

# Generating Completions for Broca’s Aphasic Sentences Using Large Language Models

Sijbren van Vaals, Yevgen Matuselych, Frank Tsiwah

**Abstract**—Broca’s aphasia is a type of aphasia characterized by non-fluent, effortful and agrammatic speech production with relatively good comprehension. Since traditional aphasia treatment methods are often time-consuming, labour-intensive, and do not reflect real-world conversations, applying natural language processing based approaches such as Large Language Models (LLMs) could potentially contribute to improving existing treatment approaches. To address this issue, we explore the use of sequence-to-sequence LLMs for completing Broca’s aphasic sentences. We first generate synthetic Broca’s aphasic data using a rule-based system designed to mirror the linguistic characteristics of Broca’s aphasic speech. Using this synthetic data (without authentic aphasic samples), we then fine-tune four pre-trained LLMs on the task of completing agrammatic sentences. We evaluate our fine-tuned models on both synthetic and authentic Broca’s aphasic data. We demonstrate LLMs’ capability for reconstructing agrammatic sentences, with the models showing improved performance with longer input utterances. Our result highlights the LLMs’ potential in advancing communication aids for individuals with Broca’s aphasia and possibly other clinical populations.

**Index Terms**—Aphasia, assistive technology, agrammatic speech, generative AI, large language models, non-fluent speech, synthetic data generation.

## I. INTRODUCTION AND BACKGROUND

Aphasia is an acquired language disorder after a focal brain injury. Typically, aphasia has long-term consequences on the communication abilities of the affected individuals, and thus, negatively impacts their quality of life [1], [2] together with their relatives’ quality of life [3]. Within the spectrum of aphasia, Broca’s aphasia, the target group for this study, is a sub-type of aphasia described as having non-fluent, effortful, telegraphic speech production, but with relatively preserved comprehension. Given the significant language fragmentation in Broca’s aphasic speech, diverse treatment approaches have been proposed [4]. While studies have demonstrated the effectiveness of these treatments, mixed findings regarding long-term maintenance of gains post-treatment [5], [6] have led

This work was (partly) supported by the PPS grant of the Ministry of Economic Affairs and Climate Policy through CLICKNL, the Top Consortium for Knowledge and Innovation (TKI) in the Creative Industries. We would like to acknowledge the support of AphasiaBank for providing us with the data used for this study. (Corresponding author: Frank Tsiwah, [f.tsiwah@rug.nl](mailto:f.tsiwah@rug.nl)).

All authors are with the Center for Language and Cognition Groningen, University of Groningen, Oude Kijk in ’t Jatstraat, 26, 9712 EK, Groningen, the Netherlands (e-mails: [s.j.van.vaals@rug.nl](mailto:s.j.van.vaals@rug.nl), [yevgen.matuselych@rug.nl](mailto:yevgen.matuselych@rug.nl), [f.tsiwah@rug.nl](mailto:f.tsiwah@rug.nl)).

to the emergence of alternative communication aids. These alternative approaches, although beneficial, present some limitations such as inefficiency, lack of personalization, and poor generalization to real-world conversations [3]. Existing communication aids for people with aphasia (PWA) include low-tech options like picture or spelling boards and high-tech devices with pre-defined messages (such as making an order for coffee), but these tools often limit expression to basic ideas or messages and can be challenging for PWA to navigate [3], [7]–[11]. Furthermore, these alternative communication aids frequently require specific skills, such as keyboard use or navigating pictograms, which can pose significant challenges for many PWA, thereby restricting their ability to engage in daily conversations [3]. Navigating pictogram sets or relying on pre-defined texts often does not fully capture the contextual dependencies inherent in real-world conversations. Consequently, developing interventions that can replicate the fluid, context-dependent language patterns observed in spontaneous interactions with familiar partners, such as spouses and caregivers, is of considerable importance. The goal of the current study is to investigate the potential of reconstructing agrammatic language into full utterances by using Large Language Models (LLMs) fine-tuned on synthetic Broca’s aphasic data<sup>1</sup>.

### A. Linguistic characteristics of Broca’s aphasic language

Broca’s aphasia is an acquired language disorder typically resulting from brain injury (e.g., stroke) or brain degeneration, which impacts an individual’s ability to produce spoken and written language, with comprehension remaining relatively intact [12]–[15]. Broca’s aphasic spontaneous speech is characterized by several linguistic deficits such as omission of bound morphemes, function words (especially determiners), difficulty in producing verbs with complex argument structure, and impairment in complex syntactic structures [12], [16]–[23]. Verbs with one argument tend to be relatively easier to produce than verbs with two or more arguments [16], [24]. Content words generally tend to be preserved.

Despite the difficulties in production, comprehension abilities in Broca’s aphasia remain relatively unaffected. This discrepancy suggests that linguistic representations may be preserved in PWA, but access to these representations is impaired [16], [17], [20]. In summary, individuals with Broca’s

<sup>1</sup>The code and data are available online: [https://github.com/sijbrenvv/Completions\\_for\\_Broca-s\\_aphasia](https://github.com/sijbrenvv/Completions_for_Broca-s_aphasia)

aphasia produce highly agrammatic speech, with one- to two-word utterances in more severe cases.

### B. AI applications in aphasiology

In recent years, Natural Language Processing (NLP) techniques have been adopted in aphasia research as well as clinical applications, particularly in analyzing connected speech production by PWA [14], [15], [25], [26]. While traditional manual analysis of patients' speech transcription is time-consuming and labor-intensive, automated NLP approaches have significantly streamlined this process [14]. Applying NLP methods not only saves time but also provides consistent and accurate evaluations, potentially aiding in the diagnosis and classification of aphasia subtypes [25], which could ultimately inform more targeted treatment strategies.

Previous research has primarily focused on using NLP techniques for the automatic detection of aphasia (with potential application in diagnostic purposes), while largely overlooking the use of these techniques for developing assistive technologies to enhance communication for PWA, their spouses and caregivers. Recent scoping reviews underscore significant gaps in the application of AI within aphasiology, particularly in its therapeutic and assistive uses [27]–[29]. One review [27] reported that 7 out of their 77 reviewed studies explored applications of AI in therapeutic or assistive applications, while most research focused on automated assessment and diagnosis of aphasia. The review highlighted a predominant reliance on supervised machine learning for classification tasks, with limited use of advanced AI techniques for personalized rehabilitation or self-management. Similarly, [29] did not identify any studies integrating AI into augmentative and alternative communication (AAC) devices specifically tailored for aphasia rehabilitation. Together, these findings emphasize a disproportionate focus on diagnostic applications over functional communication support, revealing a critical need for AI-driven innovations that address real-world conversational challenges faced by PWA and their caregivers.

At the same time, [30] recently reported an innovative approach that could potentially contribute to the advancement of assistive technology by using LLMs for reconstructing agrammatic aphasic sentences. The goal of that study was twofold. First, LLMs require large amounts of training data, while large datasets of aphasic speech do not exist. Therefore, the study [30] proposes a system that takes as an input a corpus of written texts from healthy speakers (C4 dataset [31]), and generates synthetic aphasic utterances by modifying the sentences in the corpus so that their linguistic characteristics match those reported for Broca's aphasic speech. Second, the authors then use this synthetic dataset to fine-tune a pre-trained LLM, specifically the T5 model [31], on a sentence completion task. Although the study [30] showed promising results, it presents a few methodological problems. First, the written text in the C4 corpus (used to generate synthetic aphasic utterances) does not accurately represent daily spoken language, potentially limiting the real-life applicability of the resulting system. Second, even though the study [30] aims to demonstrate the potential of synthetic datasets for developing

communication aids for PWA, the lack of testing on real aphasic utterances makes it unclear to what extent the presented approach is applicable to actual clinical scenarios. Lastly, the study's use of the C4 dataset for generating synthetic aphasic sentences and evaluating the fine-tuned LLM causes a potentially serious issue of data leakage, since the C4 dataset was included in the original training data for the T5 LLM. A data leakage could lead to artificially inflated performance metrics and even compromise the validity of the model's evaluation and its future application on real aphasic speech.

