Negative Selection Algorithms on Strings
with Efficient Training and Linear-Time Classification

Michael Elberfeld, Johannes Textor

Institut für Theoretische Informatik, Universität zu Lübeck, 23538 Lübeck, Germany

Abstract

A string-based negative selection algorithm is an immune-inspired classifier that infers a bipartitioning of a string space Σ^ℓ from a training set S containing samples from only one of the partitions. The algorithm generates a set of patterns, called "detectors", to cover regions of the string space containing none of the training samples. These patterns are then used for classification. A major problem with all existing implementations of this approach is that the detector generation step suffers from exponential worst-case time complexity. Hence, researchers have found negative selection to be of limited use for real-world problems such as network intrusion detection. Here we show that for the two most widely used kinds of detectors, the r-chunk and r-contiguous detectors based on partial matching to substrings of length r, negative selection can be implemented efficiently by avoiding generating detectors altogether: For each detector type, training set $S \subseteq \Sigma^\ell$ and parameter $r \leq \ell$ one can construct an automaton with an acceptance behaviour that is equivalent to the algorithm's classification outcome. The resulting runtime is $O(|S|\ell r|\Sigma|)$ for constructing the automaton in the training phase and $O(\ell)$ for classifying a string.

Key words: negative selection, r-chunk detectors, r-contiguous detectors, artificial immune systems, anomaly detection

1. Introduction

The adaptive immune system successfully protects vertebrate species, including us human beings, from becoming extinguished by pathogens. According to current textbook immunology, the immune system accomplishes this without actually knowing what a pathogen is. Instead, it is trained during infancy to tolerate the tissues, cells and molecules that are normal components of its host organism – the *self* – and to attack everything else – the *nonself*. While nonself includes potentially dangerous things such as viruses, bacteria and fungi, this implies that benign intruders such as a donated kidney or liver are also attacked by the immune system. The paradigm of self-nonself-discrimination is a natural source of inspiration for computer security: Computer systems and networks are also continuously attacked by worms and other malware, and a computer program that discriminates with perfect accuracy between benign and malign software cannot exist. Thus, can we benefit from transferring the immune system's "nice hack" [17] to the computer security domain?

A popular approach to designing such computer immune systems is to mimic how the immune system's *T cells* are generated and trained to detect nonself entities. This process is known as *negative selection* [14, 15]: T cell receptors are generated by random assembly of gene fragments. In an organ called the *thymus*, newborn T cells are exposed to proteins from self. Every cell whose receptor matches a self protein is destroyed. Only the cells that survive negative selection leave from the thymus and start to continuously circulate through the organism, screening for nonself entities. A *negative selection algorithm* is essentially an abstraction of this process.

The negative selection algorithms that we consider in this paper are binary classifiers operating on a string space Σ^ℓ. The classification problem is posed as follows (Figure 1): Σ^ℓ is assumed to be pre-partitioned in two pairwise disjoint subsets $\mathcal{S}$ (self) and $\mathcal{N}$ (nonself). The strings can represent, for example, data packets in a computer network

Email addresses: elberfeld@tcs.uni-luebeck.de (Michael Elberfeld), textor@tcs.uni-luebeck.de (Johannes Textor)

