two languages are aligned by taking the co-occurrence counts of each French/Bulgarian/Chinese topic with each English topic in the parallel corpus. The topics that co-occur the most frequently are considered to be aligned. Finally, the English topic models are also used to obtain the most probable words for each topic.

## 4.2 Incorporating Topic Information

### 4.2.1 Side Constraints

A simple yet effective way of integrating information at the sentence level is to use side constraints (Sennrich et al., 2016a), which consists in prepending (or appending) an extra token to the input sentence associated with the feature to which you wish to bias the translation. These feature tokens are added to the vocabulary and treated as normal tokens. It has previously been used for a range of phenomena including politeness (in the original article) and domain adaptation (Kobus et al., 2017; Caswell et al., 2019).

We therefore prepend a special token to the beginning of the source sequence to represent the topic of the current source section (cf. Figure 3). Our intuition is that the model will learn to associate the presence of each topic with related target vocabulary and in a way that provides more adapted information than topics trained at the document level.

---

<topic64> Elle a débuté en août 2013 avec “My Student Teacher”.  
 <topic91> Le 21 mars 2014, NC.A sort son single digital “Hello Baby”.  
 <topic91> Elle sort ensuite son premier mini-album le 9 avril.

---

Figure 3: An example of topics used as side constraints.

### 4.2.2 Cache-based Neural Model

Our second method draws inspiration from Louis and Webber (2014) in their use of caches to represent section information. We adapt the model proposed by Kuang et al. (2018) to the Transformer (Vaswani et al., 2017), introducing some small changes. Kuang et al. (2018)’s model uses two caches containing words relevant to the topic of the document and the preceding context respectively (termed “topic cache” and “dynamic cache”). In our version of the model, the topic refers to the learned topic of the section rather than the document.




Figure 4: An example of how the topic cache and dynamic contents are determined.

As illustrated in Figure 4, the topic cache is loaded with words from the current topic, and the dynamic cache conveys information about the recent context preceding the current sentence. For a given sentence, the topic cache is loaded with the most probable words from the topic’s distribution as learnt using the topic model, and the dynamic cache contains the set of unique content words<sup>4</sup> from previously translated sentences from the same section. The dynamic cache is updated as translation progresses to include the content words from the most recent preceding sentences. The cache is of a fixed size and if the length of the cache is exceeded, older context is removed to allow for newer context.

The caches are concatenated and passed to a neural cache model, which computes a probability distribution over the words. At timestep  $t$ , scores over cache words  $y_c$  are computed based on the current decoder state  $h_t$ , the previous output from the NMT model  $y_{<t}$  and the encoder context  $c_e$  using a feedforward network  $f_{cache}$ :

$$score(y_c|h_t, c_e, y_{<t}) = f_{cache}(h_t, c_e, y_{<t}, y_c) \quad (1)$$

The distribution over cache words is obtained by applying the softmax function:

$$p_{cache} = softmax(score(y_c|h_t, c_e, y_{<t})) \quad (2)$$

At each decoding step, this distribution is combined with  $p_{NMT}$ , the NMT model’s distribution, using linear interpolation to obtain a final distribution over target words  $y_t$ :

$$g = \sigma(f_{gate}(h_t, c_e, y_{<t})) \quad (3)$$

$$p(y_t|y_{<t}, c) = g \cdot p_{NMT} + (1 - g) \cdot p_{cache} \quad (4)$$

To adapt this method to the Transformer, we use the output of self-attention as a representation of the previously output

<sup>4</sup>A list of grammatical (stopwords) was used to obtain only content words. We choose however to include pronouns and certain auxiliaries (e.g. was, were) that could be useful indicators of tense choices (e.g. for past events).

words, the encoder-decoder attention to represent the source side context and the final feed-forward network as the current decoder state (see Figure 5).