The present study aims to address these problems by answering the following research question: To what extent can a pre-trained sequence-to-sequence LLM, fine-tuned on synthetic Broca's aphasic data, effectively reconstruct authentic agrammatic sentences produced by individuals with Broca's aphasia into complete, grammatically correct sentences? To address this research question, we first generate synthetic Broca's aphasic sentences based on spoken language sentences produced by neurotypical individuals. This is done using a rule-based system adapted from the above-mentioned study [30]. Second, we use our synthetic data to fine-tune a series of pre-trained LLMs on a sentence completion task, and then evaluate these LLMs on synthetic and authentic aphasic sentences.

## II. METHOD

We aim to assess the extent to which fragmented Broca's aphasic sentences can be reconstructed using sequence-to-sequence LLMs fine-tuned on synthetic aphasic data. Our overall methodological setup, therefore, consists of two steps: (1) synthetic data generation, and (2) fine-tuning LLMs on a sentence completion task using synthetic data.

For the synthetic data generation, we employ a tailor-made rule-based system inspired by the approach of [30]. Its rules are based on the linguistic features associated with Broca's aphasic speech, deduced from existing literature. This approach ensures that our synthetic data resembles Broca's aphasic speech in terms of its linguistic characteristics.

For the sentence completion task, we fine-tune four sequence-to-sequence LLMs from the T5 [31] and FLAN-T5 families [32] on the task of completing synthetic Broca's aphasic sentences generated by our rule-based system. We quantitatively evaluate the resulting models and then test the best-performing model on authentic Broca's aphasic sentences. Since the authentic sentences do not have ground-truth data (i.e., target completions), in this case we only assess the generated completions qualitatively.

### A. Synthetic data generation

To generate synthetic sentences, we use an adapted version of the general approach described in [30]. The system takes a natural language corpus and, using a set of predefined rules, modifies sentences so that they resemble fragmented sentences commonly produced by PWA. This approach effectively transforms each original sentence into a single synthetic aphasic sentence. Compared to the original system of [30], we use different sentence modification rates, additional features,

and rules to reduce the complexity and grammaticality of sentences.

First, we consider a natural language corpus (see Section II-C below for more details) and discard all sentences which are longer than 15 words or, following [30], contain punctuation (other than a comma or full stop) or special characters. We also extract verb and noun phrases using spaCy [33] and discard sentences with the ratio of noun phrases to verb phrases greater than 2, based on existing studies of aphasic speech [12], [34].

For the retained sentences, we use a spaCy pipeline [33] with the UDPipe English model [35] to obtain a dependency parse for each sentence and retrieve each word’s part-of-speech (PoS) tag. We then probabilistically apply a series of PoS-specific rules to each word (see Table I), with probabilities established based on [30] and our manual inspection of generated sentences. Table II shows examples of synthetic utterances next to their corresponding original utterances, and highlights how the synthetic data retains the core content of the original sentences while exhibiting the fragmented and simplified structure typical of Broca’s aphasic speech.

TABLE I  
RULES USED BY OUR SYNTHETIC DATA GENERATION SYSTEM.

POS tag	Rule	Prob.	Example
Noun	Change grammatical number	30 %	<i>table</i> → <i>tables</i>
Adjective / adverb / verb	Discard	50 %	<i>regular</i> → $\emptyset$
Pronoun (pos- sessive / demonstrative)	Replace with another pronoun of the same type	40 %	<i>this</i> → <i>those</i>
Determiner / adposition / particle	Discard	70 %	<i>the</i> → $\emptyset$
Other	Do not change	100 %	N/A

In some extreme cases, a simultaneous application of multiple rules could result in having unrealistically short utterances. Therefore, as a final step, we exclude synthetic utterances that do not adhere to either or both of the following length requirements: (1) the synthetic utterance must have at least three words, and (2) the number of words in a synthetic utterance must be within the range 25–75% of the number of words in the respective original utterance.

Ideally, our synthetic sentences should be similar to authentic aphasic sentences and different from sentences produced by neurotypical individuals. There are many possible ways to compare sentence characteristics, but what matters in our setup is that the patterns described above hold in the context of LLMs. This is why we estimate the average acceptability (or predictability) of a sample of sentences in each data set using two language models commonly used for this purpose, GPT-2 [36] and RoBERTa [37]. Predictability is estimated through the standard measure called surprisal, or negative log-probability [38], [39], which measures how easy it is on average for a language model to predict words in a sentence given their context. We apply this approach to sentences produced by neurotypical individuals, PWA, and our synthetic data generation

TABLE II  
EXAMPLES OF SYNTHETIC UTTERANCES GENERATED BY THE  
RULE-BASE SYSTEM.

Original utterance	Synthetic utterance
But we are talking just the regular light horses you know.	But we are just light horse you know.
That band had sung that very night.	Band have very night.
Which just does not sound like a very practical proposal.	Which do sound a very proposal.
I mean crumbed and in the jar.	I crumb and the jar.

system. We only consider sentences that are longer than 3 words and downsample each data set to 4,448 sentences (the size of the smallest, authentic aphasic, data set after filtering out short sentences). We expect that the language produced by neurotypical individuals should be highly predictable and, therefore, associated with lower surprisal, while both authentic and synthetic authentic data should be more irregular, and, therefore, less predictable and associated with higher surprisal.

The presented approach relies on two types of data. To produce a number of synthetic utterances sufficient for fine-tuning a LLM, our system needs a large set of utterances produced by neurotypical individuals. To evaluate our system, we additionally need a set of authentic utterances produced by PWA. We describe the relevant data sources in Section II-C below, but first we present our sentence completion models.

## B. Sentence completion

We use a synthetic dataset obtained as described above to fine-tune sequence-to-sequence LLMs on the task of sentence completion. We choose four large pretrained encoder-decoder models based on their performance in common NLP tasks, the availability of their source code, the computational resources available to us, and, for T5, its use in an earlier study that is most similar to ours [30]. Each LLM has been pretrained on large amounts of text data, and in our setup we only fine-tune it on our target task.

1) *Models*: T5-BASE and T5-LARGE [31] belong to the Text-to-Text Transfer Transformer (T5) family. These models are based on the transformer architecture [40] and leverage transfer learning through the use of task-specific prefixes that allow framing any NLP problem as a text-to-text task. The models consist of encoder and decoder components with stacked layers, self-attention mechanisms, and feed-forward networks. There are several versions of the model that differ in the number of parameters, and we use two of them: T5-BASE is the baseline version of the model with 220M parameters, and T5-LARGE has 770M parameters. While there are even larger models in the T5 family, we could not use them in this study due to the lack of appropriate computational resources. All T5 models have been pretrained on the large Colossal Clean Crawled Corpus (C4) dataset, comprising 750 GB of text crawled from the web.

FLAN-T5-BASE and FLAN-T5-XL belong to the FLAN-T5 family [32]. These models are based on the T5 models but are extended using instruction fine-tuning on a large variety

of tasks (dialog, arithmetic reasoning, etc.), which improves the models’ generalization to unseen tasks. FLAN-T5-BASE has 250M parameters, while FLAN-T5-XL has 3B parameters. Again, we could not use larger versions of the FLAN models due to our computational limitations.

2) *Fine-tuning setup*: We fine-tune each of the four models on a synthetic aphasic dataset generated as described in the previous section. During this process, the models learn to reconstruct (synthetic) fragmented Broca’s aphasic sentences. We use ChrF (CHaRacter-level F-score [41]) as the evaluation metric for updating the models’ parameters. To take into account possible influence of prefixes on the models’ performance, we fine-tune the models in two conditions: without any prefix and with prefix “Complete this sentence” (for comparison to [30]).

