

# Tokenisation is NP-Complete

Philip Whittington, Gregor Bachmann, Tiago Pimentel

Department of Computer Science, ETH Zürich

{philip.whittington, gregor.bachmann, tiago.pimentel}@inf.ethz.ch

## Abstract

In this work, we prove the NP-completeness of two variants of tokenisation, defined here as the problem of compressing a dataset to at most  $\delta$  symbols by either finding a vocabulary directly (*direct* tokenisation), or selecting a sequence of merge operations (*bottom-up* tokenisation).

## 1 Introduction

Tokenisation is at the heart of natural language processing (NLP) being the first step required to use a language model. Given a string of **characters**  $c$ , a tokeniser converts it into a string of **subwords**  $s$ . Language models are then trained to estimate distributions over subword-strings—never seeing the original character-strings. Despite its prominent role, however, much remains unknown about tokenisation. We still do not know, for instance, what makes a good tokeniser (Gowda and May, 2020; Cognetta et al., 2024): which characteristics should its produced subwords  $s$  have to be a good starting point for language modelling? If we knew this, then we could define an **objective function** with which we could evaluate tokenisers.

Another open question is how to—given such an objective function—efficiently find a tokeniser which maximises it. Byte pair encoding (BPE; Gage, 1994; Sennrich et al., 2016), for instance, is a greedy solution to find a tokeniser which maximises a text’s compression. UnigramLM (Kudo, 2018) is a heuristic method to find a tokeniser that maximises its tokenised text’s unigram log-probability. Both these methods, however, are approximate: they do not necessarily find an optimal tokeniser according to their objective function. This raises the question of whether finding such optimal tokenisers efficiently is even possible.

In this paper, we answer this question (at least partially), proving the NP-completeness of several variants of this tokenisation problem. Specifically, we focus on finding tokenisers that maximise the

**compression** of a text.<sup>1</sup> Given this objective, we then define the **tokenisation problem** as the task of finding a tokeniser which compresses a dataset to at most  $\delta$  symbols. Notably, prior work imposes different constraints on how tokenisers are defined; here we consider two variants. In **direct tokenisation**, the desired compression must be reached by choosing a vocabulary (i.e., a set of subwords) which is directly used to represent the text. In **bottom-up tokenisation**, the desired compression must be reached by finding a sequence of merge operations instead, which we apply to the input text.

We prove the NP-hardness of both of these tokenisation problems (as well as of some variants thereof) by reducing from the **max 2-satisfiability** problem.<sup>2</sup> Practically speaking, our results imply that we are unlikely to discover an efficient algorithm for the problem of finding optimal tokenisers, and that we should focus on approximate algorithms (such as BPE or UnigramLM) instead.

## 2 How to Choose a Tokeniser?

Given a tokeniser, any character-level distribution has an equivalent subword-level distribution (Pimentel and Meister, 2024; Phan et al., 2024; Giulianelli et al., 2024). Ergo, despite the distribution we may wish to language model, a sufficiently expressive model should be able to represent it exactly; this is true regardless which tokeniser is used. In theory, thus, a researcher’s choice of tokeniser should not influence their language model’s quality.

In practice, however, a bad choice of tokeniser can have undesirable effects on downstream appli-

<sup>1</sup>The compression achieved by a tokeniser correlates with downstream language modelling performance (Gallé, 2019; Zouhar et al., 2023a) and computational efficiency.

<sup>2</sup>We note two related concurrent works. Kozma and Voderholzer (2024) also prove the NP-completeness of bottom-up tokenisation; in fact, they prove something stronger: its APX-hardness. Lim et al. (2025) prove the NP-completeness of a restricted variant of direct tokenisation, in which a set of candidate tokens is previously specified.

cations. For instance, performing standard arithmetic tasks (e.g.,  $317 + 421$ ) can be difficult even for large models (Nogueira et al., 2021; Muffo et al., 2022) due to the arbitrary splitting of numbers into subwords. Indeed, simple changes in how numbers are tokenised can improve performance in such tasks (Singh and Strouse, 2024). Similar issues arise when prompting language models to count letters, where even advanced models such as GPT-4 infamously cannot correctly count the number of occurrences of the letter `r` in the word `strawberry`.<sup>3</sup>

This raises the question of how to select a good tokeniser. Ideally, we would choose the tokeniser which maximises downstream language modelling performance. Unfortunately, we do not know how to measure such performance without fully training a model, making its direct maximisation computationally infeasible. Rather, we thus optimise proxy objectives—assumed to correlate with downstream performance. Among these are unigram log-probability (Kudo, 2018), Rényi efficiency (Zouhar et al., 2023a), and compression (Gallé, 2019).

We focus on compression in this paper. Denoting our tokenisation’s **objective function** as  $\mathfrak{G}$ , we write this objective as:  $\mathfrak{G}(\mathbf{s}) = -|\mathbf{s}|$ . Improved compression leads to: (i) more efficient training and inference, due to shortened inputs;<sup>4</sup> (ii) improved downstream performance, at least to a certain extent (Gallé, 2019; Rust et al., 2021; Zouhar et al., 2023a; Goldman et al., 2024);<sup>5</sup> and (iii) fairer multilingual treatment—assuming similar compression across languages—given models’ limited context lengths and the per-token costs to use proprietary models (Petrov et al., 2023; Ahia et al., 2023).

#### Our Notation’s Colour-coding

- **Blue** for raw data (i.e., characters  $\mathbf{c} \in \Sigma^*$ );
- **Magenta** for tokeniser-specific data (i.e., subwords  $\mathbf{s} \in \mathcal{S}^*$  and mergers  $\mathbf{m} \in \mathcal{M}^*$ );
- **Orange** for functions (e.g., `tok`).

<sup>3</sup>In a recent paper, Lesci et al. (2025) show that tokenisation strongly influences language models’ outputs: a character-string  $\mathbf{c}$  is assigned up to 17 times less probability by language models if tokenised into two subwords than if represented as a single subword; a difference of 2.88 nats.

<sup>4</sup>Recent work tries to improve the computational efficiency of byte-level models (Yu et al., 2023; Pagnoni et al., 2024).

<sup>5</sup>Although, see Ali et al. (2024), who argue that compression might be a necessary but not sufficient condition for good downstream performance, and Schmidt et al. (2024), who argue that compression and downstream performance have a more complex relationship than prior work suggests.

### 3 Defining a Tokeniser

A tokeniser can be defined as a 3-tuple  $\langle \mathcal{S}, \text{tok}, \text{detok} \rangle$ , composed of a vocabulary, a tokenisation and a detokenisation function. Before defining these terms, however, we require some notation. Let  $\mathbf{c} \in \Sigma^*$  be a **character-string**,<sup>6</sup> i.e., a sequence of characters  $\mathbf{c}$  from alphabet  $\Sigma$ , which we write as  $\mathbf{c} = c_1 c_2 \dots c_{|\mathbf{c}|}$ . Further, let  $\mathcal{D} = \{\mathbf{c}_n\}_{n=1}^N$  be a dataset of character-strings.<sup>7</sup> A subword  $\mathbf{s} \in \mathcal{S}$  represents a non-empty character-string  $\mathbf{c}$  (where sequence  $\mathbf{c}$  can have length one). Finally, let  $\mathbf{s} \in \mathcal{S}^*$  be a **subword-string**. Just like a single subword, a subword-string  $\mathbf{s} = \langle s_1, s_2, \dots, s_{|\mathbf{s}|} \rangle$  represents a character-string via the concatenation of its subwords’ characters:

$$\text{concat}(\mathbf{s}) = s_1 \circ s_2 \circ \dots \circ s_{|\mathbf{s}|} \quad (1)$$

and we say that a pair of character and subword strings are equivalent if:

$$\mathbf{c} \stackrel{\circ}{=} \mathbf{s} \iff \mathbf{c} = \text{concat}(\mathbf{s}) \quad (2)$$

Given the notation above, we can now define the items in tuple  $\langle \mathcal{S}, \text{tok}, \text{detok} \rangle$ . A tokeniser’s **vocabulary** is a set of subwords  $\mathcal{S} \subset \Sigma^+$  such that  $\Sigma \subseteq \mathcal{S}$ ;<sup>8</sup> we say its size is  $|\mathcal{S}| = |\Sigma| + K$ . Further, a **detokenisation function** is defined as  $\text{detok} : \mathcal{S}^* \rightarrow \Sigma^*$  and given a subword-string it outputs the character-string it represents. This function thus is simply defined as  $\text{detok}(\mathbf{s}) \stackrel{\text{def}}{=} \text{concat}(\mathbf{s})$ .

Finally, we are left with defining a **tokenisation function**  $\text{tok} : \Sigma^* \rightarrow \mathcal{S}^*$ , which maps from character- to subword-strings. Notably, these functions always ensure the equivalence  $\mathbf{c} \stackrel{\circ}{=} \mathbf{s}$  for  $\mathbf{s} = \text{tok}(\mathbf{c})$ . Several tokenisation functions, however, are compatible with this constraint, as given a vocabulary, many subword-strings may be equivalent to the same character-string. For instance, given  $\mathcal{S} = \{a, c, t, at\}$ , the string  $\mathbf{c} = c, a, t$  could be tokenised as  $\mathbf{s} = \langle c, a, t \rangle$  or as  $\mathbf{s} = \langle c, at \rangle$ . Most researchers define tokenisation functions in one of two ways, which we term direct and bottom-up tokenisation functions; we define these next.

<sup>6</sup>We note that  $\Sigma^*$  denotes the Kleene star of  $\Sigma$  (i.e.,  $\cup_{i=0}^{\infty} \Sigma^i$ ), and  $\Sigma^+$  denotes its Kleene plus (i.e.,  $\cup_{i=1}^{\infty} \Sigma^i$ ).

<sup>7</sup>We note that we use set notation here, but our datasets are actually multisets—datasets can include the same string  $\mathbf{c}$  multiple times. We show that tokenisation is still NP-complete for datasets with no repetitions in §6.3. Further, we impose no constraint on the kind of string present in these datasets: each  $\mathbf{c}_n$  can be either a raw or pre-tokenised character-string (i.e., either a full document or a whitespace-separated word).

<sup>8</sup> $\Sigma \subseteq \mathcal{S}$  is typically enforced to guarantee that every  $\mathbf{c} \in \Sigma^*$  can be represented by at least one subword-string  $\mathbf{s} \in \mathcal{S}^*$ .

### 3.1 Direct Tokenisation Functions

In direct tokenisation, a character-string is directly replaced by an optimal subword-string. To implement this, one must thus first define what *optimal* means; this is done through an objective function  $\mathfrak{G}$  which, given a subword-string, returns a score. Given a previously chosen vocabulary  $\mathcal{S}$  (we discuss how to find  $\mathcal{S}$  in §5), a direct tokenisation function then encodes string  $\mathbf{c}$  as:

$$\text{tok}_{\diamond}[\mathcal{S}](\mathbf{c}) = \arg \max_{\mathbf{s} \in \mathcal{S}^*} \mathfrak{G}(\mathbf{s}) \quad (3)$$

s.t.  $\mathbf{s} \stackrel{\circ}{=} \mathbf{c}$

In words, given a vocabulary  $\mathcal{S}$ , function  $\text{tok}_{\diamond}$  returns the optimal subword-string  $\mathbf{s} \in \mathcal{S}^*$  which is equivalent to the input character-string  $\mathbf{c}$ . We then set  $\text{tok}(\mathbf{c}) \stackrel{\text{def}}{=} \text{tok}_{\diamond}[\mathcal{S}](\mathbf{c})$ . Different choices of  $\mathfrak{G}$  recover methods such as UnigramLM’s tokenisation function (Kudo, 2018) or PathPiece (Schmidt et al., 2024). Notably, in general, this function is not efficiently computable.<sup>9</sup>

In this paper, we are concerned with tokenisers that use compression as their objective: that is, for which  $\mathfrak{G}(\mathbf{s}) = -|\mathbf{s}|$ . In this case, we can rewrite the direct tokenisation function as:

$$\text{tok}_{\diamond}[\mathcal{S}](\mathbf{c}) = \arg \min_{\mathbf{s} \in \mathcal{S}^*} |\mathbf{s}| \quad (4)$$

s.t.  $\mathbf{s} \stackrel{\circ}{=} \mathbf{c}$

Importantly, in the case of compression, this equation can be computed efficiently (as shown in §5.1).

### 3.2 Bottom-up Tokenisation Functions

In bottom-up tokenisation, one starts with a set of character-strings, and merges their symbols bottom-up, one pair at a time.<sup>10</sup> Formally, let  $m \in \mathcal{M}$  be a **merge**, defined as a pair of subwords:  $m = \langle s_1, s_2 \rangle$ . Further, let  $\mathcal{M} \stackrel{\text{def}}{=} \Sigma^+ \times \Sigma^+$ . Now, let **merge** be a functional; given merge  $m = \langle s_1, s_2 \rangle$ , it returns a function  $\text{merge}[m] : \mathcal{S}^* \rightarrow (\mathcal{S} \cup \{s_1 \circ s_2\})^*$  which operates on string  $\mathbf{s}$  left-to-right, replacing every occurrence of  $s_1$  followed by  $s_2$  in it with subword  $s' = s_1 \circ s_2$ . E.g., given  $\mathbf{s} = \langle wo, r, ld \rangle$  and  $m = \langle wo, r \rangle$ , the output of  $\text{merge}[m](\mathbf{s})$  is  $\langle wor, ld \rangle$ .

<sup>9</sup>In fact, Geh et al. (2024) shows that it is NP-complete for  $\mathfrak{G}(\mathbf{s}) = \sum_{t=1}^{|\mathbf{s}|} \log p_{\theta}(s_t | \mathbf{s}_{<t})$ , where  $p_{\theta}$  is a language model.

<sup>10</sup>Currently, this is likely the most common tokenisation function, being used in popular tokenisers such as, e.g., GPT-4’s (OpenAI, 2023), LLaMA’s (Touvron et al., 2023a,b), and Pythia’s (Biderman et al., 2023).