One difference between the present implementation and the cache model of Kuang et al. (2018) should be noted. While they update the dynamic cache at every timestep, providing the previously translated words from the same sentence, here the dynamic cache only contains words from preceding sentences. This is because the Transformer model has better access to output at preceding timesteps through self-attention. Furthermore, as the cache size is fixed, adding words from the current translation would reduce the length of preceding context that is available.

## 5 Experiments

We apply the two methods described in Section 4 to the translation of three languages (Fr, Zh and Bg) into English. We train and test the models using the Wikipedia biography data previously described in Section 3.

### 5.1 Topic Models

We use the monolingual datasets to train LDA topic models for the four languages (Fr, Zh, Bg, En). For each language we train two topic models: one that learns topics of sections within documents, and one that learns topics of the whole documents (for comparison purposes). The hyperparameters alpha and beta of LDA are set to 0.001 and 0.01 to encourage sparsity, following (Louis and Webber, 2014). We set the number of topics to 100 for all models.

### 5.2 MT Model Training

**Pre-training and Fine-tuning** We train an out-of-domain Transformer-base model (Vaswani et al., 2017) for each language pair and use it as a basis to fine-tune all our models using our in-domain biography data. This pre-training ensures that all models are strong, having been trained on large quantities of data in addition to in-domain data. We also train an in-domain baseline by fine-tuning the pretrained model on our in-domain data, without additional information about document structure.

**Data and Preprocessing** For the out-of-domain pretrained model, we use data from WMT (Barrault et al., 2019) for Zh→En (~24.2m parallel sentences) and Fr→En (~39m parallel sentences), and data from Opus (Tiedemann, 2012) for Bg→En (~44m parallel sentences). For all in-domain data we use the data described in Section 3. The data is tokenised and BPE segmented (Sennrich et al., 2016b). For cache models, words in the trained topics are also segmented using BPE. For the side constraints models we prepend a token representing the topic of each section.

**Architecture and Settings** The parameters of the Transformer are set to standard values (Vaswani et al., 2017): the encoder and decoder have 6 stacked layers, the embedding size is 512 and the feed-forward network hidden layer dimensionality is 2048. We use the Nematus toolkit (Sennrich et al., 2017) for all models.

**Cache model parameters** For cache-based models, we fix the size of the two caches to 100 words each. The scoring feed-forward network has hidden dimensions 1000 and

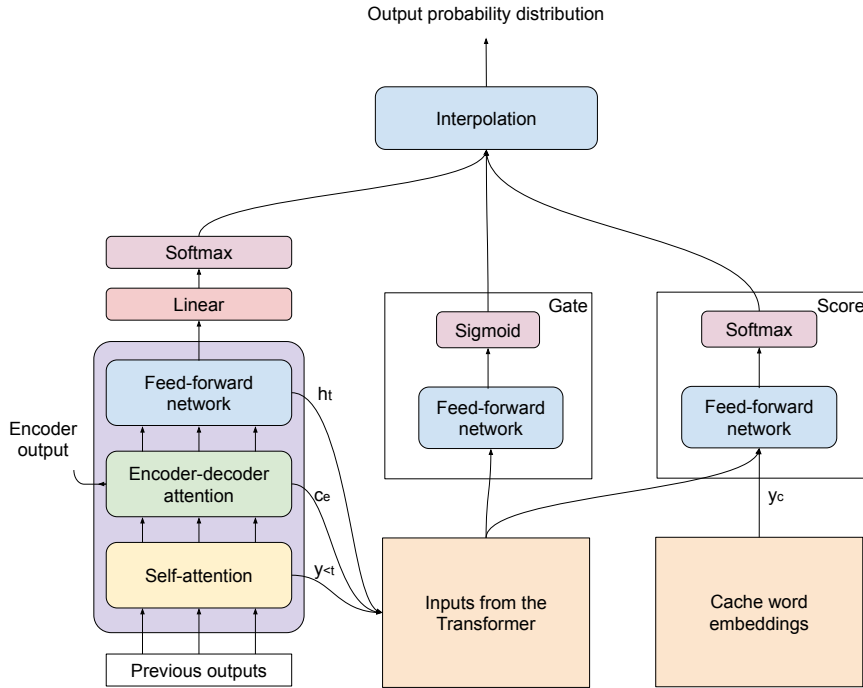


Figure 5: The neural cache model integrated with the Transformer decoder.

