

Learning from Partially Annotated Data: Example-aware Creation of Gap-filling Exercises for Language Learning

Semere Kiros Bitew*, Johannes Deleu*, A. Seza Dođruöz, Chris Develder
and Thomas Demeester

IDLab, Ghent University - imec

* Equal contribution

{semerekiros.bitew, as.dogruoz, firstname.lastname}@ugent.be

Abstract

Since performing exercises (including, e.g., practice tests) forms a crucial component of learning, and creating such exercises requires non-trivial effort from the teacher. There is a great value in automatic exercise generation in digital tools in education. In this paper, we particularly focus on automatic creation of gap-filling exercises for language learning, specifically grammar exercises. Since providing any annotation in this domain requires human expert effort, we aim to avoid it entirely and explore the task of converting existing texts into new gap-filling exercises, purely based on an example exercise, *without explicit instruction or detailed annotation* of the intended grammar topics. We contribute (i) a novel neural network architecture specifically designed for aforementioned gap-filling exercise generation task, and (ii) a real-world benchmark dataset for French grammar. We show that our model for this French grammar gap-filling exercise generation outperforms a competitive baseline classifier by 8% in F1 percentage points, achieving an average F1 score of 82%. Our model implementation and the dataset are made publicly available¹ to foster future research, thus offering a standardized evaluation and baseline solution of the proposed partially annotated data prediction task in grammar exercise creation.

1 Introduction

While digital education tools have been increasingly developed and deployed for over a decade, the e-learning sector has definitely boomed in the wake of COVID-19, even leading to a new Digital Education Action Plan from the European Commission.² As one application in e-learning, we particularly focus on language education, and specifically on the automatic generation of gap-filling grammar exercises. This type of exercises has been

shown to be very effective in language learning, with a noticeable effect of such practice tests on students progress and is generally considered as a global measure of language proficiency (Oller Jr, 1973). Furthermore, automatic generation of exercises has been shown produce relatively high quality exercises, for example, for multiple choice questions (Mitkov et al., 2006), demonstrating the potential effectiveness of reducing human effort and offering cost-effective solutions towards personalized exercise generation. In terms of technology, recent developments in natural language processing, e.g., BERT (Devlin et al., 2018), GPT-3 (Brown et al., 2020), InstructGPT (Ouyang et al., 2022), open up new opportunities for further up-scaling and improving automatic generation of such tests/exercises.

In this paper we specifically propose to generate grammar exercises from existing texts, by inducing well-chosen gaps in a given input sentence, following a set of given example exercise sentences. Further, we aim to create models that can be trained on the exercises themselves, without further annotations. The latter implies that we want to forgo a fully supervised learning setting, because such models would require each gap in the available exercises to be manually annotated with additional metadata, such as the particular exercise type, e.g., for gap-filling exercises, a suitable category such as a verb tense. Thus, we focus on converting given input texts into gap-filling exercises, by mimicking the implicit rules underlying a given example exercise, rather than by following explicit instructions such as a prescribed exercise type.

Application scenario: Consider a language teacher, who just introduced a particular grammatical topic (e.g., a new verb tense), and needs the students to practice. The grammar topic of interest may need to be practiced in combination with particular other topics (e.g., related tenses already stud-

¹<https://github.com/semerekiros/GF2/>

²<https://education.ec.europa.eu/focus-topics/digital-education/action-plan>

ied by the students). Given that gap-filling questions can be completed online and automatically assessed (Daradoumis et al., 2019), the teacher creates a new gap-filling exercise, covering these combined grammar topics. The goal of our model is then to support the automatic creation of new exercises, based on that example exercise, by transforming other texts provided by the teacher into additional gap-filling exercises that target the same linguistic topics to be practiced, without explicit instructions by the teacher of which topics the model should include. This would allow the teacher to rapidly create new training material for the students, potentially more diverse, for example, in terms of topics of the texts, their temporal relevance, or the inherent linguistic difficulty.

Learning from partially annotated data: The scenario outlined above represents a learning task in between one-shot learning (i.e., learning from one example (Wang et al., 2020) and full supervision (i.e., based on the full annotation of all examples). On the one hand, the one-shot setting considers the example exercise as a single training instance defining the nature of the prediction task by the way it was constructed by the teacher (in this case, the included grammar topics). On the other hand, the fully supervised setting would require at least explicit knowledge of all exercise instructions (i.e., gap types per exercise). Although we assume the availability of an entire corpus of such exercises, on overlapping grammar topics, we will not rely on explicit annotation of the nature of the gaps (i.e., gap type that defines the type/scope of the grammar exercise, or even just identifying the word category). Thus, we do want to learn from partially annotated examples, where the annotation is limited to just the indication of the gap and the text span that constitutes the expected answer. This basically amounts to the type of information that would be available in a one-/few-shot setting, but we aim to leverage the complete corpus to train our models.

Note that, while creating exercises, teachers are aware of the envisioned exercise type and the gap types, and such exercise type would also be communicated (e.g., as a free-text instruction) to students. Still, to keep our experiments and the gained insights transparent, we left out any exercise level instructions for our experiments.