We use the HuggingFace transformer library [42] to load and quantise each model to 4-bit precision with a batch size of 8. For efficiency, we employ HuggingFace’s Parameter-Efficient Fine-Tuning (PEFT) library [43] that incorporates Quantized Low-Rank Adaptation (QLoRA) [44] (see more details in Supplementary Materials). Using these techniques, we fine-tune each model. T5-BASE is fine-tuned for a maximum of 5 epochs with early stopping after 2 epochs with no improvement, using a learning rate of  $10^{-4}$  with half precision (FP16). The other models are fine-tuned for 1 epoch only due to overfitting. Supplementary Materials provide technical details about fine-tuning each model.

To evaluate our sentence completion models, we use two different methods. First, we compare the models’ completions of (unseen) synthetic sentences to their original counterparts (i.e., the same sentences before they are processed by our synthetic data generation system). This comparison is done using three measures: ChrF and RougeL [45] between the two sentences (mentioned above), and cosine similarity between the embeddings of the two sentences (using embeddings of [46]). While ChrF and RougeL measure the degree of surface-level (i.e., character and word, respectively) alignment between the two sentences, cosine similarity measures how close the meanings of the two sentences are. Although BLEU [47] is a popular evaluation metric in sentence comparison tasks, we found it to be unsuitable for our sentence completion task. The explanation with examples can be found in Supplementary Materials. Second, we use our models to generate completions for authentic sentences produced by PWA. In this case, we have no ground-truth data to compare our models’ completions to, and we limit our evaluation to qualitative analyses.

### C. Data sources

Considering the setup described in the two previous sections, we need two main types of data: (1) authentic utterances produced by PWA, to be used for the evaluation of our synthetic data generation system and our sentence completion models, and (2) utterances produced by neurotypical individuals, to be used as input to our synthetic data generation system, as well as for its evaluation. Ideally, the two types of data should be similar in their register or even content. We use the following data sources.

1) *Authentic aphasic data*: To obtain authentic data produced by PWA, we use AphasiaBank [48], a foundational resource in aphasia research containing interactions between PWA and clinicians. It provides speech data collected from PWA during clinical studies. We use this data for the final evaluation of our sentence completion models, where the models generate completions for these utterances.

2) *Neurotypical individuals’ data from AphasiaBank*: AphasiaBank also contains utterances produced by neurotypical individuals (i.e., the control condition in many clinical studies) in the same setting as our authentic aphasic data. Thanks to their close similarity to the aphasic data, these utterances perfectly suit our goal of evaluating the quality of our synthetic Broca’s aphasic utterances. At the same time, we decided against using this data as an input to our sentence generation system, because the content of these utterances can be very similar to that of the authentic aphasic utterances: for example, both neurotypical individuals and PWA could be describing the same picture as part of the data collection procedure. Therefore, feeding the neurotypical utterances from AphasiaBank into our sentence generation system would yield synthetic data that’s potentially too similar to our evaluation data, leading to an unfair evaluation of our sentence completion models.

3) *Neurotypical individuals’ data from SBCSAE*: The Santa Barbara Corpus of Spoken American English (SBCSAE) [49] is a compilation of naturally occurring spoken interactions containing approximately 250k word tokens. SBCSAE is comparable to AphasiaBank in that it consists of transcriptions of spoken language. Finally, the corpus is very diverse and representative, as the data was collected from individuals of varied regional origin, age, ethnic and social background, occupation, and gender. We use utterances from SBCSAE to provide input to our synthetic data generation system, as well as to evaluate this system.

The sentences in both AphasiaBank and SBCSAE include linguistic annotations, special characters, and other types of noise. We preprocess the data sources to extract only the actual word tokens and remove all annotations (i.e., pauses, actions, text within angle and square brackets), unicode errors and special characters, instances of stuttering and repeated words, etc. A full list of preprocessing routines with examples can be found in Supplementary Materials.

## III. RESULTS

In this section, we first evaluate our synthetic data generation system. We then assess the LLMs’ ability to reconstruct synthetic Broca’s aphasic sentences. Finally, we evaluate the best-performing LLM on completing authentic sentences produced by PWA.

### A. Evaluating synthetic data generation system

To evaluate our synthetic data generation system, we consider the average sentence predictability (surprisal) values shown in Table III. Note that GPT-2 and RoBERTa have very different architectures and are trained on different data sets, therefore, the surprisal values *across* the two models (columns) are not directly comparable, and we only compare the values

TABLE III  
AVERAGE SENTENCE PREDICTABILITY (SURPRISAL) ASSIGNED TO EACH DATA TYPE BY TWO PRETRAINED LLMs. WHILE THE SYNTHETIC APHASIC SENTENCES ARE NOT DIRECTLY EXTRACTED FROM ANY SPECIFIC SOURCE, THEY ARE GENERATED BASED ON THE SBSCAE DATA.

Data type	Source	GPT-2	RoBERTa
Neurotypical	SBSCAE	4.76	2.96
Neurotypical	AphasiaBank	5.43	3.50
Authentic aphasic	AphasiaBank	6.48	5.08
Synthetic aphasic	– (SBSCAE)	6.39	5.51

TABLE IV  
EVALUATION RESULTS OF THE SENTENCE COMPLETION MODELS ON THE SYNTHETIC TEST SET. MEAN AND STANDARD ERROR ARE PROVIDED FOR EACH MEASURE.

Model	Without prefix			With prefix		
	ChrF	RougeL	Cos. sim.	ChrF	RougeL	Cos. sim.
T5-BASE	56.39±0.48	0.71±0.00	0.88±0.00	56.56±0.48	0.71±0.00	0.88±0.00
FLAN-T5-BASE	49.53±0.42	0.67±0.00	0.87±0.00	50.15±0.42	0.67±0.00	0.87±0.00
T5-LARGE	53.95±0.46	0.68±0.00	0.88±0.00	54.94±0.47	0.70±0.00	0.88±0.00
FLAN-T5-XL	55.71±0.47	0.69±0.00	0.88±0.00	56.19±0.48	0.69±0.00	0.88±0.00

within each model (column). We first notice that the two data sets that contain sentences from neurotypical individuals (rows 1–2) are associated with lower surprisal values compared to the authentic aphasic data (row 3), and this result holds for the values generated by both language models, the pattern we expected. Second, our synthetic data (row 4) yields surprisal values that are comparable to the values for the authentic aphasic data (row 3) and higher than those for the neurotypical data (rows 1–2). This is an important result that suggests that our synthetic generation system produces sentences where words are harder for LLMs to predict from context – similar to authentic aphasic sentences and in contrast to neurotypical language. Upon manual inspection, our generated data also appears satisfactory, therefore, we proceed with using the synthetic Broca’s aphasic data generated by our system to fine-tune the four LLMs.

### B. Evaluating sentence completion models

1) *Synthetic data*: In this setting, we compare the agrammatic (synthetic) sentences completed by our four LLMs to the original sentences.

T5-BASE model demonstrates consistent performance on the synthetic test set regardless of whether the prefix “Complete this sentence” is used. The other models show slightly improved performance with the prefix, which suggests that they may benefit from additional contextual cues provided by the prefix.

A visual inspection of the generated completions (see examples in Table V) reveals that the models, especially T5-BASE and FLAN-T5-XL, show great capacity in addressing some of the main deficits observed in Broca’s aphasic utterances, such as word omission. While T5-BASE achieves the highest performance scores, it occasionally generates semantically incorrect completions, which highlights limitations in its ability to generate contextually appropriate sentences. FLAN-T5-XL, on the other hand, performs consistently well, often

TABLE V  
EXAMPLES OF LLMs’ COMPLETIONS FOR SYNTHETIC UTTERANCES. PREFIX USED IN ALL CASES.