	En orig.	Zh-En Zh orig.	all	En orig.	Fr-En Fr orig.	all	En orig.	Bg-En Bg orig.	all
Out-of-domain baseline	14.3	<b>11.6</b>	13.1	45.3	46.1	45.7	22.0	21.9	22.0
In-domain baseline	16.9	10.8	14.0	<b>52.4</b>	<b>51.8</b>	<b>52.1</b>	24.4	22.0	24.2
Document-level side constraints	17.1	10.8	14.1	52.0*	<b>51.8</b>	51.9	<b>24.7*</b>	23.1	24.5**
Section-level side constraints	<b>17.3**</b>	11.2*	<b>14.4***</b>	52.0	50.6***	51.3***	<b>24.7**</b>	<b>23.3</b>	<b>24.6***</b>
Document-level cache-based	16.6	10.9	13.9	52.0	51.6	51.8	24.2	22.6	24.0
Section-level cache-based	16.8	10.7	13.9	52.0	51.0**	51.5**	24.3	22.5	24.1

Table 4: Bleu score results. We distinguish between the two translation directions, depending on which language was the language of the original text (“orig.”), and also calculate the Bleu score on the concatenation of these two test sets (“all”). The highest score for each set is marked in bold and improvements that are statistically significant from those of the in-domain baseline are indicated with asterisks (\* for  $p \leq 0.05$ , \*\* for  $p \leq 0.01$  and \*\*\* $p \leq 0.001$ ).

500 and the gate feed-forward network has hidden dimensions 500 and 200, following the configurations reported in (Kuang et al., 2018). The cache word embeddings are shared with the Transformer decoder. During training we provide the real topic of the target sentence for half of the training data, and for the other half we provide the topic projected from the source, in order for the model to learn from but not be over-reliant on gold (source) topics. During training we also use the real target side sentences to load the dynamic cache. At inference time, the topic is a projection from source to target and the dynamic cache is loaded with words from previously translated sentences.

**Comparative Systems** In order to assess the usefulness of document structure, we provide contrastive experiments whereby both methods are applied to the full Wikipedia articles, disregarding section boundaries, the difference being that topic models are trained on whole documents rather than individual document sections. We will refer to the models using section boundaries as “section-level” and to the ones that do not as “document-level”.

### 5.3 Results

We report Bleu scores (Papineni et al., 2002), calculated using SacreBleu (Post, 2018) on our tests sets for each of our models in Table 4. We also provide scores for the pre-trained model (out-of-domain baseline) and the in-domain fine-tuned baseline. Improved results that are statistically significant<sup>5</sup> from the in-domain baseline results are indicated (\* for  $p \leq 0.05$ , \*\* for  $p \leq 0.01$  and \*\*\* $p \leq 0.001$ ).

**Effect of fine-tuning** Fine-tuning the out-of-domain baseline gives improved results for all language pairs, with the greatest gain being seen for Fr-En with a +6.4 increase in Bleu. The smallest difference is seen for Zh-En, where fine-tuning actually degrades performance on the side of the test set that was originally in Chinese. Even for Fr and Bg, smaller gains are seen on this side of the test set.

<sup>5</sup>Calculated using bootstrap resampling using the Moses script `bootstrap-hypothesis-difference-significance.pl` on outputs and references tokenised using the WMT standard ‘13a’ tokenisation as used in SacreBleu.