Link with related research: In broad terms, the proposed work fits within the area of automatic question generation (AQG) for the educational domain. In the field of education, creating questions manually is an arduous task that demands considerable time, training, experience, and resources from educators (Davis, 2009). As a solution to this challenge, researchers have turned towards AGQ approaches to automatically generate homework, test, and exam exercises from readily available plain text that requires little to no human calibration. In particular, educational AQG systems have been developed for generating *factoid questions* covering several subjects such as history (Al-Yahya, 2011; Papasalouros et al., 2008), general sciences (Sun et al., 2018; Stasaski and Hearst, 2017; Conejo et al., 2016), health and biomedical sciences (Pugh et al., 2016; Afzal and Mitkov, 2014), etc., as well as for *language learning* such as vocabulary or grammar exercises (Susanti et al., 2017; Hill and Simha, 2016; Goto et al., 2010). There has been some more generic recent work, however, on finding distractors for multiple choice questions across subjects and languages (Bitew et al., 2022). It is line with recent work on training deep neural networks for general-purpose question generation (Du et al., 2017), based on large training sets. There is a clear preference for two question types that allow for automated assessment, i.e., multiple-choice questions (e.g., in (Stasaski and Hearst, 2017; Pugh et al., 2016; Afzal and Mitkov, 2014; Papasalouros et al., 2008)) or gap-filling questions (as in (Hill and Simha, 2016; Malinova and Rahneva, 2016; Perez-Beltrachini et al., 2012; Goto et al., 2010)).

Our work is focused on gap-filling questions, which typically require test-takers to fill in blank spaces in a text with missing word(s) omitted by test developers. The missing words can either be chosen from a set of possible answers (i.e., closed cloze questions), or generated from scratch using hints provided in the text (i.e., open cloze questions). To generate such questions, various strategies were employed, such as deleting every *n*th word from a text (Taylor, 1953), or rationally deleting words according to specific purpose, e.g., usage of prepositions (Lee and Seneff, 2007), verbs (Sumita et al., 2005) etc. Previous studies have relied on selecting informative sentences (Slavuj et al., 2021; Pino et al., 2008) from existing corpora, such as textbooks (Agarwal and Mannem, 2011), WordNet (Pino et al., 2008), and

Example 1

- 1 Vous travaillerez beaucoup?
1 Will you work a lot?
- 2 En ne mangeant plus de bonbons, tu maigriras vite!
2 By not eating sweets, you will lose weight quickly!
- 3 J'espère que mon équipe favorite ne perdra plus aucun match.
3 I hope my favorite team won't lose any more games.
- 4 Maxime m'a promis qu'il ne mentira plus jamais.
4 Maxime promised me that he will never lie again.
- 5 Maman préparera des spaghettis ce soir.
5 Mum will make spaghetti tonight.

Example 2

A l'âge de 27 ans, le Californien David Blancarte
At the age of 27, Californian David Blancarte had
a eu un grave accident de scooter. Quand il s'est
a serious scooter accident. When he woke up
réveillé à l'hôpital, il ne sentait plus ses
in the hospital, he no longer felt his
jambes. On lui a expliqué qu'il ne pourrait
legs. It was explained to him that he couldn't
plus marcher. C' était une vraie catastrophe
walk anymore. It was a real disaster for him!
pour lui! Pendant une longue période de
During a long period of rehabilitation, he learned
revalidation, il a appris à se déplacer en chaise.
to move around in a wheelchair.
roulante. ...

Figure 1: French grammar exercise from the GF2 corpus, with English translations for convenience shown in light grey. Green spans (with solid underline) are actual gaps as selected by teachers in the dataset, red spans represent potential gaps on other grammar topics but were not marked as gaps. (Left) Isolated sentence exercise with focus on a single tense (*futur simple*); (right) full text exercise combining two tense types (*imparfait* and *passé composé*).

then using techniques such as POS tagging (Agarwal and Mannem, 2011) or term frequency analysis (Mitkov et al., 2006) to determine gap positions. More recently, Marrese-Taylor et al. (2018), have developed sequence labeling model to automate the process of generating gap-filling exercises.

Another very relevant work by Felice et al. (2022) devised a method to adapt an ELECTRA (Clark et al., 2020) model for the purpose of generating open cloze grammar exercises in English. Their approach involved classifying each individual token as either a gap or non-gap. However, there exist several notable distinctions between their approach and our own. Firstly, unlike their method that solely focused on individual tokens, we make gap decisions based on spans. This distinction is essential as our gaps can encompass multiple words, allowing for more comprehensive and contextually accurate grammar exercises. Secondly, our objective and experimental setup differ significantly. Our ultimate goal is to generate multiple versions of the same text, with each version targeting a distinct grammar aspect (e.g., future tense, prepositions of time or combinations of different

types). In contrast, their approach consistently produces exercises of the same type for a given input text (i.e., similar to our baseline model), lacking the versatility and adaptability our model offers.

We observed a tendency in generation of gap-filling questions aiming at well-defined tasks. To the best of our knowledge, none of the prior works have proposed strategies to capture common underlying structures in terms of task definition, while training on a heterogeneous set of real-world examples (e.g., covering various grammatical topics).

Key research contributions:

- We introduce the task of the example-aware prediction of suitable linguistic gaps in texts based on partially annotated data. This task is of paramount importance in the development of new gap-filling exercises.
- We present our real-world dataset of French gap-filling exercises covering unknown combinations of grammatical aspects. Our dataset called GF2 (*'Gap-Filling for Grammar in French'*) is released as a research benchmark for the introduced task.

- We propose and train a suitable neural network architecture for the task, and show that conditioning the model’s output for a given input text on an example exercise of the envisioned exercise type, leads to an increased effectiveness, compared to an example-independent baseline model. Additionally we analyse the model’s ability to disentangle elementary exercise types, without being explicitly trained to do so, and we observe that it can recognize types to some extent, especially for the most commonly occurring types in the test set.

2 Gap-filling Exercise Creation as a Span Detection Task

This section describes the particular prediction task this paper focuses on. We cast the creation of a French gap-filling exercise from an input text as a *binary span detection task*: the goal is detecting each span (i.e., consecutive sequence of tokens) that represents a correct gap. For clarity, we left out creating the ‘hint’ (e.g., the infinitive for verbs) which would make it a finalized gap-filling exercise, as it is considered less challenging and may deviate attention from the core problem of identifying the correct spans.

