

A Theme-Rewriting Approach for Generating Algebra Word Problems

Rik Koncel-Kedziorski

Ioannis Konstas

Luke Zettlemoyer

Hannaneh Hajishirzi

University of Washington

kedzior@uw.edu, {ikonstas, lsz, hannaneh}@cs.washington.edu

Abstract

Texts present coherent stories that have a particular theme or overall setting, for example science fiction or western. In this paper, we present a text generation method called *rewriting* that edits existing human-authored narratives to change their theme without changing the underlying story. We apply the approach to math word problems, where it might help students stay more engaged by quickly transforming all of their homework assignments to the theme of their favorite movie without changing the math concepts that are being taught. Our rewriting method uses a two-stage decoding process, which proposes new words from the target theme and scores the resulting stories according to a number of factors defining aspects of syntactic, semantic, and thematic coherence. Experiments demonstrate that the final stories typically represent the new theme well while still testing the original math concepts, outperforming a number of baselines. We also release a new dataset of human-authored rewrites of math word problems in several themes.

1 Introduction

Storytelling is the complex activity of expressing a plot, its events and participants in words meaningful to an audience. Automatic storytelling systems can be used for customized sport commentaries, enriching video games with personalized or dynamic plot-lines (Barros and Musse, 2007), or providing customized learning materials which meet each individual student’s needs and interests (Bartlett, 2004). In this paper, we focus on generating narrative-style

Jim walked 0.2 of a mile from school to David’s house and 0.7 of a mile from David’s house to his own house. How many miles did Jim walk in all?

Star Wars

Uncle Owen walked 0.2 of a mile from hangar to Luke Skywalker’s room and 0.7 of a mile from Luke Skywalker’s room to his own room. How many miles did Uncle Owen walk in all?

Cartoon

Finn squished 0.2 of a mile from cupboard to Melissa’s dock and 0.7 of a mile from Melissa’s dock to his own dock. How many miles did Finn squish in all?

Western

Duane strolled 0.2 of a mile from barn to Madeline’s camp and 0.7 of a mile from Madeline’s camp to his own camp. How many miles did Duane stroll in all?

Figure 1: An example story and rewrites in 3 themes.

math word problems (Figure 1) and demonstrate that it is possible to design an algorithm that can automatically change the overall theme of a text without changing its underlying story, for example to create more engaging homework that is in the theme of a student’s favorite movie.

A math word problem is a coherent story that provides the student with good clues to the correct mathematical operations between the numerical quantities described therein. However, the particular *theme* of a problem, whether it be about collecting apples or traveling distances through space, can vary significantly so long as the correlation between the story and underlying equation is maintained. Students’ success at solving a word problem is tied to their interest in the problem’s theme (Renninger

et al., 2002), and personalizing word problems increases student understanding, engagement, and performance in the problem solving process (Hart, 1996; Davis-Dorsey et al., 1991).

Motivated by this need for thematically diverse, highly coherent stories, we address the problem of *story rewriting*, or transforming human-authored stories into novel, coherent stories in a new theme. Rather than synthesizing first a story plot (McIntyre and Lapata, 2009; McIntyre and Lapata, 2010) or script (Chambers and Jurafsky, 2009; Pichotta and Mooney, 2016; Granroth-Wilding and Clark, 2016) from scratch, we instead begin from an existing story and iteratively edit it towards a thematically novel but –most crucially– semantically compatible story. This approach allows us to reuse much, but not all, of the syntactic and semantic structure of the original text, resulting in the creation of more coherent and solvable math word problems.

We define a theme to be a collection of reference texts, such as a movie script or series of books. Given a theme, the *rewrite* algorithm constructs new texts by substituting thematically appropriate words and phrases, as measured with automatic metrics over the theme text collection, for parts of the original texts. This process optimizes for a number of metrics of overall text quality, including syntactic, semantics, and discourse scores. It uses no hand-crafted templates and requires no theme-specific tuning data, making it easy to apply for new themes in practice. Tables 4–6 show example stories generated from the rewrite system.

To evaluate performance, we collected a corpus of 450 rewrites of math word problems in Star Wars and Children’s Cartoon themes via crowdsourcing.¹ Experiments with automated metrics and human evaluations demonstrate that the approach described here outperforms a number of baselines and can produce solvable problems in multiple different themes, even with no in-domain tuning.

2 Related Work

Our approach is related to the previous work in story generation (e.g., McIntyre and Lapata (2010)) and sentence rewriting (e.g., text simplification (Xu et

al., 2016)), as reviewed in this section. It has three major differences from all these approaches: First, we focus on multi-sentence stories where preserving the coherence, discourse relations, and solvability is essential. Previous work mainly focuses on rewriting single sentences. Second, we build a theme from a text corpus and show how the stories can be adapted to new themes. Third, our method leverages the human-authored story to capture the semantic skeleton and the plot of the current story, rather than synthesizing the story plot. To our knowledge, we are the first to introduce a text rewriting formulation for story generation.