Model	Synthetic utterance	Generated completion
T5-BASE	I do i thethey they did	I do not know i do not know thethey did not know.
	I voting favor of it.	I am voting in favor of it.
	I said you me better than i myself.	I said you would know me better than i did myself.
	Lenore what it.	Lenore knows what it is.
FLAN-T5-XL	They our lawns.	They are going to have to cut their lawn.
	Big or go all.	Big or go all the way.
	No you it.	No you can not do it.
	I will worry it.	I will not worry about it.
	I will him call.	I will call him up.
	It lot first.	It was a lot of fun at first.

opting for simpler but semantically plausible completions. This tendency reveals a trade-off between complexity and fluency, and suggests that reconstructing more morphosyntactically complex utterances is a challenging task for our models. More examples of generated completions for the synthetic test set are presented in Supplementary Materials.

2) *Authentic data*: In this setting, we have no ground-truth data for authentic Broca’s aphasic utterances. Therefore, for these utterances, we only present a qualitative analysis of sentence completions generated by T5-BASE and FLAN-T5-XL fine-tuned using the prefix “Complete this sentence”. We decided to focus on these two models as they achieved the highest performance on the synthetic data.

Table VI (top half) presents examples of completions generated by the two LLMs. Although the models exhibit some capability of handling Broca’s aphasic speech, important limitations remain. In particular, both models sometimes only minimally change the authentic utterance or even simply reproduce it. Also, they occasionally add negations (2.99% of the time for FLAN-T5-XL), which can change the sentence meaning to the opposite, or expand contractions wrongly (e.g., “He’s” to “He is” instead of “He has”). The occasional addition of negation reflects the general difficulty of this phenomenon for natural language processing models [50], [51].

While there are instances of grammatical completions, the models often produce outputs that do not correct the grammatical errors present in the authentic Broca’s aphasic utterances. This indicates a gap in the models’ ability to fully interpret and rectify the Broca’s aphasic speech patterns. Both models occasionally insert common conversational fillers or affirmations such as “Yeah”, “Okay” and “I do not know”, which do not meaningfully contribute to completing the utterance but reflect an attempt to provide a plausible conversational continuation. There are some utterances so agrammatic that completing them meaningfully would be challenging even for humans. In these cases, the models either do not change the utterances at all or provide completions that do not improve clarity. For example, “Peanut, b@you,.” by T5-BASE remains largely unreadable.

To estimate how well FLAN-T5-XL performs in the target task, we have manually annotated 10% of our test data, namely

TABLE VI  
 EXAMPLES OF LLMs’ COMPLETIONS FOR ALL AUTHENTIC UTTERANCES (TOP HALF) AND WITH ONE-WORD AND TWO-WORD UTTERANCES EXCLUDED (BOTTOM HALF). THE PREFIX IS USED IN ALL CASES.

Authentic aphasic utterance	Generated by T5-BASE	Generated by FLAN-T5-XL
Two slice. Oh ohman. Hold on a second. Peanut , bæ@you .. So it is youknow cinderella beautiful person youknow. Because i it nothing. Everything’s okay. Because kid in here. So that is. Nice !. And ride away because bebecause cinderella was good.	Two slices. Oh ohman. Hold on for a second. Peanut, b@you.. So it is you know cinderella is a beautiful person youknow. Because i thought it was nothing. Everything’s okay. Because there is a kid in here. So that is. Nice! And then ride away because bebecause cinderella was good.	Two slices. Oh ohman. Hold on a second. Peanut, b@you.. So it is youknow cinderella beautiful person youknow. Because i do not know it nothing. Everything is okay. Because the kid is in here. So that is what it is. Nice!. And ride away because bebecause cinderella was good.
What in the world is this ?. And he is. Oh is it hid feet. And peanut butter is. I do it. You hafta come with these. It was old time. And this is with them. Ah i have lost my. I know say. And you just.	What in the world is this? And he is. Oh is it hid under the feet. And peanut butter is not. I do not know how to do it. You hafta come with this. It was old time. And this is what is with them. Ah i have lost my. I know say. And you just.	What in the world is this? And he is not. Oh is it hid under my feet. And peanut butter is. I do not do it. You hafta come with these. It was an old time. And this is what is going on with them. Ah i have lost my mind. I know what to say. And you just just sat there.

455 sentences. We did the annotation without considering their broader context, to simulate the task that the model was faced with. We found that 196 sentences did not require any completion by the model, while 259 did. From the sentences that needed to be completed, 66 would not be possible to complete, given the lack of the broader context, while 193 could in principle be completed. Out of these, the model has successfully completed 80 sentences, simply reproduced 57 of them, and produced errors in 56 sentences (42 cases of incomplete/minimal change, 12 unnecessary additions, and 2 cases of added negation). On top of that, the model introduced unnecessary changes into the 13 (out of the 196) sentences that did not require completion.

Very short utterances might be difficult for our models to complete not only because they are often cryptic, but also because our models were not fine-tuned on this task: recall from Section II-A above that our synthetic utterances had to contain at least three words. Table VI (bottom half) presents examples of completions for longer and more complex utterances (three words or more, see Supplementary Materials for more examples). For such longer utterances, the models demonstrate a better grasp of context, producing more complete and semantically coherent responses, with less redundancy and fewer nonsensical completions. Even for uncommon phrases, the models are better at generating plausible continuations. For instance, “Oh is it hid feet” was completed to “Oh is it hid under the feet” by T5-BASE, which, despite a minor grammatical error, shows an effort to logically complete the utterance.

#### IV. DISCUSSION

In this study, we aimed to explore the potential of reconstructing Broca’s aphasic language into full utterances using

LLMs fine-tuned on synthetic Broca’s aphasic data. Using a rule-based system, we generated synthetic data based on transcribed spoken language from neurotypical individuals. Our synthetic data mirrors the characteristics of Broca’s aphasic speech. Subsequently, we fine-tuned four LLMs on this synthetic data to evaluate their ability to complete agrammatic aphasic sentences. To the best of our knowledge, this is the first study to evaluate LLMs on the task of reconstructing authentic Broca’s aphasic sentences.

While beneficial, current rehabilitation methods for Broca’s aphasia often fall short in the long-term sustainability of treatment, demand considerable time and resources and do not fully cover the diverse communication challenges faced by PWA. LLM-based approaches provide a possibility to overcome these restraints by offering ongoing and consistent support tailored to the specific needs of a patient, improving their quality of life and that of their spouses and caregivers.

##### A. Automatic agrammatic sentence completion

1) *Synthetic data*: We found that the models performed well on synthetic aphasic data, generating coherent and contextually appropriate completions. The T5-BASE and FLAN-T5-XL models showed consistent performance in the synthetic setting, with FLAN-T5-XL exhibiting a slightly better use of context and producing slightly better completions. These findings are consistent with the study of [30], who demonstrated the utility of synthetic data for training LLMs to reconstruct agrammatic aphasic sentences. However, while [30] used written text from the C4 corpus to generate synthetic aphasic utterances, our approach is based on spoken language, which better reflects conversational contexts encountered by PWA. Additionally, unlike [30] who evaluated their model solely on synthetic

data, we extended the evaluation to authentic aphasic sentences, addressing a key limitation in their study. This broader evaluation provides a more realistic assessment of the models' applicability to clinical scenarios.

*2) Authentic data:* When tested on authentic Broca's aphasic sentences, the models' performance varied. The models were able to generate reasonable completions based on visual inspections, with better results observed for longer input sentences, which provide more contextual information, resulting in improved model performance. Visual inspections indicated that FLAN-T5-XL generated more coherent and contextually relevant completions compared to T5-BASE, particularly for longer utterances. In contrast to [30] study, which did not test their approach on real aphasic data, our findings demonstrate that fine-tuned LLMs can generalize effectively to authentic agrammatic sentences produced by individuals with Broca's aphasia. This step highlights the potential and the practical relevance of such models for rehabilitation purposes. While previous AI applications in aphasia have largely focused on diagnostic tasks or static assistive technologies [3], [27]–[29], our results suggest that LLM-based systems have the potential to provide adaptive solutions for reconstructing agrammatic speech into complete sentences. Incorporating longer input utterances further improves model performance, suggesting the importance of designing systems capable of leveraging contextual information to support real-world communication needs for PWA and their caregivers.