Figure 1 shows two example gap-filling exercises, with indication of the ground truth spans in green (and with solid underline). We denote the distinguishing feature of each gap as its *gap type* (e.g., the tense *futur simple* for each of the valid tags in Example 1). An exercise typically covers multiple gap types, and the particular combination that characterizes a given exercise is called its *exercise type*. As such, many different exercise types can be constructed, and some may be unseen in the training data. For example, Example 2 (again in Fig. 1) combines three tenses (*imparfait*, *passé composé*, and *conditionnel présent*), which constitutes its exercise type. However, the same text could have been enriched with different gaps, corresponding to a different exercise type. In fact, our test set of one hundred exercises, for which we annotated gap types in terms of 12 elementary verb tenses, covers a total of 35 such composite exercise types.

Considering the lack of information regarding the exercise types for the training exercises, we further define the task we are examining more precisely. The objective is to detect the valid spans (i.e., spans that will be designated as gaps) of a given flat *input* text that mimics the same underly-

ing exercise type as an example gap-filling exercise, which we denote as the *exemplar*. This exemplar serves as an indirect reference for the model to understand the desired exercise type. By utilizing this approach, we can better inform the model about the desired exercise type while accounting for the the lack of exercise information available.

Note that our goal is working with real-world data. Our training data contains gap-filling examples following particular unknown exercise types. Moreover, teachers appear to not always select every possible span that satisfies the exercise type. We saw cases in our dataset (cf. Section 4.1), where the same verb occurring twice in the same form would be selected as a valid gap only once. Such real-world ‘inconsistencies’ contribute to the challenging nature of learning from such data without additional annotations.

3 Example-aware span detection model

This section describes our baseline model and proposed example-aware gap detection model. Figure 2 provides a schematic overview. We first detail the part indicated as *Baseline model*, inside the smaller dashed box, followed by the part that encodes the exemplar, which leads to the full model.

Baseline model: An input text \mathbf{t} , consisting of N tokens $\mathbf{t} = [t_0, t_1, \dots, t_{N-1}]$ is encoded by a transformer based masked language model (MLM), in our experiments the multilingual XLM-RoBERTa (Conneau et al., 2019). From the corresponding transformer outputs $[\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_{N-1}]$ (with $\mathbf{h}_i \in \mathbb{R}^k$, $i=0 \dots N-1$), vector representations are constructed for all possible spans inside the input sequence, up to a certain length (in our experiments 12 tokens). The goal is then to make a binary prediction in terms of valid gaps, for each of these spans. In particular, for a span $\varsigma = [t_{\text{start}}, \dots, t_{\text{end}}]$ with endpoint tokens t_{start} , t_{end} and width $|\varsigma| = (\text{end} - \text{start} + 1)$ in the input text, the corresponding span representation \mathbf{h}_{ς} is constructed as

$$\mathbf{h}_{\varsigma} = \text{FFNN}(\mathbf{h}_{\text{start}} \oplus \mathbf{h}_{\text{end}} \oplus \mathbf{h}_{|\varsigma|})$$

in which \oplus represents vector concatenation, $\mathbf{h}_{|\varsigma|}$ corresponds to a span width embedding, jointly learned with the model, and FFNN is a fully connected feed-forward model with a single hidden layer, ReLU activation, and output dimension k .

