

Combinatorial Algorithms for Subsequence Matching: A Survey

Maria Kosche, Tore Koß, Florin Manea, Stefan Siemer

Göttingen University, Germany

{maria.kosche,tore.koss,florin.manea,stefan.siemer}@cs.uni-goettingen.de

In this paper we provide an overview of a series of recent results regarding algorithms for searching for subsequences in words or for the analysis of the sets of subsequences occurring in a word.

1 Introduction

For a string $w = w_1w_2\dots w_n$, where each w_i is a single symbol from some alphabet Σ , any string $v = w_{i_1}w_{i_2}\dots w_{i_k}$ with $k \leq n$ and $1 \leq i_1 < i_2 \leq \dots < i_k \leq n$ is called a *subsequence* (also called sometimes *scattered factor* or *subword*) of w (denoted by $v \preceq w$).

The concept of subsequences is employed in many different areas of computer science. Subsequences appear in areas of theoretical computer science such as, for instance, in formal languages and logics (e. g., where they are used in relation to piecewise testable languages [54, 55, 35, 36, 37], or to define the subword order and downward closures [31, 41, 40, 60]) or in combinatorics on words, where they are used to define the notions of binomial equivalence and binomial complexity, or to introduce the notion of subword histories, [49, 23, 43, 42, 51, 47, 50]; however, subsequences are also used in more applied settings, e. g., for modelling concurrency [48, 52, 13], or in database theory (especially *event stream processing* [5, 30, 61]). Moreover, many classical algorithmic problems are based on subsequences, e. g., longest common subsequence [6] or shortest common supersequence [46], and, in particular, such problems have recently regained interest in the context of fine-grained complexity (see [11, 12, 1, 2]).

There are two main types of algorithmic problems for subsequences investigated in the literature. Firstly, the class of *matching* problems, where one has to decide whether a string v is a subsequence of a string w , i. e., whether $v \preceq w$ (the name matching comes from the fact that the string v can be seen as a pattern that has to be identified, or matched, within the string w). Secondly, the class of *analysis problems*, which are concerned with the algorithmic analysis of the sets $\text{Subseq}_k(w)$ of all length- k subsequences of a given string w . Some more concrete examples of analysis problems are the following: for given string $w \in \Sigma^+$ and integer $k \in \mathbb{N}$, we want to decide whether $\text{Subseq}_k(w) = \Sigma^k$ (the *universality* problem), or, for an additional string v , we want to decide whether $\text{Subseq}_k(v) = \text{Subseq}_k(w)$ (the *equivalence* problem). For classical subsequences (as defined above), the matching problem is trivial, while the analysis problems are well-investigated and relatively well-understood. In particular, the equivalence problem was introduced by Imre Simon in his PhD thesis [54], and was intensely studied in the combinatorial pattern matching community (see [32, 28, 56, 57, 16, 20] and the references therein), before being optimally solved in 2021 [29]. The work on these problems was extended to classes of constrained subsequences, for which very different results were obtained, by fundamentally different methods [18, 39].

In this work we overview a series of algorithmic, combinatorial, and complexity theoretic results concerning subsequences. For their original presentation, please see [7, 17, 20, 29, 38, 39, 18] and the references therein.

Going a bit more into detail, the results surveyed here cover different settings and algorithmic frameworks regarding subsequences, ranging from the classical (and well motivated) case of unrestricted subsequences to some novel ones, where constrained subsequences appear.

For the case of classical subsequences, the accent of our presentation is put on analysis problems (as matching is trivial). We survey results related to the equivalence and universality problems originally presented in [7, 17, 20, 29, 38].

For the case of constrained subsequences, we feel that some more discussion about their motivation and origin would be in order. So, the notion of constrained subsequences is rooted in the following main idea: it seems unrealistic for particular scenarios to consider occurrences of v in w where the positions of w that are matching, respectively, the first and last symbol of v (or, similarly, the positions of w matching consecutive symbols of v) are very far away from each other. It seems indeed questionable, for instance, whether considering an alignment of DNA-sequences v and w where the nucleotides of v are spread over a factor of w which is several times longer than v itself (or, alternatively, where the nucleotides of v occur in w with arbitrarily long gaps between them) is still meaningful. Similarly, but in a totally different context, when observing a computation, which is represented by a string, one might be more interested in its recent history (and the sequences of events occurring there), rather than analysing the entire computation. Or, in the same setting, one might be interested in sequences of events which occur in a computation, such that the computation executed between two events in such a sequence is constrained by some precise rules. In a similar situation, occurring this time in the context of complex event processing [5, 30, 61], it might be desirable to describe the situation that between the events of a job A , only events associated to a job B appear (e. g., due to unknown side-effects this leads to a failure of job A). In this case, we are interested in occurrences of a string v as a subsequence of a string w such that the gaps between the positions of w which correspond to the symbols of v , only contain symbols from a certain subset of the alphabet (i. e., the events associated to job B). Moreover, the fact that in many practical scenarios (including those mentioned above) one has to process streams, which, at any moment, can only be partly accessed by our algorithms, enforces even more the idea that the case where one is interested in subsequences occurring arbitrarily in a given stream (or long string) is less realistic and less useful than the case where one is interested in the subsequences occurring in bounded ranges of the respective stream/string (which can be entirely accessed and processed at any moment by our algorithms). So, wrapping this up, in practice, it makes sense to reason both about the length and the actual content of gaps induced by an occurrence of v in w , as well as about the length of the factor (or range) of w in which such an occurrence is contained.

To this end, in our overview, for the case of constrained subsequences, we will overview a series of algorithmic results for problems related to the case of subsequences in which, given strings w and v , constraints on either the factors of w in which v may occur as subsequence (called bounded range constraints, see [39]) are imposed, or constraints on the factors of w occurring between two consecutive letters of an occurrence of v in w (called gap constraints, see [18]) are imposed. In this setting, we overview results regarding the matching problem (which is no longer trivial) as well as results on analysis problems. For the original presentation of these results, as well as for a more detailed overview of the motivations for these particular classes of constrained subsequences and related work, we refer to [39, 18].

This paper is structured as follows. We first give a series of general definitions and preliminaries related to subsequences in Section 2. In Section 3, the matching problem is covered and results are presented for all three cases, the classical subsequences setting, for subsequences occurring in bounded ranges, and for subsequences with gap constraints. Section 4 covers the analysis problems in respective subsections, namely the universality problem in Section 4.1, the analysis of absent subsequences in

Section 4.2, and the equivalence problem in Section 4.3. We conclude with a section covering a series of related problems as well as directions for future work.

2 Basic Definitions

Let \mathbb{N} be the set of natural numbers, including 0. For $m, n \in \mathbb{N}$, we define the range (or interval) of natural numbers lower bounded by m and upper bounded by n as $[m : n] = \{m, m + 1, \dots, n\}$. An alphabet Σ is a non-empty finite set of symbols (called letters). A *string (or word)* is a finite sequence of letters from Σ , thus an element of the free monoid Σ^* . Let $\Sigma^+ = \Sigma^* \setminus \{\varepsilon\}$, where ε is the empty string. The *length* of a string $w \in \Sigma^*$ is denoted by $|w|$. The i^{th} letter of $w \in \Sigma^*$ is denoted by $w[i]$, for $i \in [1 : |w|]$. For $a \in \Sigma$, let $|w|_a = |\{i \in [1 : |w|] \mid w[i] = a\}|$; let $\text{alph}(w) = \{x \in \Sigma \mid |w|_x > 0\}$ be the smallest subset $S \subseteq \Sigma$ such that $w \in S^*$. For $m, n \in \mathbb{N}$, with $m \leq n$, we define the range (or factor) of w between positions m and n as $w[m : n] = w[m]w[m + 1] \dots w[n]$. If $m > n$, then $w[m : n]$ is the empty word. Also, by convention, if $m < 1$, then $w[m : n] = w[1 : n]$, and if $n > |w|$, then $w[m : n] = w[m : |w|]$. A factor $u = w[m : n]$ of w is called a *prefix* (respectively, *suffix*) of w if $m = 1$ (respectively, $n = |w|$).

The powers of a word w are defined as: $w^0 = \varepsilon$ and $w^{k+1} = ww^k$, for $k \geq 0$. Define w^ω as the right infinite word which has w^n as prefix for all $n \geq 0$. The positive integer $p \leq |w|$ is a period of a word w if w is a prefix of $w[1 : p]^\omega$.

We now recall the main notion of this paper, namely the notion of subsequence.