*3) Implications:* The demonstrated ability of sequence-to-sequence LLMs to generate reasonable completions for Broca's aphasic sentences highlights their potential in advancing communicative technologies for non-fluent aphasic patients. These models could become valuable tools in speech and language therapy by providing contextually appropriate sentence completions, potentially helping individuals with Broca's aphasia convey their thoughts more effectively.

Moreover, the study emphasises the importance of leveraging authentic data to build robust models for real-world applications. The performance improvements observed when focusing on longer utterances imply that taking into account a minimum input length could enhance the effectiveness of such models in practical settings.

Beyond the application to Broca's aphasia, the methodology used for generating synthetic data could be expanded, in order to build a comprehensive database of synthetic speech and language data not only for individuals with Broca's aphasia but also for other types of aphasia and potentially for other clinical populations with speech and language disorders.

Synthetic data also addresses ethical concerns such as patient privacy and consent, providing an anonymous alternative. It provides a cost-effective way to create large datasets, which reduces the need for expensive, time-consuming data collection [52], although it is important to keep in mind that no synthetic data can completely replace authentic human data.

## B. Limitations

The study's focus on Broca's aphasia possibly limits the generalisability of the findings to other types of non-fluent

aphasia or related language impairments, since different forms of non-fluent aphasia present distinct linguistic challenges. The study also does not fully address the variability in individual language use and recovery patterns among people with Broca's aphasia, leading to suboptimal performance for certain users.

The computational resources required for training and fine-tuning large transformer-based models are substantial. This limitation may restrict the accessibility and scalability of the proposed approaches, specifically for researchers and practitioners with limited computational resources. Moreover, we only trained and tested the models on one high-resource language, English, and future work should determine whether our proposed method can also be used in other languages.

Finally, recovering the intended meaning of agrammatic aphasic sentences without extensive context or direct interaction with the individual is inherently challenging. For instance, the fragment "He book table" could be interpreted in multiple ways such as "He put the book on the table", "He took the book from the table" or "He wants the book on the table." This variability emphasises the challenge that a sentence cannot always be accurately recovered without additional grounding such as understanding the broader conversation or using external data such as visual cues.

## C. Future work

In terms of synthetic Broca's aphasic data generation, one route for future research is the exploration of data augmentation techniques tailored to agrammatic aphasic sentences, improving generalisation capabilities of the models on real-world Broca's aphasic speech.

Another interesting direction for future work is studying the domain influence, possibly enabling the creation of task-specific models based on the patient's interests or hobbies. The incorporation of domain knowledge has demonstrated improved model performance in NLP, including machine translation [53] and NLP-based educational applications [54]. In terms of training setup, it could be promising to explore in-context learning methods with the same models, and compare the resulting scores with those reported in our study (similar to, e.g., [55]). Our preliminary experiments with FLAN-T5-XL in a ten-shot learning setup suggest that fine-tuning methods yield better performance (on ChrF and Rouge), but experimenting with a variety of prompts may be necessary in the future.

Finally, we recognize that the observed correct sentence completion rate is not sufficient for deployment as a ready-to-use communication aid in clinical settings. Nonetheless, this work represents an initial step toward building such applications, since there are currently no existing benchmarks for performance on authentic aphasic sentences. Therefore, future research could ask native English speakers to produce corresponding target completions for the authentic Broca's aphasic utterances from AphasiaBank. The availability of such ground-truth data is crucial for a quantitative evaluation of models' completions of the authentic Broca's aphasic sentences. The creation of a reliable benchmark would provide a reference for evaluating how well the sentence completion

models handle typical errors and omissions in Broca's aphasia. Moving beyond the sentence completion task, such data would contribute immensely to creating a broader benchmark for related tasks in the field of aphasiology.

## V. CONCLUSION

In summary, we conclude that sequence-to-sequence large language models exhibit the capability of reconstructing agrammatic sentences produced by individuals with Broca's aphasia. The models generate reasonable and contextually appropriate completions, denoting their potential as valuable tools in speech and language therapy. These findings support the ongoing development and refinement of machine learning models tailored to assist in communication for those affected by language impairments. The broader potential to create synthetic data for various clinical populations opens avenues for improving their communicative independence and quality of life.

