

Transforming Sequence Tagging Into A Seq2Seq Task

Karthik Raman*, Iftekhhar Naim,
Jiecao Chen, Kazuma Hashimoto, Kiran Yalasangi, Krishna Srinivasan
Google Research, Mountain View

{karthikraman, inaim, chenjiecao, kazumah, kirancy, krishnaps}@google.com

Abstract

Pretrained, large, generative language models (LMs) have had great success in a wide range of sequence tagging and structured prediction tasks. Casting a sequence tagging task as a Seq2Seq one requires deciding the formats of the input and output sequences. However, we lack a principled understanding of the trade-offs associated with these formats (such as the effect on model accuracy, sequence length, multilingual generalization, hallucination). In this paper, we rigorously study different *formats* one could use for casting input text sentences and their output labels into the **input** and **target** (*i.e.*, output) of a Seq2Seq model. Along the way, we introduce a new format, which we show to be both simpler and more effective. Additionally the new format demonstrates significant gains in the multilingual settings – both zero-shot transfer learning and joint training. Lastly, we find that the new format is more robust and almost completely devoid of *hallucination* – an issue we find common in existing formats. With well over a 1000 experiments studying 14 different formats, over 7 diverse public benchmarks – including 3 multilingual datasets spanning 7 languages – we believe our findings provide a strong empirical basis in understanding how we should tackle sequence tagging tasks.

1 Introduction

The advent of powerful generative language models (LMs) such as T5 (Raffel et al., 2019) and mT5 (Xue et al., 2021) have unlocked new opportunities. Recent works have shown great success applying these models to a wide range of sequence tagging and structured prediction tasks (Athiwaratkun et al., 2020; Du et al., 2021; Paolini et al., 2021). The key to leveraging these models is to cast the sequence tagging task as a suitable Seq2Seq task – by transforming the original input

text and output labels into the **input** and **target** (*i.e.*, output) of a Seq2Seq model.

| |
|---|
| Slot-filling example from Athiwaratkun et al. (2020) |
| Input: Add Kent James to the Disney soundtrack |
| Target: Add [Kent James <i>artist</i>] to the [Disney <i>playlist</i>] soundtrack |
| SRL example from Paolini et al. (2021) (predicate = "sold") |
| Input: The luxury auto maker last year [sold] 1,214 cars in the U.S. |
| Target: [The luxury auto maker <i>subject</i>] [last year <i>temporal</i>] sold [1,214 cars <i>object</i>] [in the U.S. <i>location</i>] |

Table 1: Examples of the **TaggedSpans** approach used in current works to transform NLP problems into the Seq2Seq setting.

While the list of NLP problems tackled in this manner are diverse – Named Entity Recognition, Co-reference resolution, Semantic Role Labeling (SRL), Slot-filling, Entity and Relation extraction to name a few – the Seq2Seq transformation used across different works is surprisingly similar (with slight variants). As seen from the examples in Table 1, existing works use the same Seq2Seq transformation – what we refer to as the **TaggedSpans** *format*. More specifically, the *target* sentence resembles the input sentence, except for tagged decorators surrounding the corresponding token spans.

While exciting results have been achieved using this *format*, there has been no principled study understanding this format choice. Towards that end, we make the following contributions in this paper:

- Based on rigorous experiments, we derive insights into the performance of the **TaggedSpans** format (as well as other Seq2Seq transformation formats) along different dimensions.
- We propose a **new Sentinel-based format** that proves to be not only more effective (achieving significantly higher accuracy across all the benchmarks) but also more efficient than the formats used in existing works.
- We perform an in-depth study understanding performance of different formats in the *multilingual learning* setting – with the new **Sentinel-based** format outperforming other

* Corresponding Author

formats by upto **30-40%** in some settings.

- We demonstrate the *robustness* of different formats to factors such as model size, sequence length and decoding strategy.
- Lastly, we show that the existing **TaggedSpans** format (as well as other similar formats) are highly susceptible to hallucinations (Maynez et al., 2020), *i.e.*, the proclivity of large LMs adding, modifying, or deleting input words while generating the output sequence. While the **TaggedSpans** had hallucinated tokens in over **50%** of examples in some settings, the proposed **Sentinel-based** format is virtually **hallucination-free** (< 0.15% hallucinations).

With well over a 1000 experiments involving 14 different formats, 7 datasets (3 multilingual), we believe this to be the most comprehensive study on this topic to-date.

2 Related Work

Several recent papers have shown the strength of the large Seq2Seq LMs for sequence tagging – in the context of Named Entity Recognition (Yan et al., 2021; Qin and Joty, 2022; Paolini et al., 2021), Slot-Labeling for Semantic Parsing (Krishnan et al., 2021; Ahmad et al., 2021; Du et al., 2021; Athiwaratkun et al., 2020), and Semantic Role Labeling (Daza and Frank, 2018). Yan et al. (2021) proposed a general Seq2Seq framework for NER tasks, both for disjoint spans and nested overlapping spans. Ahmad et al. (2021) showed that Seq2Seq models can significantly outperform the encoder-only models for slot-filling tasks. Raffel et al. (2019) converted a wide range of sequence tagging tasks (e.g., answer span extraction for Question-Answering) into Seq2Seq problems. All these methods, however, apply fairly similar text-to-text formats in that: they (1) prepend the input with a task prefix and leave rest of it mostly unchanged and (2) use the **TaggedSpans** format (or slight variants) for the *target* (output) sequence *i.e.*, copy over input tokens and interlace with the tags for the associated spans. In this paper, we systematically evaluate this approach, and demonstrate its relative strengths and weaknesses across a large collection of benchmarks.

Recent papers investigated the impact of carefully hand-crafted input prompts for large language models and showed different prompt templates can significantly change the prediction accuracy, espe-

cially in zero-shot or few-shot settings (Reynolds and McDonnell, 2021; Wang et al., 2021; Zhong et al., 2021). For example, Zhong et al. (2021) has demonstrated that different prompts describing classification labels can have large impact on the zero-shot classification accuracy. Similar nuances are likely at play for output sequence formats as well, but there hasn't been enough systematic study to understand the impact of different output format choices. We present empirical evaluations of different input/output formats for sequence tagging tasks, and how they impact sequence tagging accuracy in English-only, multilingual, and zero-shot multilingual settings.

3 The Seq2Seq Transformation

Applying pretrained generative models on these sequence labeling / structured predictions problems requires creating the *input* and *target* strings from the original input sentences and labels. The examples in Table 1 indicate how previous works have tackled this transformation. In this section, we formalize this transformation and consider different alternatives to the existing **TaggedSpans** format in how we can create the *input* and *target* strings. Note that to simplify the exposition, we assume that the original labels can be represented using the standard **BIO** (Begin-Inside-Outside) notation – as is conventional for most of these tasks.

3.1 The Tagged Spans Format

This is the approach used by existing works for transforming into a Seq2Seq task. Here the target sentence is an interlaced amalgamation of the input and the associated tag labels. Consider the example shown in Table 2. Here the "**<ARTIST>**" tag precedes the token "Kent" to indicate the start of an "ARTIST" span, and is closed with the "**</>**" marker after the token "James". In other words, salient token spans are explicitly surrounded by the associated tags ¹

An astute reader may notice that the example from Table 2 also contains tagged decorators for the "**O**" (Outside) class as well unlike existing works. This was done to simplify exposition and demonstrate differences from the other formats we study. While we do discuss dropping these Outside tags in Sec. 3.4, it is worth noting that in our empirical

¹Note that, to simplify exposition of different formats, we use a slightly different punctuation symbol for the tags than the square brackets used in the examples in Table 1. This choice does not affect performance though.

| Format | Sample Input & Target | | | | | | | |
|----------------------------------|------------------------------------|-----------------------|-----------------------|-----------------|------------------|-----------------------|-------------------------|--|
| Tagged Spans | <i>input:</i> Add | Kent | James | to | the | Disney | soundtrack | |
| | <i>target:</i> <O> Add </> | <ARTIST> Kent | James </> | <O> to </> | <O> the </> | <PLAYLIST> Disney </> | <O> soundtrack </> | |
| Input + Tag | <i>input:</i> Add | Kent | James | to | the | Disney | soundtrack | |
| | <i>target:</i> <O> Add | <ARTIST> Kent | <I-ARTIST> James | <O> to | <O> the | <PLAYLIST> Disney | <O> soundtrack | |
| Tag Only | <i>input:</i> Add | Kent | James | to | the | Disney | soundtrack | |
| | <i>target:</i> O | ARTIST | I-ARTIST | O | O | PLAYLIST | O | |
| Sentinel + Tag | <i>input:</i> <extra_id_0> Add | <extra_id_1> Kent | <extra_id_2> James | <extra_id_3> to | <extra_id_4> the | <extra_id_5> Disney | <extra_id_6> soundtrack | |
| | <i>target:</i> <extra_id_0> O | <extra_id_1> ARTIST | <extra_id_2> I-ARTIST | <extra_id_3> O | <extra_id_4> O | <extra_id_5> PLAYLIST | <extra_id_6> O | |
| Extractive Tagged Spans | <i>input:</i> Add | Kent | James | to | the | Disney | soundtrack | |
| | <i>target:</i> <O> Add | <ARTIST> Kent | James | | | <PLAYLIST> Disney | | |
| Extractive Sentinel + Tag | <i>input:</i> <extra_id_0> Add | <extra_id_1> Kent | <extra_id_2> James | <extra_id_3> to | <extra_id_4> the | <extra_id_5> Disney | <extra_id_6> soundtrack | |
| | <i>target:</i> <extra_id_1> ARTIST | <extra_id_2> I-ARTIST | | | | <extra_id_5> PLAYLIST | | |

Table 2: Example illustrating how the *input* and *target* differ across the high-level formats studied here.

analyses we found that having these Outside tag markers actually helps improve performance for the **TaggedSpans** and other formats (Table 5).

Why this format? Given how common this format is in existing works, one may suspect that this choice originates from some meticulous study. However until this paper, there have been no previous efforts exploring other choices, or understanding the relative benefits & drawbacks of this format.

In particular, one source of concern is the need to **faithfully** repeat the input as part of the output, **in addition to** identifying the correct tags for the token spans. This makes the learning task significantly harder (as evidenced by the diminished empirical performance observed in Sec. 6). Furthermore this copying of text also makes generalization in other languages harder (seen in Sec. 7). Additionally the model is harder to interpret since the log-likelihood scores are a combination of the generated input spans and their associated output tags. Lastly, this copying makes the model susceptible to hallucinations (in over 50% of examples in some settings – Sec. 9.2).

3.2 Other Formats

A slight variant of the **TaggedSpans** format is the **Input + Tag** format. As seen from the example in Table 2, the output here is an interlacing of the input tokens with the BIO labels *i.e.*, the input token is preceded by its BIO label. Unfortunately, this format shares the same drawbacks due to the need to accurately repeat the input.

To remedy the issues caused by requiring the input in the target, one solution is to drop the input tokens altogether. This format – which we call the **Tag-Only** format – is the simplest yet as it simply drops the input tokens from the **Input + Tag** format’s target. While this remedies some issues, it makes the learning problem significantly harder as the model now needs to also track token indices implicitly – a task even harder in non-English languages (as we later demonstrate empirically).

3.3 The proposed Sentinel+Tag Format

To avoid these issues with the aforementioned formats, we would like a format that:

- Avoids having to repeat the input to simplify decoding and avoid hallucination.
- Allows for easier way of tracking tokens (and indices) and associating them with their tags.
- Does so in a language-agnostic manner so as to enable multilingual generalization.

To achieve this, the **key insight** we had was *to slightly modify the input to simplify tracking tokens*. In particular, we proposed using text-independent, language-independent **Sentinel tokens** – see example in Table 2. Introducing these sentinels in the input, allows us to avoid repeating the original text tokens in the target, and instead use the sentinel tokens² associated with each input token to help learn the corresponding BIO tag. By virtue of the resulting targets being independent of the text tokens and language-agnostic, the **Sentinel+Tag** format can generalize well and be largely immune to hallucination. Furthermore, unlike the **Tag-Only** approach, the learning problem is simplified as the sentinels help uniquely ground tokens (spans) and the associate tags. Lastly, while the input is slightly longer, this is completely offset by the massive reduction in output sequence lengths – see Table 13.