**Translation direction of test sets** In fact, all language pairs and models show different results depending on the original language of the test sets - the scores are consistently higher for the En-originating test sets, than for the ones originating in Fr, Bg or Zh. This effect is particularly striking for Zh-En, where the difference reaches over 6 Bleu points for some models. One explanation for this is a difference in writing style. Text translated from another language has different properties from text originally written in the language. Since the majority of in-domain training data is text that was originally written in English, it is not surprising that the models perform worse on the test sets originating in Fr/Bg/Zh.

**Language-specific differences** Bleu scores are very different depending on the language pair. Zh-En scores appear very low (from 10.8 to 17.3 Bleu), whereas Fr-En is characterised by very high scores (from 45.3 to 52.4 BLEU). Upon closer inspection, it appears that one reason for the low scores achieved by the Zh-En models is that the references for the test sets are not always literal translations of the source, but contain many paraphrases and inexact translations (see Table 2 for examples). The very high Bleu scores for Fr-En can be explained partly by the presence of fewer deviations in formulation, possibly due to the structural similarity of the two languages. The ease of this set is also aided by the presence of a number of short sentences making up lists of works (e.g. books, musical compositions, etc.), which are relatively simple to translate.

**Side-constraints vs. cache model** The comparative results of the two methods differ by language pair. For Bg-En, the side-constraints method shows a decisive advantage over the cache model method, with significantly higher Bleu scores ( $p \leq 0.05$ ).<sup>6</sup> This is also observed for Zh-En, although to a lesser degree, and is only significant at the section level ( $p \leq 0.05$ ). The Fr-En results do not show any consistent pattern in the differences between the side-constraints and cache models. Overall it appears that the cache-based model does not offer a systematic advantage over the simpler side-constraints model, despite finer grained information being provided to the model.

**Document- and section-level models** According to Bleu, there is no systematic pattern between document- and section-level across all language pairs, and the scores depend on the method used. Whereas little difference is seen between the document- and section-level experiments for Zh-En and Bg-En, section-level information actually appears to degrade performance when compared to document-level information for Fr-En, and this for both the side-constraints and cache-based methods. This effect is statistically significant for the Fr-original test set ( $p \leq 0.05$ ).

## 6 Analysis

To gain more insight into the differences between the proposed models, we briefly analyse the cache model topics (Section 6.1), and in Section 6.2 provide a manual evaluation and some qualitative analysis of translations.

<sup>6</sup>With the exception of the Bg-orig. test set for which there is no significant difference between the two methods at the document level, which could be due to the fact that this test set is very small.

### 6.1 Analysis of Topic Models

Table 6 shows the top ten most relevant words for selected topics. These lists show clearly that the topics are coherent, containing words relevant to the same subject matter. For instance, Topic 0 from the English topic model relates to sections concerning members of musical bands, and more specifically discussing the band’s musical career. This is also true of the non-English topic models: Chinese Topic 51 is about higher education and academic career.

### 6.2 Qualitative Analysis of Translations

While BLEU scores provide a general impression of the performance of MT models, they do not give any insight into the specific strengths and weaknesses of the models. Therefore, we manually evaluate a sample of the outputs. We compare the two best performing models (document-level and section-level side constraints), randomly sampling sections from each of the test sets, and keeping the first 100 sentences from each as sets for manual evaluation.

Lang. pair	Orig.	Better	Worse	Equal	Identical
Fr-En	En	15	14	22	49
	Fr	13	16	33	38
Bg-En	En	22	17	23	38
	Bg	24	24	32	20
Zh-En	En	22	9	34	35
	Zh	9	11	65	15

Table 5: Manual evaluation of 100 sentences per test set. Comparisons are classified in terms of the number of times the section-level translation is better, worse, equal or identical to the document-level output.