Definition 1 (Subsequence). *A word v is a subsequence of length k of w (denoted $v \leq w$), where $|w| = n$, if there exist positions $1 \leq i_1 < i_2 < \dots < i_k \leq n$, such that $v = w[i_1]w[i_2] \dots w[i_k]$. The set of all subsequences of w is denoted by $\text{Subseq}(w)$.*

In the following, we will also discuss some other concepts regarding (classical) subsequences, namely subsequences with gap constraints, (partitioned into length constraints, regular constraints, or combined length and regular constraints) for the factors of w between two consecutive letters of v , and subsequences within bounded ranges (where we consider, similar to a sliding window scenario, an integer p an integer p as an upper bound on the range of a word in which a subsequence may occur).

Firstly, we introduce subsequences with gap constraints, or gapped subsequences. This presentation is based on [18]. We begin by defining the notion of gap constraints. We recall that for a string w , an embedding is a function $e : [k] \rightarrow [|w|]$ such that $i < j$ implies $e(i) < e(j)$ for all $i, j \in [k]$, and it induces the subsequence $\text{subseq}_e(w) = w[e(1)]w[e(2)] \dots w[e(k)]$ of w . For every $j \in [k - 1]$, the j^{th} gap of w induced by e is the string $\text{gap}_e(w, j) = w[e(j) + 1 .. e(j + 1) - 1]$. We say that e is the embedding of $\text{subseq}_e(w)$ in w .

Definition 2 (Gap constraints). *An ℓ -tuple of gap constraints is a tuple $gc = (C_1, C_2, \dots, C_\ell)$ with $C_i \subseteq \Sigma^*$ for every $i \in [\ell]$. For convenience, we set $gc[i] = C_i$ for every $i \in [\ell]$. We say that an embedding e satisfies a $(k - 1)$ -tuple of gap constraints gc with respect to a string w if it has the form $e : [k] \rightarrow [|w|]$, and, for every $i \in [k - 1]$, $\text{gap}_e(w, i) \in C_i$. For a $(k - 1)$ -tuple gc of gap constraints, the set $\text{Subseq}(gc, w)$ contains all subsequences of w induced by embeddings that satisfy gc , i. e., $\text{Subseq}(gc, w) = \{\text{subseq}_e(w) \mid e \text{ is an embedding that satisfies } gc \text{ w. r. t. } w\}$. The elements of $\text{Subseq}(gc, w)$ are also called the gc -subsequences of w .*

For a $(|u| - 1)$ -tuple gc of gap constraints, we write $u \leq_{gc} v$ to denote that $u \leq_e v$ for some embedding $e : [|u|] \rightarrow [|v|]$ that satisfies gc with respect to v , i. e., $u \leq_{gc} v$ means that u is a gc -subsequence of v .

We generally distinguish the following types of gap constraints:

- *regular constraints* if $C_i \in \text{REG}$ for every $i \in [k-1]$. For every $i \in [k-1]$, the regular constraint C_i is represented by a deterministic finite automaton (for short, DFA) A_i accepting it.
- *length constraints* if, for every $i \in [k-1]$, there are $L^-(i), L^+(i) \in \mathbb{N} \cup \{0, +\infty\}$ with $L^-(i) \leq L^+(i)$, such that $C_i = \{v \in \Sigma^* \mid L^-(i) \leq |v| \leq L^+(i)\}$. Length constraints are succinctly represented by pairs of numbers $(L^-(i), L^+(i))$, $i \in [k-1]$, in binary encoding.
- *reg-len constraints* if, for every $i \in [k-1]$, C_i is the conjunction of a regular constraint C'_i and a length constraint $(L^-(i), L^+(i))$, i. e., $C_i = C'_i \cap \{v \in \Sigma^* \mid L^-(i) \leq |v| \leq L^+(i)\}$. Such constraints are represented by $((L^-(i), L^+(i)), A'_i)$, where A'_i is a DFA accepting C'_i .

We move now further, and introduce the concept of p -subsequence, or subsequences occurring within bounded ranges. For this presentation, we follow [39].

Definition 3 (Bounded range constraints). 1. Let $v, w \in \Sigma^*$ with $|v| = m, |w| = n$. The string v is called a p -subsequence of w (denoted $v \leq_p w$) if there exists an embedding $e: [m] \rightarrow [n]$ such that $v = \text{subseq}_e(w)$ and $|w[e(1) : e(m)]| \leq p$, or equivalently $e(m) - e(1) \leq p - 1$. 2. For $p \in \mathbb{N}$ and $w \in \Sigma^*$, we denote the set of all p -subsequences of w by $p\text{-Subseq}(w) = \{v \in \Sigma^* \mid v \leq_p w\}$. Furthermore, for $k \in \mathbb{N}$, we denote the set of all p -subsequences of length k of w by $p\text{-Subseq}_k(w)$.

Once these main concepts introduced, we can now discuss several preliminaries which are necessary for understanding the surveyed results.

Computational Model. In general, the problems surveyed here are of algorithmic nature. The computational model used to describe the algorithms is the standard unit-cost RAM with logarithmic word size: for an input of size N , each memory word can hold $\log N$ bits. Arithmetic and bitwise operations with numbers in $[1 : N]$ are, thus, assumed to take $O(1)$ time. In all the problems, it is assumed that we are given a word w or two words w and u , with $|w| = n$ and $|v| = m$ (so the size of the input is $N = n + m$), over an alphabet $\Sigma = \{1, 2, \dots, \sigma\}$, with $2 \leq |\Sigma| = \sigma \leq n + m$. That is, the processed words are assumed to be sequences of integers (called letters or symbols), each fitting in $O(1)$ memory words. This is a common assumption in string algorithms: the input alphabet is said to be *an integer alphabet*. For more details see, e. g., [15].

The algorithmic results (upper bounds) that are surveyed here are complemented by a series of lower bounds. In those cases, the results hold already for the case of constant alphabets. That is, they hold already when the input of the problem is restricted to words over an alphabet $\Sigma = \{1, 2, \dots, \sigma\}$, with $\sigma \in O(1)$.

Complexity Hypotheses. For the series of conditional lower bounds for the time of complexity of the considered problems, we now recall some standard computational problems and complexity hypotheses regarding them, respectively, on which the proofs of lower bounds are based.

The *Satisfiability problem for formulas in conjunctive normal form*, CNF-SAT, gets as input a Boolean formula F in conjunctive normal form as a set of clauses $F = \{c_1, c_2, \dots, c_m\}$ over a set of variables $V = \{v_1, v_2, \dots, v_n\}$, i. e., for every $i \in [m]$, we have $c_i \subseteq \{v_1, \neg v_1, \dots, v_n, \neg v_n\}$. The question is whether F is satisfiable. By k -CNF-SAT, we denote the variant where $|c_i| \leq k$ for all $i \in [m]$.

The *Orthogonal Vectors problem* (OV for short) gets as input two sets A, B each containing n Boolean-vectors of dimension d , where $d \in \omega(\log n)$. The question is whether there exist two vectors $\vec{a} \in A$ and $\vec{b} \in B$ which are orthogonal, i. e., $\vec{a}[i] \cdot \vec{b}[i] = 0$ for every $i \in [d]$.

We shall use the following algorithmic hypotheses based on CNF-SAT and OV that are common for obtaining conditional lower bounds in fine-grained complexity. In the following, poly is any fixed polynomial function:

– *Exponential Time Hypothesis* (ETH) [34, 44]: 3-CNF-SAT cannot be solved in time $2^{o(n)} \text{poly}(n + m)$.

– *Strong Exponential Time Hypothesis* (SETH) [33, 59]: For every $\varepsilon > 0$ there exists k such that k -CNF-SAT cannot be decided in $O(2^{n(1-\varepsilon)} \text{poly}(n))$ time.

The following result, which essentially formulates the Orthogonal Vectors Hypothesis (OVH), can be shown (see [9, 10, 59]).

Lemma 4. *OV cannot be solved in $O(n^{2-\varepsilon} \text{poly}(d))$ time for any $\varepsilon > 0$, unless SETH fails.*

3 Matching problems

The matching problem MATCH for subsequences is to decide, given two words $u, w \in \Sigma^*$ with $|u| = m$ and $|w| = n$, whether u is a subsequence of w . In the general case, that is without further restrictions like gap constraints or bounded ranges, it is quite easy to answer: if we go left-to-right through w and greedily search for the letters $u[1]$ to $u[m]$, we answer positively if and only if we find all letters from u in w . This greedy approach obviously is correct and works in linear time $O(n)$.

Considering bounded range constraints. If we consider subsequences occurring within bounded ranges, the problem changes as follows: for u, w as above and $p \in \mathbb{N}$ with $p \geq m$, we need to decide whether u is a p -subsequence of w . Simply using the greedy approach above for each range still works, but is not optimal anymore (it has $O(np)$ complexity). However, by reading the word w left to right and maintaining an array which saves for every $i \in [1 : m]$ the length of the shortest suffix of the current range $w[t - p + 1 : t]$ containing $u[1 : i]$ (if there is any) and updating the array when we increment t (i. e., read the a new letter of the word w), we can reduce the time complexity to $O(mn)$ (see [39] for a more detailed description of the algorithm).