While questions may arise about the effect of imperfect tokenization (e.g. in non-space separated languages like Thai), our empirical study in Section 7 spanning 7 languages aims to answer this.

3.4 Variants and Simplifications

We also investigated further variants of these formats, that simplify the target texts. One such simplification pertains to the *Inside* tags. Rather than having to produce the complete tag (e.g. "I-ARTIST"), we also evaluated variants where the target simply uses the tag "I" regardless of the opened tag.

²In this work we use one sentinel per token. In particular, the sentinel we used for token number *k* is <extra_id_k> – since these are present in the mT5 vocabulary as single tokens.

| Dataset | Lang | # Train | # Valid | # Test | Token/ex | Tagged span/ex | % tokens tagged | # Tag classes | Tag Entropy |
|-------------|------|---------|---------|--------|----------|----------------|-----------------|---------------|-------------|
| mATIS | en | 4478 | 500 | 893 | 11.28 | 3.32 | 36.50 | 79 | 3.888 |
| | hi | 540 | 60 | 893 | 11.39 | 3.14 | 34.21 | 65 | 3.823 |
| | tr | 540 | 60 | 715 | 10.80 | 3.03 | 33.32 | 63 | 3.847 |
| SNIPS | en | 13084 | 700 | 700 | 9.00 | 2.60 | 51.33 | 39 | 4.857 |
| MovieTrivia | en | 7005 | 811 | 1953 | 20.31 | 2.95 | 64.74 | 12 | 2.867 |
| Movies | en | 8722 | 1019 | 2441 | 10.23 | 2.19 | 38.65 | 12 | 3.145 |
| Restaurant | en | 6845 | 789 | 1516 | 9.25 | 2.01 | 38.02 | 8 | 2.809 |
| mTOP | en | 15667 | 2235 | 2235 | 7.61 | 1.71 | 44.18 | 74 | 4.681 |
| | de | 13424 | 1815 | 3549 | 7.97 | 1.70 | 42.73 | 74 | 4.613 |
| | es | 10933 | 1527 | 2998 | 9.01 | 1.63 | 42.95 | 70 | 4.522 |
| | fr | 11814 | 1577 | 3193 | 9.03 | 1.61 | 39.08 | 73 | 4.604 |
| | hi | 10933 | 2012 | 2789 | 8.28 | 1.63 | 38.71 | 72 | 4.714 |
| | th | 10759 | 1671 | 2765 | 9.21 | 1.58 | 48.37 | 73 | 4.608 |
| mTOD | en | 30506 | 4178 | 8615 | 7.26 | 1.69 | 43.46 | 15 | 2.434 |
| | es | 3616 | 1981 | 3043 | 7.46 | 1.55 | 52.43 | 11 | 2.323 |
| | th | 2154 | 1235 | 1692 | 7.46 | 1.55 | 52.43 | 11 | 2.323 |

Table 3: Statistics for the different datasets, including split sizes as well as per-examples averages of tokens/example, tagged spans/example. A higher % *tokens tagged* value indicates more tokens being part of tagged spans. The number of tag classes and tag entropy indicate the difficulty of identifying a tag for any given span. Note that the last 5 columns are computing using only the training sets.

When this simplification is used, we refer to it as **Simplified Inside (SI)**.

Another such simplification pertains to the *Outside* tags. Rather than emitting these tags, we also evaluated variants where the Outside tag was omitted from the target. Examples for these variants can be found in Table 17 in the Appendix. We refer to this simplification as **Simplified Outside (SO)**.

Note that the rest of the paper assumes non-nested spans for simplicity. However most of these above mentioned formats can easily support nested or overlapping spans (Example: By having multiple tags for each opened span after the corresponding token / sentinel in the **Tag + Input** or **Sentinel + Tag** approaches).

3.5 Formats for Extractive Applications

The approaches discussed so far allowed us to map the predicted label to the exact span occurrence in the input string. However in some scenarios this input-alignment is not necessary. In particular, for certain tasks and applications we primarily care about extracting only the labeled tags and associated phrases. In such cases we can further simplify the target as seen in the target of the **Extractive Tagged Spans** (abbrv. **ETS**) in Table 2. This format, which is used in the T5 paper (Raffel et al., 2019), simply outputs only the labeled spans. While this simplifies the task, it comes at the cost of not being able to map spans precisely³ back to the input for non-extractive tasks.

³While some heuristic matching can be used, it is hard to do so in a multilingual-friendly manner. Additionally multiple mentions with different semantic meanings can be problematic (e.g. “who plays harry potter in the harry potter movies”)

We can analogously define **Extractive Sentinel+Tag** (abbrv. **ES+T**) formats for this extractive setting (see last row of Table 2). Note that, unlike the existing ETS approaches, the sentinel approach allows us to precisely map labels back to the original input token indices and thus can be used seamlessly for **both** extractive and non-extractive tasks. However, to simplify exposition of the paper, we separate the ES+T and ES+T **Simplified** (abbrv. **ES+T(S)**) – which drops Inside tags (see example in Table 17) – approaches and only discuss these with the other extractive formats.

4 Datasets

To gain a comprehensive understanding of how these different formats perform, we ran experiments on 7 different public datasets – 4 monolingual and 3 multilingual spanning 7 languages from different language groups. The monolingual datasets include: SNIPS (Coucke et al., 2018), MIT Movie Trivia⁴, MIT Movie corpus, and MIT Restaurant corpus. The multilingual datasets include: mATIS (Upadhyay et al., 2018) – a multilingual generalization of the ATIS dataset (Price, 1990), mTOP (Li et al., 2021), and mTOD (Schuster et al., 2019). These datasets span different domains – such as cinema, dining, travel, dialog . . . – and contain different kinds of texts – including queries, reviews, and conversations. For instance, mATIS, SNIPS, mTOP, and mTOD contain conversational queries or dialog-like texts, and involve slot labelling for semantic parsing. The MIT Movie

⁴All the MIT datasets were downloaded from <https://groups.csail.mit.edu/sls/downloads/>.

Trivia, MIT Movies, and MIT Restaurants datasets, on the other hand, are widely used for Named Entity Recognition (NER).

As seen from the statistics listed in Table 3, these datasets tend to have different properties. MIT Movie Trivia dataset tends to have longer sentences, longer tagged spans, and more named entities but fewer number of output tag categories. On the other hand, mATIS and mTOP have relatively shorter inputs but contain over 70 different output tags. The mTOD dataset has shorter inputs and fewer output categories and is the largest training set (> 30K for English), making it a relatively easier task. Finally, the high entropy of tags (as in the mTOP and SNIPS) indicates a higher degree of difficulty with fairly ambiguous categories. Together these datasets provide a broad and diverse collection of public benchmarks for evaluating different Seq2Seq formats.

5 Empirical Setting

We studied different input/output formats using the mT5 pretrained models (Xue et al., 2021), which have been used to obtain state-of-the-art results on multiple public benchmarks. In particular, we used the **Base**-sized model for the majority of our experiments as it enabled us to maximize our experimentation given our compute budget, while still providing for a very powerful initialization. However, results and trends largely stay the same for other model sizes as we discussed in Sec 9.

In all cases, we trained these models using the T5X⁵ code for 5000 steps – selecting the checkpoint with the highest validation set sequence accuracy. As recommended in the T5 paper we used the default learning rate of 0.001 for all our experiments, and used a batch size of 128. We appended EOS tokens at the ends of all input and target formats for consistency. Lastly, we used an input sequence length of 128 and target sequence length of 256, while also using packing. The impact of training sequence length on accuracy and speed is discussed further in Sec 9.

To evaluate the different Seq2Seq formats, we primarily relied on two metrics:

1. **Perfect:** This reflects the % of times, an example is parsed perfectly *i.e.*, all tagged spans are identified correctly with the right tag.
2. **F1:** Using CoNLL style evaluation of tagged spans, we report Micro F1 (Macro-F1 trends

were consistent with Micro-F1 findings).

One important aspect to note is that the outputs of all the formats are different. Hence they need to be evaluated differently. In particular, formats that repeat the input (**Tagged Spans**, **Input + Tag**) are more prone to hallucination. To try and decouple this effect where possible, we evaluated the **Tagged Spans** and **Input + Tag** format by checking the token indices of the tagged spans – rather than the generated token strings. While this leads to slightly higher metrics, it allows for a more fair and hallucination-free evaluation against approaches like **Tag-Only** and **Sentinel + Tag**.

6 Performance on English-benchmarks

The first question we looked to answer was: *How well do these Seq2Seq approaches perform compared with the previous encoder-only approaches?* To do so, as a baseline we used the previous benchmark setting mBERT (Devlin et al., 2019) model – which is an encoder-only model having a similar size as the T5 base model’s encoder. We evaluated the models on the simpler *English-only* setting to start with. The results on all 7 English benchmarks can be found in Table 4. Perhaps unsurprisingly, we find that the Seq2Seq approaches significantly outperform the encoder-only mBERT approach on all datasets. This result is consistent with prior papers (Ahmad et al., 2021; Raffel et al., 2019). All but the simplest Tag-Only Seq2Seq format outperform the mBERT baseline on all the datasets. This conclusively demonstrates the value of modeling these problems using the encoder-decoder models.

The next question we looked to answer was: *How do the different basic formats described in Sec 3 compare against each other?* In particular, we wanted to understand whether not repeating the input tokens hurts performance. As seen from the third to sixth columns in Table 4, the proposed **Sentinel+Tag** ends up outperforming all previous formats. In fact, the **Sentinel+Tag** outperforms all the other basic approaches in 6 out of the 7 benchmarks and is significantly better when averaging out the results across all runs.

These results also show that while the **Tag Only** approach is quite competitive, using either the sentinels or input tokens leads to better results. Additionally, we also find that the **Input + Tag** is actually marginally better than the more commonly used Tagged Spans approach (though this gap is too narrow to be statistically significant).

