

Span Pointer Networks for Non-Autoregressive Task-Oriented Semantic Parsing

Akshat Shrivastava Pierce Chuang Arun Babu Shrey Desai
Abhinav Arora Alexander Zotov Ahmed Aly

Facebook

{akshats, pichuang, arbabu, shreyd,
abhinavarora, azotov, ahhegazy}@fb.com

Abstract

An effective recipe for building seq2seq, non-autoregressive, task-oriented parsers to map utterances to semantic frames proceeds in three steps: encoding an utterance x , predicting a frame’s length $|y|$, and decoding a $|y|$ -sized frame with utterance and ontology tokens. Though empirically strong, these models are typically bottlenecked by length prediction, as even small inaccuracies change the syntactic and semantic characteristics of resulting frames. In our work, we propose span pointer networks, non-autoregressive parsers which shift the decoding task from text generation to span prediction; that is, when imputing utterance spans into frame slots, our model produces endpoints (e.g., $[i, j]$) as opposed to text (e.g., “6pm”). This natural quantization of the output space also provides consistency in the length prediction task, allowing our length predictor to be responsible for frame syntax and the decoder for frame syntax, creating a coarse-to-fine model. We evaluate our approach on several task-oriented semantic parsing datasets. Notably, we bridge the quality gap between non-autoregressive and autoregressive parsers, achieving 87 EM on TOPv2 (Chen et al., 2020). Furthermore, due to our more consistent gold frames, we show strong improvements in model generalization in both cross-domain and cross-lingual transfer in low-resource settings. Finally, due to our diminished output vocabulary, we observe 70% reduction in latency and 83% in memory at beam size 5 compared to prior non-autoregressive parsers.

1 Introduction

Task-oriented conversational assistants typically first employ semantic parsers to map utterances to frames (Hemphill et al., 1990; Coucke et al., 2018; Gupta et al., 2018; Rongali et al., 2020; Aghajanyan et al., 2020). Due to performance constraints in real-world deployments, recent work in

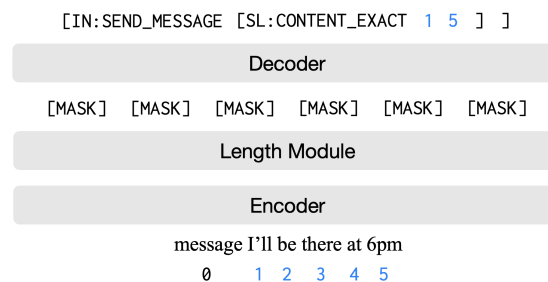


Figure 1: Illustration of a span pointer network for task-oriented semantic parsing based on seq2seq, non-autoregressive, mask-predict models (Ghazvininejad et al., 2019). In this example, we map the utterance x “message I’ll be there at 6pm” to the frame y $[IN:SEND_MESSAGE [SL:CONTENT_EXACT 1 5]]$ where $[1, 5]$ corresponds to the span “I’ll be there at 6pm”. Our model operates in three stages: (1) an encoder consumes the utterance; (2) a length module creates $|y| \times [MASK]$ tokens corresponding to the frame’s length; and (3) a decoder swaps each $[MASK]$ token with an utterance or ontology (e.g., intent or slot) token. Unlike typical methodology, we model span endpoints, which empirically simplifies the parsing task.

task-oriented semantic parsing has shifted towards building seq2seq, non-autoregressive parsers optimized for both quality and latency (Zhu et al., 2020; Babu et al., 2021). These models enforce strong independence assumptions during decoding, allowing frame components (e.g., ontology and utterance tokens) to be generated in parallel. However, the application of off-the-shelf, non-autoregressive algorithms to task-oriented semantic parsing is not trivial, often leading to brittle implementations with sub-optimal, opinionated components (Babu et al., 2021).

One popular family of seq2seq, non-autoregressive, task-oriented parsers is based on the mask-predict algorithm (Ghazvininejad et al., 2019), which operates in three steps: encoding, length prediction, and decoding. Unlike machine translation, task-oriented semantic parsing does

not benefit much from iterative refinement; therefore, frames are typically generated in one step. However, this increases the burden of upstream components, namely placing a major bottleneck on the length prediction module. Therefore, even off-by-one length errors can change the syntactic and semantic characteristics of resulting frames.

In this work, we create **span pointer networks** which model a seq2seq, non-autoregressive parsing task centered around span prediction as opposed to text generation. Figure 1 illustrates an application of our model; when parsing the utterance “message I’ll be there at 6pm”, our model produces span endpoints [1, 5] as opposed to exact text “I’ll be there at 6pm” in the appropriate leaf slot. As a result, our length module implicitly predicts the syntax of a frame (i.e., how many intents and slots does the utterance have?) and the decoder resolves leaf arguments with span endpoints (i.e., what utterance spans should be included in the frame?). This additionally creates a parallel among our approach and coarse-to-fine (Dong and Lapata, 2018) modeling as our length predictor predicts a general structure that our decoder further refines. Despite this shift in behavior, our approach is largely compatible with the typical mask-predict methodology, as we largely modify the frame representation and model architecture to be span-based.

We evaluate span pointer networks on three axes: quality, generalizability, and resources. First, we benchmark exact match (EM) on multiple task-oriented semantic parsing datasets, including TOPv2 (Chen et al., 2020) and TOP (Gupta et al., 2018). Our non-autoregressive parser is competitive, achieving 87 EM on TOPv2, and matches or exceeds autoregressive parsers in many cases. Second, we evaluate generalizability by setting up cross-domain and cross-lingual transfer experiments. Our non-autoregressive outperforms strong baselines, notably outperforming autoregressive parsers by +15 EM when averaged across 5 zero-shot multilingual settings. Third, due to the restricted decoder vocabulary, we measure resource savings, in particular, latency and memory usage. Compared to non-autoregressive models our parser achieves 70% reduction in latency and 83% reduction in memory with a beam size of 5.