Story generation has been of long interest to AI researchers (Meehan, 1976; Lebowitz, 1987; Turner, 1993; Liu and Singh, 2002; Mostafazadeh et al., 2016). Recent methods in story generation first synthesize candidate plots for a story and then compile those plots into text. Li et al. (2013) use crowdsourcing to build plot graphs. McIntyre and Lapata (2009; 2010) address story generation through the automatic deduction and reassembly of scripts (Schank and Abelson, 1977), or structured representations of events and their participants, and causal relationships involved. Leveraging the automatic script learning methods of Chambers and Jurafsky (2009), McIntyre and Lapata (2010) learn candidate entity-centered plot graphs, or possible events involving the entity and an ordering between these events, with the use of a genetic algorithm. Then plots are compiled into stories through the use of a rule-based text surface realizer (Lavoie and Rambow, 1997) and reranked using a language model.

Polozov et al. (2015) automatically generate math word problems tailored to a student’s interest using Answer Set Programming to satisfy a collection of pedagogical and narrative requirements. This method naturally produces highly coherent, personalized story problems that meet pedagogical requirements, at the expense of building the thematic ontologies and discourse constraints by hand.²

Additionally, there is related work in text simplification (Wubben et al., 2012; Kauchak, 2013; Zhu et al., 2010; Vanderwende et al., 2007; Woodsend and Lapata, 2011b; Hwang et al., 2015), sentence

¹Data and code available at <https://gitlab.cs.washington.edu/kedzior/Rewriter/>.

²According to Polozov et al. (2015) building small thematic ontologies of types, relations, and discourse tropes (100-200 entries) for each of only 3 literary settings took 1-2 person months.

compression (Filippova and Strube, 2008; Rush et al., 2015), and paraphrasing (Ganitkevitch et al., 2013; Chen and Dolan, 2011; Ganitkevitch et al., 2011). All these tasks are focused on rewriting sentences under a predefined set of constraints, such as simplicity. Different rule-based and data-driven approaches are introduced by Petersen and Ostendorf (2007), Vickrey and Koller (2008), and Siddharthan (2004). Most data-driven approaches take advantage of machine translation techniques, use source-target sentence pairs, and learn rewrite operations (Yatskar et al., 2010; Woodsend and Lapata, 2011a), or use additional external paraphrasing resources (Xu et al., 2016).

Finally, this work is related to those on automatically solving math word problems. Specific topics include number word problems (Shi et al., 2015), logic puzzle problems (Mitra and Baral, 2015), arithmetic word problems (Hosseini et al., 2014; Roy and Roth, 2015), algebra word problems (Kushman et al., 2014; Zhou et al., 2015; Koncel-Kedziorski et al., 2015a; Roy et al., 2016), and geometry word problems (Seo et al., 2015; Seo et al., 2014). Several datasets of word problems are available (Koncel-Kedziorski et al., 2016; Huang et al., 2016), though none address the need for thematic text.

3 Problem Formulation

Our system takes as input a story s and a theme t , and outputs the best rewrite s^* from generated candidates S .

A theme t is defined as a textual corpus that describes a topic or a domain. This is an intentionally broad definition that allows a variety of textual resources to serve as themes. For example, the collection of all Science Fiction stories from the Project Gutenberg can be a theme, or the script of a single movie, or a sampling of fan fiction from the Internet. This flexibility adds to the utility of our work, as varying amounts of thematic text may be available.

The generated candidate s^* is the most *thematically* fit problem that is *syntactically* and *semantically* coherent given the original problem s and the new theme t . We represent a story in terms of the words it contains, so that $s = \{w_1, w_2, \dots, w_n\}$ and




Figure 2: An overview of our method for scoring a candidate story s' given a human-authored story s and a theme t . $Syn(s'|s)$: compatibility of syntactic relations (purple arrows), $Sem_{pair}(s'|s)$: coherence of semantic relations (blue arrows), $Sem_{Lex}(s'|s)$: semantic mapping of individual words, and $Th(s'|t)$: thematicity.

$|s| = n$. The new story s' is defined as:

$$s' = \{f(w_1), f(w_2), \dots, f(w_n)\}$$

where the function $f(w) : \mathcal{V}_o \rightarrow \mathcal{V}_t^K \cup \emptyset$, rewrites a word from the vocabulary of the original problem \mathcal{V}_o to either a word, a trivial noun compound of length K (e.g., multi-word named entity) from the vocabulary of the thematic vocabulary \mathcal{V}_t , or reduces to the empty symbol, i.e., omits the input word entirely; hence the length of s' can differ from that of the original problem.

Formally, our goal is to select the candidate $s' \in S$ by maximizing a scoring function \mathcal{R} over thematic, syntactic and semantic constraints, subject to a set of parameters θ :

$$s^* = \arg \max_{s' \in S} \mathcal{R}(s'|s, t; \theta) \quad (1)$$

In order to find the best story s^* , our problem reduces to generating candidate stories s' from the space of possible rewrites of the human-authored story s in a new theme t (Section 5). Since there are exponentially many rewrites, we follow a two-stage decoding approach: first we identify only the content words w_i in the input problem, and provide for each a list of the top- k most *salient* thematic candidate words and trivial noun compounds. We then search the space by progressively introducing more rewrites in the beam, and scoring them according to \mathcal{R} (Section 4). Figure 2 shows the overview of the scoring function for a candidate sentence s' .