⁵<https://github.com/google-research/t5x>

| Dataset | mBERT (and SI) | | Tag Only | | Input + Tag | | Tagged Spans | | Sentinel + Tag | | Tag Only (SI) | | Input + Tag (SI) | | Sentinel + Tag (SI) | |
|--------------|------------------|------------------|-----------------|-----------------|-----------------|------------------------|------------------------|-----------------|-------------------------------------|-------------------------------------|-----------------|-----------------|------------------|------------------------|-------------------------------------|-------------------------------------|
| | Perfect | F1 | Perfect | F1 | Perfect | F1 | Perfect | F1 | Perfect | F1 | Perfect | F1 | Perfect | F1 | Perfect | F1 |
| mATIS(en) | 85.11 (84.88) | 93.65 (93.25) | 88.17 ± 0.37 | 93.83 ± 0.28 | 88.73 ± 0.46 | 95.40 ± 0.21 | 88.95 ± 0.43 | 95.55 ± 0.16 | 89.81 [†] ± 0.09 | 95.91 [†] ± 0.05 | 88.65 ± 0.78 | 94.77 ± 0.37 | 88.76 ± 0.56 | 95.32 ± 0.34 | 90.07 [†] ± 0.47 | 95.96 [†] ± 0.13 |
| SNIPS | 86.57 (88.29) | 93.66 (94.34) | 85.38 ± 0.54 | 92.24 ± 0.42 | 87.05 ± 0.36 | 94.30 ± 0.20 | 87.14 ± 0.31 | 94.28 ± 0.09 | 90.14 [†] ± 0.70 | 95.47 [†] ± 0.40 | 89.10 ± 0.36 | 94.73 ± 0.28 | 87.76 ± 0.78 | 94.47 ± 0.29 | 89.81 [†] ± 0.07 | 95.53 [†] ± 0.04 |
| Movie Trivia | 29.80 (29.60) | 64.69 (64.40) | 22.65 ± 3.20 | 59.10 ± 2.48 | 39.84 ± 0.47 | 72.58 ± 0.20 | 40.21 ± 0.09 | 72.19 ± 0.22 | 39.43 ± 0.68 | 72.67 ± 0.58 | 34.27 ± 0.23 | 68.41 ± 0.25 | 39.73 ± 0.30 | 72.76 ± 0.14 | 39.85 ± 0.58 | 73.01 ± 0.27 |
| Movies | 60.59 (61.41) | 80.87 (81.40) | 71.24 ± 0.18 | 86.79 ± 0.06 | 71.94 ± 0.27 | 87.26 ± 0.25 | 71.42 ± 0.36 | 86.95 ± 0.13 | 73.17 ± 0.54 | 87.79 ± 0.34 | 72.06 ± 0.58 | 87.18 ± 0.30 | 72.12 ± 0.02 | 87.47 ± 0.13 | 72.74 ± 0.95 | 87.56 ± 0.48 |
| Restaurant | 50.06 (50.79) | 70.25 (71.03) | 61.28 ± 0.19 | 78.63 ± 0.26 | 62.12 ± 0.14 | 80.14 ± 0.19 | 61.68 ± 0.69 | 79.87 ± 0.35 | 62.53 ± 0.65 | 80.02 ± 0.46 | 61.87 ± 0.19 | 79.38 ± 0.07 | 62.42 ± 0.22 | 80.42 ± 0.15 | 62.93 ± 0.30 | 80.39 ± 0.31 |
| mTOP(en) | 81.16 (81.21) | 88.48 (88.69) | 81.98 ± 0.34 | 88.50 ± 0.34 | 84.34 ± 0.45 | 90.80 ± 0.39 | 84.22 ± 0.66 | 90.82 ± 0.46 | 85.58 [†] ± 0.58 | 91.68 [†] ± 0.42 | 83.06 ± 0.48 | 89.59 ± 0.31 | 84.88 ± 0.21 | 91.20 ± 0.06 | 86.56 [†] ± 0.69 | 92.28 [†] ± 0.44 |
| mTOD(en) | 92.05 (93.32) | 95.71 (97.12) | 92.53 ± 0.10 | 95.97 ± 0.05 | 92.39 ± 0.07 | 96.01 ± 0.04 | 92.13 ± 0.08 | 95.90 ± 0.08 | 92.69 ± 0.13 | 96.21 ± 0.08 | 93.09 ± 0.06 | 96.34 ± 0.04 | 92.66 ± 0.03 | 96.17 ± 0.01 | 93.19 ± 0.17 | 96.42 ± 0.08 |
| Average | 69.33 (70.05) | 83.90 (84.32) | 71.89 ± 0.25 | 85.01 ± 0.22 | 75.20 ± 0.18 | 88.07 ± 0.11 | 75.11 ± 0.11 | 87.94 ± 0.03 | 76.19 [†] ± 0.19 | 88.54 [†] ± 0.16 | 74.59 ± 0.15 | 87.20 ± 0.11 | 75.48 ± 0.19 | 88.26 ± 0.08 | 76.45 [†] ± 0.26 | 88.73 [†] ± 0.16 |

Table 4: Results for how the different Seq2Seq formats perform on the English benchmarks. Metrics are averaged over 3 runs and reported (with standard deviation). mBERT results include both w/o and with Simplified Inside. The [†] symbol indicates 99+% significant improvement (per the z-test) against **all** non-sentinel approaches.

| Dataset | Input + Tag (SI, SO) | Tagged Spans (SO) | Sentinel + Tag (SI, SO) |
|--------------------------------|----------------------|-------------------|-------------------------|
| English-only benchmarks | | | |
| Average | 74.35 | 74.85 | 75.68 |
| (Δ vs. non-SO) | -1.13 | -0.26 | -0.77 |
| Multilingual Zero-Shot | | | |
| Average (non-en) | 21.33 | 21.05 | 44.76 |
| (Δ vs. non-SO) | -15.14 | -10.24 | -2.35 |
| Multilingual Joint | | | |
| Average (non-en) | 73.27 | 72.93 | 79.01 |
| (Δ vs. non-SO) | -1.62 | 0.33 | 0.23 |

Table 5: **Perfect** metric scores for the variant Seq2Seq formats that modify or drop the Outside tag (along with the performance difference due to this simplification).

6.1 Do variants & simplifications work?

We next tried to understand performance on making tweaks to the formats as discussed in Sec 3.4. The last 3 columns of Table 4, includes results for variants that simplify the Inside tag to just "I. Results for the mBERT baseline with the same simplification are also provided. Overall the **Sentinel + Tag (SI)** resulted in the aggregate best performance across all formats on the 7 English benchmarks. The small gap between SI and non-SI variants makes sense given the autoregressive nature of the decoder (that can generate the full Inside tags by attending to any opened Begin tag that need to be continued). We should note that the **SI** variants have one significant advantage. As seen in Table 13, these formats have significantly shorter target sequences across all datasets, resulting in faster training and inference.

Another simplification we evaluated involved the *Outside* tags. Interestingly, as seen in Table 5, results worsen when using this simplification – for all formats – in nearly all settings. This indicates that these Seq2Seq models find it easier to produce a consistent prediction for all tokens rather than having to skip tokens. More details and analysis

for this simplification can be found in Appendix A.

7 Multilingual Capabilities

A major benefit of large LMs is their generalization capabilities in the multilingual setting. In this section, we evaluated the different formats in the more challenging multilingual setting. Specifically we ran experiments on 3 benchmarks (spanning 7 different languages) in two different settings:

- **Joint Multilingual:** Models are trained by combining data for all languages. The weights for all languages are equal *i.e.*, the models see roughly the same number of examples in all languages. For simplicity, checkpoint selections happens per-language using that language’s validation set (though we did not find this choice to affect results significantly).
- **Zero-shot:** Here we train the model using only English data, and evaluate on all languages. This is the more challenging setting for evaluating multilingual generalization.

As seen in the results for the zero-shot setting in Table 6, the findings from the previous section are only further emphasized with the **Sentinel+Tag** approaches significantly outperforming all the other approaches. For instance, when looking at the *Perfect* metrics averaged across all 9 non-English test sets (2 mATIS, 5 mTOP, 2 mTOD), we find **Sentinel+Tag (SI)** format greatly outperforms all non-sentinel formats with the next best alternative (**Input+Tag**) being more than **10.5%** absolute worse (**47.1%** vs **36.5%**, a **30%** relative increase in accuracy). Furthermore it is nearly **16%** better (a **47%** increase) than the current standard *i.e.*, – **Tagged Spans (31.3 vs 47.1)**.

These gains are all the more noteworthy when compared to the performance of the **Tag Only** (no sentinels) and **Input+Tag** (input token instead of

| Dataset | Tag Only | | Tag Only (SI) | | Input + Tag | | Input + Tag (SI) | | Tagged Spans | | Sentinel + Tag | | Sentinel + Tag (SI) | |
|---------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|------------------|-----------------|-----------------|-----------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 |
| mATIS (3 langs) | 36.71 ± 3.31 | 52.79 ± 4.22 | 36.97 ± 1.09 | 53.30 ± 1.79 | 38.96 ± 1.45 | 62.62 ± 2.14 | 39.47 ± 1.22 | 63.42 ± 1.38 | 35.58 ± 1.10 | 58.55 ± 2.19 | 43.62[†] ± 1.99 | 68.54[†] ± 1.89 | 42.49 [†] ± 2.15 | 66.67 [†] ± 2.31 |
| mTOP (6 langs) | 43.41 ± 1.02 | 51.48 ± 1.06 | 44.98 ± 0.66 | 53.04 ± 0.73 | 45.60 ± 1.01 | 53.68 ± 1.02 | 46.57 ± 0.35 | 54.81 ± 0.26 | 44.00 ± 0.91 | 54.53 ± 0.89 | 57.14 [†] ± 1.06 | 69.65 [†] ± 1.08 | 59.31[†] ± 0.98 | 71.22[†] ± 1.34 |
| mTOD (3 langs) | 66.42 ± 1.41 | 74.27 ± 1.51 | 67.78 ± 1.08 | 75.42 ± 0.76 | 65.11 ± 1.18 | 73.75 ± 1.01 | 65.55 ± 0.55 | 74.74 ± 0.42 | 58.73 ± 1.84 | 68.71 ± 1.56 | 66.92 ± 1.49 | 76.42 [†] ± 0.94 | 70.07 ± 1.80 | 78.51[†] ± 1.11 |
| Average | 47.49 ± 1.09 | 57.50 ± 1.31 | 48.68 ± 0.39 | 58.70 ± 0.81 | 48.82 ± 0.37 | 60.93 ± 0.46 | 49.54 ± 0.46 | 61.94 ± 0.33 | 45.58 ± 0.57 | 59.08 ± 0.82 | 56.20 [†] ± 0.95 | 71.07 [†] ± 0.79 | 57.79[†] ± 1.38 | 71.90[†] ± 1.40 |
| Average (non-en) | 34.17 ± 1.27 | 45.75 ± 1.65 | 35.51 ± 0.41 | 47.13 ± 0.98 | 35.58 ± 0.43 | 49.89 ± 0.58 | 36.47 ± 0.63 | 51.18 ± 0.45 | 31.29 ± 0.64 | 47.41 ± 1.03 | 45.20 [†] ± 1.23 | 63.24 [†] ± 1.02 | 47.11[†] ± 1.73 | 64.25[†] ± 1.80 |

Table 6: Multilingual Results in the **Zero-shot** setting (mean and standard deviation over 3 runs). Detailed (per-language) results can be found in Table 23 in the Appendix. [†] is 98% significance against **all** non-sentinel methods.

| Dataset | Tag Only | | Tag Only (SI) | | Input + Tag | | Input + Tag (SI) | | Tagged Spans | | Sentinel + Tag | | Sentinel + Tag (SI) | |
|---------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|------------------|-----------------|-----------------|-----------------|------------------------------------|------------------------------------|------------------------------------|------------------------------|
| | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 |
| mATIS (3 langs) | 68.89 ± 1.09 | 82.90 ± 0.74 | 70.13 ± 0.46 | 83.94 ± 0.50 | 71.03 ± 0.47 | 87.22 ± 0.39 | 70.38 ± 0.54 | 86.69 ± 0.26 | 69.14 ± 0.21 | 86.53 ± 0.23 | 73.26 [†] ± 1.17 | 88.37[†] ± 0.37 | 73.43[†] ± 0.74 | 88.26 [†] ± 0.53 |
| mTOP (6 langs) | 69.43 ± 0.64 | 77.31 ± 0.58 | 73.40 ± 0.20 | 81.00 ± 0.12 | 75.21 ± 0.31 | 83.10 ± 0.26 | 76.22 ± 0.34 | 83.90 ± 0.24 | 73.12 ± 0.37 | 82.10 ± 0.33 | 81.23[†] ± 0.23 | 88.26[†] ± 0.16 | 80.74 [†] ± 0.57 | 87.85 [†] ± 0.42 |
| mTOD (3 langs) | 88.81 ± 0.14 | 92.70 ± 0.09 | 89.58 ± 0.26 | 93.38 ± 0.18 | 89.61 ± 0.22 | 93.53 ± 0.17 | 89.63 ± 0.17 | 93.50 ± 0.10 | 89.19 ± 0.22 | 93.30 ± 0.15 | 89.97[†] ± 0.08 | 93.72 ± 0.07 | 89.92 [†] ± 0.12 | 93.64 ± 0.12 |
| Average | 74.14 ± 0.18 | 82.55 ± 0.13 | 76.63 ± 0.08 | 84.83 ± 0.05 | 77.76 ± 0.09 | 86.74 ± 0.15 | 78.11 ± 0.25 | 87.00 ± 0.14 | 76.15 ± 0.12 | 86.01 ± 0.16 | 81.42[†] ± 0.38 | 89.65[†] ± 0.14 | 81.21 [†] ± 0.10 | 89.40 [†] ± 0.10 |
| Average (non-en) | 70.37 ± 0.21 | 79.56 ± 0.20 | 73.22 ± 0.04 | 82.25 ± 0.05 | 74.60 ± 0.14 | 84.56 ± 0.21 | 74.89 ± 0.32 | 84.81 ± 0.17 | 72.60 ± 0.14 | 83.65 ± 0.20 | 78.94[†] ± 0.47 | 88.08[†] ± 0.17 | 78.78 [†] ± 0.10 | 87.82 [†] ± 0.10 |

Table 7: Multilingual Results in the regular **joint multilingual** setting (mean and stdev over 3 runs). Detailed (per-language) results can be found in Appendix (Table 24).[†] is 99% significance against **all** non-sentinel methods.