- [1] M. L. Berthier, "Poststroke aphasia: epidemiology, pathophysiology and treatment," *Drugs & Aging*, vol. 22, no. 2, pp. 163–182, 2005.
- [2] K. Hilari, S. Northcott, P. Roy, J. Marshall, R. D. Wiggins, J. Chataway, and D. Ames, "Psychological distress after stroke and aphasia: the first six months," *Clinical Rehabilitation*, vol. 24, no. 2, pp. 181–190, 2010.
- [3] N. Azevedo, G. Le Dorze, G. Jarema, C. Alary Gauvreau, T. Ogourtsova, S. Beaulieu, C. Beaujard, M. Yvon, and E. Kehayia, "Understanding the experience of users of communication aids and applications through focus group discussions with people with aphasia and family members," *Frontiers in Communication*, vol. 8, 2023. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fcomm.2023.1219331>
- [4] P. Cuperus, "Aphasia therapy software: research, development, and implementation," Ph.D. dissertation, University of Groningen, [Groningen], 2023.
- [5] L. Nickels, "Therapy for naming disorders: revisiting, revising, and reviewing," *Aphasiology*, vol. 16, no. 10-11, pp. 935–979, 2002.
- [6] V. de Aguiar, R. Bastiaanse, and G. Miceli, "Improving production of treated and untreated verbs in aphasia: A meta-analysis," *Frontiers in Human Neuroscience*, vol. 10, p. 468, 2016.
- [7] R. Koul, M. Corwin, and S. Hayes, "Production of graphic symbol sentences by individuals with aphasia: Efficacy of a computer-based augmentative and alternative communication intervention," *Brain and Language*, vol. 92, no. 1, pp. 58–77, 2005.
- [8] R. Koul, M. Corwin, R. Nigam, and S. Oetzel, "Training individuals with chronic severe Broca's aphasia to produce sentences using graphic symbols: implications for AAC intervention," *Journal of Assistive Technologies*, vol. 2, no. 1, pp. 23–34, 03 2008. [Online]. Available: <https://doi.org/10.1108/17549450200800004>
- [9] A. Mooney, S. Bedrick, G. Noethe, S. Spaulding, and M. Fried-Oken, "Mobile technology to support lexical retrieval during activity retell in primary progressive aphasia," *Aphasiology*, vol. 32, no. 6, pp. 666–692, 2018.
- [10] N. Alam, S. Munjal, N. K. Panda, R. Kumar, and S. Gupta, "Efficacy of Jellow app as an adjunct to stimulation therapy in improvement in language and quality of life in patients with chronic Broca's Aphasia," *Disability and Rehabilitation: Assistive Technology*, vol. 18, pp. 596–602, 2023.
- [11] T. Chavers, C. Cheng, and R. Koul, "AAC interventions in persons with aphasia," in *Augmentative and Alternative Communication: Challenges and Solutions*, B. T. Ogletree, Ed. San Diego, CA: Plural Publishing, Incorporated, 2021, p. 141.
- [12] C. K. Thompson, L. P. Shapiro, L. Li, and L. Schendel, "Analysis of verbs and verb-argument structure: A method for quantification of aphasic language production," *Clinical aphasiology*, vol. 23, pp. 121–140, 1995.
- [13] R. Bastiaanse and C. K. Thompson, *Perspectives on agrammatism*. Psychology Press, 2012. [Online]. Available: <https://doi.org/10.4324/9780203120378>
- [14] M. Perez, D. Le, A. Romana, E. Jones, K. Licata, and E. M. Provost, "Seq2seq for automatic paraphasia detection in aphasic speech," 2023.
- [15] P. De Clercq, C. Puffay, J. Kries, H. Van Hamme, M. Vandermosten, T. Francart, and J. Vanthornhout, "Detecting post-stroke aphasia via brain responses to speech in a deep learning framework," in *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2024, pp. 1–5.
- [16] L. P. Shapiro, E. Zurif, and J. Grimshaw, "Sentence processing and the mental representation of verbs," *Cognition*, vol. 27, no. 3, pp. 219–246, 1987.
- [17] L. P. Shapiro, B. Gordon, N. Hack, and J. Killackey, "Verb-argument structure processing in Broca's and Wernicke's aphasia," *Brain and Language*, vol. 45, no. 3, pp. 423–447, 1993.
- [18] H. Goodglass and E. Kaplan, *The Assessment of Aphasia and Related Disorders*, 2nd ed. Lea & Febiger Philadelphia, 1983.
- [19] Y. Grodzinsky, "The syntactic characterization of agrammatism," *Cognition*, vol. 16, no. 2, pp. 99–120, 1984.
- [20] E. M. Saffran, R. S. Berndt, and M. F. Schwartz, "The quantitative analysis of agrammatic production: Procedure and data," *Brain and Language*, vol. 37, no. 3, pp. 440–479, 1989. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0093934X89900308>
- [21] H. Goodglass, E. Kaplan, and S. Weintraub, *BDAE: The Boston diagnostic aphasia examination*. Lippincott Williams & Wilkins Philadelphia, PA, 2001.
- [22] H. Kolk, "How language adapts to the brain," *The syntax of non-sententials: Multi-disciplinary perspectives*. John Benjamins Publishing Company, Amsterdam, the Netherlands, 2006.
- [23] M. B. Ruiters, H. H. Kolk, T. C. Rietveld, and I. Feddema, "Combining possibly reciprocally dependent linguistic parameters in the quantitative assessment of aphasic speakers' grammatical output," *Aphasiology*, vol. 27, no. 3, pp. 293–308, 2013.
- [24] C. K. Thompson, K. J. Ballard, M. E. Tait, S. Weintraub, and M. Mesulam, "Patterns of language decline in non-fluent primary progressive aphasia," *Aphasiology*, vol. 11, no. 4-5, pp. 297–321, 1997.
- [25] C. Themistocleous, B. Ficek, K. Webster, D.-B. den Ouden, A. E. Hillis, and K. Tsapkini, "Automatic subtyping of individuals with primary progressive aphasia," *Journal of Alzheimer's Disease*, vol. 79, no. 3, pp. 1185–1194, 2021.
- [26] K. C. Fraser, N. Linz, H. Lindsay, and A. König, "The importance of sharing patient-generated clinical speech and language data," in *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, K. Niederhoffer, K. Hollingshead, P. Resnik, R. Resnik, and K. Loveys, Eds. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 55–61. [Online]. Available: <https://aclanthology.org/W19-3007>
- [27] A. Adikari, N. Hernandez, D. Alahakoon, M. L. Rose, and J. E. Pierce, "From concept to practice: a scoping review of the application of AI to aphasia diagnosis and management," *Disability and Rehabilitation*, vol. 46, no. 7, pp. 1288–1297, 2024.
- [28] A. Privitera, S. Ng, A.-H. Kong, and B. Weekes, "AI and aphasia in the digital age: A critical review," *Brain Sciences*, vol. 14, no. 4, p. 383, 2024.
- [29] N. Azevedo, E. Kehayia, G. Jarema, G. L. Dorze, C. Beaujard, and M. Yvon, "How artificial intelligence (AI) is used in aphasia rehabilitation: A scoping review," *Aphasiology*, vol. 38, no. 2, pp. 305–336, 2024.
- [30] R. Misra, S. S. Mishra, and T. K. Gandhi, "Assistive completion of agrammatic aphasic sentences: Amalgamation of NLP and neurolinguistics-based synthetic dataset," in *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2023, pp. 1–4.
- [31] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of Machine Learning Research*, vol. 21, no. 140, pp. 1–67, 2020. [Online]. Available: <http://jmlr.org/papers/v21/20-074.html>
- [32] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, Y. Li, X. Wang, M. Dehghani, S. Brahma, A. Webson, S. S. Gu, Z. Dai, M. Suzgun, X. Chen, A. Chowdhery, A. Castro-Ros, M. Pellat, K. Robinson, D. Valter, S. Narang, G. Mishra, A. Yu, V. Zhao, Y. Huang, A. Dai, H. Yu, S. Petrov, E. H. Chi, J. Dean, J. Devlin, A. Roberts, D. Zhou, Q. V. Le, and J. Wei, "Scaling instruction-finetuned language models," *Journal of Machine Learning Research*, vol. 25, no. 70, pp. 1–53, 2024. [Online]. Available: <http://jmlr.org/papers/v25/23-0870.html>
- [33] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd, "spaCy:

- Industrial-strength Natural Language Processing in Python,” 2020, [using model: en\_core\_web\_sm].
- [34] K. A. Tetzloff, R. L. Utianski, J. R. Duffy, H. M. Clark, E. A. Strand, K. A. Josephs, and J. L. Whitwell, “Quantitative analysis of agrammatism in agrammatic primary progressive aphasia and dominant apraxia of speech,” *Journal of Speech, Language, and Hearing Research*, vol. 61, no. 9, pp. 2337–2346, 2018. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/30098169/>
- [35] M. Straka, J. Hajič, and J. Straková, “UDPipe: Trainable pipeline for processing CoNLL-U files performing tokenization, morphological analysis, POS tagging and parsing,” in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, N. Calzolari, K. Choukri, T. Declerck, S. Goggi, M. Grobelnik, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, and S. Piperidis, Eds. Portorož, Slovenia: European Language Resources Association (ELRA), May 2016, pp. 4290–4297. [Online]. Available: <https://aclanthology.org/L16-1680>
- [36] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [37] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “RoBERTa: A robustly optimized BERT pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [38] R. Levy, “Expectation-based syntactic comprehension,” *Cognition*, vol. 106, no. 3, pp. 1126–1177, 2008.
- [39] A. Goodkind and K. Bicknell, “Predictive power of word surprisal for reading times is a linear function of language model quality,” in *Proceedings of the 8th workshop on cognitive modeling and computational linguistics (CMCL 2018)*, 2018, pp. 10–18.
- [40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fb0d53c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fb0d53c1c4a845aa-Paper.pdf)
- [41] M. Popović, “chrF: character n-gram F-score for automatic MT evaluation,” in *Proceedings of the Tenth Workshop on Statistical Machine Translation*, O. Bojar, R. Chatterjee, C. Federmann, B. Haddow, C. Hokamp, M. Huck, V. Logacheva, and P. Pecina, Eds. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 392–395. [Online]. Available: <https://aclanthology.org/W15-3049>
- [42] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, “Transformers: State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 38–45. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>
- [43] S. Mangrulkar, S. Gugger, L. Debut, Y. Belkada, S. Paul, and B. Bossan, “Peft: State-of-the-art parameter-efficient fine-tuning methods,” <https://github.com/huggingface/peft>, 2022.
- [44] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, “Qlora: Efficient finetuning of quantized llms,” in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 10088–10115. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/1feb8787143631bdc0f2beaa62a049b-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/1feb8787143631bdc0f2beaa62a049b-Paper-Conference.pdf)
- [45] C.-Y. Lin, “ROUGE: A package for automatic evaluation of summaries,” in *Text Summarization Branches Out*. Barcelona, Spain: Association for Computational Linguistics, Jul. 2004, pp. 74–81. [Online]. Available: <https://aclanthology.org/W04-1013/>
- [46] J. Ni, G. Hernandez Abrego, N. Constant, J. Ma, K. Hall, D. Cer, and Y. Yang, “Sentence-T5: Scalable sentence encoders from pre-trained text-to-text models,” in *Findings of the Association for Computational Linguistics: ACL 2022*, S. Muresan, P. Nakov, and A. Villavicencio, Eds. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 1864–1874. [Online]. Available: <https://aclanthology.org/2022.findings-acl.146/>
- [47] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [48] B. MacWhinney, D. Fromm, M. Forbes, and A. Holland, “AphasiaBank: Methods for studying discourse,” *Aphasiology*, vol. 25, no. 11, pp. 1286–1307, 2011.
- [49] J. W. Du Bois, W. L. Chafe, C. Meyer, S. A. Thompson, R. Englebretson, and N. Martey, “Santa barbara corpus of spoken american english, parts 1-4,” *Philadelphia: Linguistic Data Consortium*, 2000-2005. [Online]. Available: <https://www.linguistics.ucsb.edu/research/santa-barbara-corpus>
- [50] A. Hosseini, S. Reddy, D. Bahdanau, R. D. Hjelm, A. Sordani, and A. Courville, “Understanding by understanding not: Modeling negation in language models,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, and Y. Zhou, Eds. Online: Association for Computational Linguistics, Jun. 2021, pp. 1301–1312. [Online]. Available: <https://aclanthology.org/2021.naacl-main.102/>
- [51] M. Rezaei and E. Blanco, “Paraphrasing in affirmative terms improves negation understanding,” in *Proceedings of the Association for Computational Linguistics (Volume 2: Short Papers)*, L.-W. Ku, A. Martins, and V. Srikumar, Eds. Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 602–615. [Online]. Available: <https://aclanthology.org/2024.acl-short.55/>
- [52] O. Melamud and C. Shivade, “Towards automatic generation of shareable synthetic clinical notes using neural language models,” in *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, A. Rumshisky, K. Roberts, S. Bethard, and T. Naumann, Eds. Minneapolis, Minnesota, USA: Association for Computational Linguistics, Jun. 2019, pp. 35–45. [Online]. Available: <https://aclanthology.org/W19-1905>
- [53] E. Khiu, H. Toossi, J. Liu, J. Li, D. Anugraha, J. Flores, L. Roman, A. S. Doğruöz, and E.-S. Lee, “Predicting machine translation performance on low-resource languages: The role of domain similarity,” in *Findings of the Association for Computational Linguistics: EACL 2024*, Y. Graham and M. Purver, Eds. St. Julian’s, Malta: Association for Computational Linguistics, Mar. 2024, pp. 1474–1486. [Online]. Available: <https://aclanthology.org/2024.findings-eacl.100>
- [54] T. Sakakini, H. Gong, J. Y. Lee, R. Schloss, J. Xiong, and S. Bhat, “Equipping educational applications with domain knowledge,” in *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, H. Yannakoudakis, E. Kochmar, C. Leacock, N. Madnani, I. Pilán, and T. Zesch, Eds. Florence, Italy: Association for Computational Linguistics, Aug. 2019, pp. 472–477. [Online]. Available: <https://aclanthology.org/W19-4448>
- [55] M. Mosbach, T. Pimentel, S. Ravfogel, D. Klakow, and Y. Elazar, “Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation,” in *Findings of the Association for Computational Linguistics: ACL 2023*, A. Rogers, J. Boyd-Graber, and N. Okazaki, Eds. Toronto, Canada: Association for Computational Linguistics, Jul. 2023, pp. 12 284–12 314. [Online]. Available: <https://aclanthology.org/2023.findings-acl.779/>
- [56] P. van Kooten, “Custom python library to unfold contractions.” 2022. [Online]. Available: <https://github.com/kootenpv/contractions>

## VI. SUPPLEMENTARY MATERIALS

### A. BLEU exclusion

Since the sentence completion task is inherently different from machine translation, we do not use the conventional BLEU [47] metric. BLEU often results in scores of 0.00 for reasonable completions because it relies on exact n-gram matches, which can be too strict for this task. Even slight variations in word choice or order can lead to a BLEU score of 0.00 if the n-grams do not align perfectly. Table VII shows a few examples of reasonable completions assigned a BLEU score of 0.00 along with the assigned ChrF and RougeL scores, demonstrating the appropriateness of ChrF and RougeL over BLEU for this task. ChrF considers character-level n-grams, providing a more nuanced evaluation by capturing partial matches and minor variations. RougeL finds the longest common subsequence between the generated and reference text, allowing us to estimate how well the generated sentence preserves word order and grammatical coherence. Thus, offering a better assessment of the quality of the generated completions. Moreover, Table VIII shows the descriptive statistics of the BLEU metric calculated on all completions of the synthetic test set by FLAN-T5-XL fine-tuned using the prefix, further indicating the unsuitability of the BLEU metric for the sentence completion task.

TABLE VII  
EXAMPLES OF REASONABLE COMPLETIONS WITH BLEU SCORES OF 0.00 AND THEIR CORRESPONDING CHRf AND ROUGEL SCORES. ALL EXAMPLES ARE TAKEN FROM FLAN-T5-XL FINE-TUNED USING THE PREFIX.

Syn. utterance	Target completion	Gen. completion	BLEU score	ChrF score	RougeL score
And run it computer.	And run it through the computer.	And run it on the computer.	0.00	63.75	0.83
Janine uses Charlie times.	Janine uses Charlie all the time.	Janine uses Charlie several times.	0.00	65.45	0.55
Without the risk of killed.	Without running the risk of getting killed.	Without the risk of being killed.	0.00	56.87	0.77
Gradually bring it.	Just gradually bring it around.	Gradually bring it up.	0.00	55.65	0.44
And then they.	And then they break.	And then they sat down.	0.00	58.39	0.67

TABLE VIII  
DESCRIPTIVE STATISTICS OF THE BLEU SCORE COMPUTED FOR FLAN-T5-XL ON THE SYNTHETIC TEST SET USING THE PREFIX "COMPLETE THIS SENTENCE".

Statistic	Score
Mean	0.24
Standard deviation	0.30
Min	0.00
Median	0.00

### B. Fine-tuning setup details

During fine-tuning, we initialise the QLoRA configuration (see Tables below for quantisation parameters and hyperparameters), prepare the model for 4-bit fine-tuning and add the LoRA adaptor using the initialised configuration.

TABLE IX  
QUANTISATION PARAMETERS IN THE BITSANDBYTES CONFIGURATION.

BitsAndBytes configuration parameters
load_in_4bit=True
bnb_4bit_use_double_quant=True
bnb_4bit_quant_type="nf4"
bnb_4bit_compute_dtype=torch.bfloat16

TABLE X  
HYPERPARAMETERS IN THE LORA CONFIGURATION.

LoRA configuration parameters
r=8
lora_alpha=32
lora_dropout=0.05
inference_mode=False
bias="none"
task_type=TaskType.SEQ_2_SEQ_LM

### C. Training object

We use the Seq2SeqTrainingArguments and Seq2SeqTrainer classes from the HuggingFace transformers library [42] to define the training arguments and trainer object respectively.

TABLE XI  
THE TRAINING ARGUMENTS (SEQ2SEQTRAININGARGUMENTS) USED TO FINE-TUNE THE MODELS.