4 Scoring Stories

The scoring function \mathcal{R} decomposes into three components, capturing aspects of syntactic compatibility, semantic coherence, and thematicity:

$$\begin{aligned} \mathcal{R}(s'|s, t; \theta) = & \alpha \times \text{Sem}(s'|s) \\ & + \beta \times \text{Syn}(s'|s) \\ & + \gamma \times \text{Th}(s'|s, t) \end{aligned} \quad (2)$$

The syntactic (Syn) and semantic (Sem) coherence components measure the coherence of the words in the new story s' , as well as their compatibility to the syntactic and semantic relations in the original story s . On the other hand, thematicity (Th) scores the relevance and importance of words in the new story with respect to theme t .

We describe each of these components and the decoding process in the following sections.

4.1 Thematicity

Recall that a theme t is defined as a collection of documents that share a common topic, such as books in the science fiction genre, or scripts of horror movies. We define thematicity of a word w' as the measure of *salience*, or how discriminative that word is to a given theme.³ For example, *robot* and *spaceship* are expected to be highly thematic with respect to Star Wars. In our setting we extend this definition to a candidate problem s' given s and t as:

$$\text{Th}(s'|s, t) = \sum_i^{|s|} \text{Sal}(w'_i, t) \quad (3)$$

where w'_i is a word from the candidate problem, and Sal is its salience score with respect to the theme. In the context of this work we argue that the thematic adaptation of the content words, i.e., nouns, verbs, named entities, and adjectives, plays the most important role in forming a new thematic problem. Therefore, we define their salience (except named entities) based on their tf-idf score over the theme t , and set it to zero for function words. Since named entities have relatively low frequencies in the theme corpus we set their salience to $1 - \frac{1}{c(w'_i)}$, where $c(w'_i)$

³We will be interchangeably referring to w' as either the word or the head of the multi-word noun compound that rewrites the equivalent word w in the original problem.

is the number of times w'_i occurs in the theme. In the example story in Figure 2 the thematicity score is derived as $\text{Sal}(\text{Luke Skywalker}) + \text{Sal}(\text{ships}) + \text{Sal}(\text{droids})$.

4.2 Syntactic compatibility

This work offers a new method for syntactic and discourse coherence based on preserving human-authored syntactic structure in generated text (hence our use of the term *rewriting*). The syntactic constructs in a document play a distinctive role in maintaining cohesion across sentences. We consider the human-authored syntax of the original story s as gold standard, and use it to score a candidate problem s' by considering how well the syntactic relations of s apply to s' .

Formally, given a dependency triple (w_i, w_j, l) from a parse of a sentence in s , we compute the likelihood for the corresponding triple (w'_i, w'_j, l) for w'_i, w'_j in s' . We define the syntactic score for all sentences in s' as:

$$\text{Syn}(s'|s) = \sum_{i,j,l|(w_i,w_j,l) \in \text{Dep}(s)} \mathcal{L}_{\text{Dep}}(w'_i, w'_j, l) \quad (4)$$

where $\text{Dep}(s)$ are the dependency parse trees for all sentences in s ; \mathcal{L}_{Dep} is a 3-gram language model over dependency triples which gives the likelihood of an arc label l being used between a pair of words (w'_i, w'_j) . For example in Figure 2, the syntactic compatibility score includes dependency likelihoods of $\mathcal{L}_{\text{Dep}}(\text{ship}, 2, \text{num})$, $\mathcal{L}_{\text{Dep}}(\text{had}, \text{ship}, \text{obj})$.

Therefore, the Syn function prefers stories s' that (a) have similar dependency structure to the original story s and (b) make use of a common syntactic configuration.

4.3 Semantic Coherence

The semantic coherence component expresses how well a candidate s' rewrites individual words and realizes the semantic relationships that exist in the human-authored story s . Ideally, we would like to preserve enough of the semantics of s in order to produce a coherent story s' , yet we are populating s' with words taken from an unrelated theme. Therefore, we model the semantics of a story s' in terms of the lexical semantics contributed by individual words as well as semantic relationships that exist between its elements. Note that the relationships can

cross the sentence boundaries, promoting discourse coherence.

We decompose semantic relations in a story into a set of local, lexical relationships between pairs of words. Specifically, we consider semantic relations for noun-noun and verb-verb pairs as provided by WordNet (Miller, 1995). Since some relations are not directly outlined in these resources (e.g., the selectional preferences of nouns with regard to their adjectival modifiers), we also consider the word-embedding similarity between words. For example in Figure 2 the semantic relationships are denoted with blue arrows between pairs of content words in the story (e.g., {Sam, dogs}, {dogs, puppies}, etc).

More formally, we define the semantic coherence of s' with respect to s as:

$$\begin{aligned} \text{Sem}(s'|s) &= \sum_i^{|s'|} \text{Sem}_{Lex}(w_i, w'_i) \\ &+ \sum_{i,j \in CW} \text{Sem}_{Pair}(\{w_i, w_j\}, \{w'_i, w'_j\}) \end{aligned} \quad (5)$$

where CW is the set of pairs of indices of content words (nouns, verbs, adjectives, and named entities) from s . We focus on the content words of the original problem, as they carry most of the semantic information. Sem_{Lex} and Sem_{Pair} functions are semantic adaptation scores for individual words and semantic relations respectively, described below.