| Dataset | Extractive TS | | Extractive S+T | | Extractive S+T(S) | |
|---------------------------|-----------------|-----------------|------------------------------|------------------------------|------------------------------------|------------------------------------|
| | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 |
| mATIS (3 langs) | 35.17 ± 0.38 | 57.56 ± 0.84 | 37.62 ± 2.87 | 64.07 [†] ± 2.13 | 42.17[†] ± 0.91 | 67.48[†] ± 0.66 |
| mTOP (6 langs) | 65.43 ± 1.16 | 48.39 ± 1.79 | 80.14 [†] ± 1.00 | 67.37 [†] ± 1.99 | 80.81[†] ± 1.01 | 68.57[†] ± 1.18 |
| mTOD (3 langs) | 52.85 ± 2.60 | 66.46 ± 2.47 | 60.11 [†] ± 1.05 | 72.74 [†] ± 0.44 | 61.47[†] ± 0.39 | 73.66[†] ± 0.26 |
| Average | 54.72 ± 0.17 | 55.20 ± 0.06 | 64.50 [†] ± 1.48 | 67.89 [†] ± 1.42 | 66.31[†] ± 0.83 | 69.57[†] ± 0.82 |
| Average (non-en) | 43.07 ± 0.20 | 42.46 ± 0.09 | 56.34 [†] ± 1.74 | 59.60 [†] ± 1.69 | 58.55[†] ± 1.05 | 61.64[†] ± 1.04 |

Table 8: **Zero-shot** results (mean and stdev over 2 runs) for **extractive** methods. Per-language results reported in Table 20.[†] indicates 99% significance.

sentinel) approaches. While those two formats have very similar performance – with **Input+Tag** being slightly better than **Tag Only** – the addition of the sentinel tokens drastically improves performance. Specifically, the language-agnostic and input-agnostic nature of these sentinel tokens enables the language-agnosticity of the decoder leading to improved generalization across languages despite the training data being only English.

As seen in Table 7, these trends largely continue in the joint multilingual training setting as well. Despite the use of training data from all languages, the **Sentinel+Tag** formats outperform all non-sentinel formats by between **4 and 9 pp** on the Perfect metric. We again find a large gap (**6.3 pp**) to the current **Tagged Spans** approach (a **8.7%** increase from **72.6** to **78.9**).

A few additional observations: (a) As on English

| Dataset | Extractive TS | | Extractive S+T | | Extractive S+T(S) | |
|---------------------------|------------------------|------------------------|------------------------------|------------------------------|------------------------------------|------------------------------------|
| | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 |
| mATIS (3 langs) | 70.01 ± 1.20 | 86.66 ± 0.48 | 72.06 ± 1.08 | 87.75 ± 0.55 | 72.94[†] ± 0.34 | 88.09[†] ± 0.31 |
| mTOP (6 langs) | 87.78 ± 0.30 | 77.17 ± 0.51 | 90.57 [†] ± 0.87 | 84.15 [†] ± 0.17 | 91.89[†] ± 0.00 | 86.02[†] ± 0.80 |
| mTOD (3 langs) | 89.02 ± 0.16 | 93.32 ± 0.03 | 87.48 ± 0.70 | 92.36 ± 0.40 | 88.58 ± 0.30 | 92.88 [†] ± 0.24 |
| Average | 83.65 ± 0.49 | 83.58 ± 0.38 | 85.17 [†] ± 0.53 | 87.10 [†] ± 0.12 | 86.33[†] ± 0.16 | 88.26[†] ± 0.54 |
| Average (non-en) | 82.25 ± 0.57 | 80.72 ± 0.43 | 84.11 [†] ± 0.42 | 85.24 [†] ± 0.00 | 85.04[†] ± 0.16 | 86.40[†] ± 0.71 |

Table 9: **Joint multilingual** results (mean and stdev over 2 runs) for **extractive** methods. Per-language results reported in Table 21.[†] indicates 99% significance.

benchmarks, the Inside simplification slightly improved performance – though the gaps are within error margins. (b) As seen in Table 5, the Outside simplification tends to hurt zeroshot performance significantly for all techniques (c) **Tagged Spans** – the current standard in published works – is significantly worse than other formats in both settings and especially in the zero-shot setting.

8 Performance of Extractive methods

As discussed in Sec 3.5, certain applications may prefer extractive formats. To verify if our findings extend to these settings we also evaluated these extractive formats on the same datasets. Tables 8 (zeroshot) and 9 (joint multilingual) summarize the results. A few key observations:

- The extractive sentinel-based approaches continue the trend of significantly outperform exist-

| Size | Average (all) [Zero-Shot] | | | | Average (non-en) [Zero-Shot] | | | | Average (all) [Joint] | | | | Average (non-en) [Joint] | | | |
|--------------|---------------------------|--------------|--------------|----------------------------------|------------------------------|--------------|--------------|----------------------------------|-----------------------|--------------|--------------|----------------------------------|--------------------------|--------------|--------------|----------------------------------|
| | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T |
| Small (300M) | 44.57 ± 0.03 | 44.96 ± 0.00 | 42.78 ± 0.05 | 46.61 [†] ± 0.01 | 30.28 ± 0.07 | 30.73 ± 0.02 | 27.96 ± 0.15 | 32.49 [†] ± 0.01 | 74.53 ± 0.06 | 75.52 ± 0.13 | 73.64 ± 0.29 | 78.55 [†] ± 0.06 | 71.07 ± 0.11 | 72.25 ± 0.20 | 70.14 ± 0.31 | 75.76 [†] ± 0.09 |
| Base (580M) | 48.68 ± 0.39 | 49.54 ± 0.46 | 45.58 ± 0.57 | 57.79 [†] ± 1.38 | 35.51 ± 0.41 | 36.47 ± 0.63 | 31.29 ± 0.64 | 47.11 [†] ± 1.73 | 76.63 ± 0.08 | 78.11 ± 0.25 | 76.15 ± 0.12 | 81.21 [†] ± 0.10 | 73.22 ± 0.04 | 74.89 ± 0.32 | 72.60 ± 0.14 | 78.78 [†] ± 0.10 |
| Large (1.2B) | 51.26 ± 0.55 | 50.27 ± 0.62 | 46.15 ± 0.28 | 59.80 [†] ± 0.89 | 38.96 ± 0.68 | 37.12 ± 0.79 | 31.76 ± 0.41 | 49.74 [†] ± 1.13 | 79.24 ± 0.11 | 80.03 ± 0.19 | 78.51 ± 0.24 | 81.37 [†] ± 0.10 | 76.24 ± 0.15 | 77.05 ± 0.19 | 75.13 ± 0.31 | 78.89 [†] ± 0.14 |
| XL (3.7B) | 54.72 ± 0.32 | 59.01 ± 0.22 | 60.27 ± 0.84 | 64.47 [†] ± 0.08 | 42.86 ± 0.39 | 48.49 ± 0.30 | 50.17 ± 1.18 | 55.70 [†] ± 0.12 | 82.44 ± 0.21 | 82.63 ± 0.07 | 81.73 ± 0.04 | 84.04 [†] ± 0.10 | 79.73 ± 0.30 | 79.97 ± 0.12 | 78.80 ± 0.05 | 81.73 [†] ± 0.10 |
| XXL | 53.27 | 57.32 | 59.97 | 65.55 | 40.95 | 46.27 | 49.84 | 57.36 | 82.13 | 82.59 | 82.20 | 83.86 | 79.35 | 79.91 | 79.34 | 81.62 |

Table 10: Impact of model size (# of params in parentheses) on performance of formats. Averaged Perfect metric scores are reported over the same 3 benchmarks (12 test sets) as Table 6, on both zero-shot and joint settings. The 4 methods compared are **T-O**: Tag Only (SI), **I+T**: Input + Tag (SI), **TS**: Tagged Spans and **S+T**: Sentinel + Tag (SI). Base results are averaged over 3 runs. XL, Large and Small were averaged over 2 runs. Due to compute limits, XXL (13B params) was run once for 2k steps (since trial runs plateaued there).[†] indicates 99% significance.

| Size | Average (all) [Zero-Shot] | | | | Average (non-en) [Zero-Shot] | | | | Average (all) [Joint] | | | | Average (non-en) [Joint] | | | |
|-------|---------------------------|--------------|-------|-------------|------------------------------|--------------|-------|-------------|-----------------------|-------------|------|-------------|--------------------------|-------------|------|-------------|
| | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T |
| Small | 29.92 | 33.88 (2.5) | 32.14 | 0.03 | 39.14 | 44.99(3.32) | 42.61 | 0.03 | 8.90 | 5.81 (0.88) | 6.19 | 0.00 | 10.82 | 7.61 (1.18) | 8.10 | 0.00 |
| Base | 29.68 | 40.60 (4.43) | 43.22 | 0.02 | 38.87 | 53.89 (5.91) | 57.38 | 0.02 | 8.53 | 5.42 (0.9) | 6.10 | 0.01 | 10.58 | 7.06 (1.15) | 7.87 | 0.01 |
| Large | 30.60 | 40.14 (1.9) | 47.78 | 0.07 | 39.67 | 53.37 (2.54) | 63.48 | 0.09 | 8.68 | 5.04 (0.87) | 5.81 | 0.00 | 10.71 | 6.52 (1.21) | 7.53 | 0.00 |
| XL | 29.92 | 24.76 (1.55) | 20.39 | 0.11 | 39.64 | 32.96 (2.08) | 27.15 | 0.14 | 3.20 | 3.81 (0.66) | 4.10 | 0.01 | 4.12 | 5.04 (0.95) | 5.42 | 0.01 |
| XXL | 32.81 | 32.64 (2.16) | 21.35 | 0.68 | 43.41 | 43.46 (2.89) | 28.34 | 0.88 | 4.91 | 3.71 (0.56) | 3.79 | 0.00 | 6.38 | 4.90 (0.75) | 4.99 | 0.00 |

Table 11: Percentage of examples with hallucinations observed for the different formats / models from Table 10. For **Input+ Tag** we also report (in parenthesis) the % of examples with hallucinations in correctly tagged spans.

ing extractive approaches across all metrics.

- When comparing Tables 6 vs 8 and Tables 7 vs 9, we find that all the Extractive formats improve on the *Perfect* metric as opposed to formats which output labels for the entire input. On further analysis we find this is due to the simpler (and shorter length) prediction task in this case.
- However, when we look at the *MicroF1* metrics, we find that extractive models are **worse** at *extracting* the correct labeled spans than models that simply label every input token. This performance drop is particularly significant for the existing (TaggedSpans) approach (3-5pp) – vs. the drop in sentinel models (1.2-2.5pp).

These results may have wider implications on how we should model extractive tasks and the formats used there. However we leave this to future work.

9 Robustness & Efficiency

Beyond performance, efficiency and robustness of a model are equally important factors in practice. We focus on measuring these in this section.

9.1 Understanding effect of model size

While experiments and results so far used *Base* sized models (to maximize experimentation given compute), we would like to understand how robust our findings are as we vary model sizes. As seen in Table 10 (and Table 25), the previously observed trends are all consistently repeated across all model sizes. In particular, we still find the **Sentinel+Tag** models performing the best across all the datasets

| Dataset | Zeroshot | | | Joint | | |
|-------------------------|----------|-------------|-------------|-------|-------|-------------|
| | E TS | E S+T | E S+T(S) | E TS | E S+T | E S+T(S) |
| mATIS (3 langs) | 32.65 | 0.41 | 0.17 | 4.09 | 0.07 | 0.06 |
| mTOP (6 langs) | 17.68 | 0.18 | 0.03 | 0.56 | 0.01 | 0.00 |
| mTOD (3 langs) | 24.59 | 0.12 | 0.12 | 0.31 | 0.11 | 0.07 |
| Average | 23.15 | 0.22 | 0.09 | 1.38 | 0.05 | 0.03 |
| Average (non-en) | 30.75 | 0.25 | 0.10 | 1.74 | 0.05 | 0.03 |

Table 12: % of examples with hallucinations for different **extractive** formats. Per-dataset results in Table 22.

for both the zero-shot and joint settings. Notably, zero-shot performance improves significantly for all datasets as model size increases. Consequently, the larger models have a smaller gap between zero-shot and joint settings. Put together, the performance of the **Sentinel+Tag** format is robust across all model sizes, datasets, and experiment settings.