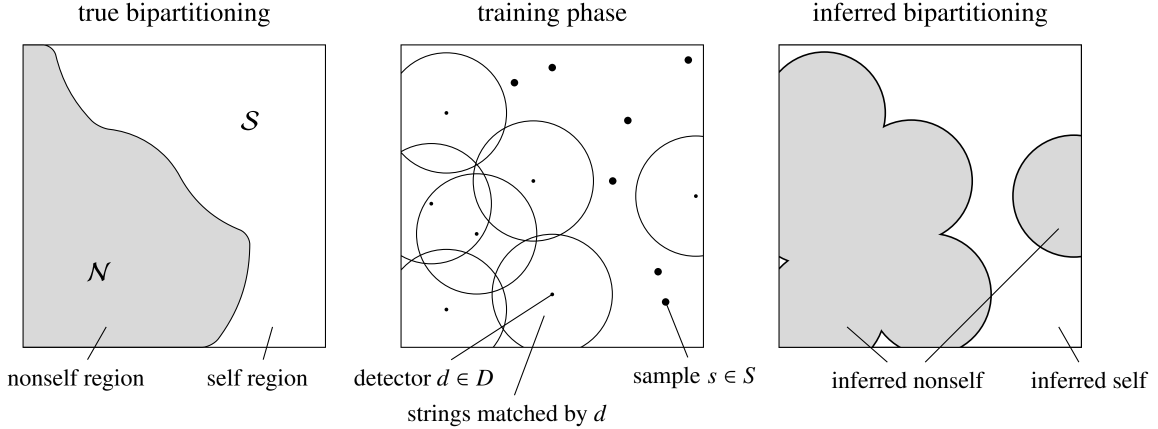

Figure 1: The classification problem solved by a negative selection algorithm. The string universe Σ^ℓ is prepartitioned in two regions $\mathcal{S}$ (self) and $\mathcal{N}$ (nonself). The classifier is given a training set $S \subseteq \mathcal{S}$ (large dots) and generates a detector set D (small dots) to cover regions of the universe containing none of the training examples (circles). This detector set induces a classification boundary that approximates the partitions $\mathcal{S}$ and $\mathcal{N}$.

[3] or sequences of system calls from UNIX processes [13], where the self and nonself partitions would correspond to "normal" and "anomalous" behaviour, respectively. The algorithm is given a sample $S \subseteq \mathcal{S}$ of self strings, called *self-set*, and a set $M \subseteq \Sigma^\ell$ of strings to classify, called *monitor set*. It then generates a set D of patterns called *detectors*. In analogy to the T cells in the immune system, this is typically done by generating the detectors randomly and discarding those that match any string in the self-set. Consequently, each string $m \in M$ is classified by labelling m as non-self if it is matched by *any* detector, and self otherwise. In particular, m is never labeled non-self if it also occurs in the self-set.

From a broader machine learning perspective, negative selection is usually described as an *anomaly detection* technique [16, 6]. The following two important properties distinguish negative selection from many well-known classifiers: (1) The training data consists of examples from only one class. Other techniques with this property include classifiers based on kernel density estimation [4, 20] and the one-class support vector machine [22]. (2) Classification is based on a *negative representation* of training data, typically on short substrings (r-grams) that do not occur in the self-set. While positive representations such as the r-gram frequency distribution used e.g. for identification of language [11] and text categorization [5] are more common in the machine learning domain, similar negative representations have been studied in string theory. For example, certain sets of non-occurring substrings *(forbidden words)* can be used to describe the complexity of a language [8].

1.1. Contribution of this paper

This paper presents two algorithms that implement negative selection with r-chunk and r-contiguous detectors by generating compressed representations of the respective detector sets, from which automata are constructed that simulate the classification outcome through their acceptance behaviour. Both algorithms use time $O(|S|\ell r|\Sigma|)$ to construct an automaton for a given self-set S and parameter r, which is equivalent to the *training phase* of the simulated negative selection algorithm. Upon construction, the automaton classifies each string in linear time $O(\ell)$. This improves upon the exponential worst-case complexity of existing algorithms, and thus removes one major obstacle for applying negative selection to real-world problems [28, 23, 27]. In comparison to our preliminary conference version [12], the algorithms presented in this paper are based on prefix trees instead of patterns. This reduces the overall runtime significantly (Table 1), generalizes to higher alphabets, and allows for a simpler and more concise presentation. In addition to the classification itself, the automata can also be used to efficiently count the detectors and, if necessary, enumerate them explicitly.

The r-chunk and r-contiguous detectors considered here are among the most common ones in the artificial immune systems literature [16]: (1) An *r-contiguous detector* is a string of length ℓ and matches all strings to which it is identical in at least r contiguous positions. (2) An *r-chunk detector* is a string of length r (or r-gram) with a position

index and matches all strings in which the r-gram occurs at that position. Figure 2 shows an example self-set $S \subseteq \{a, b\}^5$ along with the complete sets of 3-chunk and 3-contiguous detectors that do *not* match any string in S, as well as the partitioning of $\{a, b\}^5$ induced by these detector sets. The r-contiguous detectors are directly based on a model of antigen recognition by T cell receptors [21, 14], and r-chunk detectors were later introduced to achieve better results on data where different regions of the input strings have highly different meanings, such as network data packets [3].