Semantic Compatibility between words (Sem_{Lex}) is defined as:

$$\text{Sem}_{Lex}(w_i, w'_i) = \cos(w_i, w'_i) + \text{Resnik}(w_i, w'_i) \quad (6)$$

where $\cos(w_i, w'_i)$ denotes the cosine similarity between the vector space embeddings of two words w_i and w'_i , and $\text{Resnik}(w_i, w'_i)$ expresses the information content of the lowest subsumer of $\{w_i, w'_i\}$ in WordNet. For example in Figure 2, the semantic compatibility score incorporates lexical similarities $\text{Sem}_{Lex}(\text{dog}, \text{ship})$, etc.

Compatibility score between semantic relations (Sem_{Pair}) is defined by adding two components: $PairSim$ and $Analogy$ that compute how semantic relations between pairs of words are preserved in

⁴For the ease of notation, we represent the embedding of the words with w_i as well.

the new story:

$$PairSim = \cos(w_i, w_j) * \cos(w'_i, w'_j) \quad (7)$$

$$Analogy = \cos(w'_i + w_j - w_i, w'_j) \quad (8)$$

$PairSim$ preserves the similarity between pairs of words $\{w_i, w_j\}$ in s and the corresponding pair $\{w'_i, w'_j\}$ in the new story s' . Intuitively, if w_i and w_j are semantically close to each other, we would like the corresponding words to be close in the new story as well. For example in Figure 2, ‘dog’ and ‘puppy’ are similar in the original story, we expect the corresponding words ‘ship’ and ‘droid’ to be similar in the new story. The $Analogy$ function, inspired by Mikolov et al. (2013), computes the analogy of w'_j from w'_i given the relationship that holds between w_i and w_j in the vector space. For example in Figure 2, the relation between ‘Sam’ and ‘dog’ is similar to the relation between ‘Luke Skywalker’ and ‘ship’.

5 Decoding

Our decoding process begins by first identifying the content words w_i (nouns, verbs, adjectives and named entities) in the original problem s that will be considered as *initial points* for rewriting. For each of these lexical classes we extract the top- k most thematic words and trivial noun compounds from the theme t . For example, in Figure 2, candidate nouns are: ‘ships’, ‘robots’, ‘droids’, etc., and for verbs: ‘blast’, ‘soar’, ‘command’, etc. Recall that the space of candidate rewrites is large, prohibiting an exhaustive enumeration. We therefore do approximate search with a beam by considering simultaneously all possible *paths* that start at the different initial points. At each step the decoder considers an additional rewrite from the list of candidates, adds it to the existing hypothesis path, and scores it according to function \mathcal{R} (Equation 2).

All the counterpart scores are locally optimal, as they factor over each new word w'_i or pair of $\{w'_i, w'_j\}$, where w'_j is a rewrite *already* existing in the hypothesis path. At any given step we may recombine hypotheses that share the same prefix hypothesis path, and keep the top scoring one. The process terminates when there are no more rewrites left. We also experimented decoding with a variety of orderings of the text in the original problem s , including left-to-right, and head-first following the

dependency tree of each sentence and then concatenating these linearizations; we observed that considering multiple paths achieves the best performance.

6 Data Collection

For the set of human-authored stories $\{s\}$, we use a corpus of math word problems described in Koncel-Kedziorski et al. (2016). We select a subset of 150 problems targeting 5th and 6th grade levels, all of which involve a single equation in one variable. These problems have 2.7 sentences and 29.4 words on average, 12.6 of which are considered content words by our system. In order to tune and evaluate our model, we collect a corpus of human-authored rewrites produced by workers from Amazon Mechanical Turk based on two themes: Star Wars, and Adventure Time (a children’s cartoon).

We experimented with different ways of helping to define the theme for the workers, including offering automatically generated word clouds or enforcing that a response includes one of several keywords. In practice, we have found that using specific cultural elements as themes (such as famous movie or cartoon franchises) attracts workers who already have a strong knowledge of the theme, resulting in higher quality work.

To help explain the rewriting process, we show workers three examples of thematic rewrites with varying degrees of correlation to the original problems. We then show workers a random problem from the original set $\{s\}$ and a corresponding equation for that problem. We instruct the workers to “rewrite” the problem according to the theme, ensuring that their rewritten problem can be solved by the provided equation. The final dataset collection comprises of 450 human-authored rewrites. We collect 3 rewrites for 100 of the original problems for the Star Wars theme (based on the popular Star Wars sequel movies), and 3 rewrites for the rest of the 50 original problems, for the Children Cartoons Theme (CARTOON), based on the Adventure Time TV show. We keep 150 examples from the Star Wars theme for development ($STAR_{dev}$), and the rest 150 for testing ($STAR_{test}$).

We collected the $STAR_{dev}$ and CARTOON data based on workers with the “master” designation and at least 95% approval rating. Then we pro-

ceeded collecting $STAR_{test}$ by a subset of the authors of $STAR_{dev}$ who self-identify as theme experts and whose quality of work is manually confirmed.