To summarize, our contributions are: (1) We propose span pointer networks, non-autoregressive parsers which use span prediction as opposed to text generation; (2) We evaluate several aspects

of span pointer networks, conducting a series of experiments around quality, generalizability, and resources; (3) We empirically show improvements on all three axes, notably outperforming both non-autoregressive and autoregressive parsers.

2 Background and Related Work

Task-oriented semantic parsing broadly consists of mapping textual utterances to structured frames (Gupta et al., 2018; Aghajanyan et al., 2020; Chen et al., 2020; Rongali et al., 2020; Li et al., 2021). Frames are structured semantic representations of utterances, and are comprised of both ontology tokens (e.g., intents and slots) and utterance tokens. For example, we can map the utterance “message I’ll be there at 6pm” into a frame [IN:SEND_MESSAGE [SL:CONTENT_EXACT I’ll be there at 6pm]], which has several components: (1) a “send message” intent; (2) a “content exact” slot; and (3) a span “I’ll be there at 6pm”, which is an argument of the “content exact” slot. Here, intents and slots are akin to functions and arguments, respectively, in an API call.

Modern parsers are typically based on seq2seq transformers which encode utterances and decode frames. These parsers, by virtue of being seq2seq models, can be learned either in an autoregressive or non-autoregressive fashion. Autoregressive parsers generate frame components in sequence, while non-autoregressive parsers generate frame components in parallel. Both paradigms have their advantages and disadvantages, primarily trading-off latency and quality: while non-autoregressive models are optimized for inference and suitable in production settings, they are often qualitatively worse than autoregressive models due to strong independence assumptions during generation (Gu et al., 2018; Lee et al., 2018; Ghazvininejad et al., 2019; Kasai et al., 2020; Babu et al., 2021).

Our work focuses on improving the performance of one family of non-autoregressive parsers, namely those based on the mask-predict algorithm (Ghazvininejad et al., 2019). These parsers operate in three stages: (1) on the source-side, an encoder consumes an utterance x ; (2) using its representations, a length module predicts a frame’s length $|y|$ and creates $|y| \times [\text{MASK}]$ tokens; and (3) a length-conditioned decoder produces a frame y with ontology and utterance tokens. Frames are typically generated in one step as, unlike machine translation, task-oriented semantic parsing does not

benefit much from iterative refinement (Babu et al., 2021). Consequently, this increases the burden of accuracy on upstream components, especially placing a major bottleneck on the length prediction module. Even off-by-one length errors can dramatically change resulting frames, which is one of the drivers behind the quality gap between non-autoregressive and autoregressive models.

3 Span Pointer Networks

In this section, we introduce span pointer networks: seq2seq, non-autoregressive parsers which model a span-based, task-oriented semantic parsing task. Our approach is based on the mask-predict algorithm, which requires an encoder, length module, and decoder to map utterances to frames. However, our core idea is shifting from **text generation** to **span prediction** when creating leaf arguments in frames; that is, replacing text spans (e.g., “I’ll be there at 6pm”) with index spans (e.g., [1, 5]). Text generation and span prediction are *functionally* identical; though span prediction requires converting index spans to text spans post-hoc, both resolve to the same frame. However, text generation and span prediction are *technically* different, as they present distinct targets for non-autoregressive parsing. Note that our approach is different from Pasupat et al. (2019), who develop a CKY-based parser which maps spans to labels; in contrast, we explore span prediction with seq2seq modeling.

To build intuition for how these paradigms affect mask-predict parsers, in particular, consider the utterance “message I’ll be there at 6pm” and frames [IN:SEND_MESSAGE [SL:CONTENT_EXACT I’ll be there at 6pm]] and [IN:SEND_MESSAGE [SL:CONTENT_EXACT 1 5]], one with text spans (requiring text generation) and one with index spans (requiring span prediction). **For accurate parsing, a length module needs to understand both frame syntax and semantics with text generation, but only frame syntax with span prediction.** Specifically, with text generation, the length module outputs 9 [MASK] tokens, which requires it to implicitly guess the entire frame—both its syntactic structure and semantic arguments—before decoding. In contrast, with span prediction, the length module outputs 6 [MASK] tokens, which requires it to implicitly guess syntactic structure but leave semantic arguments to the decoder, which will subsequently resolve the exact span indices.

To make this argument more clear, consider the

frames [SL:ARTIST Beyonce] and SL:LOCATION Seattle, Washington]: these frames both have the same length despite being semantically distinct, so during fine-tuning, the length module learns syntax and the decoder learns semantics. We view this type of span prediction as analogous to coarse-to-fine modeling (Dong and Lapata, 2018); the length module predicts a coarse-grained frame structure and the decoder infills incomplete parts (e.g., leaf arguments) with fine-grained details.

Below, we elaborate on the core components required to implement span pointer networks: a span-based **frame representation** and the **model architecture**.

3.1 Frame Representation

Our span-based frame representation is a variant of the **decoupled** frame representation (Aghajanyan et al., 2020), following prior work in task-oriented semantic parsing (Aghajanyan et al., 2020; Li et al., 2021). This representation requires frames to mimic tree structures comprised of ontology tokens (e.g., intents and slots) and utterance tokens, where, critically, utterance tokens appear only as leaf arguments to slots. However, seq2seq parsers fine-tuned on decoupled frames typically require text generation, as leaf arguments consist of text spans. We refer to the original decoupled representation as the **canonical form** and propose two successive modifications, the **index form** and **span form** which, together, enable span prediction. A side-by-side view of these forms is depicted in Table 1, and we elaborate on the non-canonical forms below:

1. **Index Form:** We capitalize on the closed nature of the semantic parsing task, replacing utterance tokens with index tokens which instead “point” to utterance positions (Rongali et al., 2020). Here, the definition of an index is contingent on the tokenization algorithm used and requires careful preprocessing of utterances and frames. However, when used consistently, index form yields several advantages. Transformers can model index-based frames out-of-the-box by leveraging the positional embeddings of utterance tokens (Vaswani et al., 2017). Furthermore, we can significantly restrict the decoder vocabulary as the maximum sequence length (e.g., 100) is typically several orders of magnitude smaller than the size of popular subword vocabular-