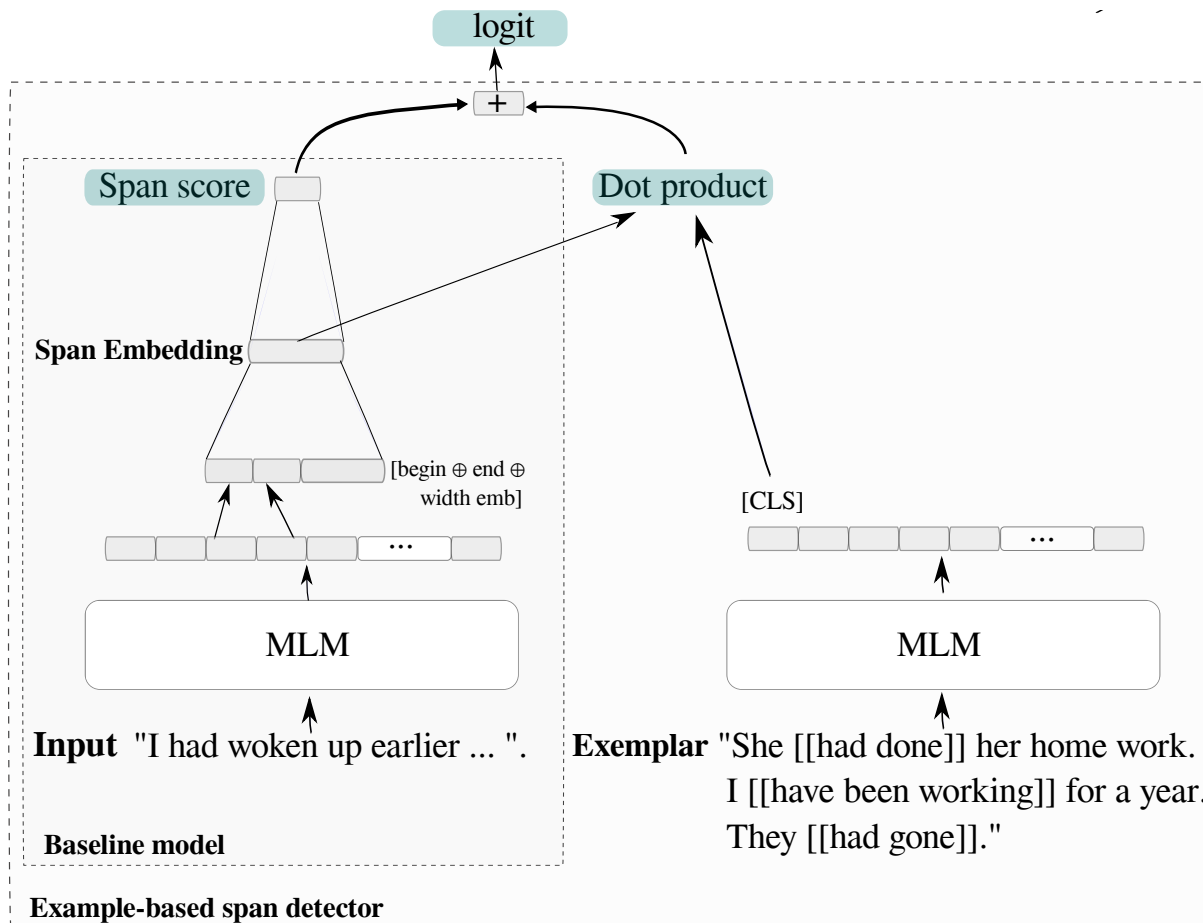


Figure 2: Example-aware gap detection model architecture. \oplus denotes concatenation. In general, the model considers all possible spans up to a maximum width, but we depict here only one span from the input for brevity.

The XLM-ROBERTa output representations $\mathbf{h}_{\text{start}}$ and \mathbf{h}_{end} of the start and end token of ς are concatenated with the span width embedding $\mathbf{h}_{|\varsigma|}$, and transformed through FFNN into the k -dimensional span representation \mathbf{h}_{ς} . The probability of span ς representing a valid gap is modeled as

$$p_{\text{base}}(\varsigma) = \sigma(\mathbf{w} \cdot \mathbf{h}_{\varsigma} + b)$$

in which the trainable parameters \mathbf{w} and b are a k -length coefficient vector and bias, respectively, σ is the sigmoid function, and \cdot represents the dot product. The baseline model is trained by minimizing the cross entropy loss between each span’s score $p_{\text{base}}(\varsigma)$ and its label (1 for valid gaps, 0 otherwise). At inference, spans are predicted as gaps as soon as $p_{\varsigma} \geq 0.5$.

Example-aware gap detection model: As shown in Fig. 2, our example-aware model is a direct extension of the baseline model which by construction makes example-unaware predictions. The same MLM that encodes the input, is now

used to also encode the exemplar, which contains the example exercise text as well as the correct gap information. The latter is added by surrounding each gap with the special tokens ‘[’ and ‘]’ (as seen in the figure). Details on how the examples are chosen, are provided in Section 4.2. The exemplar representation $\mathbf{h}_{\text{exemplar}}$ is obtained as the MLM’s [CLS] representation³.

We then quantify the compatibility of each span ς in the input text with the exemplar, through the dot product $\mathbf{h}_{\text{exemplar}} \cdot \mathbf{h}_{\varsigma}$ of their respective representations. In a direct extension of the baseline model, it leads to the proposed model for the probability $p_{\text{example-aware}}(\varsigma)$ that ς represents a valid gap:

$$p_{\text{example-aware}}(\varsigma) = \sigma(\mathbf{h}_{\varsigma} \cdot \mathbf{w} + \mathbf{h}_{\varsigma} \cdot \mathbf{h}_{\text{exemplar}} + b)$$

³[CLS] is a special token that is prepended to the input, and its corresponding output representation is pretrained to represent the entire sequence that is used for classification tasks

4 Empirical validation on real-world data

In this section, we first introduce the dataset that we will publicly release. Then, we explain how we train our models and use them for inference. Finally, we describe the strategies we adopted to evaluate the effectiveness of our models.

4.1 GF2 dataset: Gap-Fill for Grammar in French

We denote our new dataset as “Gap-Filling for Grammar in French” (GF2). It was contributed by Televic Education⁴, and gathered through its education platform assessmentQ⁵. AssessmentQ is a comprehensive online platform for interactive workforce learning and high-stakes exams. It allows teachers to compose their questions and answers for practice and assessment. As a result, the dataset is made up of a real-world set of gap-filling grammar exercise questions for French, manually created by experts. We cleaned and preprocessed the data before we could use it to train our models. First, organizational metadata information was removed. Other elements that we removed are the hints within the body of the text that could easily give away the gap positions, as well as inline instructions (if present) about the exercise type. Second, we automatically stripped off HTML tags from the documents. Our final dataset contains a total of 768 exercise documents, in which a total of 5,530 spans are tagged as gaps. The exercises were randomly split into 618 train documents, and 50 and 100 for validation and test, respectively. Table 1 summarizes GF2’s descriptive statistics.

For the validation and test exercises, we made an extra manual effort to enrich each of the existing gaps with their gap type. Our annotations reflect the fact that the data contains a mix of verb and non-verb gaps. Every gap has an associated word type attribute (e.g. adverb, adjective, verb) and in case of verbs a tense attribute. In what follows we zoom in on the verb gaps and consider the tense as the main gap type. The bottom half of Table 1 shows the frequency of occurrence for the main verb types in the development and test documents. We use these annotations to get insights into the dataset and to evaluate the properties of our models (see Section 5). Note that the examples shown in Fig. 1 are actual entries from the GF2 dataset.

⁴<https://www.televic.com/en/education>

⁵<https://www.televic-education.com/en/assessmentq>

4.2 Training and inference

Our baseline model is relatively straightforward to train. We designate all spans indicated as gaps in our training data as valid gaps, which are considered positive examples. Conversely, any spans that are not indicated as gaps are labeled as negatives. We train our model by minimizing the cross entropy loss between each span’s predicted score and its label as described in Section 3. However, training our example-aware model poses a challenge due to the lack of knowledge regarding the exercise types of the training exercises. Using one exercise as an example and another exercise of the same type as the input, along with the corresponding targets, is not therefore feasible. Instead, we make the assumption that exercises are generated by teachers who consistently follow the underlying exercise type throughout the entire exercise. As a result, we divide the training exercises into two parts: one part is used as an exemplar, and the other part serves as the actual input, for which the gaps are assumed to follow the same exercise type.