Consider now  $\mathbf{m} \in \mathcal{M}^*$ , a sequence of merges. Given a character-string  $\mathbf{c} \in \Sigma^*$ , a bottom-up tokenisation function compresses it as:

$$\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c}) = \left( \bigodot_{z=1}^{|\mathbf{m}|} \text{merge}[m_z] \right) (\mathbf{c}) \quad (5)$$

where  $\odot$  represents function composition, e.g.,  $\bigodot_{z=1}^2 \text{merge}[m_z] = \text{merge}[m_2] \odot \text{merge}[m_1]$ . Bottom-up tokenisers then set  $\text{tok} \stackrel{\text{def}}{=} \text{tok}_{\uparrow}[\mathbf{m}]$ . Further, a merge sequence  $\mathbf{m}$  is also used to set a bottom-up tokeniser’s vocabulary as:

$$\mathcal{S} = \Sigma \cup \{s_1 \circ s_2 \mid \langle s_1, s_2 \rangle \in \mathbf{m}\} \quad (6)$$

where  $|\mathbf{m}| = K$  implies this vocabulary has size  $|\mathcal{S}| = |\Sigma| + K$ , as before.

## 4 Maximum 2-Satisfiability

Our paper’s goal is to prove the NP-completeness of tokenisation. To show this, we must reduce an NP-hard problem to tokenisation in polynomial time. We will rely on the **maximum 2-satisfiability** problem (max-2-SAT) for this, whose definition we provide here. The NP-hardness of max-2-SAT was proven by Garey et al. (1974).

**Definition 1.** Let  $\mathcal{X} = \{X_j\}_{j=1}^J$  be a set of variables; each of these variables are assigned values  $x_j \in \{\mathbb{F}, \mathbb{T}\}$ , and we write  $\chi = \{x_j\}_{j=1}^J \in \{\mathbb{F}, \mathbb{T}\}^J$ . Let  $\mathcal{C} = \{(L_i^1 \vee L_i^2)\}_{i=1}^I$  be a set of clauses,<sup>11</sup> where each literal  $L$  represents either a variable  $X_j$  or its negation  $\neg X_j$ . The **max-2-SAT decision problem** requires deciding whether there exists an assignment for which at least  $\gamma$  clauses are satisfied:

$$\gamma \leq \max_{\chi \in \{\mathbb{F}, \mathbb{T}\}^J} \sum_{i=1}^I \mathbb{1}_{\chi} \{L_i^1 \vee L_i^2\} \quad (7)$$

where  $\mathbb{1}_{\chi}$  is an indicator function which evaluates the clause and returns one if the clause is satisfied by  $\chi$  and zero otherwise.

For mathematical convenience, we will write  $\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma)$  for a function which returns  $\mathbb{T}$  if its input is satisfiable under a max-2-SAT decision problem, and  $\mathbb{F}$  otherwise. As a concrete example, consider the set of variables  $\mathcal{X} = \{X_1, X_2\}$  and the set of clauses  $\mathcal{C} = \{X_1 \vee X_2, \neg X_1 \vee X_2, X_1 \vee \neg X_2, \neg X_1 \vee \neg X_2\}$ .

<sup>11</sup>max-2-SAT also allows clauses to have a single literal  $L_i$ . In this case, we can always rewrite the clause as  $(L_i \vee L_i)$  with no change to the solution of this decision problem.

The assignment  $x_1 = \text{T}$ ,  $x_2 = \text{T}$  leads to 3 clauses being satisfied, which is the optimum. For this example, we thus have that  $\text{M2S}(\mathcal{X}, \mathcal{C}, 3) = \text{T}$ , but that  $\text{M2S}(\mathcal{X}, \mathcal{C}, 4) = \text{F}$ .

## 5 Finding Optimal Direct Tokenisers

We are now left with the task of finding an optimal tokeniser. We do this by selecting either: its vocabulary in direct tokenisation, since  $\text{tok} = \text{tok}_\phi[\mathcal{S}]$ ; or its merge sequence in bottom-up tokenisation, since  $\text{tok} = \text{tok}_\uparrow[\mathbf{m}]$  and since its vocabulary is chosen according to Eq. (6). (Note that in §3, we only showed how to apply tokenisers at inference time, but not how to find them.) In this section, we focus on direct tokenisation, defining its optimisation and decision problems; we then prove its NP-completeness. The optimisation problem is defined as follows.

**Definition 2.** Given a dataset  $\mathcal{D}$  and a vocabulary size  $K$ , the **direct tokenisation optimisation problem** is to find a vocabulary  $\mathcal{S}_{\text{opt}} \subset \Sigma^+$  which maximally compresses  $\mathcal{D}$ :

$$\mathcal{S}_{\text{opt}} = \arg \min_{\mathcal{S} \subset \Sigma^+} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| \quad (8)$$

s.t.  $|\mathcal{S}| = |\Sigma| + K$

We can similarly define direct tokenisation’s decision problem.

**Definition 3.** Given a dataset  $\mathcal{D}$  and a vocabulary size  $K$ , the **direct tokenisation decision problem** requires deciding whether there exists a vocabulary  $\mathcal{S} \subset \Sigma^+$  which compresses  $\mathcal{D}$  to at most  $\delta$  symbols:

$$\delta \geq \min_{\mathcal{S} \subset \Sigma^+} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| \quad (9)$$

s.t.  $|\mathcal{S}| = |\Sigma| + K$

We write  $\text{Tok}_\phi(\mathcal{D}, K, \delta)$  for a function which returns T if a direct tokenisation decision problem with those inputs is satisfiable, and F otherwise. Note that, whenever  $|\mathcal{D}| \leq K$ , the solution to the problem above is trivial, as an optimal solution simply requires including all strings  $\mathbf{c}_n$  in vocabulary  $\mathcal{S}$ . As we show next, however, in the general case the above decision problem is NP-complete. We now state this as a theorem, which we will prove in the next two sections.

**Theorem 1.** The direct tokenisation decision problem, as in Definition 3, is NP-complete.

*Proof.* A decision problem is considered to be NP-complete if: (i) it is in NP; (ii) it is NP-hard. We prove these conditions in §5.1 and §5.2.  $\square$


### 5.1 Direct Tokenisation is in NP

A decision problem is in the nondeterministic polynomial time class (NP) if, given a **certificate** of polynomial length, one can verify that certificate in polynomial time. Specifically, a certificate usually encodes a decision problem’s solution, allowing us to verify its satisfiability. In the case of direct tokenisation, this certificate would be a vocabulary  $\mathcal{S}$  which leads a dataset  $\mathcal{D}$  to be compressed to at most  $\delta$  symbols. Verifying this certificate simply requires computing the sum in Eq. (9), i.e.:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| \quad (10)$$

**Lemma 1.** The direct tokenisation decision problem, as in Definition 3, is in NP.

*Proof.* As noted above, whenever  $|\mathcal{D}| \leq K$ , each  $\mathbf{c}_n \in \mathcal{D}$  can be included in the vocabulary  $\mathcal{S}$  and fully compressed to a single symbol; we can thus verify the problem’s satisfiability by simply checking that  $\delta \geq |\mathcal{D}|$  as this is the best reachable compression. Assuming  $K$  to be bounded by  $|\mathcal{D}|$ —and therefore polynomial in the input—we have that the certificate  $\mathcal{S}$  also has polynomial length. Given such a certificate  $\mathcal{S}$ , verifying it simply requires computing the sum in Eq. (10). In turn, computing this sum requires  $|\mathcal{D}|$  calls to function  $\text{tok}_\phi$ . It follows that, if function  $\text{tok}_\phi$  runs in polynomial time, then direct tokenisation is in NP. Luckily, this function can indeed be computed efficiently using Schmidt et al.’s (2024) PathPiece method, which runs in  $O(|\mathbf{c}|^2)$  time. This is achieved by first converting  $\mathbf{c}$  into a directed acyclic graph where nodes represent string positions  $[0, 1, \dots, |\mathbf{c}|]$  and where two nodes  $t, t'$  are connected if there exists a subword  $s \in \mathcal{S}$  for which  $\mathbf{c}_{t:t'} \stackrel{\circ}{=} s$ . E.g., for a string  $\mathbf{c}$  and  $\mathcal{S} = \Sigma \cup \{\mathbf{c}_{0:2}, \mathbf{c}_{1:t}\}$ , we build a graph:



The shortest path from node 0 to  $|c|$  in this graph then gives us  $\text{tok}_\phi[\mathcal{S}](\mathbf{c})$ . As the shortest path of a directed acyclic graph can be computed in  $O(N + V)$  time, the time complexity of finding the shortest path in this directed acyclic graph is thus bounded by  $O(|\mathbf{c}|^2)$ .  $\square$

### 5.2 Direct Tokenisation is NP-hard

We now use a reduction from max-2-SAT to prove the NP-hardness of direct tokenisation.

**Reduction 1.** Let us have an instance of the max-2-SAT decision problem as in Definition 1. To reduce this instance to an instance of the direct tokenisation decision problem, as in Definition 3, we first define an alphabet  $\Sigma = \{\odot\} \cup \{x_j^T, x_j^F\}_{j=1}^J$ . We then construct three sets of strings:

$$\mathcal{D}_1 = \{\odot x_j^T \odot\}_{j=1}^J \cup \{\odot x_j^F \odot\}_{j=1}^J \quad (11a)$$

$$\mathcal{D}_2 = \{\odot x_j^T \odot x_j^F \odot\}_{j=1}^J \quad (11b)$$

$$\mathcal{D}_3 = \{\odot L_i^1 \odot L_i^2 \odot\}_{i=1}^I \quad (11c)$$

In these strings  $L_i$  is replaced by either character  $x_j^T$  or  $x_j^F$ , depending on whether it represents  $X_j$  or  $\neg X_j$ , respectively. We then construct our dataset  $\mathcal{D}$ , and choose  $K$  and  $\delta$  as:

$$\mathcal{D} = \left( \bigcup_{j=1}^f \mathcal{D}_1 \right) \cup \left( \bigcup_{j=1}^{f'} \mathcal{D}_2 \right) \cup \mathcal{D}_3 \quad (12a)$$

$$K = J, \quad \delta = (4f + 3f')J + 5I - 2\gamma \quad (12b)$$

where we set  $f' \stackrel{\text{def}}{=} 2I + 1$  and  $f \stackrel{\text{def}}{=} 4f'J + 4I + 1$ .

We write  $\text{R1}(\mathcal{X}, \mathcal{C}, \gamma)$  to represent a function which, given an instance of max-2-SAT, returns an instance of the tokenisation problem given by our reduction (i.e.,  $\mathcal{D}, K, \delta$ ). For our reduction to be correct, we must have that:

$$\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \iff \text{Tok}_\odot(\text{R1}(\mathcal{X}, \mathcal{C}, \gamma)) \quad (13)$$

meaning that a max-2-SAT instance is satisfiable if and only if its reduced direct tokenisation instance is as well. We now set out to prove this. We start by proving the forward direction of this iff clause.

**Lemma 2.** If a max-2-SAT instance is satisfiable, then the direct tokenisation instance output by Reduction 1 is also satisfiable. Formally:

$$\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \implies \text{Tok}_\odot(\text{R1}(\mathcal{X}, \mathcal{C}, \gamma)) \quad (14)$$

*Proof sketch.* See a formal proof in App. A. Our proof works by first fixing a satisfying solution to max-2-SAT with values  $x_j^*$ . Given this solution, for each variable, we add to our vocabulary  $\mathcal{S}$  a subword  $\odot x_j^T \odot$  if  $x_j^*$  is true, or  $\odot x_j^F \odot$  if  $x_j^*$  is false. Given these subwords, strings in  $\mathcal{D}_1$  and  $\mathcal{D}_2$  occupy a total length of  $(4f + 3f')J$ . Further, since at least  $\gamma$  of the max-2-SAT clauses are satisfied by  $x_j^*$ , the strings in  $\mathcal{D}_3$  will occupy a total length smaller or equal to  $5I - 2\gamma$ . This solution to the tokenisation problem thus gives us a total length which is smaller or equal to  $\delta = (4f + 3f')J + 5I - 2\gamma$ .  $\square$

Now, we are left with proving the backward direction of the iff clause in Eq. (13). We do so in the following lemma.

**Lemma 3.** If the direct tokenisation instance output by Reduction 1 is satisfiable, the max-2-SAT instance which generated it is as well. Formally:

$$\text{Tok}_\odot(\text{R1}(\mathcal{X}, \mathcal{C}, \gamma)) \implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (15)$$

*Proof sketch.* See a formal proof in App. B. Our proof works in three steps. First, we show that any satisfying solution must only have subwords of the form  $\odot x_j^T \odot$  or  $\odot x_j^F \odot$ , since this is required to compress strings in  $\mathcal{D}_1$  to at most  $4fJ$  symbols. Second, we show that any satisfying solution must only have either subword  $\odot x_j^T \odot$  or  $\odot x_j^F \odot$  for any variable  $X_j$ ; this is required to compress strings in  $\mathcal{D}_2$  to at most  $3f'J$  symbols. Finally, we show that if a tokeniser compresses strings in  $\mathcal{D}_3$  to  $5I - 2\gamma$ , then there is an assignment  $\chi$  which satisfies at least  $\gamma$  of the original max-2-SAT problem.  $\square$

Given both lemmas above, we can now trivially prove that direct tokenisation is NP-hard.

**Lemma 4.** The direct tokenisation decision problem, as in Definition 3, is NP-hard.

*Proof.* First, it is easy to see that Reduction 1 runs in polynomial time. Second, max-2-SAT is an NP-hard problem (Garey et al., 1974). This lemma then follows trivially from Lemmas 2 and 3, which together show that an instance of the tokenisation problem generated through Reduction 1 is satisfiable if and only if the max-2-SAT instance used to produce it is also satisfiable.  $\square$

## 6 Finding Optimal Bottom-up Tokenisers

We now shift our attention to bottom-up tokenisation. We define both its optimisation and decision problems, and then prove its NP-completeness. We start with defining the optimisation problem.

**Definition 4.** Given a dataset  $\mathcal{D}$  and a vocabulary size  $K$ , the **bottom-up tokenisation optimisation problem** is to find a merge sequence  $\mathbf{m}_{\text{opt}} \in \mathcal{M}^*$  which maximally compresses  $\mathcal{D}$ :

$$\mathbf{m}_{\text{opt}} = \arg \min_{\mathbf{m} \in \mathcal{M}^*} \sum_{c \in \mathcal{D}} |\text{tok}_\uparrow[\mathbf{m}](c)| \quad (16)$$

s.t.  $|\mathbf{m}| = K$

As can be seen, this optimisation problem is similar to the direct tokenisation problem, albeit its target is to find a merge sequence instead of a vocabulary. We similarly define a decision problem.