Theorem 5. *MATCH in bounded ranges can be solved in $O(mn)$ time.*

The algorithm presented in [39] can be, in fact, seen as an algorithm in the sliding window model with window of fixed size p (see [26, 25, 27]). More precisely, it scans the stream w left to right and, when the t^{th} letter of the stream is scanned, it reports whether the window $w[t - p + 1 : t]$ contains u as a subsequence. In other words, it reports whether the string $w[t - p + 1 : t]$ is in the regular language $L_u = \{v \mid u \leq v\}$. The problem of checking whether the factors of a stream scanned by a sliding window are in a regular language was heavily investigated, see [24] and the references therein. In particular, from the results of [26] it follows that, for a constant u (i. e., u is not part of the input), the problem of checking whether the factors of a stream scanned by a sliding window are in the language L_u cannot be solved using $o(\log p)$ bits when the window size is not changing and equals p . We note that the algorithm of [39] is optimal from this point of view: if u is constant and, thus, $m \in O(1)$, it uses $O(\log p)$ bits to store the maintained data structures.

Moreover, the algorithm presented in [39] is optimal also from the time complexity point of view, unless OVH fails.

Theorem 6. *MATCH in bounded ranges cannot be solved in time $O(n^h m^g)$, where $h + g = 2 - \varepsilon$ with $\varepsilon > 0$, conditional to OVH.*

Considering gap constraints. For subsequences with gap constraints, the matching problem is to decide, for given strings u, w , and gap constraints gc with $|gc| = |u| - 1$, whether u is a gc -subsequence of w (i. e., whether $u \in \text{Subseq}(gc, w)$).

The results from [18] give us the following upper bound in the case of reg-len constraints where gc denotes the given gap constraints, $\text{states}(gc)$ denotes the total number of states of the DFAs that represent

the regular constraints, $\text{size}(gc)$ is the total size of the automata defining these constraints and $\text{nz}(gc)$ is the number of gaps which are not equal to $\{\varepsilon\}$.

Theorem 7. *MATCH with reg-len constraints can be solved in $O(|w| \text{states}(gc) + \text{size}(gc))$ time.*

The proof given in [18] is based on a dynamic programming approach, which is implemented in the respective time complexity with the help of some relatively involved data structures, and implies that when considering only length constraints or only regular constraints (and not combined reg-len constraints), the following rectangular upper bounds hold.

Corollary 8. (1). *MATCH with length constraints can be solved in $O(|w| \text{nz}(gc))$ time.*

(2). *MATCH with regular constraints can be solved in $O(|w| \text{states}(gc) + \text{size}(gc))$ time.*

When gc only consists of constraints that are $\{\varepsilon\}$ or Σ^* , respectively, the case of string matching or, respectively, subsequence matching is modelled, which can be solved in linear time. As far as length constraints are concerned, it seems that non-trivial upper bounds lead to an increase in the difficulty of the MATCH problem; a particularly efficient approach for subsequences with general length constraints is given in [8] but in the worst case it still has rectangular complexity. However, even when non-trivial length upper bounds are used, there are still some simpler particular cases. For instance, when working with strings with *don't cares* (or partial words), where each gap has a fixed length (i. e., the lower and upper bounds are the same), Match can be solved in time $O(|w| \log |p|)$ [14].

A reduction from the OV problem is given in [18], which shows that MATCH for non-trivial length or regular constraints is more difficult than MATCH for the classical subsequences scenario. The following conditional lower bounds for subsequences with length and/or regular gap constraints are obtained.

Theorem 9. *MATCH with length constraints cannot be solved in $O(|w|^h \text{nz}(gc)^\varepsilon)$ time with $h + g = 2 - \varepsilon$ for some $\varepsilon > 0$, unless OVH fails. This holds even if $|\Sigma| = 4$ and all length constraints are $(0, \ell)$ with $\ell \leq 6$.*

Corollary 10. *MATCH with regular constraints cannot be solved in $O(|w|^h \text{states}(gc_p)^\varepsilon)$ time with $h + g = 2 - \varepsilon$ for some $\varepsilon > 0$, unless OVH fails. This holds even if $|\Sigma| = 4$ and all regular constraints are expressed by constant size DFAs.*

4 Analysis problems

When considering algorithmic analysis problems related to subsequences, typical research questions are concerned with structural properties of the set of all (constrained) subsequences occurring in a word, as well as with finding minimal (w. r. t. length or w. r. t. the subsequence relation) missing subsequences of a word.

4.1 Universality

Generally speaking, the universality problem UNI is to decide, for given integer k and string $w \in \Sigma^*$ with $|w| = n$, whether the set of subsequences of length k of w equals the set Σ^k . For convenience, we will also consider in the following the complement problem, i. e., *non-universality problem* (NUNI).

Definition 11. *A word $w \in \Sigma^*$ is called k -subsequence universal (w. r. t. Σ , for short k -universal), for $k \in \mathbb{N}$, if $\text{Subseq}_k(w) = \Sigma^k$. We abbreviate 1-universal by universal. The universality-index $\iota(w)$ of $w \in \Sigma^*$ is the largest k such that w is k -universal.*

If $\iota(w) = k$ then w is ℓ -universal for all $\ell \leq k$. Notice that k -universality is always w. r. t. a given alphabet Σ : the word $abcba$ is universal for $\Sigma = \{a, b, c\}$ but it is not universal for $\Sigma \cup \{d\}$.

The notion of k -universality coincides to that of k -richness introduced in [36, 37]. We use the name *k-universality* rather than *k-richness*, as richness of words is also used with other meanings, see, e. g., [19, 45]. We recall the arch factorisation, introduced by Hebrard [32].

Definition 12 ([32]). *For $w \in \Sigma^*$ the arch factorisation of w is $w = \text{ar}_w(1) \cdots \text{ar}_w(k)r(w)$ for some $k \in \mathbb{N}_0$ where $\text{ar}_w(i)$ is universal, the last letter of $\text{ar}_w(i)$, namely $\text{ar}_w(i)[|\text{ar}_w(i)|]$, does not occur in $\text{ar}_w(i)[1 : |\text{ar}_w(i)| - 1]$ for all $i \in [1 : k]$, and $\text{alph}(r(w)) \subset \Sigma$. The words $\text{ar}_w(i)$ are called arches of w , $r(w)$ is called the rest.*

If the arch factorisation of w contains $k \in \mathbb{N}_0$ arches, then $\iota(w) = k$. The arch factorization of a word w can be computed in linear time and, as such, we could check in linear time if a given word w is k -universal (see, e. g., [7]). The following immediate theorem based on the work of Simon [55] completely characterises the set of k -subsequence universal words, based on Hebrard's arch factorisation.

Theorem 13. *The word $w \in \Sigma^*$ is k -universal if and only if there exist the words v_i , with $i \in [1 : k]$, such that $v_1 \cdots v_k = w$ and $\text{alph}(v_i) = \Sigma$ for all $i \in [1 : k]$.*

This property gives us some insight in the combinatorial structure of the subsequences occurring inside words. For instance, it can be used to directly compare two words w. r. t. their universality index, or serve as a starting point for the analysis of the set of missing subsequences of words. See, for instance, [7, 17, 22, 21, 18, 39].

The Edit Distance to k -Subsequence Universality. As a natural extension to the universality property from above, as one can do for almost every string property, we can ask how far is a word from fulfilling that property. That is, we can ask for the distance from a give string to the set of strings which fulfill that property, with respect to some string metric. More precisely in this section, we discuss how to compute the minimal number of edits we need to apply to a word w , with $|w| = n$, $\text{alph}(w) = \Sigma$, with universality index $\iota(w)$, so that it is transformed into a word with universality index k , w. r. t. the same alphabet Σ . The edits considered are insertion, deletion, substitution, and the number we want to compute can be seen as the *edit distance* between w and the set of k -universal words over Σ .

The first thing that we can see is that, if we want to obtain a k -universal word with $k > \iota(w)$, then it is enough to consider only insertions. Indeed, deleting a letter of a word can only restrict the set of subsequences of the respective word, while in this case we are interested in enriching it. Substituting a letter might make sense, but it can be simulated by an insertion: assume one wants to substitute the letter a on position i of a word w by a b . It is enough to insert a b next to position i , and the set of subsequences of w is enriched with all the words that could have appeared as subsequences of the word where a was actually replaced by b . We might, in the end, have some extra words in the set of subsequences, which would have been eliminated through the substitution, but it does not affect our goal of reaching k -universality.