Form	Utterance / Decoupled Frame
	message I'll be there at 6pm
Canonical	[IN:SEND_MESSAGE [SL:CONTENT [I'll be there at 6pm]]]
Index	[IN:SEND_MESSAGE [SL:CONTENT [1 2 3 4 5]]]
Span	[IN:SEND_MESSAGE [SL:CONTENT [1 5]]]

Table 1: **Comparison of the canonical, index, and span forms of the decoupled frame representation.** Given the utterance “message I’ll be there at 6pm”, we present several decoupled representation forms, which each convey the same information, albeit in a different manner. Our span-based, non-autoregressive parser uses the span form.

ies (e.g., 50K BPE (Sennrich et al., 2016; Liu et al., 2019) and 250K SentencePiece (Kudo and Richardson, 2018; Conneau et al., 2020)).

2. **Span Form:** From the index form, where utterance tokens are represented as index tokens, we create the span form, where the frame representation is simplified by collapsing index tokens into spans. Put simply, we modify leaf arguments to be index endpoints corresponding to utterance spans. To represent these spans, we use index tokens instead of text tokens primarily to eliminate ambiguity: if an utterance contains multiple instances of the same text token, it might be unclear which instance a span is referring to, but assigning tokens unique indices resolves this issue. In comparison to index form, one advantage of span form is that it makes the output space denser, improving the performance of the length module; for example, both the frames [IN:CREATE_CALL [SL:CONTACT John]] and [IN:CREATE_CALL [SL:CONTACT John Smith]] would share the same length class. The span form is also attractive in tail settings: it reduces the absolute length of the frame in open-text domains such as messaging since large utterance spans are naturally compressed.

These forms, as we have described them, are different “views” of the same underlying decoupled frame representation, which allows for simple interoperability. Implementation-wise, to enable span prediction, we create `to_span_form()` and `from_span_form()` functions, which we use to encode and decode gold frames, respectively.

3.2 Model Architecture

We create a seq2seq, non-autoregressive semantic parser which maps an utterance $\mathbf{x} = (x_1, \dots, x_n)$ to a frame $\mathbf{y} = (y_1, \dots, y_m)$, where the frame is preprocessed into the span-based representation

outlined above. Following Babu et al. (2021), we leverage the mask-predict algorithm (Ghazvininejad et al., 2019). We build our seq2seq model using an encoder, length module, and decoder; though, unlike prior approaches, we optimize the decoder to better support span prediction. Below, we elaborate on our model’s components and objective:

Encoder. First, we encode an utterance

$$\mathbf{h}_1, \dots, \mathbf{h}_n = \text{Encoder}(x_1, \dots, x_n)$$

where the encoder is a pre-trained transformer encoder (Vaswani et al., 2017), such as RoBERTa (Liu et al., 2019) or XLM-R (Conneau et al., 2020).

Length Module. Next, using the encoder’s hidden states, a length module uses an MLP to predict the frame’s length ℓ , subsequently creating $\ell \times$ [MASK] tokens:

$$[\text{MASK}]_1, \dots, [\text{MASK}]_\ell = \text{LengthModule}(\mathbf{h}_1, \dots, \mathbf{h}_n)$$

Decoder. Finally, using the length module’s [MASK] tokens and encoder’s hidden states, we decode a frame:

$$y_1, \dots, y_m = \text{Decoder}([\text{MASK}]_1, \dots, [\text{MASK}]_\ell; \mathbf{h}_1, \dots, \mathbf{h}_n)$$

where the decoder is a randomly-initialized transformer decoder. Because our frames are comprised of utterance and ontology tokens, we account for this structure by equipping our decoder with a pointer-generator module (See et al., 2017; Aghajanyan et al., 2020; Rongali et al., 2020); our decoder, therefore, builds a frame by either “copying” utterance tokens ($y_t \in V_{\text{copy}}$) or “generating” ontology tokens ($y_t \in V_{\text{gen}}$).

We critically adjust V_{copy} because our span-based frames consist of utterance *index* tokens as opposed to *text* tokens. Therefore, we compute the maximum index i' across all utterances, and initialize $V_{\text{copy}} = \{i : 0 \leq i \leq i'\}$. This setup naturally

supports span prediction by enabling the decoder to unambiguously reference the endpoints of utterance spans. Furthermore, because the number of index tokens (e.g., 100) is typically far smaller than the number of text tokens (e.g., 50-200K), our parser uses substantially less memory during train- and test-time.

Model Objective. Our model requires inference at two stages; the length module predicts the frame’s length ℓ and, conditioned on this length, the decoder predicts the frame \mathbf{y} . As such, we jointly optimize for two objectives—a **length loss** $\mathcal{L}_{\text{length}} = \text{NLL}(\ell^*, \ell) + \beta_1 \text{LS}(\ell)$ and **label loss** $\mathcal{L}_{\text{label}} = \text{NLL}(\mathbf{y}^*, \mathbf{y}) + \beta_2 \text{LS}(\mathbf{y})$ —for training the length module and decoder, respectively; though, both backpropagate to the encoder. For both objectives, we compute the negative log likelihood loss between the targets (ℓ^* , \mathbf{y}^*) and predictions (ℓ , \mathbf{y}) and, additionally, regularize over-confident predictions via label smoothing (Pereyra et al., 2017). Empirically, we find the length loss converges quickly, so we control its influence using a scalar mixing parameter λ_1 . We define the loss function as $\mathcal{L}_{\text{NAR}} = \mathcal{L}_{\text{label}} + \lambda_1 \mathcal{L}_{\text{length}}$.

We also integrate R3F (Aghajanyan et al., 2021), a trust-region-based algorithm which maintains the generalizability of pre-trained representations. More details are described in Appendix §B.