The results in Table 5 show that across all test sets, many sentences are translated identically by the two models. There is a further substantial number of sentences for which the two models achieve similar quality. Among the sentences that show a difference in quality between the two models, preference for either the section- or the document-level model depends on the original language of the set. Section-level models do better across the En-originating test sets, albeit to differing degrees: the effect is quite strong for Zh-En (22 better vs. 9 worse), but less strong for Bg-En (22 vs. 17) and negligible for Fr-En (15 vs. 14). For sets originating in Zh/Fr/Bg, section-level models perform similarly or worse than document-level models. This appears to show that section information is useful compared to document-level information, but only when translating from translationese (a pattern also seen in the Bleu scores). This effect also diminishes as the MT quality increases (as is the case with En-Fr, where the Bleu scores are particularly high). These results differ from those in (Louis and Weber, 2014), where they do see improvements. However they do not test on the two translation directions (as the translationese debate in MT has only emerged in recent years). Moreover, it is possible that topic information is less useful for NMT models in high resource settings (as is the case for all three directions tested), as the quality of the models is already very good.

Figure 7 shows examples of sentences which are translated



#0	English #32	#63	French #11	Bulgarian #3	Chinese #51
band	film	married	il 'he/it'	е 'is'	年 'year'
their	films	was	Guerre 'war'	й 'her'	大学 'the University'
was	director	died	Il 'He/It'	тя 'she'	教授 'professor'
album	directed	had	mondiale 'world'	деца 'children'	毕业 'graduation'
released	was	his	guerre 'War'	има 'has'	学院 'College'
they	producer	children	1940 '1940'	майка 'mother'	研究 'the study'
tour	his	wife	1944 '1944'	баща 'father'	中国 'China'
were	Film	son	puis 'then'	години 'years'	担任 'Serve'
first	He	daughter	Seconde 'Second'	Тя 'She'	研究所 'graduate School'
They	produced	They	1945 '1945'	дъщеря 'daughter'	获 'Gain'

Figure 6: The top ten words in a selection of section-level topics.

Source	Reference	Document-aware model	Section-aware model
<b>Soliste</b> internationale, Marielle Nordmann...	An international <b>soloist</b> , Marielle Nordmann...	An international classical <b>violist</b> , Marielle Nordmann...	An international <b>solo artist</b> , Marielle Nordmann...
1919年，参加了五四运动，加入少年中国学会。	In 1919, <b>he participated in the May Fourth Movement, China Youth Association.</b>	In 1919, <b>he took part in the May 4th Movement and joined the Young China Institute.</b>	<b>He joined the Boys' China Society</b> in 1919.
Два дни по-късно <b>получава първото генералско звание - генерал-майор</b> и е назначен за командващ на 2-ра армия.	On 14 September 1944 he <b>was promoted to Major General</b> and was given the command of the Bulgarian Second Army.	He was appointed commander of the 2nd Army two days later.	Two days later, he <b>received the first rank of major-general</b> and was appointed commander of the 2nd Army.

Figure 7: Illustration of the types of differences between document-aware and section-aware side constraint model outputs.

differently by the two side-constraints models (more examples can be found in Appendix A). The main differences we observe for both models are in lexical choice, under-translations of some sentences, reformulations and differences in punctuation. There does not appear to be a clear pattern in terms of the improvements or degradations seen by each model, and it is likely that the topic information is providing some domain adaptation effect, which is difficult to observe. A possible exception to this could be the first example in Table 7, where *solist* 'soloist' is translated as *violist* (i.e. viola player), despite the woman described being a harp player. This sentence's document topic is music-related, and the model may have overfit to the topic.

## 7 Conclusion and Future Work

In this paper we propose two methods to transfer into the framework of NMT Louis and Webber (2014)'s idea of exploiting document structure when translating documents with regular and predictable structure. As in their work, we use topic modelling to model document sections, under the assumption that different sections within articles display different lexical properties, and compare the integration of this information using side constraints (Sennrich et al., 2016a) to a more complex approach using cache mechanisms, adapted from (Kuang et al., 2018).

We have created three parallel corpora of Wikipedia biographies (for En-Fr, En-Bg and En-Zh), structured into sections, as well as monolingual corpora for all four languages, all of which will be made freely available.