**Definition 5.** Given a dataset  $\mathcal{D}$  and a vocabulary size  $K$ , the **bottom-up tokenisation decision problem** requires deciding whether there exists a merge sequence  $\mathbf{m} \in \mathcal{M}^*$  which compresses  $\mathcal{D}$  to at most  $\delta$  symbols:

$$\delta \geq \min_{\mathbf{m} \in \mathcal{M}^*} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| \quad (17)$$

s.t.  $|\mathbf{m}| = K$

We write  $\text{Tok}_\uparrow(\mathcal{D}, K, \delta)$  for a function which returns T if a bottom-up tokenisation decision problem with those inputs is satisfiable, and F otherwise. We spend the rest of this section showing that bottom-up tokenisers are NP-complete.

**Theorem 2.** The bottom-up tokenisation decision problem, as in Definition 5, is NP-complete.

*Proof.* We prove this in two steps below. We first prove that this problem is in NP, in §6.1. We then prove that this problem is NP-hard, in §6.2.  $\square$

### 6.1 Bottom-up Tokenisation is in NP

We can verify this using a solution, the merge sequence  $\mathbf{m} \in \mathcal{M}^*$ , as a certificate. By showing that this certificate has polynomial length and that it can be verified in polynomial time, we prove this problem is in NP. To verify this certificate, we simply need to compute the sum in Eq. (17), i.e.:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| \quad (18)$$

which we show now can be done efficiently.

**Lemma 5.** The bottom-up tokenisation decision problem, as in Definition 5, is in NP.

*Proof.* First, if  $K$  is larger than the total number of characters in  $\mathcal{D}$ , i.e.,  $\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}|$ , then this dataset can be compressed to  $|\mathcal{D}|$  by merging each string down to a single symbol; further, compressing  $\mathcal{D}$  more than that is not possible independently of  $K$ . Verifying the satisfiability of such an instance of the tokenisation problem is thus trivial, only requiring checking if  $\delta \geq |\mathcal{D}|$ . Second, if  $K$  is bounded by  $\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}|$ —and therefore polynomial in the input—the certificate  $\mathbf{m}$  has polynomial length. Given such a certificate  $\mathbf{m}$ , verifying it then simply requires computing the sum in Eq. (18). In turn, computing this sum requires  $|\mathcal{D}|$  calls to function  $\text{tok}_\uparrow$ . It follows that, if function  $\text{tok}_\uparrow$  runs in polynomial time, then bottom-up tokenisation is in NP. The computation of  $\text{tok}_\uparrow$  can be done in polynomial time following the structure described in §3.2.

For each  $\mathbf{m} = \langle s_1, s_2 \rangle$  in  $\mathbf{m}$ , scan the current  $\mathbf{c}$  and replace each occurrence of  $s_1, s_2$  by  $s'$ . This takes time  $\mathcal{O}(|\mathbf{c}|)$  for each merge. Afterwards, the resulting string can be compared against the desired size. We obtain a total runtime of  $\mathcal{O}(|\mathcal{D}| |\mathbf{c}| |\mathbf{m}|)$ .  $\square$

### 6.2 Bottom-up Tokenisation is NP-hard

As before, we use a reduction from max-2-SAT to prove bottom-up tokenisation's NP-hardness.

**Reduction 2.** Let us have an instance of the max-2-SAT decision problem as in Definition 1. To reduce this instance to an instance of the bottom-up tokenisation decision problem, as in Definition 5, we first define an alphabet  $\Sigma = \{\odot, \otimes\} \cup \{x_j^T, x_j^F\}_{j=1}^J$ . We then construct five sets of strings:

$$\begin{aligned} \mathcal{D}_1 &= \{\odot x_j^T\}_{j=1}^J \cup \{x_j^F \odot\}_{j=1}^J \cup \{x_j^T \otimes\}_{j=1}^J \\ &\quad \cup \{\odot x_j^F\}_{j=1}^J \cup \{x_j^T \otimes\}_{j=1}^J \cup \{\otimes x_j^F\}_{j=1}^J \\ \mathcal{D}_2 &= \{\odot x_j^T \otimes\}_{j=1}^J \cup \{\otimes x_j^F \odot\}_{j=1}^J \\ &\quad \cup \{\odot x_j^T \otimes\}_{j=1}^J \cup \{\otimes x_j^F \odot\}_{j=1}^J \\ \mathcal{D}_3 &= \{\odot x_j^T \otimes x_{j'}^F \odot\}_{j=1}^J \cup \{\otimes x_j^F \odot x_{j'}^T \otimes\}_{j=1}^J \\ \mathcal{D}_4 &= \{\otimes x_j^F \odot x_{j'}^T \otimes\}_{j=1}^J \cup \{\otimes x_j^F \odot x_{j'}^T \odot\}_{j=1}^J \\ \mathcal{D}_5 &= \left\{ \begin{array}{l} \odot x_j^T \otimes x_{j'}^F \odot \text{ if } L_i^1 = X_j \text{ and } L_i^2 = \neg X_{j'} \\ \odot x_{j'}^T \otimes x_j^F \odot \text{ if } L_i^1 = \neg X_j \text{ and } L_i^2 = X_{j'} \\ \otimes x_j^F \odot x_{j'}^T \otimes \text{ if } L_i^1 = \neg X_j \text{ and } L_i^2 = \neg X_{j'} \\ \otimes x_j^T \otimes x_{j'}^F \otimes \text{ if } L_i^1 = X_j \text{ and } L_i^2 = X_{j'} \end{array} \right\}_{i=1}^I \end{aligned} \quad (19)$$

We then construct our dataset  $\mathcal{D}$ , and choose  $K$  and  $\delta$  as:

$$\begin{aligned} \mathcal{D} &= \bigcup_{j=1}^f \mathcal{D}_1 \cup \bigcup_{j=1}^{f'} \mathcal{D}_2 \cup \bigcup_{j=1}^{f''} \mathcal{D}_3 \cup \bigcup_{j=1}^{f'''} \mathcal{D}_4 \cup \mathcal{D}_5 \quad (20) \\ K &= 8J, \quad \delta = (6f + 6f' + 4f'' + 4f''') J + 3I - \gamma \end{aligned}$$

where we set:

$$f''' \stackrel{\text{def}}{=} 5I, \quad f'' \stackrel{\text{def}}{=} 10f'''J + 5I \quad (21a)$$

$$f' \stackrel{\text{def}}{=} (10f'' + 10f''') J + 5I \quad (21b)$$

$$f \stackrel{\text{def}}{=} (12f' + 10f'' + 10f''') J + 5I \quad (21c)$$

As before, we write  $\text{R2}(\mathcal{X}, \mathcal{C}, \gamma)$  for a function which, given an instance of the max-2-SAT problem, returns an instance of the bottom-up tokenisation problem. For our reduction to be correct, we must have that:

$$\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \iff \text{Tok}_\uparrow(\text{R2}(\mathcal{X}, \mathcal{C}, \gamma)) \quad (22)$$

We follow the same proof strategies as before, starting by proving the forward direction of this iff statement.

**Lemma 6.** *If a max-2-SAT instance is satisfiable, then the bottom-up tokenisation instance output by Reduction 2 is also satisfiable. Formally:*

$$\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \implies \text{Tok}_\uparrow(\text{R2}(\mathcal{X}, \mathcal{C}, \gamma)) \quad (23)$$

*Proof sketch.* See a formal proof in App. C. Without loss of generality, let a satisfying solution to max-2-SAT have values  $x_j^*$ . Our proof works by first defining the three following lists of merges, which must be included in any satisfying solution to this tokenisation problem:

$$\mathbf{m}_1 = \bigcirc_{j=1}^J [\langle \otimes, x_j^F \rangle, \langle x_j^T, \otimes \rangle] \quad (24a)$$

$$\mathbf{m}_3 = \bigcirc_{j=1}^J [\langle x_j^F, \odot \rangle, \langle \odot, x_j^T \rangle] \quad (24b)$$

$$\mathbf{m}_5 = \bigcirc_{j=1}^J [\langle \odot, x_j^F \rangle, \langle x_j^T, \odot \rangle] \quad (24c)$$

We then construct two other lists of merges, which depend on the satisfying assignments to max-2-SAT:

$$\mathbf{m}_2 = \bigcirc_{j=1}^J \left[ \begin{array}{ll} \langle \odot, x_j^T \otimes \rangle & \text{if } x_j^* = \text{T} \\ \langle \otimes x_j^F, \odot \rangle & \text{else} \end{array} \right] \quad (25a)$$

$$\mathbf{m}_4 = \bigcirc_{j=1}^J \left[ \begin{array}{ll} \langle \otimes x_j^T, \odot \rangle & \text{if } x_j^* = \text{T} \\ \langle \odot, x_j^F \otimes \rangle & \text{else} \end{array} \right] \quad (25b)$$

Finally, we create a merge sequence by concatenating these lists in order:

$$\mathbf{m} = \mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4 \circ \mathbf{m}_5 \quad (26)$$

Note that we have exactly  $K = 8J$  merges in this list. Given this merge sequence, it is easy to verify that strings in  $\mathcal{D}_1$  to  $\mathcal{D}_4$  will use exactly  $(6f + 6f' + 4f'' + 4f''') J$  symbols after being tokenised. Further, since at least  $\gamma$  of the max-2-SAT's clauses are satisfied by  $x_j^*$ , the strings in  $\mathcal{D}_5$  will occupy a total length smaller or equal to  $3I - \gamma$ . This solution to the tokenisation problem thus gives us a tokeniser which will compress  $\mathcal{D}$  to at most  $\delta = (6f + 6f' + 4f'' + 4f''') J + 3I - \gamma$ .  $\square$

We now prove the backward direction of the iff clause in Eq. (22).

**Lemma 7.** *If the bottom-up tokenisation instance output by Reduction 2 is satisfiable, the max-2-SAT instance which generated it is as well. Formally:*

$$\text{Tok}_\uparrow(\text{R2}(\mathcal{X}, \mathcal{C}, \gamma)) \implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (27)$$

*Proof sketch.* See a formal proof in App. D. Our proof works in five steps. First, we show that all satisfying solutions must include merges  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ , and  $\mathbf{m}_5$  from Eq. (24), since this is required to

compress strings in  $\mathcal{D}_1$  to at most  $6fJ$  symbols. Second, we show the other merges of any satisfying solution must be of the form:

$$\mathbf{m}_j^\odot = \left\{ \langle \odot x_j^T, \odot \rangle, \langle \odot, x_j^F \odot \rangle \right. \\ \left. \langle \odot, x_j^T \odot \rangle, \langle \odot x_j^F, \odot \rangle \right\} \quad (28a)$$

$$\mathbf{m}_j^\otimes = \left\{ \langle \odot, x_j^T \otimes \rangle, \langle \otimes x_j^F, \odot \rangle \right. \\ \left. \langle \odot x_j^T, \otimes \rangle, \langle \otimes, x_j^F \odot \rangle \right\} \quad (28b)$$

this is required to compress strings in  $\mathcal{D}_2$  to at most  $6f'J$  symbols. Third, we show that any satisfying solution will have at least one merge of each set  $\mathbf{m}_j^\odot$  and one of each set  $\mathbf{m}_j^\otimes$ ; this is required to compress strings in  $\mathcal{D}_3$  to at most  $4f''J$  symbols. Fourth, we show that any satisfying solution will have—for each  $j \in \{1, \dots, J\}$ —both its merges in sets  $\mathbf{m}_j^\odot$  and  $\mathbf{m}_j^\otimes$  containing character  $x_j^T$  or character  $x_j^F$ ; this is required to compress strings in  $\mathcal{D}_4$  to at most  $4f'''J$  symbols. Finally, we show that if a tokeniser compresses strings in  $\mathcal{D}_5$  to  $3I - \gamma$ , then there is an assignment  $\chi$  which satisfies at least  $\gamma$  of the original max-2-SAT problem.  $\square$

Finally, given both lemmas above, we can now prove that bottom-up tokenisation is NP-hard.

**Lemma 8.** *The bottom-up tokenisation decision problem, as in Definition 5, is NP-hard.*

*Proof.* First, it is easy to see that Reduction 2 runs in polynomial time. Second, max-2-SAT is an NP-hard problem (Garey et al., 1974). This lemma then follows trivially from Lemmas 6 and 7.  $\square$

### 6.3 Other Definitions of Tokenisation

We now expand our discussion to consider variations of the above tokenisation problems.

**Deduped Datasets.** Our definitions of both direct and bottom-up tokenisation allow datasets  $\mathcal{D}$  to include repeated entries. It is common, however, to deduplicate datasets in NLP—thus removing repeated entries. A small change to both our reductions is enough to adapt it to this deduplicated dataset case: simply append each string in the repeated datasets (either  $\mathcal{D}_1$  and  $\mathcal{D}_2$  in Reduction 1 or  $\mathcal{D}_1$  to  $\mathcal{D}_4$  in Reduction 2) with a unique character  $\{a_y\}_{y=1}^\infty$  and increase the target compression size  $\delta$  accordingly (by  $f + f'$  or  $f + f' + f'' + f'''$ , respectively). These new characters will never be included in optimal tokenisers' solutions, and thus the previous proofs hold, with the difference that each dataset will require extra symbols once compressed.

**A Single Long String.** In the previous sections, we considered tokenisers trained on a dataset  $\mathcal{D}$ . Work on compression, however, usually considers a single long string  $\mathbf{c}$  as its input. It is easy to see that direct tokenisation is not an NP-complete problem if its input is a single long string; including this string in vocabulary  $\mathcal{S}$  already achieves optimal compression. Bottom-up tokenisation, however, is still NP-complete even when given a single string as input. As before, this can be shown with a similar strategy to Reduction 2, but where we first append each string in dataset  $\mathcal{D}$  with a unique character  $\{a_y\}_{y=1}^{\infty}$  and then concatenate all these strings. As in the deduped case above, characters  $a_y$  will never be merged by any optimal tokeniser; they will thus serve as virtual string delimiters and will not affect our proofs beyond an increase to the target compression size  $\delta$ .