Beam Search. During inference, the length module predicts the top k lengths and we decode a parse for each length in parallel. This is similar to beam search in autoregressive decoding as we get k candidates from our model during inference; however, unlike autoregressive generation, mask-predict, non-autoregressive models are only capable of generating a single beam per target length (Ghazvininejad et al., 2019).

4 Experimental Setup

Our goal is to benchmark span pointer networks on three axes: **quality**, **generalizability**, and **resources**. We benchmark quality by inspecting exact match (EM) on semantic parsing datasets, generalizability by probing cross-domain and cross-lingual performance, and resources by measuring latency and memory usage during inference.

4.1 Datasets

We experiment with the following task-oriented semantic parsing datasets: TOP (Task Oriented Pars-

ing) (Gupta et al., 2018), TOPv2 (Chen et al., 2020), and MTOP (Multilingual Task Oriented Parsing) (Li et al., 2021). TOP evaluates parsers’ abilities to produce nested frames across the event and navigation domains. In comparison, TOPv2 consists of both linear and nested frames and extends TOP to the alarm, messaging, music, navigation, timer, and weather domains. While TOP and TOPv2 consist entirely of utterances in English, MTOP provides gold translations in Spanish, French, German, Hindi, and Thai, making it useful for evaluating parsers in multilingual settings.

4.2 Evaluation

Quality. When benchmarking our parser on task-oriented semantic parsing datasets, we primarily evaluate the exact match (EM) between predicted and gold frames. Note that, because our parser produces index tokens as opposed to text tokens in leaf arguments, we post-process predicted frames, mapping index \rightarrow text tokens.

Generalizability. Recent work evaluates the generalizability of task-oriented semantic parsers when the data distribution changes (Chen et al., 2020; Li et al., 2021). Similarly, we evaluate how generalizable our parser is in both cross-domain and cross-lingual settings, largely to understand whether span prediction is a fundamentally good paradigm for non-autoregressive parsing.

We conduct cross-domain experiments on TOPv2 (Chen et al., 2020). Our experiments proceed in two stages: we fine-tune on a high-resource, source dataset, then fine-tune on a low-resource, target dataset. Following Chen et al. (2020), we use the alarm, event, messaging, music, navigation, and timer domains as the source dataset and reminder and weather as (separate) target datasets. We also compare our parser’s EM at different levels of samples per intent slot (SPIS)¹, (Chen et al., 2020), a metric used to subsample target datasets; specifically, we compute EM at 10, 25, 50, 100, and 1000 SPIS.

In addition, we conduct cross-lingual experiments on MTOP (Li et al., 2021). We primarily evaluate the zero-shot capabilities of our parser when transferred to a different language. Specifi-

¹By performing random sampling with k SPIS, we ensure that at least k samples exist for each intent/slot label, ensuring full coverage over the ontology space. Empirically, the total number of examples sampled is $10 \times k$ SPIS, so, for example, we can roughly expect 100 unique examples with 10 SPIS. See Chen et al. (2020) for more details.

cally, we create a multilingual version of our parser using XLM-R (Conneau et al., 2020) as the pre-trained encoder. Then, we train our parser on English samples and test it on non-English samples (i.e., Spanish, French, German, Hindi, or Thai).

Resources. Production-ready conversational assistants typically have stringent resource requirements; if deployed in a real-world setting, how many resources would our parser require? We focus on latency (ms) and memory consumption (mb), as these are standard metrics practitioners track during deployment. To begin, we export our model with TorchScript (Paszke et al., 2019) and dynamically quantize it to compress its weights. Our benchmark server uses a Intel Xeon CPU with 256GB RAM, and we restrict our models to run on 4 threads. For latency, we run our TorchScript model on a CPU against the TOPv2 dev split (Chen et al., 2020). And, for memory, we track the amount of memory allocated or released during the execution of the model’s operators using the PyTorch profiler.² Unless specified otherwise, we report the 99th percentile (P99) latency and for memory consumption we measure maximum memory consumption on the longest parse.

4.3 Systems for Comparison

We chiefly compare against three task-oriented semantic parsing models, which combine a seq2seq transformer with a pointer-generator-based decoder and cover both autoregressive (AR) and non-autoregressive (NAR) training:

BART (AR). BART is a seq2seq transformer combining a transformer encoder and autoregressive transformer decoder, and is pre-trained with a denoising autoencoder objective on monolingual corpora (Lewis et al., 2020). For task-oriented semantic parsing, Aghajanyan et al. (2020) shows BART achieves state-of-the-art EM on multiple datasets.

RoBERTa AR. RoBERTa is a transformer encoder pre-trained using an optimized BERT objective on monolingual corpora (Liu et al., 2019). Because transduction-based semantic parsing requires a seq2seq model, (Aghajanyan et al., 2020; Rongali et al., 2020) combine a RoBERTa encoder and an autoregressive, randomly-initialized transformer decoder (3L, 768/1024H, 16/24A).

²https://pytorch.org/tutorials/recipes/recipes/profiler_recipe.html

Model	TOPv2	TOP
Type: Autoregressive Models (Prior)		
BERT _{LARGE} [◇]	—	83.13
RoBERTa _{LARGE} [◇]	—	86.67
RoBERTa _{LARGE} [◇]	—	84.52
BART _{LARGE} [♡]	—	87.10
Type: Autoregressive Models (Ours)		
RoBERTa _{BASE}	86.62	83.17
RoBERTa _{LARGE}	86.25	82.24
BART _{BASE}	86.73	84.33
BART _{LARGE}	87.48	85.71
Type: Non-Autoregressive Models (Ours)		
RoBERTa _{BASE}	85.78	82.37
+ Span Pointer	86.93	84.45
RoBERTa _{LARGE}	86.25	83.40
+ Span Pointer	87.37	85.07

Table 2: **EM performance on TOPv2 and TOP.** We close the quality gap between non-autoregressive and autoregressive parsing; our span pointer parser matches the BART parser, despite parallel decoding. Note that, unlike prior work, we minimize the amount of hyperparameter tuning, therefore absolute EM scores are not directly comparable. [◇]Rongali et al. (2020); [♡]Aghajanyan et al. (2020)

RoBERTa NAR. Unlike RoBERTa AR, RoBERTa NAR assumes strong independence assumptions during decoding, using the mask-predict algorithm to enable non-autoregressive generation (Ghazvininejad et al., 2019). We use the framework outlined in Babu et al. (2021), creating a seq2seq transformer with a RoBERTa encoder, MLP length module, and a non-autoregressive, randomly-initialized transformer decoder. Empirically, we find RoBERTa NAR relies on the transformer encoder more, so we make the transformer decoder shallower (1L, 768/1024H, 16/24A); these observations are also consistent with the recent “deep encoder, shallow decoder” findings in efficient machine translation (Kasai et al., 2021).