9.2 Understanding effect of hallucination

One of our key motivations for exploring alternative formats, is the prevalence of hallucination. Hallucination is one of the most notorious problems plaguing Seq2Seq models and hurting practical adoption. In particular, we found that models trained using the **Tagged Spans** or **Input + Tag** formats, often resulted in outputs containing words that not present in the input (even on English examples). Even our previous metrics for the **Input + Tag** formats were generous as they glossed over some hallucinations by comparing (tagged) token indices rather than the actual texts.

We measured hallucinations for the different formats in a simple, straightforward manner. In particular, for formats that repeated the input text in

| Format | Average / 99%-ile Seq. Len (and variants) | |
|--------------------------------|---|---|
| Input Lengths | | |
| Original Input | 14.8 / 32.6 | |
| Sentinel Input | 25.2 / 55.8 | |
| Target / Output Lengths | | |
| Tag Only | 23.3 / 70.4 | With SI: 15.5 / 34.7 |
| Input + Tag | 57.9 / 142.2 | With SI: 50.0 / 109.2, SI+SO: 33.8 / 82.3 |
| Tagged Spans | 58.3 / 123.4 | With SO: 31.3 / 67.9 |
| Sentinel + Tag | 33.0 / 90.1 | With SI: 25.2 / 54.8, SI+SO: 19.8 / 45.8 |

Table 13: Sequence length (*i.e.*, the number Sentence-Piece tokens) statistics (averaged) across datasets for the different formats studied (Full stats in Table 19).

the output (**Tagged Spans** and **Input + Tag**) we measured how often the generated sentence does not match the input (*i.e.*, one of the words are either deleted, inserted or modified). For formats which do not repeat the input, we counted the number of tokens and checked if they match – for **Tag Only** this was the number of tags produced in the output, whereas for **Sentinel + Tag** we counted the generated sentinel tokens.

Table 11 shows how the rate of hallucination across different model sizes and formats. As noted earlier, we find that hallucination is highly prevalent in the **Tagged Spans** and **Input + Tag** formats due to having to repeat the input – which often leads to errors. In the zero-shot setting, we find that nearly a third of all the examples had one of more hallucinations in most models for these formats. Furthermore, these hallucinations were present in both English as well as other languages – though slightly less common in English. Even in the joint multilingual setting, we found that about 1 in every 20-25 predictions had some hallucination. Appendix E contains examples of hallucinations for different models on these datasets.

On the other hand, the **Sentinel + Tag** based models are almost completely devoid of hallucination, with either zero or near zero hallucinated instances across all model sizes in both settings across all datasets. This order(s) of magnitude difference in robustness to hallucination serves as clear evidence of the sentinel-based approach being more robust and practical.

Extractive formats: A curious reader may wonder if the shorter, more succinct extractive formats still suffer from this issue. The results can be found in Table 12. We find similar trends here too, with existing (input-repeating) models having orders of magnitude more hallucinations than the sentinel-based formats. Examples of these hallucinations can be found in Appendix F. Additionally, we also provide examples of other (non-hallucination) wins and losses of the sentinel approach in Appendix G.

9.3 Efficiency and effect of sequence length

While running experiments, we observed that **Sentinel+Tag** and **Tag-Only** models were much faster at both training and especially at inference. The key to this is the sequence lengths of these models being far shorter than those for models which need the input to be repeated like the **Input+Tag** and **Tagged Spans** approaches. In particular, as seen in Table 13, the outputs for the best performing **Sentinel+Tag (SI)** format are often less than half the length of those of the current standard – **Tagged Spans**. While the sentinel-augmented inputs are longer (about 10 tokens on average), this difference was much smaller than the difference in the output lengths (33 tokens on average). This trend holds across all datasets, languages and aggregating functions (average, median, 99th percentile ..). Given the quadratic complexity of sequence length, such efficient gains are invaluable in practical scenarios. Consequently even though some simplifications – like the Simplified Outside variant of the **Sentinel+Tag** – did not improve performance, being 25-30% the length of current formats can translate to **10x+** speedup in training and inference.

Additionally, as detailed in Appendix D and Table 18 we also verified that the new format is robust to changes in training sequence lengths.

10 Conclusion

In this paper, we rigorously studied different input and output sequence *formats* for casting Sequence Tagging tasks as Seq2Seq problems. Using experiments across 7 public benchmarks, we found that the formats of the input and output sequences can significantly impact performance. To remedy the shortcomings of the existing formats, we also introduced a new format that uses sentinel tokens. Along with its variants, this new format proved to not only be simpler and more computationally efficient but also more accurate. The sentinel-based formats significantly outperform all other formats when it comes to multilingual generalization, with 30+% relative improvement in accuracy. While current formats are plagued by hallucination in a large percentage of examples, the new format rectifies this issue and is virtually hallucination free.

11 Acknowledgements

We sincerely thank Michael Bendersky, Emily Pitler, Slav Petrov and our anonymous reviewers for their valuable feedback.

12 Limitations

There are two notable limitations we would like to discuss of our work – specifically the proposed sentinel-based approach.

A Need for additional preprocessing and post-processing: The sentinel approach requires that *sentinel* tokens be inserted into the input, so as to simplify decoding (*i.e.*, enable the model to skip generating the original input tokens). While this leads to a host of gains as mentioned in the paper, it would be remiss of us to not point out that this approach comes with a drawback. Namely, the need for additional pre- and post-processing. Inserting the sentinels into the input required some (potentially lightweight) tokenization or lexing of the input sentence. While this may not be an issue for English and other popularly studied languages, this can be non-trivial for non-space separated languages like Thai and Lao or agglutinative languages like Finnish and German. While we tried to study the impact of this in Appendix B and found no performance loss, we should note that this is a function of the quality of the pre-processing tokenizer / lexer. If tokens are often accidentally split across tag boundaries by this pre-processing step, then this may hurt performance. Without a more detailed study on these languages (which is somewhat challenged by availability of suitable datasets in these languages), it is hard to quantify the effect this need for pre-processing may require in general.

From a practical viewpoint, this pre-processing also adds an additional (albeit likely lightweight) step before the model, and potentially an additional post-processing step (joining input tokens and predicted tags) based on the output data format required by downstream steps. While we found the computation savings from the sentinel format more than enough to counteract any additional pre-processing, this is a function of application and specific system specifications and hence worth noting for practitioners.

B Efficiency for long texts on (short) extractive tasks: Due to the insertion of the sentinel tags we increase the sequence length of our inputs (by about 10 tokens on average in our studied datasets per Table 10). This is more

than made up for by the larger decrease in output sequence length on the datasets we studied. However, for extractive tasks we can simply only output extracted spans and their associated tasks. In our running example (Table 2), this may perhaps look like: "ARTIST: Kent James, PLAYLIST: Disney". While a non-sentinel approach for such an output still suffers many of same issues as the non-sentinel formats studied in our work (namely lower performance, worse multilingual generalization, frequent hallucination ..) the efficiency gains from the sentinel approach are far less clear. In particular, when we have such an extractive task with **long** inputs (think 500+ sequence length) but very short output (*i.e.*, very few extractions), then the sentinel approach would likely be slower (at training and inference) than a non-sentinel extractive output format.

While this is not a problem setting we work on or focused on in this paper, we fully recognize that a non-trivial fraction of the community may have interest in such a setting. Hence we find this important to point out.

13 Ethics and Responsible NLP

This paper tried to provide empirical insights into how sequence tagging NLP tasks should be handled in the Seq2Seq regime given the increased prevalence of such pretrained models. Along the way it proposed a new **Sentinel**-based approach to the problem. Given the nature of the work and findings, we could not think of any explicit *risks* associated with this work or the new format. Instead we could postulate about potential benefits these new **Sentinel**-based formats could bring about with wider adoption:

- **Better multilingual generalization:** Our results demonstrate the potential to improve performance of common NLP tasks for low-resource languages (in addition to gains on high-resource languages).
- **More trustworthy NLP models:** By nearly entirely eliminating hallucinations, the new format has promise to increase fidelity of NLP models.
- **Reducing compute:** As discussed in Section 9.3, the new approach could potentially

lead to an order of magnitude reduction in training / inference time versus the current status quo for these problems. Furthermore, given concerns surrounding massive compute models, we verified our findings (like increased multilinguality and near-zero hallucination) hold true for even the "Small" and "Base" sized pretrained mT5 models (see Tables 10 and 11).

- **Possible benefits to privacy:** While not discussed in detail in the paper, one notable benefit of adding the sentinel tokens is that the output of the model no longer contains any input token. This is very amenable to privacy-preserving / privacy-focused NLP applications and potentially unlocks new opportunities for these kinds of models in more privacy sensitive settings.

References

- Wasi Ahmad, Jianfeng Chi, Tu Le, Thomas Norton, Yuan Tian, and Kai-Wei Chang. 2021. [Intent classification and slot filling for privacy policies](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4402–4417, Online. Association for Computational Linguistics.
- Ben Athiwaratkun, Cicero Nogueira dos Santos, Jason Krone, and Bing Xiang. 2020. [Augmented natural language for generative sequence labeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 375–385, Online. Association for Computational Linguistics.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.
- Angel Daza and Anette Frank. 2018. [A sequence-to-sequence model for semantic role labeling](#). In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 207–216, Melbourne, Australia. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2021. All nlp tasks are generation tasks: A general pretraining framework. *arXiv preprint arXiv:2103.10360*.
- Jitin Krishnan, Antonios Anastasopoulos, Hemant Purohit, and Huzefa Rangwala. 2021. [Multilingual code-switching for zero-shot cross-lingual intent prediction and slot filling](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 211–223, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Haoran Li, Abhinav Arora, Shuohui Chen, Anchit Gupta, Sonal Gupta, and Yashar Mehdad. 2021. [MTOP: A comprehensive multilingual task-oriented semantic parsing benchmark](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2950–2962, Online. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. Structured prediction as translation between augmented natural languages. *arXiv preprint arXiv:2101.05779*.
- Patti Price. 1990. Evaluation of spoken language systems: The atis domain. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Chengwei Qin and Shafiq Joty. 2022. [LFPT5: A unified framework for lifelong few-shot language learning based on prompt tuning of t5](#). In *International Conference on Learning Representations*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Laria Reynolds and Kyle McDonell. 2021. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. [Cross-lingual transfer learning](#)

- for multilingual task oriented dialog. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3795–3805, Minneapolis, Minnesota. Association for Computational Linguistics.
- Shyam Upadhyay, Manaal Faruqui, Gokhan Tür, Hakkani-Tür Dilek, and Larry Heck. 2018. (almost) zero-shot cross-lingual spoken language understanding. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6034–6038. IEEE.
- Zirui Wang, Adams Wei Yu, Orhan Firat, and Yuan Cao. 2021. Towards zero-label language learning. *CoRR*, abs/2109.09193.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. A unified generative framework for various NER subtasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5808–5822, Online. Association for Computational Linguistics.
- Ruiqi Zhong, Kristy Lee, Zheng Zhang, and Dan Klein. 2021. Adapting language models for zero-shot learning by meta-tuning on dataset and prompt collections. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2856–2878, Punta Cana, Dominican Republic. Association for Computational Linguistics.

A Experiments with Other Format Variants

In addition to the four primary formats and their variants with Simplified-Inside tags, we show examples of several other variants that we studied in Table 17.

For example, we experimented with the Simplified-Output (SO) variant, which omits all the Outside tags. The results in Table 14 shows that the models do not benefit from skipping any tokens or tags. While these variants result in shorter output sequences (and therefore more efficient computation), the accuracy of these methods were lower than that of the Sentinel + Tag (and its Simplified-Inside variant) as shown in Table 14.

| Dataset | Input + Tag (SI, SO) | Tagged Spans (SO) | Sentinel + Tag (SI, SO) |
|------------------------|----------------------|-------------------|-------------------------|
| mATIS(en) | 88.80 | 89.03 | 90.15 |
| SNIPS | 88.71 | 88.14 | 90.43 |
| MovieTrivia | 37.33 | 37.63 | 35.59 |
| Movies | 70.95 | 71.61 | 71.08 |
| Restaurant | 59.50 | 60.82 | 62.80 |
| mTOP(en) | 82.65 | 84.09 | 86.34 |
| mTOD(en) | 92.50 | 92.61 | 93.36 |
| Average | 74.35 | 74.85 | 75.68 |
| (Δ vs. non-SO) | -1.13 | -0.26 | -0.77 |

Table 14: Per-dataset **Perfect** metric scores for the variant formats that modify or drop the Outside tag.