If we want to obtain a word with universality index k , for $k < \iota(w)$, then it is enough to consider only deletions. Assume that we have a sequence of edits that transforms the word w into a word w' with universality index k . Now, remove all the insertions of letters from that sequence. The word w'' we obtain by executing this new sequence of operations clearly fulfils $\iota(w'') \leq \iota(w')$. Further, in the new sequence, replace all substitutions with deletions. We obtain a word w''' with a set of subsequences strictly included in the one of w'' , so with $\iota(w''') \leq \iota(w'')$. As each deletion changes the universality index by at most 1, it is clear that (a prefix of) this new sequence of deletions witnesses a shorter sequence of edits which transforms w into a word of universality index k .

So, to increase the universality index of a word it is enough to use insertions and to decrease the universality index of a word it is enough to use deletions. Nevertheless, one might be interested in what happens if we only use substitutions. In this way, we can both decrease and increase the universality index of a word. Moreover, one can see the minimal number of substitutions needed to transform w into a k -universal word as the Hamming distance between w and the set of k -universal words. In the following we list all of the resulting theorems individually. All of the results are achieved by a dynamic programming approach combined with a sophisticated analysis of the combinatorial properties of k -universal words and some new specialized data structures. See the full proofs and algorithms in [17].

Theorem 14. *Let w be a word, with $|w| = n$, $\text{alph}(w) = \Sigma$, and $\Sigma = \{1, 2, \dots, \sigma\}$. Let $k \geq \iota(w)$ be an integer. We can compute the minimal number of insertions needed to apply to w in order to obtain a k -universal word (w. r. t. Σ) in $O(nk)$ time if $k \leq n$ and $O(T(n, \sigma, k))$ time otherwise, where $T(n, \sigma, k)$ is the time needed to compute the number $k\sigma - n$.*

Theorem 15. *Let w be a word, with $|w| = n$, $\text{alph}(w) = \Sigma$, and $\Sigma = \{1, 2, \dots, \sigma\}$. Let k be an integer with $k \leq \iota(w) \leq \lfloor \frac{n}{\sigma} \rfloor$. We can compute in $O(nk)$ time the minimal number of deletions needed to obtain a word of universality index k (w. r. t. Σ) from w .*

Theorem 16. *Let w be a word, with $|w| = n$, $\text{alph}(w) = \Sigma$, and $\Sigma = \{1, 2, \dots, \sigma\}$. Let k be an integer $0 \leq k \leq \lfloor \frac{n}{\sigma} \rfloor$. We can compute the minimal number of substitutions needed to apply to w in order to obtain a k -universal word (w. r. t. Σ) in $O(nk)$ time.*

Universality and bounded range constraints. After looking into the (unrestricted) subsequence universality of a word and in particular the edit distance between a word and the set of k -universal words, we will now discuss the case of subsequences occurring in bounded ranges or words. In this case, the universality problem UNI asks to decide for given word w , alphabet Σ , and integers k and p , with $|w| = n$ and $k \leq p \leq n$, whether $p\text{-Subseq}_k(w) \neq \Sigma^k$.

Surprisingly, in this case, we get an intractability result, complemented by a fine-grained lower bound (see [39]). For convenience, this computational hardness result is stated for NUNI.

Theorem 17. *NUNI for bounded ranges is NP-hard and cannot be solved in subexponential time $2^{o(k)}$ poly(k, n) unless ETH fails.*

The proof for this result involves the reduction from the related problem NUNI for partial words, which asks to decide, for given list of partial words $S = \{w_1, \dots, w_k\}$ over $\{0, 1\}$, where every partial word has same length L , whether there exists a word $v \in \{0, 1\}^L$ such that v is not compatible with any of the partial words in S (while in this context, two partial words u and v of the same length are *compatible* if, for all $i \in [|u|]$, we have that either $u[i] = v[i]$ or at least one of $u[i]$ or $v[i]$ is undefined).

Universality and gap constraints. For the case of subsequences with gap constraints, [18] presents a series of results starting with a brute force upper bound which can be derived from the results for MATCH for subsequences with gap constraints.

Theorem 18. (1) *The problem UNI for subsequences with length (or reg-len) constraints can be solved in time $O(|\Sigma|^k \text{nz}(gc)\ell)$ (respectively, $O(|\Sigma|^k \text{states}(gc)\ell)$), where $\ell = \max\{|w|, |w'|\}$.*
 (2) *For the case of a fixed alphabet Σ (i. e., $|\Sigma| \in O(1)$), the problem UNI_Σ with length (or reg-len) constraints can be solved in time $2^{O(k)} \text{nz}(gc)\ell$ (respectively, $2^{O(k)} \text{states}(gc)\ell$), where $\ell = \max\{|w|, |w'|\}$.*

At the same time, the following lower bound (again, given for NUNI) shows that it is unlikely for significantly faster algorithms to exist.

Theorem 19. *For every fixed alphabet Σ with $|\Sigma| \geq 3$, NUNI_Σ with length constraints is NP-complete, even if all length constraints are (1,5). Moreover,*

- *it cannot be solved in subexponential time $2^{o(k)} \text{poly}(|w|, k)$ (unless ETH fails),*
- *it cannot be solved in time $O(2^{k(1-\varepsilon)} \text{poly}(|w|, k))$ (unless SETH fails).*

For a fixed alphabet Σ with $|\Sigma| = 2$, NUNI_Σ with length constraints is NP-complete even if each length constraint is (0,0) or (3,9) (meaning that each gap is either empty or has length between 3 and 9).

4.2 Absent Subsequences

In the previous sections we surveyed a series of results related to deciding whether a string contains as subsequences all strings of length up to k . Now, we focus on understanding the strings which do not occur as subsequences of a given input string.

So, in this subsection we summarize a series of algorithmic and complexity results related to decision problems concerning shortest and, respectively, minimal absent subsequences. Once more, we begin with the classical case, and then discuss the case of subsequences occurring within bounded ranges.

We begin with several definitions.

Definition 20 (– Absent subsequences). *A word v is an absent subsequence of w if v is not a subsequence of w . An absent subsequence v of w is a minimal absent subsequence (for short, MAS) of w if every proper subsequence of v is a subsequence of w . We will denote the set of all MAS of w by $\text{MAS}(w)$. An absent subsequence v of w is a shortest absent subsequence (for short, SAS) of w if $|v| \leq |v'|$ for any other absent subsequence v' of w . We will denote the set of all SAS of w by $\text{SAS}(w)$.*

Absent subsequences in words. Note that, in general, any shortest absent subsequence of a word w has length $\iota(w) + 1$, where $\iota(w)$ is the universality index of w . This already establishes a connection to the results presented in previous section. Moreover, we notice that we can easily find at least one SAS (and therefore an MAS) of w from its arch factorisation: for $1 \leq i \leq \iota(w)$ let $u_i = \text{ar}_w(i)[|\text{ar}_w(i)|]$ be the last letter of the i^{th} arch of w and $u_{\iota(w)+1}$ be any letter not occurring in $r(w)$, then $u = u_1 \cdots u_{\iota(w)+1}$ is an SAS of w . By refining this approach we can, for given $w \in \Sigma^*$, build in linear time a data structure allowing us to identify a succinct representation of a SAS of any factor $w[i : j]$ of w in constant time (and effectively output this SAS in time proportional to its length).

Theorem 21 ([38]). *For a word w of length n we can construct in $O(n)$ time data structures allowing us to answer in $O(1)$ time queries $\text{sasRange}(i, j)$: “return a representation of an SAS of $w[i : j]$ ”.*

Given $u, w \in \Sigma^*$ the problem to check whether u is an SAS or, respectively, an MAS of w is decidable in linear time (see [38]).

Theorem 22. *Given a word w of length n and a word u of length m , we can test in $O(n)$ time whether u is an SAS or MAS of w .*

In both cases, we check trivially whether u is absent from w or not. In the case we want to decide whether u is an SAS of w we simply check that $|u| = \iota(w) + 1$. When checking whether u is an MAS we calculate the shortest prefixes of w containing $u[1 : 1], u[1 : 2], \dots, u[1 : m - 1]$ respectively, as well as the shortest suffixes of w containing $u[m : m], u[m - 1 : m], \dots, u[2 : m]$ respectively. Now, u is an MAS of w if and only if for every $1 \leq i \leq m$ the shortest prefix containing $u[1 : i - 1]$ and the shortest suffix containing $u[i + 1 : m]$ do not overlap.

For the analysis of the set $\text{SAS}(w)$ we can construct in linear time data structures, visualized by a tree called SAS-tree, which encodes $\text{SAS}(w)$.