**A Hybrid Approach.** Finally, the last variant we consider is a hybrid between direct and bottom-up tokenisation, where we find a merge sequence  $\mathbf{m}$  which—when we extract a vocabulary from it as  $\mathcal{S} = \Sigma \cup \{s_1 \circ s_2 \mid \langle s_1, s_2 \rangle \in \mathbf{m}\}$ —optimally compresses a dataset  $\mathcal{D}$  using the direct tokenisation function in Eq. (4). We can easily prove the NP-hardness of this tokenisation variant by relying on Reduction 2; as our proof in Lemma 8 did not make use of the order of merges in  $\mathbf{m}$ , only of the subwords composed by it, this lemma’s proof strategy can be similarly applied to this hybrid variant.

## 7 Tokenisation and Compression

The variants of tokenisation that we consider here—with compression as their objective function—are closely related to the field of dictionary compression. In both fields, we wish to reduce the size of an input ( $\mathbf{c}$  or  $\mathcal{D}$ ) by exploiting repetitive elements. In fact, the most popular tokenisation algorithm to date, BPE, was originally proposed as a compression algorithm (Gage, 1994) and has only somewhat recently been ported into NLP to find tokenisers (by Sennrich et al., 2016).

Not surprisingly, prior work has also considered, from a theoretical perspective, the compression tokenisers achieve. Zouhar et al. (2023b), for instance, analyse bottom-up tokenisation and prove an approximation bound on the compression achieved by the tokenisers found using BPE. More recently, Kozma and Voderholzer (2024) also analyses bottom-up tokenisation, proving a tighter bound on this compression achieved by BPE.

A popular dictionary compression method, the **straight-line program** (SLP; Kieffer and Yang, 2000; Charikar et al., 2005), can be used to illustrate the similarities and differences between tokenisers and compressors.<sup>12</sup> Given a string  $\mathbf{c}$ , an SLP describes a context-free grammar from which  $\mathbf{c}$  can be uniquely derived. Formally, an SLP in Chomsky normal form (CNF) is a set of rules of form  $X \rightarrow a$  or  $X \rightarrow AB$ , where  $X, A, B$  are called nonterminals and  $a$  is a terminal.<sup>13</sup> Starting from a special nonterminal  $S$ , applying these rules exhaustively—until only terminals are left—produces exactly the desired string  $\mathbf{c}$ . Notably, given a string  $\mathbf{c}$ , it is NP-complete to find the smallest SLP which generates it (Charikar et al., 2005).

On the one hand, SLPs in CNF are closely linked to bottom-up tokenisation; each of its rules expands to two nonterminals, and thus corresponds to a merge. However, while SLPs must find the minimum number of merges (or rules) to fully compress a string into a single symbol, bottom-up tokenisers must maximally compress the string given a fixed number of merges. On the other hand, SLPs which are not in CNF (that is, for which other context-free production rules are allowed, as long as the decoding stays unique) are closely linked to direct tokenisation. In this case, a direct tokeniser could be converted into an SLP with depth two; this grammar has a start rule  $S \rightarrow \mathbf{s}$ , and a rule from each subword to its characters  $s \rightarrow \mathbf{c}$ . Again, while SLPs must find a minimal grammar representing the string, direct tokenisers must minimise the size of rule  $S \rightarrow \mathbf{s}$  given a fixed number of rules  $s \rightarrow \mathbf{c}$ .

The paragraphs above highlights two important differences between tokenisers and compressors. First, tokenisers aim to reduce only the size of the resulting tokenised text (i.e.,  $|\mathbf{s}|$ ), whereas compressors also consider the size of the compression information (e.g., considering the size required to store  $\mathcal{S}$ , which would be  $\sum_{s \in \mathcal{S}} |\text{detok}(s)|$ ). This is because tokenisers must create shorter inputs for NLP algorithms, while compressors must make information compact. Second, tokenisers and compressors have different optimisation parameters. Compression algorithms always compress a string to the best extent possible (e.g., for SLPs, until a single

<sup>12</sup>See Lohrey (2012) for an overview of straight-line programs, and Kempa and Prezza (2018); Kociumaka et al. (2023) for a more detailed overview of compression in general.

<sup>13</sup>Although not originally defined that way, SLP’s grammars are typically assumed to be in CNF, for simplicity. This does not make a big difference for compression, but will be important for our purposes.



nonterminal is reached), whereas tokenisation algorithms are given a maximum vocabulary size (i.e.,  $K$ ) and find tokenisers which only compress their input as much as possible until this limit is reached.

## 8 Conclusion

In this work, we proved the NP-completeness of two variants of tokenisation. These results underline that finding optimal tokenisers most likely will remain a difficult quest and that research should focus on approximate algorithms instead. Regarding those, there is potential both in improving the analysis of currently used algorithms, such as BPE, as well as in designing new ones.

## Limitations

While we prove the NP-completeness of multiple variants of the tokenisation problem—which is an important part of modern language modelling pipelines—we must note a few limitations in our work. First, we only prove NP-completeness of tokenisation with compression as its objective. This is a popular objective function, frequently used to judge the quality of tokenisers (e.g., Liu et al., 2025); however, it is not perfectly correlated with downstream language modelling performance, as discussed in §2.<sup>14</sup> Investigating the complexity of tokenisation under other objective functions is important. Second, our proofs do not assume a fixed alphabet size, so for fixed alphabets tokenisation might not be NP-complete. Tokenisers are frequently run at the byte level, for which specialised, more efficient algorithms might exist. Finally, while we investigated the complexity of the tokenisation problem for two types of tokenisation functions, similar results for other variants (with other tokenisation functions) remain open; we believe this would be exciting future work.

## Acknowledgments

We would like to thank Sotiris Anagnostidis, Hans-Joachim Böckenhauer, Clara Meister, Pietro Lesci, Anton Schäfer, Franz Nowak, and Marius Mosbach for their feedback on earlier versions of this manuscript. We would also like to thank Amit

<sup>14</sup>Relatedly, in practice, state-of-the-art language models still achieve very good performance (at least in English), despite suboptimal tokenisers. These tokenisers thus do not prevent good performance; rather they seem to make the training process less efficient. Studying the relationship between tokenisation choices and training efficiency, as in Schäfer et al. (2024) and Tao et al. (2024), is important future work.

Moryossef for discussions about the NP-hardness proof we present here. We further thank our three anonymous reviewers for their feedback; and one of them in particular, for suggesting we change our colour palette to make it more colour-blind friendly. Finally, we also thank the VMI of ETH Zürich for organising a scientific workshop that allowed the authors to meet and exchange ideas leading to this paper.

## References

- Orevaoghene Ahia, Sachin Kumar, Hila Gonen, Jungo Kasai, David Mortensen, Noah Smith, and Yulia Tsvetkov. 2023. [Do all languages cost the same? Tokenization in the era of commercial language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9904–9923, Singapore. Association for Computational Linguistics.
- Mehdi Ali, Michael Fromm, Klaudia Thellmann, Richard Rutmann, Max Lübbering, Johannes Leveling, Katrin Klug, Jan Ebert, Niclas Doll, Jasper Buschhoff, Charvi Jain, Alexander Weber, Lena Jurkschat, Hammam Abdelwahab, Chelsea John, Pedro Ortiz Suarez, Malte Ostendorff, Samuel Weinbach, Rafet Sifa, Stefan Kesselheim, and Nicolas Flores-Herr. 2024. [Tokenizer choice for LLM training: Negligible or crucial?](#) In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3907–3924, Mexico City, Mexico. Association for Computational Linguistics.
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*.
- M. Charikar, E. Lehman, Ding Liu, R. Panigrahy, M. Prabhakaran, A. Sahai, and A. Shelat. 2005. [The smallest grammar problem](#). *IEEE Transactions on Information Theory*, 51(7):2554–2576.
- Marco Cognition, Vilém Zouhar, Sangwhan Moon, and Naoaki Okazaki. 2024. [Two counterexamples to tokenization and the noiseless channel](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16897–16906, Torino, Italia. ELRA and ICCL.
- Philip Gage. 1994. [A new algorithm for data compression](#). *C Users Journal*, 12(2):23–38.
- Matthias Gallé. 2019. [Investigating the effectiveness of BPE: The power of shorter sequences](#). In *Proceedings of the 2019 Conference on Empirical Methods*

- in *Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1375–1381, Hong Kong, China. Association for Computational Linguistics.
- M. R. Garey, D. S. Johnson, and L. Stockmeyer. 1974. [Some simplified NP-complete problems](#). In *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing*, STOC '74, page 47–63, New York, NY, USA. Association for Computing Machinery.
- Renato Geh, Honghua Zhang, Kareem Ahmed, Benjie Wang, and Guy Van Den Broeck. 2024. [Where is the signal in tokenization space?](#) In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3966–3979, Miami, Florida, USA. Association for Computational Linguistics.
- Mario Giulianelli, Luca Malagutti, Juan Luis Gastaldi, Brian DuSell, Tim Vieira, and Ryan Cotterell. 2024. [On the proper treatment of tokenization in psycholinguistics](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Miami, Florida, USA. Association for Computational Linguistics.
- Omer Goldman, Avi Caciularu, Matan Eyal, Kris Cao, Idan Szpektor, and Reut Tsarfaty. 2024. [Unpacking tokenization: Evaluating text compression and its correlation with model performance](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2274–2286, Bangkok, Thailand. Association for Computational Linguistics.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Dominik Kempa and Nicola Prezza. 2018. [At the roots of dictionary compression: string attractors](#). In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, page 827–840, New York, NY, USA. Association for Computing Machinery.
- J.C. Kieffer and En-Hui Yang. 2000. [Grammar-based codes: a new class of universal lossless source codes](#). *IEEE Transactions on Information Theory*, 46(3):737–754.
- Tomasz Kociumaka, Gonzalo Navarro, and Nicola Prezza. 2023. [Toward a definitive compressibility measure for repetitive sequences](#). *IEEE Transactions on Information Theory*, 69(4):2074–2092.
- László Kozma and Johannes Voderholzer. 2024. [Theoretical analysis of byte-pair encoding](#). *Preprint*, arXiv:2411.08671.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Pietro Lesci, Clara Meister, Thomas Hofmann, Andreas Vlachos, and Tiago Pimentel. 2025. [Causal estimation of tokenisation bias](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Jia Peng Lim, Davin Choo, and Hady W. Lauw. 2025. [A partition cover approach to tokenization](#). *Preprint*, arXiv:2501.06246.
- Alisa Liu, Jonathan Hayase, Valentin Hofmann, Sewoong Oh, Noah A. Smith, and Yejin Choi. 2025. [SuperBPE: Space travel for language models](#). *Preprint*, arXiv:2503.13423.
- Markus Lohrey. 2012. [Algorithmics on SLP-compressed strings: A survey](#). *Groups - Complexity - Cryptology*, 4(2):241–299.
- Matteo Muffo, Aldo Cocco, and Enrico Bertino. 2022. [Evaluating transformer language models on arithmetic operations using number decomposition](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 291–297, Marseille, France. European Language Resources Association.
- Rodrigo Nogueira, Zhiying Jiang, and Jimmy Lin. 2021. [Investigating the limitations of transformers with simple arithmetic tasks](#). *Preprint*, arXiv:2102.13019.
- OpenAI. 2023. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Artidoro Pagnoni, Ram Pasunuru, Pedro Rodriguez, John Nguyen, Benjamin Muller, Margaret Li, Chunting Zhou, Lili Yu, Jason Weston, Luke Zettlemoyer, Gargi Ghosh, Mike Lewis, Ari Holtzman, and Sriniwasan Iyer. 2024. [Byte latent transformer: Patches scale better than tokens](#). *Preprint*, arXiv:2412.09871.
- Aleksandar Petrov, Emanuele La Malfa, Philip Torr, and Adel Bibi. 2023. [Language model tokenizers introduce unfairness between languages](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Buu Phan, Marton Havasi, Matthew Muckley, and Karen Ullrich. 2024. [Understanding and mitigating tokenization bias in language models](#). *Preprint*, arXiv:2406.16829.
- Tiago Pimentel and Clara Meister. 2024. [How to compute the probability of a word](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Miami, Florida, USA. Association for Computational Linguistics.

- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? On the monolingual performance of multilingual language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135, Online. Association for Computational Linguistics.
- Anton Schäfer, Thomas Hofmann, Imanol Schlag, and Tiago Pimentel. 2024. [On the effect of \(near\) duplicate subwords in language modelling](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9580–9597, Bangkok, Thailand. Association for Computational Linguistics.
- Craig W. Schmidt, Varshini Reddy, Haoran Zhang, Alec Alameddine, Omri Uzan, Yuval Pinter, and Chris Tanner. 2024. [Tokenization is more than compression](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 678–702, Miami, Florida, USA. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Aaditya K. Singh and DJ Strouse. 2024. [Tokenization counts: the impact of tokenization on arithmetic in frontier LLMs](#). *Preprint*, arXiv:2402.14903.
- Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muenighoff, Zhongwei Wan, Ping Luo, Min Lin, and Ngai Wong. 2024. [Scaling laws with vocabulary: Larger models deserve larger vocabularies](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [LLaMA: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Lili Yu, Daniel Simig, Colin Flaherty, Armen Aghajanyan, Luke Zettlemoyer, and Mike Lewis. 2023. [MEGABYTE: Predicting million-byte sequences with multiscale transformers](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Vilém Zouhar, Clara Meister, Juan Gastaldi, Li Du, Mrinmaya Sachan, and Ryan Cotterell. 2023a. [Tokenization and the noiseless channel](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5184–5207, Toronto, Canada. Association for Computational Linguistics.
- Vilém Zouhar, Clara Meister, Juan Gastaldi, Li Du, Tim Vieira, Mrinmaya Sachan, and Ryan Cotterell. 2023b. [A formal perspective on byte-pair encoding](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 598–614, Toronto, Canada. Association for Computational Linguistics.