B Sentinels in complex languages

In languages like Thai and Lao, sentences can often be composed of just one or two "words" *i.e.*, because unlike English their words are not separated by spaces. Similarly agglutinative languages like Finnish (and to a far lesser extent German) may join multiple affixes / morphemes to form very long "words".

This brings up the question of how we should insert sentinels for inputs in these languages, since we may be actually altering the underlying token representations. For example, say our input was the German phrase *schweres werkzeug* (heavy tool) and assume it was tokenized as "schweres" + "werk" + "zeug". When we create the sentinel input, we could either use: "<S> schweres <S> werk <S> zeug" or as "<S> schweres <S> werk <S>zeug" (with "<S>" representing the sentinel tokens). The two texts would be tokenized differently by the SentencePiece model (SPM) due to the lack of a space between the sentinel and "zeug" – since

SPM treats continued tokens differently than word starts.

To understand if these models are robust to that we ran experiments on the German and Thai datasets. As seen in Table 15, while there are slight differences, the models are quite robust to the tokenization and how / where the sentinels tokens are introduced – another evidence of the robustness of the sentinel based approach.

| Dataset | Space inserted | | No space | |
|--------------------|-------------------------|-------------------------|------------------|-------------------------|
| | No SI | SI | No SI | SI |
| Zeroshot | | | | |
| mTOP(de) | 60.27 \pm 1.40 | 65.15 \pm 0.88 | 60.66 \pm 1.77 | 62.89 \pm 1.63 |
| mTOP(th) | 43.30 \pm 0.75 | 43.45 \pm 1.16 | 44.41 \pm 1.45 | 44.61 \pm 2.58 |
| Joint Multilingual | | | | |
| mTOP(de) | 79.47 \pm 0.45 | 78.77 \pm 0.47 | 78.44 \pm 0.66 | 77.77 \pm 0.48 |
| mTOP(th) | 77.58 \pm 0.14 | 77.32 \pm 0.84 | 76.19 \pm 0.61 | 75.36 \pm 0.64 |

Table 15: Scores for the **Perfect** metric with and without spaces added pre-tokenization of the input of the **Sentinel + Tag** approaches.

C Robustness of decoder distribution

While the results so far have focused on the top-1 prediction, in many applications we want a robust and meaningful top- K prediction set. Thus we investigated the robustness of the decoder's output distribution. Table 16 shows the % of the test examples whose top- K predictions (produced via beam search with $K = 5$) include the ground-truth label.

| Method | mATIS(en) | SNIPS | MovieTrivia | Movies | Restaurant |
|---------------------------------|--------------|--------------|--------------|--------------|--------------|
| Tag Only | 92.16 | 96.71 | 47.47 | 88.82 | 85.03 |
| Input + Tag | 90.03 | 96.29 | 48.85 | 90.45 | 86.61 |
| Tagged Spans | 88.24 | 96.86 | 69.53 | 88.24 | 87.40 |
| Sentinel + Tag | 93.84 | 98.29 | 67.95 | 90.82 | 86.28 |
| Simplified Inside Based Formats | | | | | |
| Tag Only (SI) | 92.39 | 96.57 | 62.98 | 89.88 | 84.23 |
| Input + Tag (SI) | 91.94 | 97.14 | 66.72 | 90.66 | 87.34 |
| Sentinel + Tag (SI) | 94.40 | 98.71 | 72.20 | 90.78 | 87.53 |

Table 16: % of examples with any of the top-5 predictions containing the ground-truth label.

As expected, the observed scores are significantly higher than the perfect accuracy scores (for the Top-1) seen in Table 4. Also, consistent with previous results, the **Sentinel + Tag (SI)** performed the best on average. In general, this analysis encourages us to think about how we could leverage these instances (where the correct prediction is in the top- K but not the top-1) to further improve performance. We leave this for future work.

| Format | Sample Target |
|---|--|
| Simplified Inside Based Formats | |
| Tag Only (SI) | target: O ARTIST I O O PLAYLIST O |
| Tag + Input (SI) | target: <O> Add <ARTIST> Kent <I> James <O> to <O> the <PLAYLIST> Disney <O> soundtrack |
| Sentinel + Tag (SI) | target: <extra_id_0> O <extra_id_1> ARTIST <extra_id_2> I <extra_id_3> O <extra_id_4> O <extra_id_5> PLAYLIST <extra_id_6> O |
| Simplified Outside Based Formats | |
| Tag + Input (SI, SO) | target: Add <ARTIST> Kent <I> James to the <PLAYLIST> Disney soundtrack |
| Tagged Spans (SO) | target: Add <ARTIST> Kent James </> to the <PLAYLIST> Disney </> soundtrack |
| Sentinel + Tag (SI, SO) | target: <extra_id_0> <extra_id_1> ARTIST <extra_id_2> I <extra_id_3> <extra_id_4> <extra_id_5> PLAYLIST <extra_id_6> |
| Extractive Sentinel + Tag (ES+T) Based Formats | |
| ES+T | target: <extra_id_1> ARTIST <extra_id_2> I-ARTIST <extra_id_5> PLAYLIST |
| ES+T (Simplified) | target: <extra_id_1> ARTIST <extra_id_2> <extra_id_5> PLAYLIST |

Table 17: Examples for other format variants we studied. For all the non-sentinel approaches, examples, the *input* is the original input utterance "Add Kent James to the Disney soundtrack". For sentinel-based approaches, the *input* is "<extra_id_0> Add <extra_id_1> Kent <extra_id_2> James <extra_id_3> to <extra_id_4> the <extra_id_5> Disney <extra_id_6> soundtrack".

| Size | Average (all) [ZS] | | | | Average (non-en) [ZS] | | | | Average (all) [J] | | | | Average (non-en) [J] | | | |
|------------------|--------------------|-------|-------|--------------|-----------------------|-------|-------|--------------|-------------------|-------|-------|--------------|----------------------|-------|-------|--------------|
| | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T |
| In: 128 Out: 256 | 48.68 | 49.54 | 45.58 | 57.79 | 35.51 | 36.47 | 31.29 | 47.11 | 76.63 | 78.11 | 76.15 | 81.21 | 73.22 | 74.89 | 72.60 | 78.78 |
| In: 128 Out: 128 | 48.79 | 49.55 | 46.15 | 55.35 | 35.59 | 36.65 | 32.26 | 44.21 | 75.95 | 75.08 | 75.22 | 79.89 | 72.51 | 71.83 | 71.69 | 77.22 |
| In: 64 Out: 64 | 47.01 | 48.52 | 45.23 | 54.89 | 33.84 | 35.91 | 32.06 | 43.82 | 72.91 | 74.92 | 71.11 | 79.95 | 69.45 | 71.96 | 67.66 | 77.61 |

Table 18: Results demonstrating impact of changing training (input and output / target) sequence lengths for Base models in a similar setup as Table 10.

| Format | mATIS | | | SNIPS | MovieTrivia | Movies | Restaurant | mTOP | | | | | | mTOD | | |
|-------------------------------|----------|----------|----------|----------|-------------|----------|------------|----------|----------|----------|----------|----------|----------|---------|----------|----------|
| | en | hi | tr | | | | | en | de | es | fr | hi | th | en | es | th |
| Input Formats | | | | | | | | | | | | | | | | |
| Original Input | 16 / 37 | 24 / 54 | 20 / 42 | 12 / 26 | 28 / 56 | 13 / 32 | 12 / 27 | 9 / 23 | 11 / 28 | 12 / 31 | 13 / 33 | 15 / 37 | 10 / 24 | 9 / 20 | 11 / 24 | 13 / 28 |
| Sentinel Input | 28 / 61 | 36 / 83 | 31 / 64 | 21 / 44 | 49 / 95 | 24 / 55 | 21 / 48 | 17 / 43 | 20 / 51 | 23 / 56 | 23 / 56 | 24 / 58 | 24 / 60 | 16 / 34 | 19 / 40 | 21 / 44 |
| Sentinel Input (NoSpace) | | | | | | | | 17 / 41 | 19 / 49 | 22 / 54 | 23 / 54 | 24 / 57 | 20 / 49 | | | |
| Target Formats | | | | | | | | | | | | | | | | |
| Tag Only | 31 / 76 | 30 / 76 | 28 / 68 | 20 / 54 | 50 / 120 | 18 / 49 | 14 / 38 | 19 / 74 | 19 / 79 | 22 / 85 | 21 / 83 | 19 / 72 | 25 / 91 | 14 / 47 | 17 / 55 | 18 / 59 |
| Tag Only (SI) | 26 / 59 | 26 / 59 | 24 / 53 | 13 / 28 | 21 / 42 | 13 / 32 | 10 / 22 | 12 / 33 | 13 / 33 | 13 / 32 | 13 / 33 | 13 / 33 | 10 / 22 | 10 / 21 | 10 / 21 | |
| Input + Tag | 71 / 157 | 78 / 177 | 70 / 156 | 50 / 111 | 120 / 249 | 52 / 124 | 46 / 103 | 43 / 129 | 47 / 144 | 53 / 159 | 53 / 154 | 51 / 144 | 58 / 167 | 37 / 89 | 43 / 104 | 47 / 108 |
| Input + Tag (SI) | 66 / 144 | 74 / 161 | 67 / 148 | 43 / 87 | 91 / 176 | 47 / 108 | 41 / 92 | 37 / 86 | 40 / 98 | 45 / 106 | 45 / 104 | 45 / 105 | 47 / 111 | 33 / 67 | 36 / 75 | 39 / 79 |
| Input + Tag (SI, SO) | 45 / 101 | 51 / 116 | 45 / 103 | 29 / 64 | 69 / 145 | 28 / 67 | 24 / 55 | 24 / 70 | 26 / 80 | 29 / 87 | 29 / 83 | 30 / 83 | 32 / 91 | 21 / 50 | 25 / 61 | 26 / 61 |
| Tagged Spans | 85 / 187 | 93 / 202 | 85 / 194 | 50 / 108 | 80 / 167 | 58 / 138 | 52 / 117 | 44 / 90 | 47 / 100 | 51 / 102 | 54 / 111 | 53 / 109 | 51 / 103 | 40 / 81 | 39 / 80 | 43 / 86 |
| Tagged Spans (SO) | 49 / 109 | 55 / 120 | 49 / 111 | 29 / 59 | 44 / 82 | 27 / 65 | 24 / 47 | 22 / 54 | 25 / 60 | 26 / 58 | 26 / 60 | 28 / 65 | 27 / 61 | 20 / 49 | 21 / 43 | 22 / 44 |
| Sentinel + Tag | 42 / 98 | 41 / 99 | 38 / 90 | 29 / 70 | 71 / 156 | 29 / 71 | 23 / 57 | 26 / 91 | 27 / 98 | 31 / 105 | 30 / 103 | 27 / 90 | 34 / 113 | 21 / 60 | 24 / 69 | 26 / 72 |
| Sentinel + Tag (SI) | 38 / 83 | 37 / 82 | 35 / 74 | 22 / 45 | 42 / 82 | 23 / 55 | 19 / 43 | 20 / 47 | 21 / 50 | 22 / 52 | 22 / 52 | 21 / 50 | 23 / 53 | 17 / 35 | 17 / 37 | 18 / 37 |
| Sentinel + Tag (SI, SO) | 30 / 68 | 29 / 70 | 28 / 63 | 18 / 37 | 34 / 69 | 17 / 41 | 13 / 30 | 16 / 42 | 16 / 43 | 17 / 45 | 17 / 44 | 16 / 42 | 18 / 46 | 13 / 30 | 14 / 32 | 14 / 31 |
| Extractive Sentinel + Tag | 28 / 74 | 26 / 75 | 24 / 67 | 20 / 59 | 56 / 145 | 16 / 48 | 12 / 40 | 18 / 83 | 18 / 84 | 21 / 93 | 19 / 91 | 17 / 79 | 25 / 102 | 13 / 52 | 17 / 64 | 17 / 65 |
| Extractive Sentinel + Tag (S) | 22 / 55 | 21 / 52 | 20 / 51 | 11 / 26 | 17 / 36 | 9 / 24 | 6 / 15 | 10 / 31 | 10 / 32 | 10 / 29 | 9 / 28 | 9 / 31 | 10 / 30 | 7 / 23 | 8 / 21 | 7 / 20 |

Table 19: Detailed per-dataset / language statistics for the Mean and 99%th percentile of sequence length (*i.e.*, the number SentencePiece tokens) for the different *input* and *target* formats studied.