Because span pointer networks build on top of RoBERTa NAR, subsequently modifying the frame representation and model architecture to be span-based, we denote our parser as RoBERTa NAR + Span Pointer in our experiments.

5 Results and Discussion

5.1 Quality

Table 2 shows EM results on TOPv2 and TOP. We see that, with our span pointer formulation, we improve upon prior non-autoregressive parsers

	Weather Domain (SPIS)						Reminder Domain (SPIS)					
	10	25	50	100	1000	Avg	10	25	50	100	1000	Avg
Type: Autoregressive Models												
RoBERTa AR	69.71	74.90	77.02	78.69	86.36	77.34	49.38	56.98	62.18	69.17	78.48	63.24
BART AR	73.34	73.55	76.58	79.16	86.25	77.78	49.75	63.31	69.60	72.01	80.82	67.10
Type: Non-Autoregressive Models												
RoBERTa NAR	59.01	72.12	73.41	78.48	87.42	74.09	33.90	40.19	49.87	54.76	76.71	51.09
+ Span Pointer	72.03	74.74	74.85	78.14	88.47	77.65	48.27	60.55	68.11	73.19	80.44	66.11

Table 3: **Cross-domain generalizability experiments on TOPv2, comparing EM with SPIS.** We perform high-resource fine-tuning on multiple source domains, then low-resource fine-tuning on two target domains, weather and reminder, separately. Each target domain consists of multiple subsets for fine-tuning; these are created by randomly sampling k samples per intent slot (SPIS) (Chen et al., 2020).

by +1-2 EM. These gains are consistent across pre-trained encoders, particularly RoBERTa_{BASE}, indicating real-world applicability where smaller model architectures are preferred. We also compare against strong autoregressive parsers. Though non-autoregressive parsers typically underperform autoregressive parsers, often requiring additional tricks like distillation (Gu et al., 2018; Ghazvininejad et al., 2019; Zhou et al., 2020), we see that span pointer networks largely close this quality gap. The non-autoregressive RoBERTa_{LARGE} span pointer parser is comparable to the autoregressive BART_{LARGE} parser, the current state-of-the-art parser, despite being non-autoregressive.

5.2 Generalizability

We investigate our parser’s generalizability: how does it perform when the data distribution suddenly changes? We experiment with two types of distribution shifts, setting up cross-domain and cross-lingual experiments.

For cross-domain experiments, we conduct high-resource fine-tuning on multiple source domains, then low-resource fine-tuning on two target domains: weather and reminder. Table 3 lists individual EM scores at 10, 25, 50, 100, and 1000 SPIS as well as an average EM score across all SPIS values. Although there is a large EM gap between RoBERTa AR and RoBERTa NAR, our span-based, non-autoregressive parser closes this gap, matching or exceeding the EM both the RoBERTa and BART autoregressive parsers in most cases. Our parser notably does well on reminder, which is a more challenging domain given its large ontology and high compositionality (Chen et al., 2020).

For cross-lingual experiments, we perform zero-shot evaluations where a multilingual version of

our parser (using XLM-R as the pre-trained encoder) is trained on English samples and tested on non-English samples. Table 4 shows EM scores for five such zero-shot settings: English \rightarrow Spanish, French, German, Hindi, and Thai. Our span-based, non-autoregressive parser consistently outperforms both non-autoregressive and autoregressive baselines, achieving +14-15 average EM. We attribute the improvement in our model to be due to the length predictor indirectly predicting frame syntax which is language-agnostic. For example, [IN:GET_WEATHER [SL:LOCATION entertainment center]] and [IN:GET_WEATHER [SL:LOCATION centro de entretenimiento]] are English-Spanish parallel samples; using the canonical form, the length discrepancy makes it challenging for mask-predict models to learn the association, but by using the span form, our model is able to seamlessly bridge this gap. Our findings are consistent with prior work showing that non-autoregressive modeling can be beneficial in zero-shot multilingual settings (Zhu et al., 2020).

5.3 Resources

Our results above indicate that the RoBERTa_{BASE} NAR + Span Pointer parser is qualitatively strong, despite its smaller transformer encoder, but we have not yet investigated its resource requirements.

Using the base variants of BART AR, RoBERTa AR, and RoBERTa NAR with variable beam sizes ($k = 1$ and 5), we quantize each model to compress its weights, then benchmark latency and memory on a CPU-based server.³ Table 5 shows these results; from here, we make a couple of observations.

³We use 1 decoder layer for RoBERTa AR and RoBERTa NAR to ensure fair comparison, though note RoBERTa AR typically requires 3 decoder layers for best quality.