7 Experiments

7.1 Setup

Implementation Details We pre-process the themes using the Stanford CoreNLP tools (Manning et al., 2014) for tokenization, Named Entity Recognition (Finkel et al., 2005), and dependency parsing (Chen and Manning, 2014). For calculating salience scores, we use the ScriptBase dataset of movie scripts (Gorinski and Lapata, 2015). The Star Wars theme is constructed from the available script, roughly 7300 words. The Cartoon theme is constructed from fan-authored scripts of the first 10 episodes of the show (Springfield, 2016) totaling 1370 words. Since our thematic options are taken from arbitrary text, we use the lists of offensive terms published by The Racial Slur database (Database, 2016) and FrontGate Media (Media, 2016) to filter out offensive content. To prohibit overgeneration, we forbid the transformation of stop words or math-specific words (Survivors, 2013; Koncel-Kedziorski et al., 2015b).

For syntactic compatibility score Syn (Equation 4) we use the English Fiction subset of the Google Syntactic N-grams corpus (Goldberg and Orwant, 2013) and train a 3-gram language model using KenLM (Heafield, 2011). For Sem_{Lex} , $PairSim$ and $Analogy$ (Equations 6-8) we use the pretrained word embeddings of Levy and Goldberg (2014). These embeddings are trained using dependency contexts rather than windows of adjacent words, allowing them to capture functional word similarity. Finally, we tune the parameters of our model (Equation 2) on the development set $STAR_{dev}$ and pick those values⁵ that maximize METEOR score (Denkowski and Lavie, 2014) against 3 human references.

Evaluation We compare two ablated configurations of our method against our full model (FULL): -SYN that only uses semantic and thematicity components and does not incorporate the syntactic compatibility score, -SEM replaces the semantic coher-

⁵We set $\alpha = 0.1$, $\beta = 0.1$ and $\gamma = 1$

| Model | STAR _{dev} | STAR _{test} | CARTOON |
|-------|---------------------|----------------------|---------|
| FULL | 31.82 | 29.16 | 32.08 |
| -SEM | 28.72 | 25.55 | 27.55 |
| -SYN | 31.92 | 29.14 | 32.04 |

Table 1: METEOR results for different configuration of our model on STAR_{dev}, STAR_{test} and CARTOON datasets.

ence score with the simpler $\cos(w_i, w'_i)$, effectively rewriting only single words, and not pairs. We refrained from ablating the thematicity score as it is the core part of our model that drives the rewriting process into a new theme.

We evaluate our method using an automatic metric, and via eliciting human judgments on Amazon Mechanical Turk. For automatic evaluation, we compute the METEOR score, comparing the output of each model for a given problem and theme to the 3 human rewrites we collected, on STAR_{dev}, STAR_{test} and CARTOON. METEOR is a recall-oriented metric, widely used in the MT community; the additional stemming, synonym and paraphrase matching modules make it more applicable for our use, given the nature of our rewriting task.⁶

For human evaluation, we conduct pairwise comparison tests, pairing FULL against a human rewrite (HUMAN), FULL against -SYN, and FULL against -SEM. Participants were given a short description of the theme, and the output of each system. For each test we asked 40 subjects to select which problem they preferred over 5 pairs of outputs; we obtained a total of 200 (5x40) responses for STAR_{test} and CARTOON.

In order to better understand the strengths and weaknesses of the generated stories, we conducted a more detailed human evaluation. 8 participants were presented with the output of the three automatic systems, human rewrites (HUMAN), and a theme. The participants were asked to rate the stories across three dimensions: coherence (how coherent is the text of the problem?), solvability (can elementary school students solve it?), and thematicity (how well does the problem express them?) on a scale from 1 to 5. We collected ratings over 16 outputs from

⁶The average METEOR score comparing 1 annotator against the other 2 is 0.26, indicating that there are diverse correct strategies for solving the rewriting problem.

| Model | STAR _{test} | CARTOON |
|-------|----------------------|---------|
| FULL | 65.0 | 57.9 |
| -SYN | 35.0 | 42.1 |
| FULL | 68.8 | 69.4 |
| -SEM | 31.2 | 30.6 |
| FULL | 17.9 | 10.0 |
| HUMAN | 82.1 | 90.0 |

Table 2: Human evaluation results on pairwise comparisons between FULL and -SYN, and FULL and HUMAN, on STAR_{test} and CARTOON datasets.

| Model | Thematicity | Coherence | Solvability |
|-------|-------------|-----------|-------------|
| HUMAN | 3.7 | 3.175 | 4.025 |
| FULL | 3.7 | 3.025 | 3.9 |
| -SYN | 3.375 | 3.075 | 3.825 |
| -SEM | 3.325 | 2.65 | 3.7 |

Table 3: Human evaluation results for FULL, -SYN, -SEM and HUMAN on thematicity, coherence and solvability on STAR_{test}.

STAR_{test}, resulting in 128 responses.

7.2 Results

Table 1 reports METEOR; we notice that removing the semantic coherence scores in -SEM hurts the performance compared to FULL; this confirms our claim that semantic compatibility is crucial for building coherent stories. On the other hand, -SYN performs similarly to FULL. Closer inspection of the -SYN system’s output reveals a greater diversity in thematic elements as a result of the relaxed syntactic compatibility constraints. Hence it is more likely to have greater overlap with any of the reference rewrites, resulting in higher METEOR scores.

However, a pairwise comparison between FULL and -SYN (Table 2) reveals that human subjects consistently prefer the output of FULL instead of -SYN both for STAR_{test} and CARTOON. Table 2 also reports that HUMAN outperforms the output of the FULL model, and a pairwise comparison of FULL and -SEM which yields a result in line with the METEOR scores.