Theorem 23 ([38]). *Given a word w of length n with universality index k , we can construct in $O(n)$ time data structures allowing us to perform the following tasks:*

1. *We can check in $O(k)$ time if a word u of length $k + 1$ is an SAS of w .*
2. *We can compute in $O(k)$ time the lexicographically smallest SAS of w .*
3. *We can efficiently enumerate (i. e., with polynomial delay) all the SAS of w .*

For minimal absent subsequences the problem becomes more complicated but we can still construct data structures encoding $\text{MAS}(w)$, visualized by a directed acyclic graph called MAS-DAG, in $O(n^2\sigma)$ time. It is worth noting here that the lexicographic smallest MAS of w is $a^{|w|_a+1}$ where a is the lexicographic smallest letter of Σ which occurs in w .

Theorem 24 ([38]). *For a word w , we can construct in $O(n^2\sigma)$ time data structures allowing us to efficiently perform the following tasks:*

1. *We can check in $O(m)$ time if a word u of length m is an MAS of w .*
2. *We can compute in polynomial time the longest MAS of w .*
3. *We can check in polynomial time for a given length ℓ if there exists an MAS of length ℓ of w .*
4. *We can efficiently enumerate (with polynomial delay) all the MAS of w .*

In the end of this subsection we present a result allowing us to check whether a word u can be extended to an MAS, that is checking whether there is an MAS of w having u as a prefix, and if possible calculates the shortest such MAS.

Corollary 25. *For a word w of length n , we can construct in $O(n\sigma)$ time data structures allowing us to answer $\text{masExt}(u)$ queries: for a subsequence u of w , decide whether there exists an MAS uv of w , and, if yes, construct such an MAS uv of minimal length. The time needed to answer a query is $O(|v| + |u|)$.*

Absent subsequences and bounded range constraints. We continue by considering the case when bounded range restrictions are added in the study of absent subsequences. That is, we are only interested in subsequences not occurring in any factor of fixed length p of w (but, which may occur in longer factors). We call such a sequence *absent p -subsequence*. Similarly we define the notions of *shortest absent p -subsequences* p -SAS and *minimal absent p -subsequences* p -MAS.

Definition 26 (– Absent p -subsequence). *The word v is an absent p -subsequence of w if $v \notin p\text{-Subseq}(w)$. We also say v is p -absent from w . The word v is a p -SAS (shortest absent p -subsequence) of w if v is an absent p -subsequence of w of minimal length. The word v is a p -MAS (minimal absent p -subsequence) of w if v is an absent p -subsequence of w but all subsequences of v are p -subsequences of w .*

Adding the bounded range restriction to absent subsequences complicates significantly some of the algorithmic tasks which were efficiently solved in the original setting. In particular, checking whether u is not a p -SAS of w is NP-hard and cannot be computed in subexponential time (conditional to ETH).

Theorem 27. *Deciding whether v is not a p -SAS of w is NP-hard and cannot be solved in subexponential time $2^{o(k)} \text{poly}(k, n, m)$ unless ETH fails.*

We denote by $p\text{MAS}$ the decision problem to check for given strings $v, w \in \Sigma^*$ whether v is a p -MAS of w . $p\text{MAS}$ is still decidable in polynomial time $O(|u||w|)$, which is also optimal unless OVH fails. An optimal algorithm is given in [39].

Theorem 28. *$p\text{MAS}$ can be solved in time $O(nm)$, where $|v| = m, |w| = n$.*

Theorem 29. *pMAS cannot be solved in time $O(n^h m^g)$ where $h + g = 2 - \varepsilon$ with $\varepsilon > 0$, unless OVH fails.*

Similarly to the case of Theorem 5, the algorithm proposed in [39] can be seen as working in the sliding window model, with window of fixed size p . If, as in the case of the discussion following Theorem 5, we assume u (and m) to be constant, we obtain a linear time algorithm. However, its space complexity, measured in memory words, is $O(p)$ (as we need to keep track, in this case, of entire content of the window). In fact, when m is constant, it is easy to obtain a linear time algorithm using $O(1)$ memory words (more precisely, $O(\log p)$ bits of space) for this problem: simply try to match u and all its subsequences of length $(m - 1)$ in w simultaneously, using the algorithm from Theorem 5. Clearly, u is a p-MAS if and only if u is not a subsequence of w , but all its subsequences of length $m - 1$ are. However, the constant hidden by the O -notation in the complexity of this algorithm is proportional with m^2 . It remains open whether there exists a (sliding window) algorithm for pMAS both running in $O(mn)$ time (which is optimal, conditional to OVH) and using only $O(\log p)$ bits (which is also optimal for sliding window algorithms, see [26]).

Complementing the discussion above, one can show that it is possible to construct in linear time, for words u, w and integer $p \in \mathbb{N}$, a string w' such that deciding whether u is a p-MAS of w' is equivalent to deciding whether u is a p -subsequence of w , so solving pSubSeqMatch for the input words u and w . Hence, the lower bound from Theorem 6 carries over, and the algorithm announced in Theorem 28 is optimal (conditional to OVH) from the time complexity point of view.

Interestingly, the study of absent subsequences was not considered yet for the case of subsequences with gap constraints.

4.3 Equivalence

The equivalence problem for subsequences EQUI is to decide, for given strings $v, w \in \Sigma^*$ with $|v| = m$ and $|w| = n$ as well as an integer k , whether the sets of subsequences of length at most k of v equals the respective set of w , $\text{Subseq}_{\leq k}(v) = \text{Subseq}_{\leq k}(w)$.

EQUI, and its maximization variant in which one looks for the largest k for which EQUI with inputs v, w, k is true, were among the most studied problems in relation to subsequences. In particular, Hebrard [32] presented the aforementioned maximization problem as computing a similarity measure between strings and mentions a solution of Simon [53] for this problem which runs in $O(|\Sigma|nm)$ (the same solution is mentioned in [28]). Hebrard improves this (see [32]) in the case when Σ is a binary alphabet: given two bitstrings w and v , one can find the maximum k for which $\text{Subseq}_{\leq k}(v) = \text{Subseq}_{\leq k}(w)$ in linear time. However, the problem of finding optimal algorithms for both EQUI and its maximization variant, in the case of general alphabets, was left open in [53, 32] as the methods used in the latter paper for binary strings did not seem to scale up. In [28], Garel approaches the maximization problem and presents an algorithm based on finite automata, running in $O(|\Sigma|n)$, which computes all *distinguishing words* u of minimum length, i. e., words which are factors of only one of the words w and v from the problem's statement. Several further improvements on the aforementioned results were reported in [16, 57]. Also, in an extended abstract from 2003 [56], Simon presented another algorithm based on finite automata solving this maximization problem, which runs in $O(|\Sigma|n)$, and he conjectures that it can be implemented in $O(|\Sigma| + n)$. Unfortunately, the last claim was only insufficiently substantiated, and obtaining an algorithm with the claimed complexity remained open (in fact, Simon announced that a detailed description of this algorithm will follow shortly, but we were not able to find it in the literature).

Further, in [20], a novel approach to efficiently solving EQUI was introduced. This idea was to compute, for the two given words v and w and the given number k , their shortlex forms: the words which have the same set of subsequences of length at most k as v and w , respectively, and are also lexicographically

smallest among all words with the respective property. Clearly, $\text{Subseq}_{\leq k}(v) = \text{Subseq}_{\leq k}(w)$ if and only if the shortlex forms of v and w for k coincide.

The shortlex form of a word w of length n over Σ was computed in $O(|\Sigma|n)$ time in [20], so EQUI was also solved in $O(|\Sigma|n)$. A more efficient implementation of the ideas introduced in [20] was presented in [7]: the shortlex form of a word of length n over Σ can be computed in linear time $O(n)$, so EQUI can be solved in optimal linear time. By binary searching for the smallest k for which EQUI with inputs v, w, k is true, gives an $O(n \log n)$ time solution for the corresponding optimization problem.

Later, Gawrychowski et al. [29] solved this optimization problem (finding the maximum k such that $\text{Subseq}_{\leq k}(v) = \text{Subseq}_{\leq k}(w)$) in optimal linear time, as well. However, to achieve this result a novel data structure, the Simon-Tree, was introduced. A node of depth k in the Simon-Tree of a word w corresponds to a maximal interval $[i : j]$ (called k -block) such that for all $\ell, \ell' \in [i : j]$ it holds that $\text{Subseq}_k(w[\ell : n]) = \text{Subseq}_k(w[\ell' : n])$.