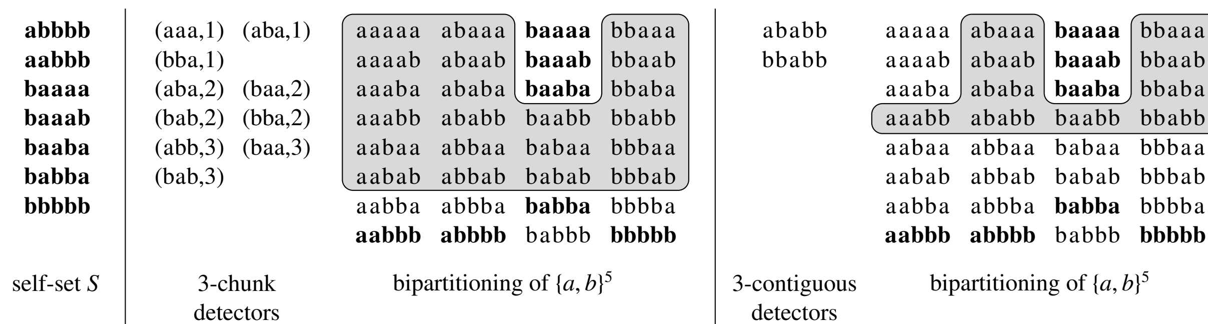

Figure 2: An example self-set $S \subseteq \{a, b\}^5$ along with all 3-chunk detectors and 3-contiguous detectors that do not match any string in S is shown. For both detector types, the bipartitionings of the shape space $\{a, b\}^5$ are illustrated with strings that are classified as nonself having a gray background and strings that are classified as self having a white background. Bold strings are members of the self-set. The *generalization region* of the negative selection classifier consists of the strings that are classified as self although they do not occur in the self-set. These strings are also called "holes" in the negative selection literature [16, 25].

1.2. Related Work on String-Based Negative Selection

The question whether negative selection with r-contiguous and r-chunk detectors can be implemented in polynomial worst-case time was open for several years. The complexity issues caused by the verbatim abstraction of negative selection as performed by the immune system are two-fold: On one hand, if the self partition is only a small fraction of Σ^ℓ, then there is an exponential number of potential detectors, and it is unclear how many of these have to be generated to achieve an acceptable detection rate. The early work of D'haeseleer and others [10, 9] addressed these problems by proving lower bounds on the number of required detectors, and presenting algorithms that generate detectors by a structured exhaustive search. However, these algorithms still have a runtime exponential in r. Similar algorithms and heuristics were later proposed by Wierzchoń [30], Ayara et al. [2], and Stibor et al. [26]. In an effort to clarify the computational complexity of negative selection, Stibor and coworkers studied the associated decision problem [24, 25]: Given a self-set S, can an r-contiguous detector be generated that does not match any string in S? It was recently suggested that this decision problem might be $\mathcal{NP}$-complete [28], although a completeness proof was not shown. The ongoing difficulties led some in the field to conclude that negative selection is computationally too expensive for real-world datasets [23, 1]. This issue was settled by the preliminary version of the present paper [12]. Most recently, Liśkiewicz and Textor discussed the idea of negative selection without explicit detector generation from a learning theoretical perspective [19].

1.3. Organization of This Paper

We start out by defining the formal underpinnings of our algorithms in the upcoming section. Afterwards, in Section 3, we sketch the construction of an automaton consisting of prefix trees and failure links that can be used to simulate negative selection with r-chunk detectors. This rather straightforward construction is used as a basis for the more involved one in Section 4, where we transform the automaton into one that allows linear-time classification with respect to r-contiguous detectors.

2. Preliminaries

In this section, we define the formal background of our work. First we review some basic terms related to strings and pattern matching techniques like automata. Then we define r-chunk detectors, r-contiguous detectors, and the corresponding classification approaches.

r-chunk detector-based algorithms	asymptotic runtime training phase	asymptotic runtime classification phase
Stibor et al. [26]	$(2^r + \|S\|)(\ell - r + 1)$	$\|D\|\ell$
Elberfeld, Textor [12]	$\|S\|(\ell - r + 1)r^2$	$\|S\|\ell^2 r$
Present paper	$\|S\|\ell r$	ℓ

r-contiguous detector-based algorithms	asymptotic runtime training phase	asymptotic runtime classification phase
D'haeseleer et al. [10] (linear)	$(2^r + \|S\|)(\ell - r)$	$\|D\|\ell$
D'haeseleer et al. [10] (greedy)	$2^r\|S\|(\ell - r)$	$\|D\|\ell$
Wierzchón [30]	$2^r(\|D\|(\ell - r) + \|S\|)$	$\|D\|\ell$
Elberfeld, Textor [12]	$\|S\|^3\ell^3 r^3$	$\|S\|^2\ell^3 r^3$
Present paper	$\|S\|\ell r$	ℓ

Table 1: Comparison of our results with the runtimes of previously published algorithms. All runtimes are given for a binary alphabet ($|\Sigma| = 2$) since not all algorithms are applicable to arbitrary alphabets. The parameter $|D|$, the desired number of detectors, is only applicable to algorithms that generate detectors explicitly – our algorithms produce the results that would be obtained with the maximal number of generated detectors.