Table 3 shows the results of the detailed comparison of Thematicity, Coherence, and Solvability. This table clearly shows the strong contribution of the semantic component of our system. The specific contribution of the syntactic component is to pro-

| |
|---|
| Star Wars |
| s_1 . Wendy bought 4 new chairs and 4 new tables for her house. If she spent 6 minutes on each piece furniture putting it together, how many minutes did it take her to finish? |
| s'_1 . Leia bought 4 new ships and 4 new guns for her room. If she spent 6 minutes on each wasteland weapon putting it together, how many minutes did it take her to terminate? |
| s_2 . My car gets 20 miles per gallon of gas. How many miles can I drive on 5 gallons of gas? |
| s'_2 . My cruiser gets 20 miles per gallon of light. How many miles can I drive on 5 gallons of light? |
| s_3 . Tyler had 15 dogs. Each dog had 5 puppies. How many puppies does Tyler now have? |
| s'_3 . Biggs had 15 creatures. Each creature had 5 creatures. How many creatures does Biggs now have? |

Table 4: Examples of the original stories s_i and rewritten math word problems s'_i in Star War theme.

duce overall more solvable and thematically satisfying problems, although it can slightly affect coherence especially when automatic parses fail. Finally, the overall high ratings for human-authored stories across all three dimensions, confirm the high quality of the crowd-sourced stories.

7.3 Qualitative Examples

Table 4–6 shows some problems generated by our method. Recall that since our system needs no annotated thematic training data, we can easily generate from any theme where thematic text is available. To demonstrate this fact, we include generated examples in a Western theme from novels from the Project Gutenberg corpus. Many of the results of our system are very legible, with only minor agreement errors. Coherent, thematic semantic relations are evident in problems such as s'_1 , where ships, guns, and weapons combine to effect the Star Wars theme; this is also evident in s'_5 , where people with western sounding names like Kurt and Madeline trade in cigarettes, an old-fashioned pre-cursor to e-cigarettes.

In some cases, semantic inconsistencies result in weird sounding problems, such as in s'_6 where the main character receives “wheat of grub”. But because of the syntactic compatibility component, our model scores this candidate higher because of the

| |
|---|
| Cartoon |
| s_7 . Dave was helping the cafeteria workers pick up lunch trays, but he could only carry 9 trays at a time. If he had to pick up 17 trays from one table and 55 trays from another, how many trips will he make? |
| s'_7 . Finn was helping the cupboard men pick up candy bottles, but he could only carry 9 bottles at a time. If he had to pick up 17 bottles from one ring and 55 bottles from another, how many swords will he make? |
| s_8 . If books came from all the 4 continents that Bryan had been into and he collected 122 books per continent, how many books does he have from all 4 continents combined? |
| s'_8 . If dances came from all the 4 mountains that Finn had been into and he collected 122 dances per mountain, how many dances does he have from all 4 mountains combined? |
| s_9 . A bucket contains 3 gallons of water. If Derek adds 6.8 gallons more, how many gallons will there be in all? |
| s'_9 . A bottle makes 3 gallons of serum. If Finn adds 6.8 gallons more, how many gallons will there be in all? |

Table 5: Examples of the original stories s_i and rewritten math word problems s'_i in Cartoon theme.

connection between “wheat” and “graze”.

Semantic incoherence is less of a problem in the cartoon theme, where absurd interactions between characters are expected. However, a difficulty for our system is demonstrated in s'_7 , where the physical entity “swords” is substituted for the nominalization of an event “trips”. Improvements to the semantic coherence component could resolve such issues.

Table 7 shows some instances where the rewrite algorithm produces unusable results. An example of under-generation is s'_{10} . Here, too many words are left untouched, resulting in both ungrammaticality and semantic incoherence. In s'_{11} , we witness some limitations of using word vectors. The rare word “Ferris” is not close to anything in the Star Wars theme, and is thus mapped almost arbitrarily to “int” (movie script shorthand for an interior shot). Better treatment of noun compounds and the use of phrase vectors would reduce such errors.

8 Conclusion

We formalized the problem of story rewriting as automatically changing the theme of a text without

| |
|---|
| Western |
| s_4 . Christians father and the senior ranger gathered firewood as they walked towards the lake in the park ... |
| s'_4 . Christian 's partner and the lone sheriff harvested barley as they strolled towards the hip in the orchard ... |
| s_5 . Sally had 27 cards. Dan gave her 41 new cards. Sally bought 20 cards. How many cards does Sally have now? |
| s'_5 . Madeline had 27 cigarettes. Kurt gave her 41 new cigarettes. Madeline bought 20 cigarettes. How many cigarettes does madeline have now? |
| s_6 . For Halloween Megan received 11 pieces of candy from neighbors and 5 pieces from her older sister. If she only ate 8 pieces a day, how long would the candy last her? |
| s'_6 . For Halloween Madeline received 11 wheat of grub from proprietors and 5 wheat from her nameless partner. If she only grazed 8 wheat a day, how long would the grub last her? |

Table 6: Examples of the original stories s_i and rewritten math word problems s'_i in Western theme.

altering the underlying story and developed an approach for rewriting algebra word problems where the rewriting model optimized for a number of measures of overall text coherence. Experiments on a newly gathered dataset demonstrated our model can produce themed texts that are usually solvable.