D Effect of Sequence Length on Performance

The results in Table 18 help us understand the robustness of different formats to changes in sequence lengths. In general, we found that that all the formats were fairly robust to reductions in training sequence length, with previous trends among formats holding. In most cases, the performance dropped by about 1-2pp, with the biggest drop observed for the **Tagged Spans** approach on the Joint setting (about **-5pp** across metrics). Thus we can conclude that while setting more appropriate sequence lengths can increase performance, failure to do so does not hurt performance significantly.

E Illustrative Hallucination Examples

Section 9.2 discussed the increased robustness of the **Sentinel+Tag** model.

To help understand what hallucinations would

look like, we provide some examples below using the running example of "Add Kent James to the Disney soundtrack". The *expected Target* (using the **Tagged Spans** format) is "<O> Add </> <ARTIST> Kent James </> <O> to </> <O> the </> <PLAYLIST> Disney </> <O> soundtrack </>".

Examples of predictions containing *hallucinations* would be:

- "<O> Add </> <ARTIST> Kent Jackson </> <O> to </> <O> the </> <PLAYLIST> Disney </> <O> soundtrack </>" changes the name of the artist from "Kent James" to "Kent Jackson".
- "<O> Add </> <ARTIST> Kent James </> <O> to </> <O> the </> <O> soundtrack </>" drops the "Disney" token altogether.
- "<O> Add </> <ARTIST> Kent James </> <O> to </> <O> the </> <PLAYLIST> Walt

| Dataset | Language | Extractive TS | | Extractive S+T | | Extractive S+T(S) | |
|------------------|--------------|---------------|---------------|----------------|---------------------|---------------------|---------------------|
| | | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 |
| mATIS | en | 89.59 | 95.60 | 88.13 | 95.39 | 89.03 | 95.59 |
| | hi | ± 0.11 ± 0.03 | ± 0.11 ± 0.03 | ± 1.12 ± 0.35 | ± 1.12 ± 0.35 | ± 0.11 ± 0.01 | ± 0.11 ± 0.01 |
| | tr | 7.39 ± 0.90 | 30.33 ± 2.70 | 22.28 ± 6.72 | 57.60 ± 3.06 | 32.87 ± 2.30 | 63.70 ± 1.50 |
| mTOD | en | 92.23 | 96.07 | 91.62 | 95.72 | 91.45 | 95.64 |
| | es | ± 0.34 ± 0.16 | ± 0.34 ± 0.16 | ± 0.59 ± 0.32 | ± 0.59 ± 0.32 | ± 0.06 ± 0.01 | ± 0.06 ± 0.01 |
| | th | 45.20 ± 5.64 | 60.74 ± 4.97 | 59.00 ± 1.00 | 72.16 ± 2.16 | 62.72 ± 1.50 | 75.57 ± 0.99 |
| mTOP | en | 87.22 | 88.60 | 87.24 | 87.20 | 88.36 | 88.88 |
| | de | ± 0.19 ± 0.17 | ± 0.19 ± 0.17 | ± 1.63 ± 1.81 | ± 1.63 ± 1.81 | ± 0.72 ± 0.47 | ± 0.72 ± 0.47 |
| | es | 73.68 ± 1.01 | 59.91 ± 2.37 | 80.70 ± 1.55 | 71.28 ± 3.60 | 81.50 ± 1.93 | 73.29 ± 2.06 |
| | fr | 83.79 ± 0.43 | 55.52 ± 1.48 | 89.39 ± 0.13 | 68.70 ± 1.76 | 89.44 ± 0.78 | 69.26 ± 2.71 |
| | hi | 63.14 ± 1.85 | 51.49 ± 3.27 | 80.02 ± 1.47 | 64.05 ± 3.15 | 81.01 ± 1.39 | 67.29 ± 2.08 |
| | th | 44.25 ± 1.47 | 33.24 ± 3.25 | 43.64 ± 1.22 | 32.02 ± 1.64 | 44.69 ± 1.27 | 34.16 ± 2.13 |
| | Average | 54.72 ± 0.17 | 55.20 ± 0.06 | 64.50 ± 1.48 | 67.89 ± 1.42 | 66.31 ± 0.83 | 69.57 ± 0.82 |
| Average (non-en) | 43.07 ± 0.20 | 42.46 ± 0.09 | 56.34 ± 1.74 | 59.60 ± 1.69 | 58.55 ± 1.05 | 61.64 ± 1.04 | |

Table 20: Detailed per-language **extractive** results for **Zero-shot** setting (mean and stdev over 2 runs).

| Dataset | Language | Extractive TS | | Extractive S+T | | Extractive S+T(S) | |
|------------------|--------------|---------------|---------------|----------------|---------------------|---------------------|---------------------|
| | | Perfect | Mic F1 | Perfect | Mic F1 | Perfect | Mic F1 |
| mATIS | en | 89.42 | 95.71 | 87.23 | 94.98 | 89.25 | 95.65 |
| | hi | ± 0.06 ± 0.10 | ± 0.06 ± 0.10 | ± 1.57 ± 0.43 | ± 1.57 ± 0.43 | ± 0.34 ± 0.12 | ± 0.34 ± 0.12 |
| | tr | 58.17 ± 2.74 | 81.21 ± 1.38 | 64.95 ± 0.90 | 84.39 ± 0.45 | 65.73 ± 1.46 | 84.62 ± 1.00 |
| mTOD | en | 91.46 | 95.69 | 90.08 | 94.85 | 91.06 | 95.31 |
| | es | ± 0.15 ± 0.09 | ± 0.15 ± 0.09 | ± 0.64 ± 0.38 | ± 0.64 ± 0.38 | ± 0.15 ± 0.08 | ± 0.15 ± 0.08 |
| | th | 84.60 ± 0.18 | 89.89 ± 0.07 | 83.96 ± 0.43 | 89.58 ± 0.31 | 84.97 ± 0.35 | 90.04 ± 0.15 |
| mTOP | en | 82.66 | 85.07 | 87.76 | 88.24 | 90.24 | 90.56 |
| | de | ± 0.70 ± 0.69 | ± 0.70 ± 0.69 | ± 1.64 ± 1.40 | ± 1.64 ± 1.40 | ± 0.05 ± 0.09 | ± 0.05 ± 0.09 |
| | es | 87.31 ± 0.38 | 83.61 ± 0.66 | 90.73 ± 0.87 | 86.71 ± 1.20 | 92.05 ± 0.08 | 88.85 ± 0.08 |
| | fr | 92.43 ± 0.00 | 81.57 ± 0.31 | 94.31 ± 0.72 | 84.41 ± 1.73 | 95.15 ± 0.15 | 86.47 ± 0.40 |
| | hi | 87.97 ± 0.25 | 82.09 ± 0.42 | 89.59 ± 0.92 | 83.05 ± 1.18 | 90.82 ± 0.19 | 84.89 ± 0.47 |
| | th | 77.02 ± 0.47 | 79.40 ± 0.55 | 81.19 ± 1.13 | 82.42 ± 1.16 | 83.24 ± 0.13 | 84.42 ± 0.00 |
| | Average | 83.65 ± 0.49 | 83.58 ± 0.38 | 85.17 ± 0.53 | 87.10 ± 0.12 | 86.33 ± 0.16 | 88.26 ± 0.54 |
| Average (non-en) | 82.25 ± 0.57 | 80.72 ± 0.43 | 84.11 ± 0.42 | 85.24 ± 0.00 | 85.04 ± 0.16 | 86.40 ± 0.71 | |

Table 21: Per-language **extractive** results for **joint multilingual** setting (mean and stdev over 2 runs).

Disney </> <O> soundtrack </>" adds a spurious token "Walt".

The above examples can also be trivially modified for the **Input + Tag** format. In particular, the first example of the above three hallucinations is also an example of why the metrics for the **Input + Tag** format (as in Table 23 and Table 24) may be a tad generous. Since the token indices would not

| Dataset | Lang | Zeroshot | | | Joint | | |
|------------------|--------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | | E TS | E S+T | E S+T(S) | E TS | E S+T | E S+T(S) |
| mATIS | en | 0.56 | 0.11 | 0.00 | 0.34 | 0.00 | 0.11 |
| | hi | ± 0.11 ± 0.11 | ± 0.11 ± 0.11 | ± 0.00 ± 0.00 | ± 0.00 ± 0.00 | ± 0.00 ± 0.00 | ± 0.11 ± 0.11 |
| | tr | 79.84 ± 3.47 | 0.62 ± 0.39 | 0.17 ± 0.06 | 9.69 ± 1.06 | 0.06 ± 0.06 | 0.00 ± 0.00 |
| mTOD | en | 17.55 | 0.49 | 0.35 | 2.24 | 0.14 | 0.07 |
| | es | ± 1.89 ± 0.49 | ± 0.21 ± 0.21 | ± 0.14 ± 0.14 | ± 0.14 ± 0.14 | ± 0.00 ± 0.00 | ± 0.07 ± 0.07 |
| | th | 0.04 ± 0.03 | 0.01 ± 0.01 | 0.01 ± 0.01 | 0.03 ± 0.01 | 0.11 ± 0.02 | 0.00 ± 0.00 |
| mTOP | en | 0.49 | 0.27 | 0.09 | 0.57 | 0.00 | 0.00 |
| | de | ± 0.15 ± 0.23 | ± 0.23 ± 0.02 | ± 0.02 ± 0.03 | ± 0.36 ± 0.23 | ± 0.00 ± 0.00 | ± 0.00 ± 0.00 |
| | es | 9.85 ± 0.97 | 0.39 ± 0.25 | 0.03 ± 0.03 | 0.23 ± 0.08 | 0.00 ± 0.00 | 0.00 ± 0.00 |
| | fr | 5.44 ± 0.40 | 0.17 ± 0.13 | 0.00 ± 0.00 | 0.22 ± 0.12 | 0.03 ± 0.03 | 0.00 ± 0.00 |
| | hi | 9.10 ± 1.46 | 0.25 ± 0.22 | 0.03 ± 0.00 | 0.17 ± 0.08 | 0.02 ± 0.02 | 0.00 ± 0.00 |
| | th | 14.34 ± 1.43 | 0.02 ± 0.02 | 0.00 ± 0.00 | 2.17 ± 0.05 | 0.00 ± 0.00 | 0.00 ± 0.00 |
| | Average | 66.89 ± 0.60 | 0.00 ± 0.00 | 0.00 ± 0.00 | 0.02 ± 0.02 | 0.00 ± 0.00 | 0.00 ± 0.00 |
| Average (non-en) | 23.15 ± 0.61 | 0.22 ± 0.16 | 0.09 ± 0.01 | 1.38 ± 0.16 | 0.05 ± 0.01 | 0.03 ± 0.02 | |

Table 22: % of hallucination per dataset (mean and stdev over 2 runs) for **extractive** formats.

have caught the hallucination in the generated token text ("Kent Jackson" instead of "Kent James"), the model's performance drops when looking at the actual text as shown in the metrics in parentheses in Table 11.

F Real Hallucination Examples (for Extractive formats)

To help us understand the kinds of hallucinations we observe in these models, below are different sets of examples of hallucinations we see in the different formats. While these examples are for the shorter (more succinct) extractive formats they are also indicative of the hallucinations we see for the longer formats that label every input token. Note though that the ordering is not indicative of the prevalence of the different types of hallucinations but largely for the sake of exposition.