Definition 30. *The Simon-Tree T_w associated to the word w , with $|w| = n$, is an ordered rooted tree. The nodes of depth k represent k -blocks of w , for $0 \leq k \leq n$, and are defined recursively.*

- *The root corresponds to the 0-block of the word w , i. e., the interval $[1 : n]$.*
- *For $k > 1$ and for a node a of depth $k - 1$, which represents a $(k - 1)$ -block $[i : j]$ with $i < j$, the children of a are exactly the blocks of the partition of $[i : j]$ in k -blocks, ordered decreasingly (right-to-left) by their starting position.*
- *For $k > 1$, each node of depth $k - 1$ which represents a singleton- $(k - 1)$ -block is a leaf.*

The Simon-Tree T_w can be constructed in linear time $O(n)$. Furthermore, for words $v, w \in \Sigma^*$ of length $|v| = m$ and $|w| = n$, Gawrychowski et al. ([29]) give a linear time algorithm to connect nodes of the Simon-Trees T_v, T_w . Two nodes $[i, j]$ in T_v (of depth k) and $[i', j']$ in T_w (also of depth k) become connected if and only if $\text{Subseq}_{\leq k}(v[\ell : m]) = \text{Subseq}_{\leq k}(w[\ell' : n])$ for all $\ell \in [i : j]$ and $\ell' \in [i', j']$. The maximal k such that $\text{Subseq}_k(v) = \text{Subseq}_k(w)$ now equals the depth of the deepest connected nodes $[1, x]$ in T_v and $[1, y]$ in T_w . The following result holds.

Theorem 31 ([29]). *Given two strings v and w , with $n = |w| \geq |v|$, the largest k for which EQUI with input v, w, k is answered positively can be computed in $O(n)$ time.*

Equivalence and bounded range constraints. Considering the problem EQUI for subsequences occurring within bounded ranges leads again to a surprising intractability result. Once more, for convenience, we give this result for the complement problem, i. e., non-equivalence problem NEQUI. We are given two words $v, w \in \Sigma^*$, and two numbers $k, p \in \mathbb{N}$, and we want to decide whether $p\text{-Subseq}_k(v) \neq p\text{-Subseq}_k(w)$. The hardness result obtained in [39] (and corresponding conditional lower bound) is the following.

Theorem 32. *NEQUI with a bounded range constraint is NP-hard and cannot be solved in subexponential time $2^{o(k)} \text{poly}(k, n, m)$ unless ETH fails.*

Equivalence and gap constraints. Finally, similarly to the case of subsequences occurring in bounded ranges, for the case of subsequences with gap constraints, one can adapt the results from [18], related to UNI, to get the following results.

Firstly, some algorithmic upper bounds.

Theorem 33. (1) The problem EQUI for subsequences with length (or reg-len) constraints can be solved in time $O(|\Sigma|^k \text{nz}(gc)\ell)$ (respectively, $O(|\Sigma|^k \text{states}(gc)\ell)$), where $\ell = \max\{|w|, |w'|\}$.

(2) For fixed alphabet Σ , the problem EQUI_Σ with length (or reg-len) constraints can be solved in time $2^{O(k)} \text{nz}(gc)\ell$ (respectively, $2^{O(k)} \text{states}(gc)\ell$), where $\ell = \max\{|w|, |w'|\}$.

Secondly, a couple of intractability results, doubled by conditional lower bounds.

Theorem 34. For every fixed alphabet Σ with $|\Sigma| \geq 3$, NEQUI_Σ with length constraints is NP-complete, even if all length constraints are (1, 5). Moreover,

- it cannot be solved in subexponential time $2^{o(k)} \text{poly}(|w|, k)$ (unless ETH fails),
- it cannot be solved in time $O(2^{k(1-\varepsilon)} \text{poly}(|w|, k))$ (unless SETH fails).

For every fixed alphabet Σ with $|\Sigma| = 2$, NEQUI_Σ with length constraints are NP-complete even if each length constraint is (0, 0) or (3, 9).

5 Conclusions, Related Problems, and Future Work

In this paper, we overviewed a series of recent algorithmic and complexity theoretic results related to the matching and analysis problems for (constrained) subsequences.

An interesting problem related to the study of analysis problems for classical subsequences is the *containment problem* (denoted CON for short), which consists in deciding whether $\text{Subseq}_k(w) \subseteq \text{Subseq}_k(v)$ for given strings $w, v \in \Sigma^*$ and integer k . This problem can be solved in polynomial time by an automata theoretic approach, as follows. We start with our input words w and v . For w we construct A_w , the subsequence-automaton [16] which accepts all the subsequences of w . Assume $|w| = n$. Then, A_w is a deterministic finite automaton, which has $n + 2$ states $\{0, \dots, n, n + 1\}$. The initial state of this automaton is 0, and all states i , with $i \in [n]$ are final; the state $n + 1$ is an error state. The transition are defined as follows: for all $i \in [n], k \in [n - i]$, and $a \in \Sigma$, we have a transition from state i to $i + k$, labelled with a , if and only if $w[i + k] = a$ and $w[i + 1..i + k - 1]$ does not contain the letter a . Moreover, for all $i \in [n]$ and $a \in \Sigma$, we have a transition from state i to state $n + 1$, labelled with letter a , if and only if $w[i + 1..n]$ does not contain a . For state $n + 1$ we have loop-transitions for all letters $a \in \Sigma$. It is straightforward that A_w accepts exactly the non-empty subsequences of w and can be constructed in $O(n|\Sigma|)$ time. For the word v , with $|v| = m$, we construct the automaton A_v , as above, and then modify it to obtain the automaton B_v by simply making the state $m + 1$ the single final state. Clearly, B_v accepts all strings which are not subsequences of v . It can be constructed in $O(m|\Sigma|)$ time. Now, we can observe that $\text{Subseq}_k(w) \not\subseteq \text{Subseq}_k(v)$ if and only if there exists a word of length at most k accepted by A_w which is also accepted by B_v . This can be checked in $O(nm|\Sigma|)$ time by simply computing the shortest word in the intersection of the language accepted by A_w with the language accepted by B_v . However, it remains an interesting open problem whether a linear time algorithm exists for CON, as it exists for EQUI.

When considering subsequences with gap constraints, it seems interesting to see if the polynomiality of the matching problem is preserved under adding *gap length equalities* to the gap constraints, i. e., constraints of the form $|\text{gap}_i| = |\text{gap}_j|$ which are satisfied by an embedding e with respect to w if $|\text{gap}_e(w, i)| = |\text{gap}_e(w, j)|$. Such length equality constraints (and more complex ones, e. g., described by linear inequalities) are of interest in the theory of string solving [3]. Unfortunately, the matching problem becomes immediately NP-hard, as it can be shown (see [18]) by adapting the NP-completeness proof for matching patterns with variables from [4].

Finally, with respect to EQUI, we can consider the following modified setting, both in the classical case and in the case of constrained subsequences (in bounded ranges, or with gap constraints). The idea

is to consider the sets of subsequences with multiplicities. We present this here for subsequences with gap constraints, following [18], but this can be trivially adapted to the case of classical subsequences (see [23]), as well as to the case of subsequences with bounded-range constraints. So, we consider the gc -subsequences of the sets $Subseq(gc, w)$ with multiplicities. For example, for $w_1 = abba$ and $w_2 = abab$, we have $Subseq(gc_2, w_1) = Subseq(gc_2, w_2) = \{aa, ab, ba, bb\}$ with $gc_2 = (\Sigma^*)$. There is exactly one way of embedding aa and bb into both w_1 and w_2 . On the other hand, ab can be embedded into w_1 in two different ways and into w_2 in three different ways. More precisely, the sets of gc_2 -subsequences of w_1 and w_2 with multiplicities are $\{(aa, 1), (ab, 2), (ba, 2), (bb, 1)\}$ and $\{(aa, 1), (ab, 3), (ba, 1), (bb, 1)\}$, respectively. So, we can now formalise this setting. For strings u and v , and a $(|u| - 1)$ -tuple gc of gap constraints, we denote by $|u|_{v, gc}$ the number of distinct embeddings $e : |u| \rightarrow |v|$ that satisfy gc and $v \preceq_e u$. For example, $|bbaa|_{ba, gc_2} = 4$, as $u[1]u[3] = u[2]u[3] = u[1]u[4] = u[2]u[4] = ba$. For any $(k - 1)$ -tuple gc of gap constraints, we define the function $\Psi_{gc}(\cdot) : \Sigma^* \rightarrow \mathbb{N}^{\Sigma^k}$ by $\Psi_{gc}(w)[p] = |w|_{p, gc}$ for every $p \in \Sigma^k$.

The *equivalence problem with multiplicities* is to decide, for a given $(k - 1)$ -tuple gc of gap constraints, and strings $w, v \in \Sigma^*$, whether $\Psi_{gc}(w) = \Psi_{gc}(v)$. Note that for the case $gc = (\Sigma^*, \dots, \Sigma^*)$ this is called the k -binomial equivalence, and was studied in the area of combinatorics on words (see, e. g., [49, 42, 43, 23]). By an automata theoretic approach one can show that the equivalence with multiplicities problem can be decided in polynomial time (in stark contrast to the NP-completeness of the case without multiplicities).