Future work could improve the thematicity and solvability components by incorporating domain-specific and commonsense knowledge, leveraging information extraction. Additionally, neural network architectures (e.g., LSTMs, seq2seq) can be trained to rewrite coherently with less reliance on brittle syntactic parses. Additionally, we plan to study rewriting in other domains such as children's short stories and extend the model to generate math word problems directly from equations. Finally, we intend to incorporate the generated problems in educational technology and tutoring systems.

Acknowledgments

This research was supported by the NSF (IIS 1616112), Allen Institute for AI (66-9175), Allen Distinguished Investigator Award, DARPA (FA8750-13-2-0008) and a Google research faculty

| |
|---|
| Poor Rewrites |
| s_{10} . It rained 0.9 inches on Monday. On Tuesday, it rained 0.7 inches less than on Monday. How much did it rain on Tuesday? |
| s'_{10} . It blasted 0.9 inches on Monday. On Tuesday, it blasted 0.7 inches less than on Monday. How much did it light on Tuesday? |
| s_{11} . The Ferris wheel in Paradise Park has 14 seats. Each seat can hold 6 people. How many people can ride the Ferris wheel at the same time? |
| s'_{11} . The int grab in chewbacca mesa has 14 areas. Each area can hold 6 troops. How many troops can ride the int grab at the same time? |

Table 7: Examples of the original stories s_i and poorer rewrites s'_i in the Star Wars theme.

award. We thank the anonymous reviewers for their helpful comments.

References

- Leandro Motta Barros and Soraia Raupp Musse. 2007. Planning algorithms for interactive storytelling. *Computers in Entertainment (CIE)*, 5(1):4.
- Lora Bartlett. 2004. Expanding teacher work roles: a resource for retention or a recipe for overwork? *Journal of Education Policy*, 19(5):565–582.
- Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised Learning of Narrative Schemas and Their Participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 602–610. Association for Computational Linguistics.
- David Chen and William B. Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Danqi Chen and Christopher Manning. 2014. A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 740–750, Doha, Qatar, October. Association for Computational Linguistics.
- The Racial Slur Database. 2016. The racial slur database.
- Judy Davis-Dorsey, Steven M Ross, and Gary R Morrison. 1991. The role of rewording and context personalization in the solving of mathematical word problems. *Journal of Educational Psychology*, 83(1):61.
- Michael Denkowski and Alon Lavie. 2014. Meteor Universal: Language Specific Translation Evaluation for

- Any Target Language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*.
- K. Filippova and M. Strube. 2008. Dependency tree based sentence compression. In *Proceedings of the Fifth International Natural Language Generation Conference (INLG)*.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, ACL '05*, pages 363–370, Stroudsburg, PA, USA. Association for Computational Linguistics.
- J. Ganitkevitch, C. Callison-Burch, C. Napoles, and B. Van Durme. 2011. Learning sentential paraphrases from bilingual parallel corpora for text-to-text generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. PPDB: The paraphrase database. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT)*, pages 758–764, Atlanta, Georgia, June. Association for Computational Linguistics.
- Yoav Goldberg and Jon Orwant. 2013. A Dataset of syntactic-Ngrams over Time from a Very Large Corpus of English Books. In *Second Joint Conference on Lexical and Computational Semantics (* SEM)*, volume 1, pages 241–247.
- Philip John Gorinski and Mirella Lapata. 2015. Movie script summarization as graph-based scene extraction. In *Human Language Technologies: The 2015 Annual Conference of the North American Chapter of the ACL*.
- Mark Granroth-Wilding and Stephen Clark. 2016. What happens next? event prediction using a compositional neural network model. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*, Phoenix, Arizona.
- Janis M Hart. 1996. The Effect of Personalized Word Problems. *Teaching Children Mathematics*, 2(8):504–505.
- Kenneth Heafield. 2011. KenLM: Faster and smaller language model queries. In *Proceedings of the Sixth Workshop on Statistical Machine Translation, WMT '11*, pages 187–197, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. Learning to Solve Arithmetic Word Problems with Verb Categorization. In *EMNLP*, pages 523–533.
- Danqing Huang, Shuming Shi, Chin-Yew Lin, Jian Yin, and Wei-Ying Ma. 2016. How well do computers solve math word problems? large-scale dataset construction and evaluation. In *Proceedings of the 2016 North American Chapter of the ACL (NAACL HLT)*.
- William Hwang, Hannaneh Hajishirzi, Mari Ostendorf, and Wei Wu. 2015. Aligning Sentences from Standard Wikipedia to Simple Wikipedia. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- David Kauchak. 2013. Improving text simplification language modeling using unsimplified text data. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*.
- Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Ang. 2015a. Parsing Algebraic Word Problems into Equations. *TACL*, 3.
- Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Ang. 2015b. Parsing algebraic word problems into equations. *Transactions of the Association for Computational Linguistics*, 3:585–597.
- Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. 2016. MAWPS: A Math Word Problem Repository. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT)*.
- Nate Kushman, Yoav Artzi, Luke Zettlemoyer, and Regina Barzilay. 2014. Learning to Automatically Solve Algebra Word Problems. In *ACL*, pages 271–281.
- Benoit Lavoie and Owen Rambow. 1997. A Fast and Portable Realizer for Text Generation Systems. In *Proceedings of the fifth conference on Applied natural language processing*, pages 265–268. Association for Computational Linguistics.
- Michael Lebowitz. 1987. Planning Stories. In *Proceedings of the cognitive science society, Hillsdale*, pages 234–242.
- Omer Levy and Yoav Goldberg. 2014. Dependency-Based Word Embeddings. In *ACL*, pages 302–308.
- Boyang Li, Stephen Lee-Urban, George Johnston, and Mark O. Riedl. 2013. Story generation with crowd-sourced plot graphs. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*.
- Hugo Liu and Push Singh. 2002. MAKEBELIEVE: Using Commonsense Knowledge to Generate Stories. In *AAAI/IAAI*, pages 957–958.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of the Conference of the Association for Computational Linguistics: System Demonstrations (ACL)*, pages 55–60.