Our experiments using this data show that there are no consistent gains to be seen across all language directions for a particular model type, and using section-level information

as opposed to document-level information does not systematically improve MT quality. These results, which are different from results found by (Louis and Webber, 2014) for SMT, suggest that while there are circumstances in which providing section (or document) topic information does help NMT through domain adaptation (shown by gains for certain subsets of the test set), in high-resource scenarios such as the ones tested here, this information is not systematically useful. A direction to look into in the future is how this type of information could help low-resource MT, as a way of injecting monolingual topic information in a more light-weight and efficient manner than most current techniques (Lample and Conneau, 2019; Lample et al., 2018).

## 8 Acknowledgments

We would like to thank Bonnie Webber, Annie Louis and Christian Hardmeier for their invaluable help and feedback on this project. This work was supported by funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825299 (GoURMET). It was also supported by the UK Engineering and Physical Sciences Research Council (EPSRC) fellowship grant EP/S001271/1 (MTStretch).

## 9 Bibliographical References

Barrault, L., Bojar, O., Costa-jussà, M. R., Federmann, C., Fishel, M., Graham, Y., Haddow, B., Huck, M., Koehn, P., Malmasi, S., Monz, C., Müller, M., Pal, S., Post, M., and Zampieri, M. (2019). Findings of the 2019 conference on machine translation (WMT19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy.

- Bawden, R. (2018). *Going beyond the sentence: Contextual Machine Translation of Dialogue*. Ph.D. thesis, LIMSI, CNRS, Université Paris-Sud, Université Paris-Saclay, Orsay, France.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- Caswell, I., Chelba, C., and Grangier, D. (2019). Tagged back-translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, Florence, Italy.
- Graham, Y., Haddow, B., and Koehn, P. (2019). Translationese in machine translation evaluation. *CoRR*, abs/1906.09833.
- Hardmeier, C. (2014). *Discourse in Statistical Machine Translation*. PhD Thesis, Uppsala University, Department of Linguistics and Philology, Uppsala, Sweden.
- Kobus, C., Crego, J., and Senellart, J. (2017). Domain control for neural machine translation. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 372–378, Varna, Bulgaria.
- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., and Herbst, E. (2007). Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic.
- Kuang, S., Xiong, D., Luo, W., and Zhou, G. (2018). Modeling coherence for neural machine translation with dynamic and topic caches. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 596–606, Santa Fe, New Mexico, USA.
- Lample, G. and Conneau, A. (2019). Cross-lingual language model pretraining. *CoRR*, abs/1901.07291.
- Lample, G., Conneau, A., Denoyer, L., and Ranzato, M. (2018). Unsupervised machine translation using monolingual corpora only. In *International Conference on Learning Representations (ICLR)*.
- Louis, A. and Webber, B. (2014). Structured and unstructured cache models for SMT domain adaptation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 155–163, Gothenburg, Sweden.
- Maruf, S., Martins, A. F. T., and Haffari, G. (2019). Selective attention for context-aware neural machine translation. *CoRR*, abs/1903.08788.
- Miculicich, L., Ram, D., Pappas, N., and Henderson, J. (2018). Document-level neural machine translation with hierarchical attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, pages 311–318, Stroudsburg, PA, USA.
- Post, M. (2018). A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Sennrich, R. and Volk, M. (2011). Iterative, MT-based sentence alignment of parallel texts. In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011)*, pages 175–182, Riga, Latvia.
- Sennrich, R., Haddow, B., and Birch, A. (2016a). Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40, San Diego, California.
- Sennrich, R., Haddow, B., and Birch, A. (2016b). Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Sennrich, R., Firat, O., Cho, K., Birch, A., Haddow, B., Hirschler, J., Junczys-Dowmunt, M., Läubli, S., Miceli Barone, A. V., Mokry, J., and Nadejde, M. (2017). Nematus: a toolkit for neural machine translation. In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain.
- Varga, D., Halácsy, P., Kornai, A., Nagy, V., Németh, L., and Trón, V. (2005). Parallel corpora for medium density languages. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 590–596, Borovets, Bulgaria.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Voita, E., Serdyukov, P., Sennrich, R., and Titov, I. (2018). Context-aware neural machine translation learns anaphora resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia.
- Wang, L. (2019). *Discourse-Aware Neural Machine Translation*. Ph.D. thesis, Dublin City University, Dublin, Ireland.
- Zhang, M. and Toral, A. (2019). The effect of translationese in machine translation test sets. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 73–81, Florence, Italy.
- Zhang, J., Luan, H., Sun, M., Zhai, F., Xu, J., Zhang, M., and Liu, Y. (2018). Improving the transformer translation model with document-level context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 533–542, Brussels, Belgium.