F.1 (Extractive) Tagged Spans

Incorrectly copied or added or modified words / phrases: This is fairly common even on English datasets.

Dataset: ATIS

Input: what is mci

Label: <airport_code> mci

Prediction: <airport_code> mco

| Size | Average (all) [Zero-Shot] | | | | Average (non-en) [Zero-Shot] | | | | Average (all) [Joint] | | | | Average (non-en) [Joint] | | | |
|-----------------|---------------------------|-----------------|-----------------|-------------------------------------|------------------------------|-----------------|-----------------|-------------------------------------|-----------------------|-----------------|-----------------|-------------------------------------|--------------------------|-----------------|-----------------|-------------------------------------|
| | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T | T-O | I+T | TS | S+T |
| Small (300M) | 55.60 ± 0.13 | 57.84 ± 0.19 | 56.74 ± 0.04 | 62.65 [†] ± 0.11 | 43.15 ± 0.15 | 45.93 ± 0.26 | 44.55 ± 0.00 | 52.09 [†] ± 0.16 | 83.06 ± 0.09 | 85.04 ± 0.15 | 84.03 ± 0.16 | 87.53 [†] ± 0.09 | 80.36 ± 0.10 | 82.66 ± 0.19 | 81.49 ± 0.17 | 85.72 [†] ± 0.11 |
| Base (580M) | 58.70 ± 0.81 | 61.94 ± 0.33 | 59.08 ± 0.82 | 71.90 [†] ± 1.40 | 47.13 ± 0.98 | 51.18 ± 0.45 | 47.41 ± 1.03 | 64.25 [†] ± 1.80 | 84.83 ± 0.05 | 87.00 ± 0.14 | 86.01 ± 0.16 | 89.40 [†] ± 0.10 | 82.25 ± 0.05 | 84.81 ± 0.17 | 83.65 ± 0.20 | 87.82 [†] ± 0.10 |
| Large (1.2B) | 59.96 ± 0.59 | 61.10 ± 1.29 | 58.65 ± 0.46 | 73.25 [†] ± 0.85 | 48.94 ± 0.68 | 49.88 ± 1.70 | 46.67 ± 0.61 | 66.07 [†] ± 1.09 | 86.78 ± 0.10 | 88.36 ± 0.13 | 87.73 ± 0.09 | 89.28 [†] ± 0.07 | 84.58 ± 0.13 | 86.42 ± 0.14 | 85.54 ± 0.10 | 87.62 [†] ± 0.14 |
| XL (3.7B) | 63.10 ± 0.55 | 70.64 ± 0.16 | 72.75 ± 0.84 | 77.96 [†] ± 0.10 | 52.50 ± 0.72 | 62.44 ± 0.20 | 65.23 ± 1.13 | 72.18 [†] ± 0.12 | 89.39 ± 0.21 | 90.06 ± 0.07 | 89.77 ± 0.04 | 91.20 [†] ± 0.10 | 87.50 ± 0.13 | 88.31 ± 0.08 | 87.94 ± 0.07 | 89.79 [†] ± 0.12 |
| XXL | 61.34 | 68.66 | 73.12 | 78.27 | 50.16 | 59.76 | 65.72 | 72.69 | 89.06 | 90.11 | 90.13 | 91.21 | 87.04 | 88.38 | 88.36 | 89.84 |

Table 25: Impact of model size (# of params in parentheses) on performance of formats. Averaged **MicroF1** scores are reported over the same 3 benchmarks (12 test sets) as Table 6, on both zero-shot and joint settings. The 4 methods compared are **T-O**: Tag Only (SI), **I+T**: Input + Tag (SI), **TS**: Tagged Spans and **S+T**: Sentinel + Tag (SI). Base results are averaged over 3 runs. XL, Large and Small were averaged over 2 runs. Due to compute limits, XXL (13B params) was run once for 2k steps (since trial runs plateaued there).[†] indicates 99% significance.

| |
|---|
| <p>Input: Will the temps be freezing tonight Target: <WEATHER_ATTRIBUTE> freezing <DATE_TIME> tonight E TS: <DATE_TIME> tonight Sentinel Target = E S+T(S): <extra_id_4> WEATHER_ATTRIBUTE <extra_id_5> DATE_TIME</p> |
| <p>Input: tonights projected low Target: <DATE_TIME> tonights E TS: <MUSIC_ALBUM_TITLE> tonights projected low Sentinel Target = E S+T(S): <extra_id_0> DATE_TIME</p> |
| <p>Input: what are some ways to cook chicken in the crock pot Target: <RECIPES_INCLUDED_INGREDIENT> chicken <RECIPES_COOKING_METHOD> crock pot E TS: <RECIPES_DISH> chicken in the crock pot Sentinel Target = E S+T(S): <extra_id_6> RECIPES_INCLUDED_INGREDIENT <extra_id_9> RECIPES_COOKING_METHOD <extra_id_10></p> |
| <p>Input: What is a low carb food Target: <RECIPES_QUALIFIER_NUTRITION> low <RECIPES_TYPE_NUTRITION> carb E TS: <RECIPES_QUALIFIER_NUTRITION> low <RECIPES_TYPE_NUTRITION> carb Sentinel Target = E S+T(S): <extra_id_3> RECIPES_QUALIFIER_NUTRITION <extra_id_4> RECIPES_TYPE_NUTRITION</p> |
| <p>Input: Read the top story on CNN.com Target: <NEWS_REFERENCE> top <NEWS_TYPE> story <NEWS_SOURCE> CNN.com E TS: Sentinel Target = E S+T(S): <extra_id_2> NEWS_REFERENCE <extra_id_3> NEWS_TYPE <extra_id_5> NEWS_SOURCE</p> |
| <p>Input: set timer for 28 minutes for cake to be done Target: <METHOD_TIMER> timer <DATE_TIME> for 28 minutes E TS: <METHOD_TIMER> timer <DATE_TIME> for 28 minutes <TIMER_NAME> cake Sentinel Target = E S+T(S): <extra_id_1> METHOD_TIMER <extra_id_2> DATE_TIME <extra_id_3> <extra_id_4></p> |
| <p>Input: what is the latest new york post headlines Target: <DATE_TIME> latest <NEWS_SOURCE> new york post <NEWS_TYPE> headlines E TS: <DATE_TIME> latest <NEWS_TOPIC> new york post <NEWS_TYPE> headlines Sentinel Target = E S+T(S): <extra_id_3> DATE_TIME <extra_id_4> NEWS_SOURCE <extra_id_5> <extra_id_6> <extra_id_7> NEWS_TYPE</p> |

Table 26: Examples of wins for **Extractive Sentinel+Tag (S)** vs the **Extractive TaggedSpans** on mTOP.

Worryingly, this is very common in all languages (including English).

Dataset: ATIS

Input: list airports

Label:

Prediction: <city_name> dallas

Dataset: mTOD

Input: will it pour ?

Label: <weather/attribute> pour

Prediction: <weather/attribute> para

Dataset: mTOD

Input: set reoccurring alarms

Label:

Prediction: <reminder/noun> reminders

Translated words / phrases: An interesting pattern we observe on some non-English datasets is where a word is accidentally translated or transliterated.

Dataset: mTOD

Input: apagar todas mi alarmas hoy

Label: <datetime> hoy

Prediction: <datetime> today

Dataset: mTOP

Input: Wie hoch ist die Temperatur

Label:

Prediction: <RECIPES_ATTRIBUTE> temperature

| |
|--|
| <p>Input: are there clear skies tonight Target = E TS: <WEATHER_ATTRIBUTE> clear <DATE_TIME> tonight Sentinel Target: <extra_id_2> WEATHER_ATTRIBUTE <extra_id_4> DATE_TIME E S+T(S): <extra_id_3> WEATHER_ATTRIBUTE <extra_id_4> DATE_TIME</p> |
| <p>Input: Can you play metal radio from Spotify for me Target = E TS: <MUSIC_GENRE> metal <MUSIC_TYPE> radio <MUSIC_PROVIDER_NAME> Spotify Sentinel Target: <extra_id_3> MUSIC_GENRE <extra_id_4> MUSIC_TYPE <extra_id_6> MUSIC_PROVIDER_NAME E S+T(S): <extra_id_3> MUSIC_PLAYLIST_TITLE <extra_id_4> MUSIC_TYPE <extra_id_6> MUSIC_PROVIDER_NAME</p> |
| <p>Input: Are there any severe weather advisories for the Pacific Northwest Target = E TS: <LOCATION> Pacific Northwest Sentinel Target: <extra_id_8> LOCATION <extra_id_9> E S+T(S): <extra_id_7> LOCATION <extra_id_8> <extra_id_9></p> |
| <p>Input: Any science related news Target = E TS: <NEWS_CATEGORY> science <NEWS_TYPE> news Sentinel Target: <extra_id_1> NEWS_CATEGORY <extra_id_3> NEWS_TYPE E S+T(S): <extra_id_1> NEWS_CATEGORY <extra_id_2> <extra_id_3> NEWS_TYPE</p> |
| <p>Input: i want the current headlines across US Target = E TS: <DATE_TIME> the current <NEWS_TYPE> headlines <NEWS_TOPIC> US Sentinel Target: <extra_id_2> DATE_TIME <extra_id_3> <extra_id_4> NEWS_TYPE <extra_id_6> NEWS_TOPIC E S+T(S): <extra_id_3> DATE_TIME <extra_id_4> NEWS_TYPE <extra_id_6> NEWS_TOPIC</p> |
| <p>Input: is there an update on the news story about Bill Clinton Target = E TS: <NEWS_TYPE> update <NEWS_TYPE> news story <NEWS_TOPIC> Bill Clinton Sentinel Target: <extra_id_3> NEWS_TYPE <extra_id_6> NEWS_TYPE <extra_id_7> <extra_id_9> NEWS_TOPIC <extra_id_10> E S+T(S): <extra_id_3> NEWS_TYPE <extra_id_6> NEWS_TYPE <extra_id_7> <extra_id_9> NEWS_TOPIC <extra_id_10></p> |

Table 27: Examples of losses for **Extractive Sentinel+Tag (S)** vs the **Extractive TaggedSpans** on mTOP.

Dataset: mTOP

Input: quel est le pourcentage de chances de pluie pour aujourd’hui

Label: <WEATHER_ATTRIBUTE> pluie <DATE_TIME> pour aujourd’hui

Prediction: <WEATHER_ATTRIBUTE> pluie <DATE_TIME> aujourd’hui

Label: <extra_id_2>

USER_ATTENDEE_EVENT

Prediction: <extra_id_5>

USER_ATTENDEE_EVENT

Model expects further tokenization: Sometimes the models expects the input to have been tokenized more than it was

F.2 (Extractive) Sentinel+Tag approach

Extra token / tag: The most frequent (albeit rare) observed is where an extra token is added to the output.

Dataset: mTOP

Input: <extra_id_0> Set

Label: <

Prediction: <extra_id_1> METHOD_TIMER

Dataset: MIT-Movies

Input: <extra_id_0> who <extra_id_1> stars <extra_id_2> in <extra_id_3> the <extra_id_4> movie <extra_id_5> titled <extra_id_6> happythankyoumoreplease

Label: <extra_id_6> TITLE

Prediction: <extra_id_6> TITLE <extra_id_7> <extra_id_8>

Dataset: mTOP

Input: <extra_id_0> Invite

Label:

Prediction: <extra_id_1> CONTACT_ADDED

Wrong sentinel number: Occasionally (and somewhat surprisingly) the model copies the wrong sentinel token index.

Dataset: mTOP

Input: <extra_id_0> get <extra_id_1> events <extra_id_2> I’m <extra_id_3> going <extra_id_4> to

G Examples of Wins and Losses

To help give a sense of the wins and losses we observed using the sentinel approach, we compared the **Sentinel+Tag** approach vs the current standard **TaggedSpans** approach on the largest dataset: mTOP. To make this analysis easier, we analyzed the extractive versions of the models as they are more succinct. The examples are provided in Tables 26 and 27.

In general we find that the sentinel model seems to be generally better at identifying and labeling the spans in the text. However we also find exam-

ples where occasionally the sentinel model adds or shortens a multi-word span (e.g. "Pacific Northwest" extended to "the Pacific Northwest", "science" extended to "science related", "the current" shortened to "current").