Hyperparameter	Value for T5-base	Value for all other models
output_dir	checkpoints_path	checkpoints_path
learning_rate	1e-4	1e-4
per_device_train_batch_size	batch_size	batch_size
per_device_eval_batch_size	batch_size	batch_size
num_train_epochs	5	1
weight_decay	0.01	0.01
predict_with_generate	True	True
evaluation_strategy	“epoch”	“epoch”
save_strategy	“epoch”	“epoch”
load_best_model_at_end	True	True
fp16	True	True
gradient_accumulation_steps	2	2
logging_dir	checkpoints_path + “/logs”	checkpoints_path + “/logs”

TABLE XII  
THE TRAINER OBJECT DEFINEMENT (SEQ2SEQTRAINER) USED TO FINE-TUNE THE SENTENCE COMPLETION MODELS.

Hyperparameter	Value
model	model
args	training_args
train_dataset	tokenized_train_dataset
eval_dataset	tokenized_valid_dataset
tokenizer	tokenizer
data_collator	data_collator
compute_metrics	compute_metrics (function)
callbacks	[EarlyStoppingCallback(early_stopping_patience=2)]

D. Data pre-processing

TABLE XIII  
EXAMPLES OF UNPROCESSED AND PRE-PROCESSED DATA OBTAINED FROM APHASIABANK AND SBSCSAE IN CHAT FORMAT.

Corpus	Unprocessed text	Pre-processed
AphasiaBank (Aphasic)	it's alright . [+ exc] 360360_360780	It is alright.
AphasiaBank (Aphasic)	&=imit:running and goes +/. 453922_454482	And goes.
AphasiaBank (Aphasic)	+”” &+b I wanna come . 727610_728460	I want to come.
AphasiaBank (Aphasic)	that he [/] he said . 804035_805805	That he said.
SBSCSAE	(.) So you don't need to go (.) borrow equipment from anybody 0.9210 to +/.	So you do not need to go borrow equipment from anybody to.
SBSCSAE	&=in &=lengthened (.) and (.) &=tsk (.) I don't know . 76950_79570	And i do not know.
SBSCSAE	you know 231970_232310 and you'd put the +/. 232310_232940	You know and you would put the.
SBSCSAE	&=in you'd have +/. 232940_233690	You would have.

The pre-processing steps include replacing multiple consecutive spaces by a single space, unfolding contractions using the contractions library [56] and the removal of the following features:

- Instances of trailing off pauses such as “(..)”, “(...)” and “+...”;
- Filler pauses including “&-um”, “&-uh”, “&-mm” and “&-like”;
- Actions described within the text, indicated by sequences such as “&=imit:runnin”;
- All Unicode errors present in the text;
- Sequences beginning with two special characters followed by text;
- Text within square brackets “[ ]” and angle brackets “ $\langle \rangle$ ”;
- Strings starting with “&+” or “U” followed by text as well as those beginning with characters such as “\*”, “?” and “!”;
- Arbitrary length strings of Xs;
- Strings starting with special characters followed by one or more numbers;
- Any remaining special characters such as “‡”, “+/.”, “@o” and “=@”;
- Instances of stuttering and repetitive words.

The effect of these pre-processing steps can be observed in Table XIII, which presents examples of the data before and after pre-processing. This cleaning process ensures that the resulting dataset is as free from noise and inconsistencies as possible.

## E. Synthetic example completions

TABLE XIV  
ADDITIONAL EXAMPLES OF GENERATED COMPLETIONS ON THE SYNTHETIC TEST SET.

Model	Synthetic aphasic utterance	Generated completion
T5-base-cts	It those i went out and coffees.	It was all those i went out and got coffee.
T5-base-cts	They been it a while.	They have been doing it for a while.
T5-base-cts	Person meaning who.	Person meaning who.
T5-base-cts	Make it like we are it.	Make it look like we are going to do it.
T5-base-cts	They so fast they shape containers.	They are so fast they can shape containers.
T5-base-cts	And i called talk him.	And i called to talk to him.
T5-base-cts	Well top guys lip.	Well top guys lip.
T5-base-cts	Oh i it they.	Oh i thought it was they.
T5-base-cts	I see venice high school.	I see venice high school.
T5-base-cts	It great in rooms we.	It was great in the room we were in.
FLAN-T5-XL-cts	And they control.	And they are in control.
FLAN-T5-XL-cts	I a lot	I have a lot of work to do.
FLAN-T5-XL-cts	I a rake but i it.	I have a rake but i don't use it.
FLAN-T5-XL-cts	There an orders to how i it.	There is an order to how i want to do it.
FLAN-T5-XL-cts	I worked my money.	I worked hard for my money.
FLAN-T5-XL-cts	I he going down.	I think he is going down.
FLAN-T5-XL-cts	This is you.	This is for you.
FLAN-T5-XL-cts	You are a aheads me son.	You are going to be ahead of me son.
FLAN-T5-XL-cts	And he good at making out budget.	And he is good at making out a budget.
FLAN-T5-XL-cts	Oh he is be you know closet everything.	Oh he is going to be you know in the closet with everything.

## F. Authentic example completions

TABLE XV  
ADDITIONAL EXAMPLES OF GENERATED COMPLETIONS FOR ALL THE AUTHENTIC DATA. THE PREFIX IS USED IN ALL CASES.

Authentic aphasic utterance	Generated by T5-base	Generated by FLAN-T5-XL
A dog chase a up to a tree.	A dog chases a dog up to a tree.	A dog chases a ball up to a tree.
Leg shirt leg pants leg.	Leg shirt leg pants leg.	Leg shirt leg pants leg.
I was think of going yeah.	I was thinking of going yeah.	I was thinking of going to the store yeah.
High that was a the firetruck.	High that was a firetruck.	High that was a firetruck.
Alright fine.	Alright fine.	Alright fine.
The prince the.	The prince of the.	The prince is the.
And get some the bread.	And get some of the bread.	And get some of the bread.
I did not know words.	I did not know how to say.	I did not know what the words were.
And i did not even what i was wrong.	And i did not even know what i was doing wrong.	And i did not even know what i was doing was wrong.
Well i do not have to do that.	Well i do not have to do that.	Well i do not have to do that.
My firstnamer my daughter.	My firstname is my daughter.	My firstname is my daughter.
But i can listen to the books.	But i can listen to the books.	But i can listen to the books.
Okay they will tell me restaurants.	Okay they will tell me about the restaurant.	Okay they will tell me about the restaurants.
She tells me lotsof words.	She tells me lots of words.	She tells me lotsof words.
And then about six sixth years i called her.	And then about six sixth years later i called her.	And then about six or six and a half years later i called her.

TABLE XVI  
ADDITIONAL EXAMPLES OF GENERATED COMPLETIONS FOR THE AUTHENTIC DATA WITH ONE-WORD AND TWO-WORD UTTERANCES REMOVED.

Authentic aphasic utterance	T5-base-cts completion	FLAN-T5-XL-cts completion
The this one ?.	Then the next one?	Then this is the one?.
Oh it is ?.	Oh it is ok?	Oh it is not?.
And they went kiss and go.	And they went to kiss and go.	And they went to kiss and go.
And there's one two.	And there's one of them.	And there's one or two.
Then the this andjelly over that.	Then the next time the andjelly over that.	Then the jar is put on this andjelly over that.
You in er in.	You are in er in.	You are in er in.
And that is that is the worst part.	And that is that is the worst part.	And that is that is the worst part.
It is one major youknow wne.	It is one major youknow wn.	It is one major youknow wn.
Imean that is that is just.	Imean that is that is just.	I mean that is just a little bit.
Stroke numb youknow i burst.	Stroke is numb you know i burst.	Stroke numb youknow i burst.
So mary'shospital surgery.	So mary's hospital surgery.	So mary'shospital had surgery.
Friend support cousin family.	Friend support cousin family.	Friend support cousin family.
Say beautiful so manage.	Say beautiful so manageable.	Say beautiful so you can manage.
So finally police.	So finally the police.	So finally the police.
Well it two am stroke.	Well it was two am stroke.	Well it is two am having a stroke.