To this end, we first segment each document in the training set into a list of sentences, along with their corresponding target gap positions. We create a new (exemplar, input) training pair by sampling one sentence to be used as the input, and uniformly sampling one up to m sentences from the remaining sentences within the same document to be used as the exemplar. The exemplar is constructed by concatenating these sampled sentences, with the addition of special symbols denoting the gap locations. (See Appendix A for details.) These are the positive training examples that encourage the model to correctly learn predicting example-aware gaps. However, to facilitate efficient learning, it is crucial to also provide negative examples on which the model should not predict gaps. To create such negative training instances, a sentence is sampled as input from the considered document, but its span targets are set to zero (no gaps), and the negative exemplar is composed as before (including indicating the gaps), but by sampling sentences from a randomly selected *other* training exercise. There is risk of incidentally creating false negative training examples, if the exemplar gaps correspond with left-out gaps in the input. However, negative exemplars appeared important for obtaining a suitable model.

We determine the optimal proportion of negative to positive instances for training our models by em-

Table 1: Statistics of the FG2 dataset and breakdown into key verb tenses (gap types) in the validation and test split. For the train split we only know gap spans, not their types, since they are not labelled.

	Train	Dev	Test
# Documents	618	50	100
# Sentences	4786	378	707
# Gaps	4518	365	647
Subjonctif Présent (SPR)	UNK	1	28
Passé Composé (participe passé) (PCP)	UNK	31	8
Passé Composé (PC)	UNK	84	108
Imparfait (IM)	UNK	8	46
Conditionnel Présent (CPR)	UNK	23	92
Passé Récent (PR)	UNK	0	12
Futur Proche (FP)	UNK	1	9
Futur Simple (FS)	UNK	8	49
Indicatif Présent (IP)	UNK	126	144
Conditionnel Passé (CPA)	UNK	0	3
Impératif (IMP)	UNK	12	26
Plus-que-parfait (PQ)	UNK	0	1

ploying a fine-tuning approach utilizing the macro F1 score as the evaluation metric on the validation set. This increases the impact of the rarer gap types in the metric, and therefore in the final model, which we considered important for practical use. Other choices could have been made, however. Ultimately, the final model is trained on the union of the training and validation splits, using the optimal proportion determined via the fine-tuning process.

During inference, we use our trained model to predict the gap positions for an input text that is implicitly conditioned on the target exercise type through the exemplar.

Implementation and training details: We implement our models using pytorch and Huggingface. We initialize our MLM encoders with xlm-roberta-base. To avoid extensive hyperparameter tuning, we made the following choices; a learning rate of 2e-5 in combination with the robust Adam optimizer. We use a batch size of 16 and train our models for 30 epochs. We consider all spans up to a maximum length 12 and we set k , the number of sentences per exemplar to 3.

4.3 Evaluation setup

In order to assess and analyze the performance of the baseline and the example-aware model, we design two evaluation strategies that look at different effectiveness aspects.

Binary gap prediction evaluation: the primary objective of our model is to mimic the real-world setting where gap labels are not given. We measure how well our models predict gap positions (i.e., gap or no-gap decisions for all input spans). To do this, we split up each of the exercise documents in our test into two parts that are roughly the same size, given that by assumption they then represent the same exercise type. We calculate the automated metrics by using one half as the exemplar and the second as the input text to our model. We repeat this process by exchanging the roles of the parts. It is worth noting that we excluded one-sentence test documents (i.e., because they can not be chunked into two parts), which amount to 16% of the total test documents. However, since most of the excluded sentences (i.e., one-line documents) only had one gap, we only removed 2.7% of the total gaps in the test set.

Gap type disentangling evaluation: The goal of the second evaluation setting is to analyze how well the model has learned to disentangle individual gap types, despite not being explicitly trained to do so. This analysis is based on the assumption that a model that scores high on that aspect, would be stronger in dealing with new or rare exercise types. Potentially even at creating new combinations of existing exercises. This is an aspect we plan to study further when designing more advanced models in future research. To this end, we construct

Table 2: Tense disentangling ability in terms of precision, recall, and F1 (in %) on the test set, as reported for each key verb tenses (with on the right their support, i.e., number of occurrences). We also show the macro F1 score for the static baseline (*baseline*) and our proposed example-aware gap prediction (*ours*).