	Zero-Shot Evaluation						Avg
	en	en→es	en→fr	en→de	en→hi	en→th	
XLM-R _{BASE} NAR	78.3	35.2	32.2	23.6	18.1	16.7	25.2
+ Span Pointer	83.0	51.2	51.4	42.0	29.6	27.3	40.3
XLM-R _{LARGE} NAR	80.5	50.9	51.5	38.7	31.6	22.8	39.1
+ Span Pointer	84.5	60.4	63.1	56.2	41.2	41.7	52.5
XLM-R _{LARGE} AR [◇]	83.9	50.3	43.9	42.3	30.9	26.7	38.8

Table 4: **Cross-lingual generalizability experiments on MTOP, comparing EM across zero-shot setups.** We perform zero-shot experiments where we fine-tune a parser on English (en), then evaluate it a non-English language—Spanish (es), French (fr), German (de), Hindi (hi), and Thai (th)—without fine-tuning. Average EM (Avg) is taken over the five non-English languages. [◇]Li et al. (2021)

Model	EM (NQ / Q) ↑	# Params ↓	Latency (ms) ↓		Memory (mb) ↓	
			$k = 1$	$k = 5$	$k = 1$	$k = 5$
BART _{BASE}	87.00 / 84.67	221M	1,143	2,131	93	437
RoBERTa _{BASE} AR	86.51 / 86.26	183M	1,154	3,983	582	2,390
RoBERTa _{BASE} NAR	85.78 / 85.60	142M	149	680	52	211
+ Span Pointer (Index Form)	86.12 / 86.05	134M	145	199	23	35
+ Span Pointer (Span Form)	86.99 / 86.80	134M	134	208	22	42

Table 5: **Latency and memory benchmarking on TOPv2.** We report EM (non-quantized / quantized) in addition to, number of parameters, latency (ms) and memory (mb) at variable beam sizes ($k \in \{1, 5\}$). Our span pointer parser achieves the best quantized EM and resource reductions over both autoregressive and non-autoregressive baselines.

First, our quantized parser achieves both higher EM and lower latency/memory compared to autoregressive BART and RoBERTa parsers. When compared to BART, our parser cuts latency by $8.5\times$ and $10\times$ at $k = 1$ and 5 , respectively, which is largely due to high parallelism during generation. Second, we also improve upon RoBERTa NAR, the base non-autoregressive parser, which we attribute to the variants of the decoupled frame representation we explore. Specifically, index form targets memory by avoiding the need to store large subword vocabularies and span form targets latency by minimizing the size of leaf arguments in frames. Overall, with $k = 5$, our parser cuts latency by $3.2\times$ and memory by $4.9\times$, indicating its usability in real-world settings.

6 Analysis

Having evaluated span pointer networks on several axes, we now turn towards understanding the driving factors behind its performance.

6.1 Model Ablations

Our parser critically relies on three components: frame representation, model architecture, and fine-tuning. In Table 6, we present a series of ablation experiments to isolate the contribution of each com-

Model (RoBERTa _{BASE})	TOPv2	TOP
Span Pointer	86.99	84.74
Ablation: Representation + Architecture		
- Span Form	85.82	82.99
- Index Form*	85.45	82.50
Ablation: Fine-tuning		
- R3F	86.73	83.64

Table 6: Model ablation experiments, comparing EM scores when the representation + architecture and fine-tuning components are modified. *To isolate the index form’s contribution, we use the canonical frame representation (§3.1) and standard non-autoregressive architecture (§3.2).

ponent. First, we consider our span pointer network without the span and index forms (§3.1). Here, to remove index form, specifically, we also have to undo our architectural changes (§3.2), making the resulting model identical to RoBERTa NAR. Our complete parser achieves higher EM, suggesting our representation and architecture changes, together, have a positive impact on performance. Second, we consider our span pointer network with and without R3F fine-tuning. While our parser achieves better EM with R3F, especially on TOP, these results indicate our method is not entirely

contingent on better fine-tuning.

6.2 Span Prediction and Target Length

	Canonical	Span
# Length Classes	47	20
Mean Frame Length per Frame	1.94	1.0
Mean Frame Length	10.50	9.71
Max Frame Length	62	58

Table 7: Length distribution statistics of TOPv2 (Chen et al., 2020) frames when comparing canonical and span forms.

In Table 7 we show the length characteristics of both the canonical and span form on TOPv2 (Chen et al., 2020). Span form leads to a tighter and more consistent distribution as it only takes 1 length class to represent each unique frame, leading to many fewer length classes as well.

6.3 Frame Syntax vs. Semantics

	Decoupled Frame	
	Syn Only	Syn + Sem
Type: Base Models		
BART	88.29	86.10
RoBERTa AR	88.11	86.29
RoBERTa NAR	88.53	84.63
Type: Large Models		
BART	89.01	87.47
RoBERTa AR	88.40	85.94
RoBERTa NAR	89.32	85.66

Table 8: EM performance in “syntax only” (Syn Only) and “syntax + semantics” (Syn + Sem) settings. Here, “syntax only” and “syntax + semantics” refer to decoupled frames without and with leaf arguments, respectively. In the “syntax only” setting, the non-autoregressive parser outperforms, suggesting the generation of leaf arguments is a major bottleneck.

Our initial motivation for modifying non-autoregressive, mask-predict parsers stems from the argument that, when shifting to span prediction from text generation, a length module only needs to predict frame syntax as opposed to both frame syntax and semantics. We make the implicit assumption that frame syntax is *easier* to learn than frame semantics; we now devise an experiment to test this hypothesis.

We refer to a regular decoupled frame [IN:SEND_MESSAGE [SL:CONTENT [I’ll be there at 6pm]]] as “syntax + semantics”,

while a decoupled frame without leaf arguments [IN:SEND_MESSAGE [SL:CONTENT]] as “syntax only”. Here, the “syntax only” frame mimics a constituency tree, representing coarse-grained *structure* rather than fine-grained *meaning*. Using TOPv2, we create two training sets, one with “syntax only” frames and one with “syntax + semantics” frames. Table 8 shows EM scores when fine-tuning BART, RoBERTa AR, and RoBERTa NAR parsers on these training sets.

In the “syntax only” setting, RoBERTa NAR slightly outperforms both BART and RoBERTa AR, while in the “syntax + semantics” setting, it consistently lags behind. Recall that, operationally, the main difference between these settings is that, in “syntax only”, RoBERTa NAR’s length module is only responsible for frame structure, while in “syntax + semantics”, it is responsible for both frame structure and arguments. Our results suggest that the main difficulty in non-autoregressive modeling is handling leaf arguments, so simplifications of frame semantics (i.e., our proposed span form) are likely to improve quality.