- Neil McIntyre and Mirella Lapata. 2009. Learning to Tell Tales: A Data-driven Approach to Story Generation. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-Volume 1*, pages 217–225. Association for Computational Linguistics.
- Neil McIntyre and Mirella Lapata. 2010. Plot Induction and Evolutionary Search for Story Generation. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1562–1572. Association for Computational Linguistics.
- FrontGate Media. 2016. Terms to block.
- James Richard Meehan. 1976. The Metanovel: Writing Stories by Computer. Technical report, DTIC Document.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781*.
- George A Miller. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11):39–41.
- Arindam Mitra and Chitta Baral. 2015. Learning to automatically solve logic grid puzzles. In *EMNLP*.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and evaluation framework for deeper understanding of commonsense stories. In *Proceedings of the 2016 North American Chapter of the ACL (NAACL HLT)*.
- Sarah Petersen and Mari Ostendorf. 2007. Text simplification for language learners: A corpus analysis. In *Proceedings of the Speech and Language Technology in Education Workshop (SLaTE)*.
- Karl Pichotta and Raymond J. Mooney. 2016. Learning statistical scripts with LSTM recurrent neural networks. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*, Phoenix, Arizona.
- Oleksandr Polozov, Eleanor ORourke, Adam M Smith, Luke Zettlemoyer, Sumit Gulwani, and Zoran Popovic. 2015. Personalized Mathematical Word Problem Generation. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*. To appear.
- KA Renninger, L Ewen, and AK Lasher. 2002. Individual interest as context in expository text and mathematical word problems. *Learning and Instruction*, 12(4):467–490.
- Subhro Roy and Dan Roth. 2015. Solving General Arithmetic Word Problems. In *EMNLP*.
- Subhro Roy, Shyam Upadhyay, and Dan Roth. 2016. Equation parsing : Mapping sentences to grounded equations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Roger C Schank and Robert P Abelson. 1977. *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum.
- Min Joon Seo, Hannaneh Hajishirzi, Ali Farhadi, and Oren Etzioni. 2014. Diagram Understanding in Geometry Questions. In *AAAI*.
- Minjoon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, and Clint Malcolm. 2015. Solving Geometry Problems: Combining Text and Diagram Interpretation. In *EMNLP*.
- Shuming Shi, Yuehui Wang, Chin-Yew Lin, Xiaojiang Liu, and Yong Rui. 2015. Automatically Solving Number Word Problems by Semantic Parsing and Reasoning. In *EMNLP*.
- Advait Siddharthan. 2004. Syntactic simplification and text cohesion. *Research on Language and Computation*, 4(1):77–109.
- Springfield. 2016. Adventure time with finn & jake episode scripts.
- Ladder Survivors. 2013. Key words for math problems.
- Scott R Turner. 1993. Minstrel: A Computer Model of Creativity and Storytelling.
- Lucy Vanderwende, Hisami Suzuki, Chris Brockett, and Ani Nenkova. 2007. Beyond sumbasic: Task-focused summarization with sentence simplification and lexical expansion. *Information Processing and Management*.
- David Vickrey and Daphne Koller. 2008. Sentence simplification for semantic role labeling. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*, pages 344–352.
- Kristian Woodsend and Mirella Lapata. 2011a. Learning to simplify sentences with quasi-synchronous grammar and integer programming. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kristian Woodsend and Mirella Lapata. 2011b. Wikisimple: Automatic simplification of wikipedia articles. In *Proceedings of the Association for Advancement of Artificial Intelligence Conference on Artificial Intelligence (AAAI)*, pages 927–932, San Francisco, CA.
- Sander Wubben, Antal Van Den Bosch, and Emiel Krahmer. 2012. Sentence simplification by monolingual

- machine translation. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*, pages 1015–1024.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of Association of Computational Linguistics*.
- Mark Yatskar, Bo Pang, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. 2010. For the sake of simplicity: Unsupervised extraction of lexical simplifications from wikipedia. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT)*.
- Lipu Zhou, Shuaixiang Dai, and Liwei Chen. 2015. Learn to Solve Algebra Word Problems Using Quadratic Programming. In *EMNLP*.
- Zheming Zhu, Delphine Bernhard, and Iryna Gurevych. 2010. A monolingual tree-based translation model for sentence simplification. In *Proceedings of the International Conference on Computational Linguistics (COLING)*.