Tenses	<i>Baseline</i>			<i>Ours</i>			Support
	P	R	F1	P	R	F1	
SPR	5.0 \pm 0.3	78.6 \pm 8.9	9.4 \pm 0.6	7.5 \pm 0.2	81.0 \pm 12.5	13.7 \pm 0.4	28
PCP	0.1 \pm 0.1	4.2 \pm 6.3	0.2 \pm 0.3	12.6 \pm 4.1	62.5 \pm 12.5	20.7 \pm 6.2	8
PC	21.3 \pm 1.2	86.4 \pm 3.7	34.2 \pm 1.8	64 \pm 9.4	86.1 \pm 1.9	73.1 \pm 5.5	108
IM	9.3 \pm 0.4	88.4 \pm 3.7	16.2 \pm 0.8	12.0 \pm 2.5	78.3 \pm 10.9	20.9 \pm 3.9	46
CPR	19.9 \pm 0.5	94.5 \pm 2.9	32.8 \pm 0.8	28.3 \pm 2.9	92.4 \pm 4.7	43.2 \pm 3.1	92
PR	2.7 \pm 0.1	100.0 \pm 0.0	5.3 \pm 0.1	9.7 \pm 2.0	100.0 \pm 0.0	17.7 \pm 3.3	12
FP	1.6 \pm 0.0	77.7 \pm 0.0	3.1 \pm 0.1	6.0 \pm 0.9	77.8 \pm 0.0	11.1 \pm 1.5	9
FS	9.9 \pm 0.3	88.5 \pm 1.7	17.8 \pm 0.5	13.6 \pm 1.1	84.4 \pm 10	23.3 \pm 1.7	49
IP	24.6 \pm 1.2	75.0 \pm 4.3	37.1 \pm 1.9	32.0 \pm 1.4	66.2 \pm 11.9	42.9 \pm 2.4	144
CPA	0.1 \pm 0.1	11.1 \pm 16	0.2 \pm 0.3	0	0	0	3
IMP	5.2 \pm 0.3	88.5 \pm 2.2	9.9 \pm 0.5	16.8 \pm 1.7	84.6 \pm 3.9	25.3 \pm 2.1	26
PQ	0.2 \pm 0.0	100.0 \pm 0.0	0.5 \pm 0.0	0.6 \pm 0.1	100 \pm 0.0	1.2 \pm 0.2	1
Macro F1		13.9			24.4		

a small set of 12 exemplars, one for each of the key verb tenses, by randomly selecting them from the original data and subsequently removing them from the train/validation/test splits. Each exemplar comprises multiple sentences, all of which are homogeneously annotated with the same intended verb type, which will serve as the desired homogeneous exercise type. We evaluate our model on every sentence of the test set, by prompting it with each of these 12 fixed exemplars. Based on the gap types we annotated on the test set, we can then compute the precision, recall and F1 score for each of these 12 tenses.

5 Experimental Results

In this section, we provide evidence of the effectiveness of our proposed model by reporting and discussing the experimental results. Table 3 summarizes the binary gap prediction evaluation of the baseline vs. the example-aware model on the test set. We report our results as the mean and standard deviation over five runs, each using a different random seed for model training. The proposed example-aware model (denoted as *ours*) consistently outperforms the example-unaware *baseline* on all metrics. In general, there is an absolute gain of 8 percentage points in F1 for the proposed model in comparison with the baseline, achieving an average F1 score of 82.4%. This confirms our intention when designing the model, that providing exam-

ple exercises leads to an increased effectiveness in terms of predicting gap positions compared to the static baseline model.

Table 3: Overall binary gap prediction in terms of precision, recall, and F1 (in %) on the test set. Results shown for the static baseline (*baseline*) and our proposed example-aware gap prediction (*ours*).

	Precision	Recall	F1
<i>Baseline</i>	74.87 \pm 2.44	73.11 \pm 2.00	73.92 \pm 0.49
<i>Ours</i>	84.30 \pm 1.70	80.74 \pm 1.80	82.40 \pm 0.20

In Table 2, we show the evaluation of our models in their ability to disentangle the 12 main verb types. We observe that for the tenses with relatively higher support, the example-aware model outperforms the baseline with certainty as demonstrated by the individual F1 scores.

The overall macro F1 score for the example-aware model stands at 24.4%, which is low in absolute value, but considerably higher than the baseline’s macro F1 score of 13.9%. We observe that the proposed model is able to recognize verb types such as passé composé (PC), imparfait (IM), and conditionnel présent (CPR) to some extent with F1 scores of 73%, 43%, and 42%, respectively. However, the low overall scores are not unexpected, because the models are not trained to recognize gap types. Furthermore, some tenses are either

very rare (e.g., PQ, CPA, PCP) as indicated by their support, or may appear mainly in combination with other exercise types. This makes achieving a better resolution in disentangling gap types without any explicit gap labels during training an inherently difficult task.

6 Conclusion

In this paper, we introduced a new task within the general challenge of training models to automatically create new exercises for use in education, based on existing exercises and without requiring additional manual annotations.

In particular, we introduced a dataset and associated prediction task, focusing on detecting gaps within a given input text, without knowledge of the exact exercise type, by only relying on an example exercise. We proposed an example-aware neural network model designed for this task, and compared it with a baseline model that does not take into account any example of the desired exercise type. We found that our example-aware model outperforms the baseline model not only in predicting gaps, but also in disentangling gap types despite not being explicitly trained on that task. Our real-world GF2 dataset of French gap-filling exercises will be publicly released together with the code to reproduce the presented empirical results.

The presented work fits with our pursuit towards supporting personalized learning experiences by either suggesting existing or generating new exercises that are tailored to students' needs. Teachers could also benefit from an increased efficiency in creating new exercises. For example, they could make many and diverse drill and practice exercises on chunks of text based on existing standard exercise types without having to provide extra meta-data information such as instructions. We hope our benchmark dataset and task will spark new research in the CL and Educational NLP community.

Limitations

We identify two limitations of the current work and make suggestions for future directions. First, while our proposed method is language-agnostic in principle, our evaluation is limited to our French benchmark dataset. Expanding our approach to encompass other languages would bring new and interesting challenges for further investigation. Second, despite topic diversity within our exercise documents (e.g., the first example in Fig. 1 consists of

independent sentences, while the second is a coherent text centered around the same topic.), it would be interesting to quantify the degree of topical bias introduced during our training process and its impact on our binary task evaluation. For future work, we first aim to adapt seq2seq models for our task particularly text-to-text models such as T5 (Raffel et al., 2020). There is also potential to explore different prompting strategies for large language models (LLMs), when generating gap-filling grammar exercises. For instance, the utilization of chain-of-thought prompting (Wei et al., 2022), which involves generating intermediate steps before producing the final response, could be explored for generating grammar exercises. Additionally, an interesting future study would involve investigating the number of example demonstrations that LLMs require in order to accurately mimic example gap exercises.