2.1. *Strings, Substrings and Languages*

An *alphabet* Σ is a nonempty and finite set of *symbols*. A *string* $s \in \Sigma^*$ is a sequence of symbols from Σ, and its length is denoted by $|s|$. Given an index $i \in \{1, \dots, |s|\}$, then $s[i]$ is the symbol at position i in s. Given two indices i and j, whenever $j \geq i$, then $s[i \dots j]$ is the *substring* of s with length $j - i + 1$ that starts at position i and whenever $j < i$, then $s[i \dots j]$ is the empty string. If $i = 1$, then $s[i \dots j]$ is a *prefix* of s and, if $j = |s|$, then $s[i \dots j]$ is a *suffix* of s. Given a string $s \in \Sigma^\ell$ and another string $d \in \Sigma^j$ with $1 \leq j \leq \ell$ and an index $i \in \{1, \dots, \ell - j + 1\}$, we say that d *occurs in s at position i* if $s[i \dots i + j - 1] = d$.

A set of strings $S \subseteq \Sigma^*$ is called a *language*. For two indices i and j, we define $S[i \dots j] = \{s[i \dots j] \mid s \in S\}$. We say that S *avoids a string d at position i* if d occurs in no $s \in S$ at position i. Alternatively, we say that S *avoids the tuple* (d, i).

2.2. *Prefix Trees, Prefix DAGs, and Automata*

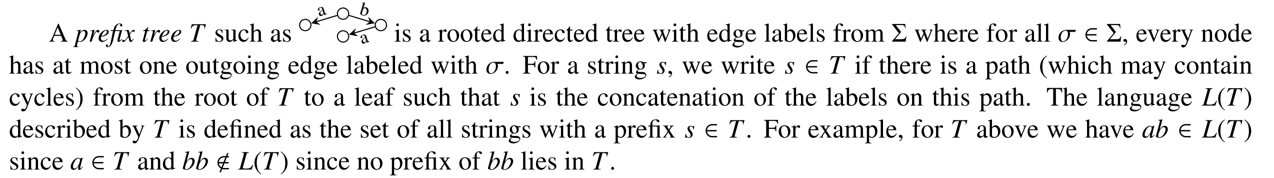

A *prefix tree* T such as [diagram] is a rooted directed tree with edge labels from Σ where for all $\sigma \in \Sigma$, every node has at most one outgoing edge labeled with σ. For a string s, we write $s \in T$ if there is a path (which may contain cycles) from the root of T to a leaf such that s is the concatenation of the labels on this path. The language $L(T)$ described by T is defined as the set of all strings with a prefix $s \in T$. For example, for T above we have $ab \in L(T)$ since $a \in T$ and $bb \notin L(T)$ since no prefix of bb lies in T.

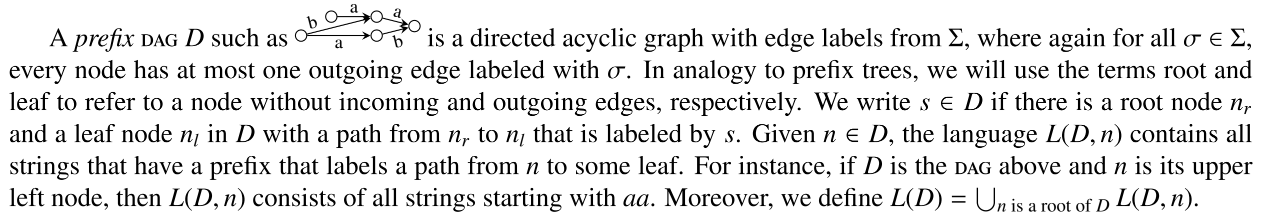

A *prefix* DAG D such as [diagram] is a directed acyclic graph with edge labels from Σ, where again for all $\sigma \in \Sigma$, every node has at most one outgoing edge labeled with σ. In analogy to prefix trees, we will use the terms root and leaf to refer to a node without incoming and outgoing edges, respectively. We write $s \in D$ if there is a root node n_r and a leaf node n_l in D with a path from n_r to n_l that is labeled by s. Given $n \in D$, the language $L(D, n)$ contains all strings that have a prefix that labels a path from n to some leaf. For instance, if D is the DAG above and n is its upper left node, then $L(D, n)$ consists of all strings starting with aa. Moreover, we define $L(D) = \bigcup_{n \text{ is a root of } D} L(D, n)$.