7 Conclusion

In this work, we present span pointer networks for task-oriented semantic parsing: seq2seq, non-autoregressive models which focus on span prediction as opposed to text generation. Our approach requires creating a span-based frame representation and model architecture, which, together, enable simple and consistent non-autoregressive modeling. We benchmark our parser on three axes—quality, generalizability, and resources—and demonstrate real-world applicability by improving upon both non-autoregressive and autoregressive baselines. Future work can extend our approach to work with discontinuous spans, as is important in session-based (Aghajanyan et al., 2020) and free word order (Li et al., 2021) modeling.

References

- Armen Aghajanyan, Jean Maillard, Akshat Shrivastava, Keith Diedrick, Michael Haeger, Haoran Li, Yashar Mehdad, Veselin Stoyanov, Anuj Kumar, Mike Lewis, and Sonal Gupta. 2020. Conversational Semantic Parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta.

2021. Better Fine-tuning by Reducing Representational Collapse. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Ahmed Aly, Kushal Lakhota, Shicong Zhao, Mrinal Mohit, Barlas Oguz, Abhinav Arora, Sonal Gupta, Christopher Dewan, Stef Nelson-Lindall, and Rushin Shah. 2018. PyText: A Seamless Path from NLP Research to Production. *arXiv preprint arXiv:1812.08729*.
- Arun Babu, Akshat Shrivastava, Armen Aghajanyan, Ahmed Aly, Angela Fan, and Marjan Ghazvininejad. 2021. Non-Autoregressive Semantic Parsing for Compositional Task-Oriented Dialog. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Xilun Chen, Ashish Ghoshal, Yashar Mehdad, Luke Zettlemoyer, and Sonal Gupta. 2020. Low-Resource Domain Adaptation for Compositional Task-Oriented Semantic Parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- 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*.
- Li Dong and Mirella Lapata. 2018. Coarse-to-Fine Decoding for Neural Semantic Parsing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. Mask-Predict: Parallel Decoding of Conditional Masked Language Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2018. Non-Autoregressive Neural Machine Translation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis. 2018. Semantic Parsing for Task Oriented Dialog using Hierarchical Representations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS Spoken Language Systems Pilot Corpus. In *Proceedings of the Workshop on Speech and Natural Language*.
- Pavel Izmailov, Dmitry Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2018. Averaging Weights Leads to Wider Optima and Better Generalization. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Jungo Kasai, James Cross, Marjan Ghazvininejad, and Jiatao Gu. 2020. Non-Autoregressive Machine Translation with Disentangled Context Transformer. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Jungo Kasai, Nikolaos Pappas, Hao Peng, James Cross, and Noah A. Smith. 2021. Deep Encoder, Shallow Decoder: Reevaluating the Speed-Quality Tradeoff in Machine Translation. In *Proceedings of the International Conference on Learning Representations*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Taku Kudo and John Richardson. 2018. SentencePiece: A Simple and Language Independent Subword Tokenizer and Detokenizer for Neural Text Processing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*.
- Jason D. Lee, Elman Mansimov, and Kyunghyun Cho. 2018. Deterministic Non-Autoregressive Neural Sequence Modeling by Iterative Refinement. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pretraining for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Haoran Li, Abhinav Arora, Shuohui Chen, Anchit Gupta, Sonal Gupta, and Yashar Mehdad. 2021. MTOP: A Comprehensive Multilingual Task-Oriented Semantic Parsing Benchmark. *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A Fast, Extensible

- Toolkit for Sequence Modeling. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL): Demonstrations*.
- Panupong Pasupat, Sonal Gupta, Karishma Mandyam, Rushin Shah, Mike Lewis, and Luke Zettlemoyer. 2019. Span-based Hierarchical Semantic Parsing for Task-Oriented Dialog. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Proceedings of the Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. 2017. Regularizing Neural Networks by Penalizing Confident Output Distributions. In *Proceedings of the International Conference on Learning Representations (ICLR): Workshop Track*.
- Subendhu Rongali, Luca Soldaini, Emilio Monti, and Wael Hamza. 2020. Don't Parse, Generate! A Sequence to Sequence Architecture for Task-Oriented Semantic Parsing.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. 2020. Large Batch Optimization for Deep Learning: Training BERT in 76 minutes. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Chunting Zhou, Graham Neubig, and Jiatao Gu. 2020. Understanding knowledge distillation in non-autoregressive machine translation. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Qile Zhu, Haidar Khan, Saleh Soltan, Stephen Rawls, and Wael Hamza. 2020. Don't Parse, Insert: Multilingual Semantic Parsing with Insertion Based Decoding. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL)*.

A Training Details

Implementation. For RoBERTa AR and RoBERTa NAR, we reference the open-source implementation of Babu et al. (2021) in PyText (Aly et al., 2018). For BART AR, we reference its implementation in fairseq (Ott et al., 2019). For experimentation on TOPv2 (Chen et al., 2020), we use 8 16GB GPUs, due to its large dataset size, and for TOP (Gupta et al., 2018), we use 1 16GB GPU. For BART_{LARGE}, we use 1 32GB GPU.

Hyperparameters. We determine model hyperparameters by sweeping across the pre-defined ranges in Appendix D. For optimization of RoBERTa and XLM-R, specifically, we use Adam (Kingma and Ba, 2015) and learning rate schedulers; autoregressive models use exponential learning decay and non-autoregressive models use decay on plateau, following the guidance in Babu et al. (2021). For optimization of BART, we use stochastic weight averaging (SWA) (Izmailov et al., 2018) and LAMB (You et al., 2020) following Aghajanyan et al. (2020).

Though fine-tuning with R3F typically yields the best results on generation tasks (Aghajanyan et al., 2021), we treat it as a hyperparameter, and report the max performance both with and without R3F. We also perform an ablation in Appendix B.1 to demonstrate its importance.