Ethics Statement

In this research, we posit that the dataset and models introduced are of low-risk in terms of potential harm to individuals. The dataset used is a curated selection of existing educational content enriched with meta-data, and we are confident that our compilation of the dataset has not introduced any additional ethical risks. However, it is crucial to emphasize the need for accountability and the establishment of clear guidelines for the deployment of grammar generation models, such as the ones benchmarked in this paper, for educational purposes.

It should be noted that our models are derived from general-purpose neural language encoders that have been trained on real-world data, which may contain biases or discriminatory content (Bommasani et al., 2021). As a result, our models may have inherited some of these biases and could potentially base their prediction on such biased information. Therefore, it is imperative for educators and researchers to thoroughly consider these ethical issues and ensure that the generated grammar questions align with educational goals and do not perpetuate harmful biases.

Educators should retain the final authority in accepting or modifying grammar question suggestions generated by such models, keeping their educational goals in mind (e.g., in terms of formative and especially summative assessment). In practice, these models are designed to enhance teachers' effi-

ciency in preparing teaching materials, rather than replacing teachers in any way. An important benefit of using AI-supported question generation with increased efficiency is the potential for personalized approaches towards students.

Acknowledgements

This work was funded by VLAIO (‘Flanders Innovation & Entrepreneurship’) in Flanders, Belgium, through the *imec-icon* project AIDA (‘AI-Driven e-Assessment’). This research also received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme. We would like to thank the AIDA partners Televic Education and WeZooz Academy for contributing data and use cases.

References

- Naveed Afzal and Ruslan Mitkov. 2014. Automatic generation of multiple choice questions using dependency-based semantic relations. *Soft Computing*, 18(7):1269–1281.
- Manish Agarwal and Prashanth Mannem. 2011. Automatic gap-fill question generation from text books. In *Proceedings of the sixth workshop on innovative use of NLP for building educational applications*, pages 56–64.
- Maha Al-Yahya. 2011. Ontoque: a question generation engine for educational assesment based on domain ontologies. In *2011 IEEE 11th International Conference on Advanced Learning Technologies*, pages 393–395. IEEE.
- Semere Kiros Bitew, Amir Hadifar, Lucas Sterckx, Johannes Deleu, Chris Develder, and Thomas De-meester. 2022. Learning to reuse distractors to support multiple choice question generation in education. *IEEE Transactions on Learning Technologies*.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosse-lut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Ricardo Conejo, Eduardo Guzmán, and Monica Trella. 2016. The siette automatic assessment environment. *International Journal of Artificial Intelligence in Education*, 26(1):270–292.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Thanasis Daradoumis, Joan Manuel Marquès Puig, Marta Arguedas, and Laura Calvet Liñan. 2019. Analyzing students’ perceptions to improve the design of an automated assessment tool in online distributed programming. *Computers & Education*, 128:159–170.
- Barbara Gross Davis. 2009. *Tools for teaching*. John Wiley & Sons.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. [Learning to ask: Neural question generation for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.
- Mariano Felice, Shiva Taslimipoor, and Paula Buttery. 2022. [Constructing open cloze tests using generation and discrimination capabilities of transformers](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1263–1273, Dublin, Ireland. Association for Computational Linguistics.
- Takuya Goto, Tomoko Kojiri, Toyohide Watanabe, Tomoharu Iwata, and Takeshi Yamada. 2010. Automatic generation system of multiple-choice cloze questions and its evaluation. *Knowledge Management & E-Learning: An International Journal*, 2(3):210–224.
- Jennifer Hill and Rahul Simha. 2016. Automatic generation of context-based fill-in-the-blank exercises using co-occurrence likelihoods and google n-grams. In *Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 23–30.
- John Lee and Stephanie Seneff. 2007. Automatic generation of cloze items for prepositions. In *Eighth Annual Conference of the International Speech Communication Association*.
- Anna Malinova and Olga Rahneva. 2016. Automatic generation of english language test questions using mathematica. In *CBU International Conference Proceedings*, volume 4, pages 906–909.

- Edison Marrese-Taylor, Ai Nakajima, Yutaka Matsuo, and Ono Yuichi. 2018. Learning to automatically generate fill-in-the-blank quizzes. *arXiv preprint arXiv:1806.04524*.
- Ruslan Mitkov, Ha Le An, and Nikiforos Karamanis. 2006. A computer-aided environment for generating multiple-choice test items. *Natural language engineering*, 12(2):177–194.
- John W Oller Jr. 1973. Cloze tests of second language proficiency and what they measure 1. *Language learning*, 23(1):105–118.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Andreas Papasalouros, Konstantinos Kanaris, and Konstantinos Kotis. 2008. Automatic generation of multiple choice questions from domain ontologies. *e-Learning*, 1:427–434.
- Laura Perez-Beltrachini, Claire Gardent, and German Kruszewski. 2012. Generating grammar exercises. In *The 7th Workshop on Innovative Use of NLP for Building Educational Applications, NAACL-HLT Workshop 2012*, pages 147–157.
- Juan Pino, Michael Heilman, and Maxine Eskenazi. 2008. A selection strategy to improve cloze question quality. In *Proceedings of the Workshop on Intelligent Tutoring Systems for Ill-Defined Domains. 9th International Conference on Intelligent Tutoring Systems, Montreal, Canada*, pages 22–32.
- Debra Pugh, Andre De Champlain, Mark Gierl, Hollis Lai, and Claire Touchie. 2016. Using cognitive models to develop quality multiple-choice questions. *Medical teacher*, 38(8):838–843.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Vanja Slavuj, L Nacinovic Prskalo, and M Brkic Bakaric. 2021. Automatic generation of language exercises based on a universal methodology: An analysis of possibilities. *Bulletin of the Transilvania University of Brasov. Series IV: Philology and Cultural Studies*, pages 29–48.
- Katherine Stasaski and Marti A Hearst. 2017. Multiple choice question generation utilizing an ontology. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 303–312.
- Eiichiro Sumita, Fumiaki Sugaya, and Seiichi Yamamoto. 2005. Measuring non-native speakers’ proficiency of english by using a test with automatically-generated fill-in-the-blank questions. In *Proceedings of the second workshop on Building Educational Applications Using NLP*, pages 61–68.
- Bo Sun, Yunzong Zhu, Yongkang Xiao, Rong Xiao, and Yungang Wei. 2018. Automatic question tagging with deep neural networks. *IEEE Transactions on Learning Technologies*, 12(1):29–43.
- Yuni Susanti, Takenobu Tokunaga, Hitoshi Nishikawa, and Hiroyuki Obari. 2017. Evaluation of automatically generated english vocabulary questions. *Research and practice in technology enhanced learning*, 12(1):1–21.
- Wilson L Taylor. 1953. “cloze procedure”: A new tool for measuring readability. *Journalism quarterly*, 30(4):415–433.
- Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. 2020. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.