The idea, firstly introduced in [23] and then used in [18], is the following. We first construct, for the first input word w with $|w| = n$ and a $(k - 1)$ -tuple gc of gap constraints, a non-deterministic finite automaton $A_{w, gc}$ that accepts exactly the gapped subsequences $p \in Subseq(gc, w)$ and has exactly $\Psi_{gc}(w)[p]$ accepting paths labelled with the subsequence p of w . Then, we will use the same construction for the second input word v , to obtain $A_{v, gc}$. Finally, we use the algorithm of [58] to test whether $A_{w, gc}$ and $A_{v, gc}$ are path equivalent, i. e., for each word p , the number of accepting paths of $A_{w, gc}$ labelled with p equals the number of accepting paths of $A_{v, gc}$ labelled with p . If this algorithm returns a positive answer, then we can conclude that $\Psi_{gc}(w) = \Psi_{gc}(v)$. Otherwise, we conclude that $\Psi_{gc}(w) \neq \Psi_{gc}(v)$.

While the algorithm discussed above related to the *equivalence problem with multiplicities* runs in polynomial time $O(\max\{|w|, |v|\}^4 k^4 + \text{size}(gc))$ for reg-len constraints, it would be interesting to see if faster algorithms exist. Moreover, the *containment problem with multiplicities* (i. e., deciding $\Psi_{gc}(w)[p] \leq \Psi_{gc}(v)[p]$ for all $p \in \Sigma^k$) seems to be more difficult. To our knowledge, whether the case of classical subsequences can be solved in polynomial time remains open.

While the discussion above highlights some clear open problems related to the topics covered in this paper, one could also extend this research by considering, for instance, other classes of constrained subsequences and investigating the matching and analysis problems for those classes as well.

References

- [1] Amir Abboud, Arturs Backurs & Virginia Vassilevska Williams (2015): *Tight Hardness Results for LCS and Other Sequence Similarity Measures*. In: *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pp. 59–78, doi:10.1109/FOCS.2015.14.
- [2] Amir Abboud, Virginia Vassilevska Williams & Oren Weimann (2014): *Consequences of Faster Alignment of Sequences*. In: *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pp. 39–51, doi:10.1007/978-3-662-43948-7_4.
- [3] Roberto Amadini (2021): *A survey on string constraint solving*. *ACM Computing Surveys (CSUR)* 55(1), pp. 1–38, doi:10.1145/3484198.

- [4] Dana Angluin (1980): *Finding Patterns Common to a Set of Strings*. *J. Comput. Syst. Sci.* 21(1), pp. 46–62, doi:10.1016/0022-0000(80)90041-0.
- [5] Alexander Artikis, Alessandro Margara, Martín Ugarte, Stijn Vansummeren & Matthias Weidlich (2017): *Complex Event Recognition Languages: Tutorial*. In: *Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems, DEBS 2017, Barcelona, Spain, June 19-23, 2017*, pp. 7–10, doi:10.1145/3093742.3095106.
- [6] Ricardo A. Baeza-Yates (1991): *Searching Subsequences*. *Theor. Comput. Sci.* 78(2), pp. 363–376, doi:10.1016/0304-3975(91)90358-9.
- [7] Laura Barker, Pamela Fleischmann, Katharina Harwardt, Florin Manea & Dirk Nowotka (2020): *Scattered Factor-Universality of Words*. In Natasa Jonoska & Dmytro Savchuk, editors: *Proc. DLT 2020, Lecture Notes in Computer Science* 12086, pp. 14–28, doi:10.1007/978-3-030-48516-0_2.
- [8] Philip Bille, Inge Li Gørtz, Hjalte Wedel Vildhøj & David Kofoed Wind (2012): *String matching with variable length gaps*. *Theor. Comput. Sci.* 443, pp. 25–34, doi:10.1016/j.tcs.2012.03.029.
- [9] Karl Bringmann (2014): *Why Walking the Dog Takes Time: Frechet Distance Has No Strongly Subquadratic Algorithms Unless SETH Fails*. In: *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pp. 661–670, doi:10.1109/FOCS.2014.76.
- [10] Karl Bringmann (2019): *Fine-Grained Complexity Theory (Tutorial)*. In: *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, pp. 4:1–4:7, doi:10.4230/LIPIcs.STACS.2019.4.
- [11] Karl Bringmann & Bhaskar Ray Chaudhury (2018): *Sketching, Streaming, and Fine-Grained Complexity of (Weighted) LCS*. In: *Proc. FSTTCS 2018, LIPIcs* 122, pp. 40:1–40:16, doi:10.4230/LIPIcs.FSTTCS.2018.40.
- [12] Karl Bringmann & Marvin Künnemann (2018): *Multivariate Fine-Grained Complexity of Longest Common Subsequence*. In: *Proc. SODA 2018*, pp. 1216–1235, doi:10.1137/1.9781611975031.79.
- [13] Sam Buss & Michael Soltys (2014): *Unshuffling a square is NP-hard*. *J. Comput. Syst. Sci.* 80(4), pp. 766–776, doi:10.1016/j.jcss.2013.11.002.
- [14] Peter Clifford & Raphaël Clifford (2007): *Simple deterministic wildcard matching*. *Inf. Process. Lett.* 101(2), pp. 53–54, doi:10.1016/j.ipl.2006.08.002.
- [15] Maxime Crochemore, Christophe Hancart & Thierry Lecroq (2007): *Algorithms on strings*. Cambridge University Press, doi:10.1017/CBO9780511546853.
- [16] Maxime Crochemore, Borivoj Melichar & Zdenek Troníček (2003): *Directed acyclic subsequence graph — Overview*. *J. Discrete Algorithms* 1(3-4), pp. 255–280, doi:10.1016/S1570-8667(03)00029-7.
- [17] Joel D. Day, Pamela Fleischmann, Maria Kosche, Tore Koß, Florin Manea & Stefan Siemer (2021): *The Edit Distance to k -Subsequence Universality*. In: *STACS, LIPIcs* 187, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 25:1–25:19, doi:10.4230/LIPIcs.STACS.2021.25.
- [18] Joel D. Day, Maria Kosche, Florin Manea & Markus L. Schmid (2022): *Subsequences With Gap Constraints: Complexity Bounds for Matching and Analysis Problems*. *CoRR* abs/2206.13896, doi:10.48550/ARXIV.2206.13896.
- [19] Xavier Droubay, Jacques Justin & Giuseppe Pirillo (2001): *Episturmian words and some constructions of de Luca and Rauzy*. *Theor. Comput. Sci.* 255(1-2), pp. 539–553, doi:10.1016/S0304-3975(99)00320-5.
- [20] Lukas Fleischer & Manfred Kufleitner (2018): *Testing Simon’s congruence*. In: *Proc. MFCS 2018, LIPIcs* 117, pp. 62:1–62:13, doi:10.4230/LIPIcs.MFCS.2018.62.
- [21] Pamela Fleischmann, Sebastian Bernhard Germann & Dirk Nowotka (2021): *Scattered Factor Universality - The Power of the Remainder*. *CoRR* abs/2104.09063, doi:10.48550/ARXIV.2104.09063. To appear in *Proc. DCFS 2022*.
- [22] Pamela Fleischmann, Lukas Haschke, Annika Huch, Annika Mayrock & Dirk Nowotka (2022): *m -Nearly k -Universal Words - Investigating Simon Congruence*. *CoRR* abs/2202.07981, doi:10.48550/ARXIV.2202.07981.