## A Proof of Lemma 2

**Lemma 2.** *If a max-2-SAT instance is satisfiable, then the direct tokenisation instance output by Reduction 1 is also satisfiable. Formally:*

$$\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \implies \text{Tok}_\phi(\text{R1}(\mathcal{X}, \mathcal{C}, \gamma)) \quad (14)$$

*Proof.* First, note that if  $\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma)$ , then we have that Eq. (7) holds:  $\gamma \leq \max_{\chi \in \{\text{F}, \text{T}\}^J} \sum_{i=1}^I \mathbb{1}\{L_i^1 \vee L_i^2\}$ . Now, without loss of generality, let a satisfying solution have values  $x_j^*$ . In this case, for each variable  $X_j$ , we construct token  $\circledast x_j^{\text{T}} \circledast$  if  $x_j^*$  is true, or  $\circledast x_j^{\text{F}} \circledast$  if  $x_j^*$  is false. This gives us a total of  $J$  new tokens, so satisfies the  $|\mathcal{S}| = |\Sigma| + K$  condition. Now we just need to count the symbols output by this solution to see if Eq. (9) is satisfied, since any given tokenisation  $\text{tok}(\cdot, \mathcal{S})$  will provide an upper bound on the optimal tokenisation in terms of compression:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| \geq \min_{\mathcal{S}' \subseteq \Sigma^+} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}'](\mathbf{c})| \quad (29)$$

s.t.  $|\mathcal{S}'| = |\Sigma| + K$

For each pair of strings  $\circledast x_j^{\text{T}} \circledast$  and  $\circledast x_j^{\text{F}} \circledast$  in  $\mathcal{D}_1$ , one is compressed into a single subword while the other is kept as originally—using 3 symbols. We thus have that the strings in  $\mathcal{D}_1$  will occupy a total of  $(1 + 3)J$  characters, and:

$$\sum_{\mathbf{c} \in (\cup_{j=1}^J \mathcal{D}_1)} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| = 4fJ \quad (30)$$

Similarly, for each string in  $\mathcal{D}_2$  of form  $\circledast x_j^{\text{T}} \circledast \circledast x_j^{\text{F}} \circledast$ , we have that either token  $\circledast x_j^{\text{T}} \circledast$  or  $\circledast x_j^{\text{F}} \circledast$  exists. So each of these strings is compressed from 5 into 3 symbols. We thus have:

$$\sum_{\mathbf{c} \in (\cup_{j=1}^{f'} \mathcal{D}_2)} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| = 3f'J \quad (31)$$

Finally, we have strings in  $\mathcal{D}_3$  of form  $\circledast L_i^1 \circledast \circledast L_i^2 \circledast$ . These strings will be compressed into 3 symbols if  $\circledast L_i^1 \circledast$  or  $\circledast L_i^2 \circledast$  (or both) exist, and kept with 5 symbols otherwise. We thus have:

$$\sum_{\mathbf{c} \in \mathcal{D}_3} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| = \sum_{i=1}^I \left( 5 - 2 \mathbb{1} \left\{ \begin{array}{l} \circledast L_i^1 \circledast \in \mathcal{S} \\ \text{or} \\ \circledast L_i^2 \circledast \in \mathcal{S} \end{array} \right\} \right) \quad (32a)$$

$$= 5I - 2 \sum_{i=1}^I \mathbb{1} \left\{ \begin{array}{l} \circledast x_j^{\text{T}} \circledast \in \mathcal{S} \text{ and } L_i^1 = X_j \\ \text{or} \\ \circledast x_j^{\text{F}} \circledast \in \mathcal{S} \text{ and } L_i^1 = \neg X_j \\ \text{or} \\ \circledast x_{j'}^{\text{T}} \circledast \in \mathcal{S} \text{ and } L_i^2 = X_{j'} \\ \text{or} \\ \circledast x_{j'}^{\text{F}} \circledast \in \mathcal{S} \text{ and } L_i^2 = \neg X_{j'} \end{array} \right\} \quad (32b)$$

$$= 5I - 2 \sum_{i=1}^I \mathbb{1}_{\chi^*} \{L_i^1 \vee L_i^2\} \quad (32c)$$

$$\leq 5I - 2\gamma \quad (32d)$$

where, by construction, we have that a subword  $\circledast L_i \circledast \in \mathcal{S}$  if and only if its associated variable ( $X_j$  or  $\neg X_j$ ) is true. Summing together the lengths in Eqs. (30) to (32), we get that

$$\sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})| \leq \delta = (4f + 3f')J + 5I - 2\gamma \quad (33)$$

which concludes the proof.  $\square$

## B Proof of Lemma 3

**Lemma 3.** *If the direct tokenisation instance output by Reduction 1 is satisfiable, the max-2-SAT instance which generated it is as well. Formally:*

$$\text{Tok}_\phi(\text{R1}(\mathcal{X}, \mathcal{C}, \gamma)) \implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (15)$$

*Proof.* First, note that the dataset  $\mathcal{D}$  output by Reduction 1 has a total number of characters:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}| = (6f + 5f')J + 5I \quad (34)$$

Further, let:

$$\text{toklen}(\mathcal{D}, \mathcal{S}) \stackrel{\text{def}}{=} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\phi[\mathcal{S}](\mathbf{c})|, \quad \mathcal{S}_0 = \Sigma \cup \bigcup_{j=1}^J \{\circledast x_j^T \circledast, \circledast x_j^F \circledast\} \quad (35)$$

The maximum number of symbols in this dataset after compression is set to  $\delta = (4f + 3f')J + 5I - 2\gamma$ . This means that, to satisfy this objective, there must exist a vocabulary whose tokeniser compresses the text by at least  $(2f + 2f')J + 2\gamma$  symbols. We now prove this lemma in three steps: ① we show that any solution which compresses the text by at least  $2fJ$  symbols must only have nontrivial subwords<sup>15</sup> of the form  $\circledast x_j^T \circledast$  or  $\circledast x_j^F \circledast$ ; ② we show that any solution which compresses the text by at least  $(2f + 2f')J$  symbols must only have either subword  $\circledast x_j^T \circledast$  or  $\circledast x_j^F \circledast$  for any variable  $X_j$ ; ③ we show that any solution which compresses the text by at least  $(2f + 2f')J + 2\gamma$  symbols must be produced by a max-2-SAT instance which has at least  $\gamma$  clauses that are simultaneously satisfiable.  $\square$

**LemmaProofStep 1.** (Step ①). *Any solution which compresses the text by at least  $2fJ$  symbols must only have nontrivial subwords of the form  $\circledast x_j^T \circledast$  or  $\circledast x_j^F \circledast$ , i.e.,:*

$$\left( \text{toklen}(\mathcal{D}, \mathcal{S}) \leq \underbrace{(4f + 5f')J + 5I}_{\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}| - 2fJ} \right) \implies \mathcal{S} \subset \mathcal{S}_0 \quad (36)$$

*Proof.* First, given a solution with  $\mathcal{S} \subset \mathcal{S}_0$ , each subword  $s \in \mathcal{S}_{\setminus \Sigma}$ , where  $\mathcal{S}_{\setminus \Sigma} \stackrel{\text{def}}{=} \mathcal{S} \setminus \Sigma$ , will replace at least  $f$  strings in  $\mathcal{D}_1$ —i.e., with form  $\circledast x_j^T \circledast$  or  $\circledast x_j^F \circledast$ —for a single subword, thus saving  $2f$  characters. Since we have  $|\mathcal{S}_{\setminus \Sigma}| = K = J$  tokens, we save exactly  $2fJ$  symbols in  $\mathcal{D}_1$ :

$$\mathcal{S} \subset \mathcal{S}_0 \implies \left( \text{toklen}(\mathcal{D}_1, \mathcal{S}') = \underbrace{4fJ}_{\sum_{\mathbf{c} \in \mathcal{D}_1} |\mathbf{c}| - 2fJ} \right) \quad (37)$$

Note now that any solution  $\mathcal{S}'$  for which  $\mathcal{S}' \not\subset \mathcal{S}_0$  has at least one nontrivial subword which is not of the form  $\circledast x_j^T \circledast$  or  $\circledast x_j^F \circledast$ ; this subword  $s \notin \mathcal{S}_0$  will thus not compress strings in  $\mathcal{D}_1$  by  $2f$  symbols, but by at most  $f$  symbols:

$$\mathcal{S}' \not\subset \mathcal{S}_0 \implies \left( \text{toklen}(\mathcal{D}_1, \mathcal{S}') \geq \underbrace{4f(J-1) + 5f}_{\sum_{\mathbf{c} \in \mathcal{D}_1} |\mathbf{c}| - 2fJ + f} \right) \quad (38)$$

Even if this new subword were able to fully compress strings in  $\mathcal{D}_2$  and  $\mathcal{D}_3$  to a single symbol each, it would reach a compression of at most  $4f'J + 4I$ . Since by design  $f = 4f'J + 4I + 1$ , we get that:

$$\mathcal{S}' \not\subset \mathcal{S}_0 \implies \left( \text{toklen}(\mathcal{D}, \mathcal{S}') \geq 4fJ + f + f'J + I > (4f + 5f')J + 5I \right) \quad (39)$$

which concludes this step of the proof.  $\square$

<sup>15</sup>We define nontrivial subwords as subwords with more than one character. Remember that by definition  $\Sigma \subseteq \mathcal{S}$ , so all characters are always included in tokenisers' vocabularies. Also note that  $|\mathcal{S}| = |\Sigma| + K$ , so those trivial subwords are not counted towards vocabulary size  $K$ .

**LemmaProofStep 2.** (Step ②). Any solution which compresses the text by at least  $(2f + 2f')J$  symbols must only have either subword  $\otimes x_j^T \otimes$  or  $\otimes x_j^F \otimes$  for any variable  $X_j$ , i.e.,:

$$\left( \text{toklen}(\mathcal{D}, \mathcal{S}) \leq \underbrace{(4f + 3f')J + 5I}_{\sum_{c \in \mathcal{D}} |c| - (2f + 2f')J} \right) \implies \forall_{j \in \{1, \dots, J\}} |\mathcal{S} \cap \{\otimes x_j^T \otimes, \otimes x_j^F \otimes\}| = 1 \quad (40)$$

*Proof.* In this step of the proof, we show that satisfying solutions must have one and only one of subwords  $\otimes x_j^T \otimes$  and  $\otimes x_j^F \otimes$  for any variable  $X_j$ . As before, it's easy to see that a solution of the form described achieves at least  $(2f + 2f')J$  symbol compression. Each subword of form  $\otimes x_j^T \otimes$  or  $\otimes x_j^F \otimes$  saves exactly  $2f$  characters in the strings in  $\mathcal{D}_1$ . Further, because we always have either subword  $\otimes x_j^T \otimes$  or  $\otimes x_j^F \otimes$  for each value of  $j$ , we also get  $2f'$  compression in the strings in  $\mathcal{D}_2$ :

$$\begin{aligned} \forall_{j \in \{1, \dots, J\}} |\mathcal{S} \cap \{\otimes x_j^T \otimes, \otimes x_j^F \otimes\}| = 1 \\ \implies \left( \text{toklen}(\mathcal{D}_1, \mathcal{S}) = \underbrace{4fJ}_{\sum_{c \in \mathcal{D}_1} |c| - 2fJ} \right) \text{ and } \left( \text{toklen}(\mathcal{D}_2, \mathcal{S}) = \underbrace{3f'J}_{\sum_{c \in \mathcal{D}_2} |c| - 2f'J} \right) \end{aligned} \quad (41)$$

Now note that this is not true if both  $\otimes x_j^T \otimes$  and  $\otimes x_j^F \otimes$  exist for a single  $j$ ; in this case, only one of the tokens can be applied to  $\otimes x_j^T \otimes x_j^F \otimes$ , and thus both tokens together lead to a benefit of 2 instead of 4. If both  $\otimes x_j^T \otimes$  and  $\otimes x_j^F \otimes$  exist for any token  $X_j$ , this implies that neither of  $\otimes x_{j'}^T \otimes$  and  $\otimes x_{j'}^F \otimes$  exists for some other  $X_{j'}$ , resulting in an uncompressed string. Then, we get at most a compression of  $2fJ + 2f'(J - 1) + 4I$ :

$$\exists_{j \in \{1, \dots, J\}} |\mathcal{S}' \cap \{\otimes x_j^T \otimes, \otimes x_j^F \otimes\}| \neq 1 \implies \left( \text{toklen}(\mathcal{D}, \mathcal{S}') \geq \underbrace{(4f + 3f')J + 2f' + I}_{\sum_{c \in \mathcal{D}} |c| - (2f + 2f')J + 2f' - 4I} \right) \quad (42)$$

By construction  $f' = 2I + 1$ , which leads to:

$$\exists_{j \in \{1, \dots, J\}} |\mathcal{S}' \cap \{\otimes x_j^T \otimes, \otimes x_j^F \otimes\}| \neq 1 \implies \left( \text{toklen}(\mathcal{D}, \mathcal{S}') > (4f + 3f')J + 5I \right) \quad (43)$$

This concludes the proof.  $\square$

**LemmaProofStep 3.** (Step ③). Any instance of the tokenisation problem with a solution which compresses the text by at least  $(2f + 2f')J + 2\gamma$  symbols must be produced by a max-2-SAT problem with at least  $\gamma$  satisfied clauses, i.e.,:

$$\left( \text{toklen}(\mathcal{D}, \mathcal{S}) \leq \underbrace{(4f + 3f')J + 5I - 2\gamma}_{\sum_{c \in \mathcal{D}} |c| - (2f + 2f')J + 2\gamma} \right) \implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (44)$$

*Proof.* Finally, we now know any solution with this compression must have—for any variable  $X_j$ —either subword  $\otimes x_j^T \otimes$  or  $\otimes x_j^F \otimes$ . We can thus create a bijection  $\text{Conv}_{\mathcal{S} \rightarrow \chi}$  between the set of possible vocabularies respecting this condition, and the set of T/F assignments to SAT variables  $\chi$ :

$$\text{Conv}_{\mathcal{S} \rightarrow \chi}(\mathcal{S}) = \left\{ \begin{array}{ll} \text{T} & \text{if } \otimes x_j^T \otimes \in \mathcal{S} \\ \text{F} & \text{if } \otimes x_j^F \otimes \in \mathcal{S} \end{array} \right\}_{j=1}^J \quad (45)$$