## 10 Language Resource References

Tiedemann, J. (2012). Parallel data, tools and interfaces in opus. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, may. European Language Resources Association (ELRA).

## A Examples from document-level and section-level side-constraints models

Source	Reference	Document-aware model	Section-aware model
<b>Soliste</b> internationale, Marielle Nordmann...	An international soloist, Marielle Nordmann...	An international classical <b>violinist</b> , Marielle Nordmann...	An international <b>solo artist</b> , Marielle Nordmann...
...cité par l'abbé Abel Gaborit. (dans "Musica et Memoria")	...mentioned by abbot Abel Gaborit. (in "Musica et Memoria")	...quoted by Abel Gaborit. (in "Musica and Memoria")	...quoted by Abel Gaborit.
Entre 1960 et 1978, elle a <b>formé</b> le Trio Nordmann avec...	Between 1960 and 1978, she <b>led</b> the Nordmann Trio with...	Between 1960 and 1978 she <b>formed</b> the Nordmann Trio...	Between 1960 and 1978 she <b>trained</b> the Nordmann Trio
相较于斯科特探险队的不幸，阿蒙森的探险比较平顺。	In contrast to the misfortunes of Scott's team, Amundsen's trek proved <b>relatively smooth and uneventful</b> .	Compared with the Scott expedition's misfortune, Amundsen's expedition was <b>smoothness</b> .	Compared with the Scott expedition's misfortune, Amundsen's expedition was <b>relatively smooth</b> .
要求抛开关于自己处境的幻想，也就是要求抛开那需要幻想的处境。	To call <b>on them to give up their illusions about their condition</b> is to call on them to give up a condition that requires illusions.	To ask for an <b>abandoned fantasy about your situation</b> , that is, a situation that requires fantasy.	To ask for an <b>abandonment of the illusion of one's own situation</b> , that is, to ask for an abandonment of the situation that requires fantasy.
1919年，参加了五四运动，加入少年中国学会。	In 1919, he <b>participated in the May Fourth Movement, China Youth Association</b> .	In 1919, he <b>took part in the May 4th Movement and joined the Young China Institute</b> .	<b>He joined the Boys' China Society</b> in 1919.
Два дни по-късно <b>получава първото генералско звание - генерал-майор</b> и е назначен за командващ на 2-ра армия.	On 14 September 1944 he <b>was promoted to Major General and</b> was given the command of the Bulgarian Second Army.	He was appointed commander of the 2nd Army two days later.	Two days later, he <b>received the first rank of major-general and</b> was appointed commander of the 2nd Army.
...е канадски геолог, изследовател на Канада, <b>спортист</b> .	... <b>was</b> a Canadian geologist, explorer <b>and athlete</b> .	... <b>was</b> a Canadian geologist, explorer, <b>and sportsman</b> .	... <b>is</b> a Canadian geologist, explorer of Canada.
По-късно започва да се занимава със строителство и <b>дърворезба</b> .	Later he began to deal with the particular construction with <b>wood carving</b> , shaping the profession quickly.	He later became involved in construction and <b>forestry</b> .	He later became involved in construction and <b>woodwork</b> .

Figure 8: Illustration of the types of differences between document-aware and section-aware side constraint model outputs.