- [23] Dominik D. Freydenberger, Pawel Gawrychowski, Juhani Karhumäki, Florin Manea & Wojciech Rytter (2015): *Testing k -binomial equivalence*. CoRR abs/1509.00622, pp. 239–248, doi:10.48550/ARXIV.1509.00622. *Multidisciplinary Creativity*, a collection of papers dedicated to G. Păun 65th birthday.
- [24] Moses Ganardi (2019): *Language recognition in the sliding window model*. Ph.D. thesis, University of Siegen, Germany.
- [25] Moses Ganardi, Danny Hucke, Daniel König, Markus Lohrey & Konstantinos Mamouras (2018): *Automata Theory on Sliding Windows*. In: *STACS, LIPIcs 96*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 31:1–31:14, doi:10.4230/LIPIcs.STACS.2018.31.
- [26] Moses Ganardi, Danny Hucke & Markus Lohrey (2016): *Querying Regular Languages over Sliding Windows*. In: *FSTTCS, LIPIcs 65*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 18:1–18:14, doi:10.4230/LIPIcs.FSTTCS.2016.18.
- [27] Moses Ganardi, Danny Hucke, Markus Lohrey & Tatiana Starikovskaya (2019): *Sliding Window Property Testing for Regular Languages*. In: *ISAAC, LIPIcs 149*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 6:1–6:13, doi:10.4230/LIPIcs.ISAAC.2019.6.
- [28] Emmanuelle Garel (1993): *Minimal Separators of Two Words*. In: *Proc. CPM 1993, Lecture Notes in Computer Science 684*, pp. 35–53, doi:10.1007/BFb0029795.
- [29] Pawel Gawrychowski, Maria Kosche, Tore Koß, Florin Manea & Stefan Siemer (2021): *Efficiently Testing Simon’s Congruence*. In Markus Bläser & Benjamin Monmege, editors: *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference), LIPIcs 187*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 34:1–34:18, doi:10.4230/LIPIcs.STACS.2021.34.
- [30] Nikos Giatrakos, Elias Alevizos, Alexander Artikis, Antonios Deligiannakis & Minos N. Garofalakis (2020): *Complex event recognition in the Big Data era: a survey*. *VLDB J.* 29(1), pp. 313–352, doi:10.1007/s00778-019-00557-w.
- [31] Simon Halfon, Philippe Schnoebelen & Georg Zetsche (2017): *Decidability, complexity, and expressiveness of first-order logic over the subword ordering*. In: *Proc. LICS 2017*, pp. 1–12, doi:10.5555/3329995.3330076.
- [32] Jean-Jacques Hebrard (1991): *An algorithm for distinguishing efficiently bit-strings by their subsequences*. *Theor. Comput. Sci.* 82(1), pp. 35–49, doi:10.1016/0304-3975(91)90170-7.
- [33] Russell Impagliazzo & Ramamohan Paturi (2001): *On the Complexity of k -SAT*. *J. Comput. Syst. Sci.* 62(2), pp. 367–375, doi:10.1006/jcss.2000.1727.
- [34] Russell Impagliazzo, Ramamohan Paturi & Francis Zane (2001): *Which Problems Have Strongly Exponential Complexity?* *J. Comput. Syst. Sci.* 63(4), pp. 512–530, doi:10.1006/jcss.2001.1774.
- [35] Prateek Karandikar, Manfred Kufleitner & Philippe Schnoebelen (2015): *On the index of Simon’s congruence for piecewise testability*. *Inf. Process. Lett.* 115(4), pp. 515–519, doi:10.1016/j.ipl.2014.11.008.
- [36] Prateek Karandikar & Philippe Schnoebelen (2016): *The Height of Piecewise-Testable Languages with Applications in Logical Complexity*. In: *Proc. CSL 2016, LIPIcs 62*, pp. 37:1–37:22, doi:10.4230/LIPIcs.CSL.2016.37.
- [37] Prateek Karandikar & Philippe Schnoebelen (2019): *The height of piecewise-testable languages and the complexity of the logic of subwords*. *Log. Methods Comput. Sci.* 15(2), doi:10.23638/LMCS-15(2:6)2019.
- [38] Maria Kosche, Tore Koß, Florin Manea & Stefan Siemer (2021): *Absent Subsequences in Words*. In Paul C. Bell, Patrick Totzke & Igor Potapov, editors: *Reachability Problems*, Springer International Publishing, Cham, pp. 115–131, doi:10.1007/978-3-030-89716-1_8.
- [39] Maria Kosche, Tore Koß, Florin Manea & Viktoriya Pak (2022): *Subsequences in Bounded Ranges: Matching and Analysis Problems*. CoRR abs/2207.09201, doi:10.48550/ARXIV.2207.09201.
- [40] Dietrich Kuske (2020): *The Subtrace Order and Counting First-Order Logic*. In: *Proc. CSR 2020, Lecture Notes in Computer Science 12159*, pp. 289–302, doi:10.1007/978-3-030-50026-9_21.

- [41] Dietrich Kuske & Georg Zetsche (2019): *Languages Ordered by the Subword Order*. In: *Proc. FOSSACS 2019, Lecture Notes in Computer Science* 11425, pp. 348–364, doi:10.1007/978-3-030-17127-8_20.
- [42] Marie Lejeune, Julien Leroy & Michel Rigo (2019): *Computing the k -binomial Complexity of the Thue-Morse Word*. In: *Proc. DLT 2019, Lecture Notes in Computer Science* 11647, pp. 278–291, doi:10.1007/978-3-030-24886-4_21.
- [43] Julien Leroy, Michel Rigo & Manon Stipulanti (2017): *Generalized Pascal triangle for binomial coefficients of words*. *Electron. J. Combin.* 24(1.44), p. 36 pp., doi:10.1016/j.aam.2016.04.006.
- [44] Daniel Lokshantov, Dániel Marx & Saket Saurabh (2011): *Lower bounds based on the Exponential Time Hypothesis*. *Bull. EATCS* 105, pp. 41–72, doi:10.1007/978-3-319-21275-3_14. Available at <http://eatcs.org/beatcs/index.php/beatcs/article/view/92>.
- [45] Aldo de Luca, Amy Glen & Luca Q. Zamboni (2008): *Rich, Sturmian, and trapezoidal words*. *Theor. Comput. Sci.* 407(1-3), pp. 569–573, doi:10.1016/j.tcs.2008.06.009.
- [46] David Maier (1978): *The Complexity of Some Problems on Subsequences and Supersequences*. *J. ACM* 25(2), pp. 322–336, doi:10.1145/322063.322075.
- [47] Alexandru Mateescu, Arto Salomaa & Sheng Yu (2004): *Subword Histories and Parikh Matrices*. *J. Comput. Syst. Sci.* 68(1), pp. 1–21, doi:10.1016/j.jcss.2003.04.001.
- [48] William E. Riddle (1979): *An Approach to Software System Modelling and Analysis*. *Comput. Lang.* 4(1), pp. 49–66, doi:10.1016/0096-0551(79)90009-2.
- [49] Michel Rigo & Pavel Salimov (2015): *Another generalization of abelian equivalence: Binomial complexity of infinite words*. *Theor. Comput. Sci.* 601, pp. 47–57, doi:10.1016/j.tcs.2015.07.025.
- [50] Arto Salomaa (2005): *Connections Between Subwords and Certain Matrix Mappings*. *Theoret. Comput. Sci.* 340(2), pp. 188–203, doi:10.1016/j.tcs.2005.03.024.
- [51] Shinnosuke Seki (2012): *Absoluteness of subword inequality is undecidable*. *Theor. Comput. Sci.* 418, pp. 116–120, doi:10.1016/j.tcs.2011.10.017.
- [52] Alan C. Shaw (1978): *Software Descriptions with Flow Expressions*. *IEEE Trans. Software Eng.* 4(3), pp. 242–254, doi:10.1109/TSE.1978.231501.
- [53] Imre Simon: *An Algorithm to Distinguish Words efficiently by their Subwords*. unpublished.
- [54] Imre Simon (1972): *Hierarchies of events with dot-depth one*. Ph.D. thesis.
- [55] Imre Simon (1975): *Piecewise testable events*. In: *Autom. Theor. Form. Lang., 2nd GI Conf., LNCS* 33, pp. 214–222, doi:10.1007/3-540-07407-4_23.
- [56] Imre Simon (2003): *Words distinguished by their subwords (extended Abstract)*. In: *Proc. WORDS 2003, TUCS General Publication* 27, pp. 6–13.
- [57] Zdenek Troníček (2002): *Common Subsequence Automaton*. In: *Proc. CIAA 2002 (Revised Papers), Lecture Notes in Computer Science* 2608, pp. 270–275, doi:10.1007/3-540-44977-9_28.
- [58] Wen-Guey Tzeng (1992): *A Polynomial-Time Algorithm for the Equivalence of Probabilistic Automata*. *SIAM J. Comput.* 21(2), pp. 216–227, doi:10.1137/0221017.
- [59] Virginia Vassilevska Williams (2015): *Hardness of Easy Problems: Basing Hardness on Popular Conjectures such as the Strong Exponential Time Hypothesis (Invited Talk)*. In: *10th International Symposium on Parameterized and Exact Computation, IPEC 2015, September 16-18, 2015, Patras, Greece*, pp. 17–29, doi:10.4230/LIPIcs.IPEC.2015.17.
- [60] Georg Zetsche (2016): *The Complexity of Downward Closure Comparisons*. In: *Proc. ICALP 2016, LIPIcs* 55, pp. 123:1–123:14, doi:10.4230/LIPIcs.ICALP.2016.123.
- [61] Haopeng Zhang, Yanlei Diao & Neil Immerman (2014): *On complexity and optimization of expensive queries in complex event processing*. In: *International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014*, pp. 217–228, doi:10.1145/2588555.2593671.