Further, note that vocabularies of this form (as shown in Eq. (41)) lead to exactly  $(2f + 2f')J$  symbols being compressed in  $\mathcal{D}_1$  and  $\mathcal{D}_2$ . To achieve the target compression, a solution must thus compress  $\mathcal{D}_3$  by at least  $2\gamma$  symbols. Now note that for any string  $\otimes L_i^1 \otimes L_i^2 \otimes$  in  $\mathcal{D}_3$  we have three compression options:  $\otimes L_i^1 \otimes$  will be compressed, saving 2 symbols;  $\otimes L_i^2 \otimes$  will be compressed, also saving 2 symbols; or nothing will be compressed. More specifically,  $\otimes L_i^1 \otimes$  can be compressed if  $L_i^1$  represents  $X_j$  and subword  $\otimes x_j^T \otimes$  exists, or if  $L_i^1$  represents  $\neg X_j$  and subword  $\otimes x_j^F \otimes$  exists; the same is true for

$\odot L_i^2 \odot$ . They cannot both be compressed, however, as there is only one symbol  $\odot$  between the literals. We thus get a compression of 2 symbols for each of these strings if at least one of its literals has an associated subword in  $\mathcal{S}$ . Note that whenever a string  $\odot L_i^1 \odot L_i^2 \odot$  is compressed by 2 symbols using vocabulary  $\mathcal{S}$ , the max-2-SAT disjunction  $L_i^1 \vee L_i^2$  will also be satisfied by assignment  $\chi = \text{Conv}_{\mathcal{S} \rightarrow \chi}(\mathcal{S})$ ; similarly, whenever this string suffers no compression (i.e., having a compression of zero), the max-2-SAT disjunction will not be satisfied. As our condition assumes a compression of at least  $2\gamma$  symbols, we know that we have at least  $\gamma$  strings for which a literal has an associated subword. We can thus write:

$$2\gamma \leq \max_{\mathcal{S} \subset \Sigma^+} \sum_{\mathbf{c} \in \mathcal{D}_3} |\mathbf{c}| - |\text{tok}_{\uparrow}[\mathcal{S}](\mathbf{c})| \quad (46a)$$

$$\text{s.t. } |\mathcal{S}| = |\Sigma| + J \text{ and } \forall_{j \in \{1, \dots, J\}} |\mathcal{S} \cap \{\odot x_j^T \odot, \odot x_j^F \odot\}| = 1$$

$$= \max_{\mathcal{S} \subset \Sigma^+} \sum_{\odot L_i^1 \odot L_i^2 \odot \in \mathcal{D}_3} 2 \mathbb{1} \left\{ \begin{array}{l} \odot L_i^1 \odot \in \mathcal{S} \\ \text{or} \\ \odot L_i^2 \odot \in \mathcal{S} \end{array} \right\} \quad (46b)$$

$$\text{s.t. } |\mathcal{S}| = |\Sigma| + J \text{ and } \forall_{j \in \{1, \dots, J\}} |\mathcal{S} \cap \{\odot x_j^T \odot, \odot x_j^F \odot\}| = 1$$

$$= \max_{\chi \in \{0,1\}^J} \sum_{i=1}^I 2 \mathbb{1}_{\chi} \{L_i^1 \vee L_i^2\} \quad (46c)$$

$$\implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (46d)$$

We thus know that, if a satisfying tokenisation solution exists, then the associated max-2-SAT problem will also be satisfiable. This concludes the proof.  $\square$

## C Proof of Lemma 6

**Lemma 6.** *If a max-2-SAT instance is satisfiable, then the bottom-up tokenisation instance output by Reduction 2 is also satisfiable. Formally:*

$$\text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \implies \text{Tok}_{\uparrow}(\text{R2}(\mathcal{X}, \mathcal{C}, \gamma)) \quad (23)$$

*Proof.* Our proof starts by first defining the three following lists of merges, which will be included in any satisfying solution to the tokenisation problem:

$$\mathbf{m}_1 = \bigcirc_{j=1}^J [\langle \otimes, x_j^F \rangle, \langle x_j^T, \otimes \rangle], \quad \mathbf{m}_3 = \bigcirc_{j=1}^J [\langle x_j^F, \odot \rangle, \langle \odot, x_j^T \rangle], \quad \mathbf{m}_5 = \bigcirc_{j=1}^J [\langle \odot, x_j^F \rangle, \langle x_j^T, \odot \rangle] \quad (47)$$

Now, without loss of generality, let a satisfying solution to max-2-SAT have values  $x_j^*$ . We then construct two other lists of merges, which depend on this max-2-SAT solution:

$$\mathbf{m}_2 = \bigcirc_{j=1}^J \left[ \begin{array}{l} \langle \odot, x_j^T \otimes \rangle \text{ if } x_j^* = \text{T} \\ \langle \otimes x_j^F, \odot \rangle \text{ else} \end{array} \right], \quad \mathbf{m}_4 = \bigcirc_{j=1}^J \left[ \begin{array}{l} \langle \odot x_j^T, \odot \rangle \text{ if } x_j^* = \text{T} \\ \langle \odot, x_j^F \odot \rangle \text{ else} \end{array} \right] \quad (48)$$

In words, we create merges  $\langle \odot, x_j^T \otimes \rangle$  and  $\langle \odot x_j^T, \odot \rangle$  if  $x_j^*$  is true, or  $\langle \otimes x_j^F, \odot \rangle$  and  $\langle \odot, x_j^F \odot \rangle$  if  $x_j^*$  is false. We then create a merge sequence by concatenating these lists in order:

$$\mathbf{m} = \mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4 \circ \mathbf{m}_5 \quad (49)$$

This gives us a total of  $|\mathbf{m}| = K = 8J$  merges. Now we just need to count the symbols output by this solution to see if Eq. (17) is satisfied, since any given tokenisation  $\text{tok}_{\uparrow}[\mathbf{m}]$  will provide an upper bound on the optimal tokenisation in terms of compression:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})| \geq \min_{\mathbf{m}' \in \mathcal{M}^*} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_{\uparrow}[\mathbf{m}'](\mathbf{c})| \quad (50)$$

$$\text{s.t. } |\mathbf{m}'| = K$$

$\mathbf{c}$	$\text{tok}_\uparrow[\mathbf{m}_1](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4 \circ \mathbf{m}_5](\mathbf{c})$	$ \text{tok}_\uparrow[\mathbf{m}](\mathbf{c}) $
$\langle \odot, x_j^T \rangle$	$\cdot$	$\langle \odot x_j^T \rangle$	$\cdot$	1
$\langle x_j^F, \odot \rangle$	$\cdot$	$\langle x_j^F \odot \rangle$	$\cdot$	1
$\langle x_j^T, \odot \rangle$	$\cdot$	$\cdot$	$\langle x_j^T \odot \rangle$	1
$\langle \odot, x_j^F \rangle$	$\cdot$	$\cdot$	$\langle \odot x_j^F \rangle$	1
$\langle x_j^T, \otimes \rangle$	$\langle x_j^T \otimes \rangle$	$\cdot$	$\cdot$	1
$\langle \otimes, x_j^F \rangle$	$\langle \otimes x_j^F \rangle$	$\cdot$	$\cdot$	1

Table 1: Example of applying  $\mathbf{m}$  in  $\mathcal{D}_1$  of bottom-up tokenisation problem obtained from Reduction 2. The dot symbol  $\cdot$  denotes the string not changing under the given merge.

$\mathbf{c}$	$\text{tok}_\uparrow[\mathbf{m}_1](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2](\mathbf{c})$		$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4](\mathbf{c})$		$ \text{tok}_\uparrow[\mathbf{m}](\mathbf{c}) $	
		$x_j^* = \text{T}$	$x_j^* = \text{F}$		$x_j^* = \text{T}$	$x_j^* = \text{F}$	$x_j^* = \text{T}$	$x_j^* = \text{F}$
$\langle \odot, x_j^T, \odot \rangle$	$\cdot$	$\cdot$	$\cdot$	$\langle \odot x_j^T, \odot \rangle$	$\langle \odot x_j^T \odot \rangle$	$\langle \odot x_j^T, \odot \rangle$	1	2
$\langle \odot, x_j^F, \odot \rangle$	$\cdot$	$\cdot$	$\cdot$	$\langle \odot, x_j^F \odot \rangle$	$\langle \odot, x_j^F \odot \rangle$	$\langle \odot x_j^F \odot \rangle$	2	1
$\langle \odot, x_j^T, \otimes \rangle$	$\langle \odot, x_j^T \otimes \rangle$	$\langle \odot x_j^T \otimes \rangle$	$\langle \odot, x_j^T \otimes \rangle$	$\cdot$	$\cdot$	$\cdot$	1	2
$\langle \otimes, x_j^F, \odot \rangle$	$\langle \otimes x_j^F, \odot \rangle$	$\langle \otimes x_j^F, \odot \rangle$	$\langle \otimes x_j^F \odot \rangle$	$\cdot$	$\cdot$	$\cdot$	2	1

Table 2: Example of applying  $\mathbf{m}$  in  $\mathcal{D}_2$  of bottom-up tokenisation problem obtained from Reduction 2. The dot symbol  $\cdot$  denotes the string not changing under the given merge.  $\mathbf{m}_5$  is not depicted as it does not affect this dataset.

By applying the merges  $\mathbf{m}$ , each string in  $\mathcal{D}_1$  will be compressed into a single subword; note that  $\mathbf{m}_2$  and  $\mathbf{m}_4$  will have no effect on these strings. This is easy to see by applying merges sequentially to these strings, as displayed in Tab. 1. leading to:

$$\sum_{\mathbf{c} \in (\cup_{j=1}^f \mathcal{D}_1)} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| = 6fJ \quad (51)$$

For each pair of strings  $\odot x_j^T \odot$  and  $\odot x_j^F \odot$  in  $\mathcal{D}_2$ , one is compressed into a single subword while the other is only compressed to two subwords—the one with  $x_j^T$  is compressed to a single symbol if  $x_j^* = \text{T}$  and the one with  $x_j^F$  otherwise. The same is true for each pair of strings  $\odot x_j^T \otimes$  and  $\otimes x_j^F \odot$ , also in  $\mathcal{D}_2$ . This is displayed in Tab. 2. We thus have that, for each variable  $X_j$ , the strings in  $\mathcal{D}_2$  will occupy a total of  $(1 + 2 + 1 + 2)J$  characters, and:

$$\sum_{\mathbf{c} \in (\cup_{j=1}^f \mathcal{D}_2)} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| = 6f'J \quad (52)$$

Similarly, each string in  $\mathcal{D}_3$  and  $\mathcal{D}_4$  will be compressed into only 2 symbols after this tokeniser is applied to it. We thus have:

$$\sum_{\mathbf{c} \in (\cup_{j=1}^{f''} \mathcal{D}_3)} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| = 4f''J, \quad \sum_{\mathbf{c} \in (\cup_{j=1}^{f'''} \mathcal{D}_4)} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| = 4f'''J \quad (53)$$

Finally, we have the strings in  $\mathcal{D}_5$ . These strings are constructed such that they will be compressed into 2 symbols if either  $L_i^1$  or  $L_i^2$  evaluates to T, and kept with 3 symbols otherwise; see Tab. 4 for a detailed simulation of why this is the case. We thus have:

$$\sum_{\mathbf{c} \in \mathcal{D}_5} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| = \sum_{i=1}^I \left( 3 - \mathbb{1} \left\{ \begin{array}{l} L_i^1 = X_j \quad \text{and} \quad \langle \odot, x_j^T \otimes \rangle, \langle \odot x_j^T, \odot \rangle \in \mathbf{m} \\ \text{or} \\ L_i^1 = \neg X_j \quad \text{and} \quad \langle \otimes x_j^F, \odot \rangle, \langle \odot, x_j^F \odot \rangle \in \mathbf{m} \\ \text{or} \\ L_i^2 = X_{j'} \quad \text{and} \quad \langle \odot, x_j^T \otimes \rangle, \langle \odot x_j^T, \odot \rangle \in \mathbf{m} \\ \text{or} \\ L_i^2 = \neg X_{j'} \quad \text{and} \quad \langle \otimes x_j^F, \odot \rangle, \langle \odot, x_j^F \odot \rangle \in \mathbf{m} \end{array} \right. \right) \quad (54a)$$

$$= 3I - \sum_{i=1}^I \mathbb{1}_{\mathcal{X}^*} \{L_i^1 \vee L_i^2\} \quad (54b)$$

$$\leq 3I - \gamma \quad (54c)$$



$\mathcal{D}$	$\mathbf{c}$	$\text{tok}_\uparrow[\mathbf{m}_1](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2](\mathbf{c})$		$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3](\mathbf{c})$		$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4](\mathbf{c})$		$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4 \circ \mathbf{m}_5](\mathbf{c})$		$ \text{tok}_\uparrow[\mathbf{m}](\mathbf{c}) $
			$x_j^* = \text{T}$	$x_j^* = \text{F}$	$x_j^* = \text{T}$	$x_j^* = \text{F}$	$x_j^* = \text{T}$	$x_j^* = \text{F}$	$x_j^* = \text{T}$	$x_j^* = \text{F}$	
$\mathcal{D}_3$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
$\mathcal{D}_3$	$\langle \otimes, x_j^{\text{F}}, \otimes, x_j^{\text{T}}, \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	2
$\mathcal{D}_4$	$\langle \otimes, x_j^{\text{F}}, \otimes, x_j^{\text{T}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
$\mathcal{D}_4$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	2

Table 3: Example of applying  $\mathbf{m}$  in  $\mathcal{D}_3$  and  $\mathcal{D}_4$  of the bottom-up tokenisation problem obtained from Reduction 2. The dot symbol  $\cdot$  denotes the string not changing under the given merge.