A Training details

In this section we detail our training procedure. As depicted in Fig. 3, we first split training exercises into list of sentences, along with their corresponding gap position indications. In order to create new (input, exemplar) pair, we sample 1 sentence from the sentence list to be used as our *input* text, and we uniformly sample 1 up to m (we set $m = 3$) sentences from the remaining sentence list to be used as our exemplar. We form our exemplar by concatenating all the sampled sentences with gap positions indicated by special tokens “[[” and “]]”. Then our model is trained by minimizing the binary cross entropy (BCE) loss between predicted gaps and their target labels (1 for valid gaps, and 0 otherwise).

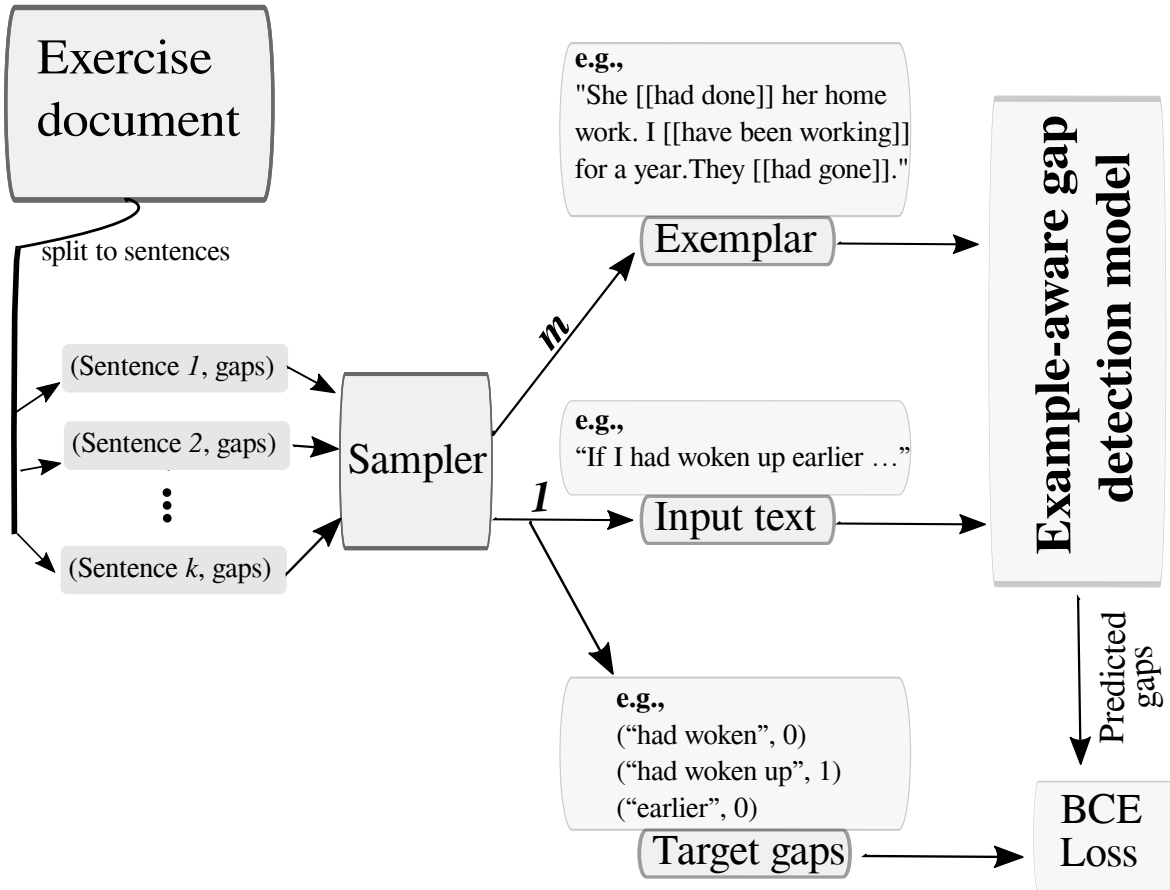


Figure 3: Training procedure of our example-aware gap detection model. First, we split exercise documents into list of sentences. Then we create (input, exemplar) training pairs that will be used by our model. We use one sentence as an input, while the exemplar is made up of sentences that are uniformly sampled from the remaining sentences. The exemplar is constructed by concatenating the m sampled sentences. The special symbols “[[” and “]]” in the exemplar indicate the gap positions. Binary cross entropy (BCE) loss is used to train our models.