B Better Fine-tuning with R3F

We also integrate R3F (Aghajanyan et al., 2021) into our proposed Span Pointer Networks, a trust-region-based algorithm which maintains the generalizability of pre-trained representations. Following Aghajanyan et al. (2021), we create auxiliary terms for both objectives

$$\begin{aligned}\mathcal{L}_{\text{R3F-length}} &= \text{KL}_S(\text{LengthModule}(x) || \\ &\quad \text{LengthModule}(x+z)) \\ \mathcal{L}_{\text{R3F-label}} &= \text{KL}_S(\text{Decoder}(x) || \\ &\quad \text{Decoder}(x+z))\end{aligned}$$

where KL_S represents the symmetric Kullback-Leibler divergence between a regular and noised input and $z \sim \mathcal{U}(-\sigma, \sigma)$. Then, we add these terms to the length and label losses, respectively, resulting in a smoother objective:

$$\mathcal{L}_{\text{R3F-NAR}} = \mathcal{L}_{\text{NAR}} + \lambda_2 \mathcal{L}_{\text{R3F-length}} + \lambda_3 \mathcal{L}_{\text{R3F-label}}$$

B.1 R3F Ablations

In Table 9 we show the impact of leveraging R3F (Aghajanyan et al., 2021) across our baseline model architectures. We see consistent improvements across generation strategies (autoregressive and non-autoregressive) for our BART and RoBERTa-based parsers.

Model	TOPv2	TOP
Type: Autoregressive Models		
RoBERTa _{BASE}	86.29	82.08
+ R3F	86.62	83.17
RoBERTa _{LARGE}	85.94	82.64
+ R3F	86.25	82.24
BART _{BASE}	86.10	83.62
+ R3F	86.73	84.33
BART _{LARGE}	87.48	85.53
+ R3F	86.96	85.71
Type: Non-Autoregressive Models		
RoBERTa _{BASE}	85.13	82.06
+ R3F	85.78	82.37
RoBERTa _{LARGE}	85.93	82.57
+ R3F	86.25	83.40

Table 9: EM scores of autoregressive and non-autoregressive parsers when leveraging R3F-based fine-tuning (Aghajanyan et al., 2021).

C Autoregressive Span Pointer Networks

Model	TOPv2	TOP
Type: Autoregressive Models		
RoBERTa _{BASE}	86.62	83.17
+ Span Pointer	86.72	84.54
RoBERTa _{LARGE}	86.25	82.24
+ Span Pointer	86.86	84.66
Type: Non-Autoregressive Models		
RoBERTa _{BASE}	85.78	82.37
+ Span Pointer	86.93	84.45
RoBERTa _{LARGE}	86.25	83.40
+ Span Pointer	87.37	85.07

Table 10: EM scores for auto-regressive Span Pointer Networks

In table 10 we present the results of autoregressive span pointer networks. We see that our formulation of leveraging span prediction helps in the autoregressive setting as well as the non-autoregressive setting. However, we see larger improvements in non-autoregressive span pointer networks, even surpassing the auto-regressive variants despite being non-autoregressive.

D Hyperparameters

We specify the details of our hyperparameter sweeps in two tables below. Table 13 specifies the auto-regressive hyperparameters for BART and RoBERTa models. Table 14 we specify the non-autoregressive hyperparameters for RoBERTa and XLM-R models used in the paper. For the auto-regressive and non-autoregressive models, the models were hand tuned for initial parameters and the final numbers reported were based on a hyperparameter sweep with a total of 36 runs. For BART (Lewis et al., 2020) based models, we keep optimization parameter consistent with Aghajanyan et al. (2020) rather than tuning the model ourselves. For our non-autoregressive models, both the baseline (Babu et al., 2021) and our proposed span pointer parser, we leverage the exact same hyperparameter sweep to ensure a fair comparison.

Parameter	Value
R3F λ	Swept over
Noise	Uniform
ϵ	0.00001

Table 11: R3F constant hyperparameters.

Parameter	Value
SWA LR	0.0002
Start Step	18K
Frequency	230

Table 12: Stochastic Weight Averaging (SWA) hyperparameters used for BART fine-tuning.

Parameter	Autoregressive Models			
	BART _{BASE}	BART _{LARGE}	RoBERTa _{BASE}	RoBERTa _{LARGE}
Optimizer	SWA-Lamb	SWA-Lamb	Adam	Adam
Learning Rate Scheduler	Exp-LR($\gamma=0.95$)	Exp-LR($\gamma=0.95$)	Exp-LR($\gamma=[0.5, 0.99]$)	Exp-LR($\gamma=[0.5, 0.99]$)
Learning Rate	[5e-7, 5e-3]	[5e-7, 5e-3]	[0.000001, 0.001]	[0.000001, 0.001]
Batch Size	16	16	{4,8,16}	{4,8}
R3F λ	0.01	0.01	0.01	0.01
# GPU	1	1	{1,8}	{1,8}
GPU Memory 32GB	16GB	16GB	16GB	16GB

Table 13: Hyperparameter values for autoregressive model architectures.

Parameter	Non-Autoregressive Models		
	RoBERTa _{BASE}	RoBERTa _{LARGE}	XLM-R _{LARGE}
Optimizer		Adam	
Learning Rate Scheduler		ReduceLR	
Learning Rate		[0.0002, 0.000002]	
Batch Size	{32, 64}	{8, 16}	16
Length Loss λ_1		[0.1, 1.0]	
R3F $\lambda_{\text{label}} (\lambda_2)$		0.001	
R3F $\lambda_{\text{length}} (\lambda_3)$		0.01	
$\beta_1(\mathcal{L}_{\text{labels}})$		[0, 0.2]	
$\beta_2(\mathcal{L}_{\text{length}})$		[0, 0.5]	
# GPU		{1,8}	
GPU Memory	16GB	16GB	32GB

Table 14: Hyperparameter values for non-autoregressive model architectures.