Assignment	Condition	$\mathbf{c}$	$\text{tok}_\uparrow[\mathbf{m}_1](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3](\mathbf{c})$	$\text{tok}_\uparrow[\mathbf{m}_1 \circ \mathbf{m}_2 \circ \mathbf{m}_3 \circ \mathbf{m}_4](\mathbf{c})$	$ \text{tok}_\uparrow[\mathbf{m}](\mathbf{c}) $
$L_i^1 = X_j$ and $L_i^2 = \neg X_j$	$x_j^* = \text{T} \wedge x_j^* = \text{T}$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
	$x_j^* = \text{F} \wedge x_j^* = \text{T}$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	3
	$x_j^* = \text{T} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
	$x_j^* = \text{F} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
$L_i^1 = \neg X_j$ and $L_i^2 = X_j$	$x_j^* = \text{T} \wedge x_j^* = \text{T}$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
	$x_j^* = \text{F} \wedge x_j^* = \text{T}$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{F}}, \otimes \rangle$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
	$x_j^* = \text{T} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	3
	$x_j^* = \text{F} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	$\langle \otimes x_j^{\text{T}}, \otimes, x_j^{\text{F}} \otimes \rangle$	2
$L_i^1 = \neg X_j$ and $L_i^2 = \neg X_j$	$x_j^* = \text{T} \wedge x_j^* = \text{T}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\cdot$	3
	$x_j^* = \text{F} \wedge x_j^* = \text{T}$	$\langle \otimes, x_j^{\text{F}}, \otimes, x_j^{\text{T}}, \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	2
	$x_j^* = \text{T} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	2
	$x_j^* = \text{F} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes x_j^{\text{F}}, \otimes, x_j^{\text{T}} \otimes \rangle$	2
$L_i^1 = X_j$ and $L_i^2 = X_j$	$x_j^* = \text{T} \wedge x_j^* = \text{T}$	$\cdot$	$\cdot$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	2
	$x_j^* = \text{F} \wedge x_j^* = \text{T}$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}}, \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\cdot$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	2
	$x_j^* = \text{T} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	2
	$x_j^* = \text{F} \wedge x_j^* = \text{F}$	$\cdot$	$\cdot$	$\cdot$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	$\langle \otimes, x_j^{\text{T}}, \otimes, x_j^{\text{T}} \otimes \rangle$	3

Table 4: Example of applying  $\mathbf{m}$  in  $\mathcal{D}_5$  of the bottom-up tokenisation problem obtained from Reduction 2. The dot symbol  $\cdot$  denotes the string not changing under the given merge.  $\mathbf{m}_5$  is not depicted as it does not affect this dataset.

where, by construction, we have a merge in our sequence (e.g.,  $\langle \otimes, x_j^{\text{T}} \otimes \rangle$  or  $\langle \otimes x_j^{\text{F}}, \otimes \rangle$ ) if and only if its value is in a satisfying assignment (e.g.,  $x_j^* = \text{T}$  or  $x_j^* = \text{F}$  respectively). Summing together the lengths in Eqs. (51) to (54), we get that:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| \leq \delta = (6f + 6f' + 4f'' + 4f''')J + 3I - \gamma \quad (55)$$

which concludes the proof.  $\square$

## D Proof of Lemma 7

**Lemma 7.** *If the bottom-up tokenisation instance output by Reduction 2 is satisfiable, the max-2-SAT instance which generated it is as well. Formally:*

$$\text{Tok}_\uparrow(\text{R2}(\mathcal{X}, \mathcal{C}, \gamma)) \implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (27)$$

*Proof.* First, note that:

$$\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}| = (12f + 12f' + 10f'' + 10f''')J + 5I \quad (56)$$

Further, let:

$$\text{toklen}(\mathcal{D}, \mathbf{m}) \stackrel{\text{def}}{=} \sum_{\mathbf{c} \in \mathcal{D}} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})| \quad (57)$$

$$\begin{aligned} \mathbf{m}_1 &= \bigcirc_{j=1}^J [\langle \otimes, x_j^{\text{F}} \rangle, \langle x_j^{\text{T}}, \otimes \rangle], & \mathbf{m}_3 &= \bigcirc_{j=1}^J [\langle x_j^{\text{F}}, \otimes \rangle, \langle \otimes, x_j^{\text{T}} \rangle], & \mathbf{m}_5 &= \bigcirc_{j=1}^J [\langle \otimes, x_j^{\text{F}} \rangle, \langle x_j^{\text{T}}, \otimes \rangle] \\ \mathbf{m}_j^\otimes &= \left\{ \langle \otimes x_j^{\text{T}}, \otimes \rangle, \langle \otimes, x_j^{\text{F}} \otimes \rangle \right\}, & \mathbf{m}_j^\otimes &= \left\{ \langle \otimes, x_j^{\text{T}} \otimes \rangle, \langle \otimes x_j^{\text{F}}, \otimes \rangle \right\} \\ \mathbf{m}_j^\otimes &= \left\{ \langle \otimes, x_j^{\text{T}} \otimes \rangle, \langle \otimes x_j^{\text{F}}, \otimes \rangle \right\}, & \mathbf{m}_j^\otimes &= \left\{ \langle \otimes x_j^{\text{T}}, \otimes \rangle, \langle \otimes, x_j^{\text{F}} \otimes \rangle \right\} \\ \mathbf{m}_j^\text{T} &= \left\{ \langle \otimes x_j^{\text{T}}, \otimes \rangle, \langle \otimes, x_j^{\text{T}} \otimes \rangle \right\}, & \mathbf{m}_j^\text{F} &= \left\{ \langle \otimes, x_j^{\text{F}} \otimes \rangle, \langle \otimes x_j^{\text{F}}, \otimes \rangle \right\} \\ \mathbf{m}_j^\text{T} &= \left\{ \langle \otimes, x_j^{\text{T}} \otimes \rangle, \langle \otimes x_j^{\text{T}}, \otimes \rangle \right\}, & \mathbf{m}_j^\text{F} &= \left\{ \langle \otimes x_j^{\text{F}}, \otimes \rangle, \langle \otimes, x_j^{\text{F}} \otimes \rangle \right\} \end{aligned}$$

The maximum number of symbols in this dataset after compression is set to  $\delta = (6f+6f'+4f''+4f''') J + 3I - \gamma$ . This means that to satisfy this objective, there must exist a vocabulary whose tokeniser compresses the text by at least  $(6f + 6f' + 6f'' + 6f''') J + 2I + \gamma$  symbols. We now prove this lemma in five steps: ① we show that any solution which compresses the text by at least  $6fJ$  symbols must include all merges in  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ , and  $\mathbf{m}_5$ ; ② we show that any solution which compresses the text by at least  $(6f + 6f')J$  symbols must only include either merges in  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ ,  $\mathbf{m}_5$ , or in either  $\mathbf{m}_j^\circledast$  or  $\mathbf{m}_j^\otimes$ ; ③ we show that any solution which compresses the text by at least  $(6f + 6f' + 6f'')J$  symbols must include, for each  $j \in \{1, \dots, J\}$ , exactly one merge in set  $\mathbf{m}_j^\circledast$  and one in set  $\mathbf{m}_j^\otimes$ ; ④ we show that any solution which compresses the text by at least  $(6f + 6f' + 6f'' + 6f''')J$  symbols must include, for each  $j \in \{1, \dots, J\}$ , exactly two merges in either set  $\mathbf{m}_j^T$  or in set  $\mathbf{m}_j^F$ ; ⑤ we show that any solution which compresses the text by at least  $(6f + 6f' + 6f'' + 6f''')J + 2I + \gamma$  symbols must be produced by a max-2-SAT problem with at least  $\gamma$  satisfied clauses.  $\square$

**LemmaProofStep 1.** (Step ①). *Any solution which compresses the text by at least  $6fJ$  symbols must include all merges in  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ , and  $\mathbf{m}_5$ , i.e.,:*

$$\left( \text{toklen}(\mathcal{D}, \mathbf{m}) \leq \underbrace{(6fJ + (12f' + 10f'' + 10f''')J + 5I)}_{\sum_{c \in \mathcal{D}} |c| - 6fJ} \right) \quad (58)$$

$$\implies \underbrace{\bigcirc_{j=1}^J [\langle \otimes, x_j^F \rangle, \langle x_j^T, \otimes \rangle]}_{\mathbf{m}_1} \subset \mathbf{m}, \quad \underbrace{\bigcirc_{j=1}^J [\langle x_j^F, \otimes \rangle, \langle \otimes, x_j^T \rangle]}_{\mathbf{m}_3} \subset \mathbf{m}, \quad \underbrace{\bigcirc_{j=1}^J [\langle \otimes, x_j^F \rangle, \langle x_j^T, \otimes \rangle]}_{\mathbf{m}_5} \subset \mathbf{m}$$

*Proof.* We prove this statement by contradiction. Assume that one of the subwords above is not present in the tokeniser's merge sequence  $\mathbf{m}$ . In that case, the strings in  $\mathcal{D}_1$  which contain this character string will not be compressed, and will thus still be represented with 2 symbols. There will thus be at most  $6J - 1$  strings in  $\mathcal{D}_1$  represented with a single symbol, and at least one represented with two symbols. The minimum length achievable would thus be:

$$\text{toklen}(\mathcal{D}, \mathbf{m}) = \underbrace{\sum_{c \in \bigcup_{i=1}^f \mathcal{D}_1} |\text{tok}_\uparrow[\mathbf{m}](c)|}_{\geq (6J-1)f+2f} + \underbrace{\sum_{c \in \mathcal{D} \setminus (\bigcup_{i=1}^f \mathcal{D}_1)} |\text{tok}_\uparrow[\mathbf{m}](c)|}_{>0} \quad (59a)$$

$$> (6J + 1)f \quad \text{By construction } f = (12f' + 10f'' + 10f''')J + 5I \quad (59b)$$

$$= (6f + 12f' + 10f'' + 10f''')J + 5I \quad (59c)$$

which contradicts the proof's statement.  $\square$

**LemmaProofStep 2.** (Step ②). *Any solution which compresses the text by at least  $(6f + 6f')J$  symbols must only include either merges in  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ ,  $\mathbf{m}_5$ , or in either  $\mathbf{m}_j^\circledast$  or  $\mathbf{m}_j^\otimes$ , i.e.,:*

$$\left( \text{toklen}(\mathcal{D}, \mathcal{S}) \leq \underbrace{(6f + 6f')J + (10f'' + 10f''')J + 5I}_{\sum_{c \in \mathcal{D}} |c| - (6f+6f')J} \right) \quad (60)$$

$$\implies \mathbf{m} \setminus (\mathbf{m}_1 \circ \mathbf{m}_3 \circ \mathbf{m}_5) \subseteq \underbrace{\left\{ \langle \otimes, x_j^T \otimes \rangle, \langle \otimes x_j^F, \otimes \rangle, \langle \otimes x_j^T, \otimes \rangle, \langle \otimes, x_j^F \otimes \rangle \right\}_{j=1}^J}_{\bigcup_{j=1}^J (\mathbf{m}_j^\circledast \cup \mathbf{m}_j^\otimes)}$$

*Proof.* We again prove this statement by contradiction. Assume that  $\mathbf{m}$  has all merges  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ ,  $\mathbf{m}_5$ , but one of its other merges is in neither of the sets  $\mathbf{m}_j^\circledast$  and  $\mathbf{m}_j^\otimes$ . This means that at least one of the sets  $\mathbf{m}_j^\circledast \cup \mathbf{m}_j^\otimes$  will not have at least two merges in the solution; this is because there are  $J$  such sets  $(\mathbf{m}_j^\circledast \cup \mathbf{m}_j^\otimes)$ , which—coupled together with the  $6J$  already selected merges in  $\mathbf{m}_1$ ,  $\mathbf{m}_3$ ,  $\mathbf{m}_5$ —would amount to the maximum of  $8J$  merges. In that case, the strings (e.g.,  $\otimes x_j^T \otimes$ ,  $\otimes x_j^F \otimes$ ,  $\otimes x_j^T \otimes$  and  $\otimes x_j^F \otimes$ ) in  $\mathcal{D}_2$  containing the characters this absent merge represents will not be fully compressed to a single symbol,

being represented with 2 symbols instead. Out of the  $4J$  strings in  $\mathcal{D}_2$  then, there will thus be at most  $2J - 1$  represented with a single symbol, and at least  $2J + 1$  represented with two symbols—resulting in a total of at least  $2J - 1 + 2(2J + 1) = 6J + 1$  symbols. The minimum length achievable would thus be:

$$\text{toklen}(\mathcal{D}, \mathbf{m}) = \underbrace{\sum_{\mathbf{c} \in \bigcup_{j=1}^f \mathcal{D}_1} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})|}_{=6fJ} + \underbrace{\sum_{\mathbf{c} \in \bigcup_{j=1}^{f'} \mathcal{D}_2} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})|}_{\geq(6J+1)f'} + \underbrace{\sum_{\mathbf{c} \in \mathcal{D} \setminus (\mathcal{D}_1 \cup \mathcal{D}_2)} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})|}_{>0} \quad (61a)$$

$$> 6fJ + (6J + 1)f' \quad \text{By construction } f' = (10f'' + 10f''')J + 5I \quad (61b)$$

$$= (6f + 6f' + 10f'' + 10f''')J + 5I \quad (61c)$$

which contradicts the proof's statement.  $\square$

**LemmaProofStep 3.** (Step ③). *Any solution which compresses the text by at least  $(6f + 6f' + 6f'')J$  symbols must include all merges in  $\mathbf{m}_1, \mathbf{m}_3, \mathbf{m}_5$ , and, for each  $j \in \{1, \dots, J\}$ , exactly one merge in set  $\mathbf{m}_j^{\odot}$  and one in set  $\mathbf{m}_j^{\otimes}$ , i.e.,:*

$$\left( \text{toklen}(\mathcal{D}, \mathbf{m}) \leq \underbrace{(6f + 6f' + 4f'')J + 10f'''J + 5I}_{\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}| - (6f + 6f' + 6f'')J} \right) \quad (62)$$

$$\implies \forall_{j \in \{1, \dots, J\}} \left| \underbrace{\mathbf{m} \cap \left\{ \begin{array}{l} \langle \otimes x_j^T, \odot \rangle, \langle \odot, x_j^F \otimes \rangle \\ \langle \odot, x_j^T \otimes \rangle, \langle \otimes x_j^F, \odot \rangle \end{array} \right\}}_{\mathbf{m}_j^{\odot}} \right| = 1 \quad \text{and} \quad \left| \underbrace{\mathbf{m} \cap \left\{ \begin{array}{l} \langle \odot, x_j^T \otimes \rangle, \langle \otimes x_j^F, \odot \rangle \\ \langle \otimes x_j^T, \otimes \rangle, \langle \otimes, x_j^F \otimes \rangle \end{array} \right\}}_{\mathbf{m}_j^{\otimes}} \right| = 1$$

*Proof.* We again prove this statement by contradiction. First, assume that  $\mathbf{m}$  contains all the merges in  $\mathbf{m}_1, \mathbf{m}_3, \mathbf{m}_5$ ; further, assume all its other merges are contained in sets  $\mathbf{m}_j^{\odot}$  and  $\mathbf{m}_j^{\otimes}$ . Note now that, if any merge in  $\mathbf{m}_j^{\otimes}$  is in the selected merges  $\mathbf{m}$ , the string  $\otimes x_j^F \otimes x_j^T \otimes$  in  $\mathcal{D}_3$  will be compressed to 2 symbols (e.g.,  $\langle \otimes x_j^F, \otimes x_j^T \otimes \rangle$ ); if none of these merges is present, however, this string will only be compressed to 3 symbols (e.g.,  $\langle \otimes x_j^F, \odot, x_j^T \otimes \rangle$ ). The same is true for strings  $\odot x_j^T \otimes x_j^F \odot$  and merges in  $\mathbf{m}_j^{\odot}$ . Now, assume the contradictory case: for a value of  $j \in \{1, \dots, J\}$ ,  $\mathbf{m}$  does not satisfy the condition above. As, by construction, our solution has  $K = 8J$  merges, and because  $|\mathbf{m}_1 \circ \mathbf{m}_3 \circ \mathbf{m}_5| = 6J$ , we know that we have  $2J$  merges in sets  $\mathbf{m}_j^{\odot}$  and  $\mathbf{m}_j^{\otimes}$ . As there are exactly  $2J$  such sets, if the condition above does not hold, at least one of these sets must have no merge present in  $\mathbf{m}$ . In that case, the strings in  $\mathcal{D}_3$  which contain the character string represented by these absent merges will be compressed to three symbols, while others will be compressed to two symbols. There will thus be at most  $2J - 1$  strings in  $\mathcal{D}_3$  represented with two symbols, and at least one represented with three symbols. The minimum length achievable would thus be:

$$\text{toklen}(\mathcal{D}, \mathbf{m}) = \underbrace{\sum_{\mathbf{c} \in \bigcup_{j=1}^f \mathcal{D}_1 \cup \bigcup_{j=1}^{f'} \mathcal{D}_2} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})|}_{=(6f+6f')J} + \underbrace{\sum_{\mathbf{c} \in \bigcup_{j=1}^{f''} \mathcal{D}_3} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})|}_{\geq(2J-1)2f''+3f''} + \underbrace{\sum_{\mathbf{c} \in \bigcup_{j=1}^{f'''} \mathcal{D}_4 \cup \mathcal{D}_5} |\text{tok}_{\uparrow}[\mathbf{m}](\mathbf{c})|}_{>0} \quad (63a)$$

$$> (6f + 6f')J + (4J + 1)f'' \quad \text{By construction } f'' = 10f'''J + 5I \quad (63b)$$

$$= (6f + 6f' + 4f'' + 10f''')J + 5I \quad (63c)$$

which contradicts the proof's statement.  $\square$

**LemmaProofStep 4.** (Step ④). *Any solution which compresses the text by at least  $(6f + 6f' + 6f'' + 6f''')J$  symbols must include all merges in  $\mathbf{m}_1, \mathbf{m}_3, \mathbf{m}_5$ , and, for each  $j \in \{1, \dots, J\}$ , exactly one merge*

in set  $\mathbf{m}_j^\circledast$  and one in set  $\mathbf{m}_j^\otimes$ , such that either both these merges are in  $\mathbf{m}_j^T$  or both are in  $\mathbf{m}_j^F$ , i.e.,:

$$\left( \text{toklen}(\mathcal{D}, \mathbf{m}) \leq \underbrace{(6f + 6f' + 4f'' + 4f''')J + 5I}_{\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}| - (6f + 6f' + 6f'' + 6f''')J} \right) \quad (64)$$

$$\implies \forall_{j \in \{1, \dots, J\}} \left| \mathbf{m} \cap \underbrace{\left\{ \langle \circledast x_j^T, \circ \rangle, \langle \circ, x_j^T \otimes \rangle, \right\}}_{\mathbf{m}_j^T} \right| = 2 \text{ or } \left| \mathbf{m} \cap \underbrace{\left\{ \langle \circ, x_j^F \circ \rangle, \langle \otimes x_j^F, \circ \rangle, \right\}}_{\mathbf{m}_j^F} \right| = 2$$

*Proof.* First, note that the conditions of the step of our proof are stricter than previous ones, so we assume the conditions of steps ① to ③ hold—i.e.,  $\mathbf{m}$  contains all merges in  $\mathbf{m}_1, \mathbf{m}_3, \mathbf{m}_5$ ; further, it has one and only one merge from each set  $\mathbf{m}_j^\circledast$  and  $\mathbf{m}_j^\otimes$ . (Note that  $\mathbf{m}_j^\circledast \cup \mathbf{m}_j^\otimes = \mathbf{m}_j^T \cup \mathbf{m}_j^F$ , and that the just-mentioned condition implies  $|\mathbf{m} \cap (\mathbf{m}_j^T \cup \mathbf{m}_j^F)| = 2$ .) We now again prove this statement by contradiction. Consider now the case:

$$\left| \mathbf{m} \cap \underbrace{\left\{ \langle \circledast x_j^T, \circ \rangle, \langle \circ, x_j^T \otimes \rangle, \right\}}_{\mathbf{m}_j^T} \right| = 2 \text{ or } \left| \mathbf{m} \cap \underbrace{\left\{ \langle \circ, x_j^F \circ \rangle, \langle \otimes x_j^F, \circ \rangle, \right\}}_{\mathbf{m}_j^F} \right| = 2 \quad (65)$$

If this is true, then strings  $\circledast x_j^F \circ x_j^T \otimes$  and  $\otimes x_j^F \circ x_j^T \circ$  in  $\mathcal{D}_4$  will be compressed to 2 symbols each (e.g., to  $\langle \circledast x_j^F, \circ x_j^T \otimes \rangle$  and  $\langle \otimes x_j^F, \circ x_j^T \circ \rangle$  or  $\langle \circledast x_j^F \circ, x_j^T \otimes \rangle$  and  $\langle \otimes x_j^F \circ, x_j^T \circ \rangle$ ); if this condition is false, however, one of these strings will only be compressed to 3 symbols (e.g., to  $\langle \circledast x_j^F, \circ x_j^T \otimes \rangle$  and  $\langle \otimes x_j^F, \circ, x_j^T \circ \rangle$ ). Now, assume the contradictory case: for a value of  $j \in \{1, \dots, J\}$ ,  $\mathbf{m}$  does not satisfy the condition above. In that case, the strings in  $\mathcal{D}_4$  for which the condition does not hold will be compressed to  $3 + 2$  symbols, while others will be compressed to  $2 + 2$  symbols. There will thus be at most  $2J - 1$  strings in  $\mathcal{D}_4$  represented with two symbols, and at least one represented with three symbols. The minimum length achievable would thus be:

$$\begin{aligned} \text{toklen}(\mathcal{D}, \mathbf{m}) &= \underbrace{\sum_{\mathbf{c} \in \bigcup_{i=1}^f \mathcal{D}_1 \cup \bigcup_{i=1}^{f'} \mathcal{D}_2 \cup \bigcup_{i=1}^{f''} \mathcal{D}_3} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})|}_{=(6f+6f'+4f'')J} + \underbrace{\sum_{\mathbf{c} \in \bigcup_{i=1}^{f'''} \mathcal{D}_4} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})|}_{\geq (2J-1)2f''' + 3f'''} + \underbrace{\sum_{\mathbf{c} \in \mathcal{D}_5} |\text{tok}_\uparrow[\mathbf{m}](\mathbf{c})|}_{>0} \quad (66a) \\ &> (6f + 6f' + 4f'')J + (4J + 1)f''' \quad \text{By construction } f''' = 5I \quad (66b) \\ &= (6f + 6f' + 4f'' + 4f''')J + 5I \quad (66c) \end{aligned}$$

which contradicts the proof's statement.  $\square$

**LemmaProofStep 5.** (Step ⑤). *Any instance of the tokenisation problem with a solution which compresses the text by at least  $(6f + 6f' + 6f'' + 6f''')J + 2I + \gamma$  symbols must be produced by a max-2-SAT problem with at least  $\gamma$  satisfied clauses, i.e.,:*

$$\left( \text{toklen}(\mathcal{D}, \mathcal{S}) \leq \underbrace{(6f + 6f' + 4f'' + 4f''')J + 3I - \gamma}_{\sum_{\mathbf{c} \in \mathcal{D}} |\mathbf{c}| - (6f + 6f' + 6f'' + 6f''')J - 2I - \gamma} \right) \implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma)$$

*Proof.* Finally, we now know any solution with this compression must have—for any variable  $X_j$ —either two merges in  $\mathbf{m}_j^T$  or in  $\mathbf{m}_j^F$  (and never both). We can thus create a bijection  $\text{Conv}_{\mathbf{m} \rightarrow \chi}$  between the set of possible merge sequences respecting this condition, and the set of T/F assignments to SAT variables  $\chi$ :

$$\text{Conv}_{\mathbf{m} \rightarrow \chi}(\mathbf{m}) = \left\{ \begin{array}{ll} \text{T} & \text{if } |\mathbf{m} \cap \mathbf{m}_j^T| = 2 \\ \text{F} & \text{if } |\mathbf{m} \cap \mathbf{m}_j^F| = 2 \end{array} \right\}_{j=1}^J \quad (67)$$

Further, note that merge sequences of this form (as shown in Eq. (41)) lead to exactly  $(6f + 6f' + 6f'' + 6f''')J$  symbols being compressed in datasets  $\mathcal{D}_1$  to  $\mathcal{D}_4$ . To achieve the target compression, a solution

must thus compress  $\mathcal{D}_5$  by at least  $2I + \gamma$  symbols. Now note that for any string in  $\mathcal{D}_5$ , e.g.,  $\odot x_j^T \odot x_{j'}^F \odot$ , we have three compression options:  $\odot x_j^T \odot$  and  $x_{j'}^F \odot$  will be compressed, saving 3 symbols;  $\odot x_j^T$  and  $\odot x_{j'}^F \odot$  will be compressed, also saving 3 symbols; or only  $\odot x_j^T$  and  $x_{j'}^F \odot$  will be compressed saving only 2 symbols. More specifically,  $\odot x_j^T \odot$  will be compressed to a single symbol if merge  $\langle \odot, x_j^T \odot \rangle$  exists; similarly,  $\odot x_{j'}^F \odot$  will be compressed to a single symbol if merge  $\langle \odot x_{j'}^F, \odot \rangle$  exists. They cannot both be compressed, however, as there is only one symbol  $\odot$  between the literals. We thus get a reduction of 3 symbols for each of these strings if at least one of its literals has an associated merge in  $\mathbf{m}$ . Note thus that whenever a string  $\odot x_j^T \odot x_{j'}^F \odot$  is compressed by 3 symbols using merges  $\mathbf{m}$ , the max-2-SAT disjunction  $X_j \vee \neg X_{j'}$  will also be satisfied by assignment  $\chi = \text{Conv}_{\mathbf{m} \rightarrow \chi}(\mathbf{m})$ ; similarly, whenever this string is only compressed by two symbols, the max-2-SAT disjunction will not be satisfied. A similar logic applies to all potential strings in  $\mathcal{D}_5$ :  $\odot x_j^T \odot x_{j'}^F \odot$ ,  $\odot x_{j'}^T \odot x_j^F \odot$ ,  $\otimes x_j^F \otimes x_{j'}^F \otimes$ , and  $\otimes x_j^T \otimes x_{j'}^T \otimes$ . As our condition assumes a compression of at least  $2I + \gamma$  symbols, we know that we have at least  $\gamma$  strings for which a literal has an associated merge. Now, let  $\mathbf{m}^* \in \mathcal{M}^*$  be a valid solution to  $\text{Tok}_{\uparrow}(\text{R2}(\mathcal{X}, \mathcal{C}, \gamma))$  and  $\chi^* = \text{Conv}_{\mathbf{m} \rightarrow \chi}(\mathbf{m}^*)$  the equivalent max-2-SAT assignment. We can thus write:

$$2I + \gamma \leq \sum_{\mathbf{c} \in \mathcal{D}_5} |\mathbf{c}| - |\text{tok}_{\uparrow}[\mathbf{m}^*](\mathbf{c})| \quad (68a)$$

$$= 2I + \sum_{\mathbf{c} \in \mathcal{D}_5} \mathbb{1} \left\{ \begin{array}{l} \left( \begin{array}{l} (x_j^T \in \mathbf{c}) \text{ and } (|\mathbf{m}_j^T \cap \mathbf{m}^*| = 2) \\ \text{or} \\ (x_{j'}^F \in \mathbf{c}) \text{ and } (|\mathbf{m}_{j'}^F \cap \mathbf{m}^*| = 2) \end{array} \right) \\ \text{or} \\ \left( \begin{array}{l} (x_{j'}^T \in \mathbf{c}) \text{ and } (|\mathbf{m}_{j'}^T \cap \mathbf{m}^*| = 2) \\ \text{or} \\ (x_j^F \in \mathbf{c}) \text{ and } (|\mathbf{m}_j^F \cap \mathbf{m}^*| = 2) \end{array} \right) \end{array} \right\} \quad (68b)$$

$$= 2I + \sum_{i=1}^I \mathbb{1}_{\chi^*} \{L_i^1 \vee L_i^2\} \quad (68c)$$

$$\implies \text{M2S}(\mathcal{X}, \mathcal{C}, \gamma) \quad (68d)$$

where  $\mathbb{1}_{\chi^*} \{L_i^1 \vee L_i^2\}$  evaluates  $L_i^1 \vee L_i^2$  using assignments  $\chi^*$ . We thus know that, if a satisfying tokenisation solution exists, then the associated max-2-SAT problem will also be satisfiable. This concludes the proof.  $\square$