We will construct finite automata to decide the membership of strings in languages. Formally, a *finite automaton* is a tuple $M = (Q, q_\text{i}, Q_\text{a}, \Sigma, \Delta)$, where Q is a *set of states* with a distinguished *initial state* $q_\text{i} \in Q$, $Q_\text{a} \subseteq Q$ the set of *accepting states*, Σ the *alphabet* of M, and $\Delta \subseteq Q \times \Sigma \times Q$ the *transition relation*. Furthermore, we assume that the transition relation is *unambiguous*: for every $q \in Q$ and every $\sigma \in \Sigma$ there is at most one $q' \in Q$ with $(q, \sigma, q') \in \Delta$. It is common to represent the transition relation as a graph with nodes Q (with the initial state and the accepting states highlighted properly) and labeled edges (with a σ-labeled edge from q to q' if $(q, \sigma, q') \in \Delta$.) An automaton M is said to *accept* a string s if its graph contains a path from q_i to some $q \in Q_\text{a}$ whose concatenated edge labels equal s. The language $L(M)$ contains all strings accepted by M. Note that every prefix DAG D can be turned into a

finite automaton M with $L(D) = L(M)$. For a more detailed discussion of automata-based string processing, we refer to the textbook of Crochemore, Hancart and Lecroq [7].

2.3. Detectors and Self-Nonself-Discrimination

We fix an alphabet Σ, a string length ℓ, a *self-set* $S \subseteq \Sigma^\ell$, and a matching parameter $r \in \{1, \dots, \ell\}$.

Definition 2.1 (r-chunk detector). An *r-chunk detector* (d, i) is a tuple of a string $d \in \Sigma^r$ and an index $i \in \{1, \dots, \ell - r + 1\}$. It *matches* a string s if d occurs in s at position i.

The *set of r-chunk detectors for S*, denoted by $\textsc{chunk}(S, r)$, contains exactly the r-chunk detectors (d, i) that do *not* match any string in S. Let $m \in \Sigma^\ell$. The string m is *nonself with respect to S and r-chunk detectors* if m matches an r-chunk detector from $\textsc{chunk}(S, r)$ and *self*, otherwise. The set $\textsc{chunk-nonself}(S, r)$ contains exactly the strings of length ℓ over Σ that are nonself with respect to S and r-chunk detectors.

Definition 2.2 (r-contiguous detector). An *r-contiguous detector* is a string $d \in \Sigma^\ell$. It *matches* a string $s \in \Sigma^\ell$ if there is an index $i \in \{1, \dots, \ell - r + 1\}$ where $d[i \dots i + r - 1]$ occurs in s.

Similarly to the chunk detector case, we define the *set of r-contiguous detectors for S and r*, $\textsc{cont}(S, r)$, as the set of all r-contiguous detectors that do not match any string in S. Let $m \in \Sigma^\ell$. The string m is *nonself with respect to S and r-contiguous detectors* if m matches an r-contiguous detector from $\textsc{cont}(S, r)$ and *self*, otherwise. The set $\textsc{cont-nonself}(S, r)$ contains exactly the strings that are nonself with respect to S and r-contiguous detectors.

Figure 2 from the introduction shows an example of a self-set S, the corresponding detector sets $\textsc{chunk}(S, 3)$ and $\textsc{cont}(S, 3)$, and the corresponding partitions of the shape space into self and nonself.

3. Negative Selection with Chunk Detectors

In this section, we discuss how to construct automata for $\textsc{chunk-nonself}(S, r)$. The construction is a rather straightforward combination of two standard string processing tools: prefix trees and failure links. We present the construction here for sake of completeness and because we use it as a building block for the more intricate one in the next section.

Theorem 3.1. *There exists an algorithm that, given any $S \subseteq \Sigma^\ell$ and $r \in \{1, \dots, \ell\}$, constructs a finite automaton M with $L(M) \cap \Sigma^\ell = \textsc{chunk-nonself}(S, r)$ in time $O(|S|\ell r|\Sigma|)$.*

Proof. By definition, a string $m \in \Sigma^\ell$ lies in the set $\textsc{chunk-nonself}(S, r)$ exactly if S avoids $(m[i \dots i + r - 1], i)$ for some index $i \in \{1, \dots, \ell - r + 1\}$. Therefore, to classify m in time $O(\ell r)$, it suffices to construct, for every position $i \in \{1, \dots, \ell - r + 1\}$ independently, a prefix tree T_i with $L(T_i) \cap \Sigma^r = \Sigma^r \setminus S[i, \dots, i + r - 1]$. The prefix tree T_i can be constructed as follows: Start with an empty prefix tree and insert every $s \in S[i, \dots, i + r - 1]$ into it. Next, for every non-leaf node n and every $\sigma \in \Sigma$ where no edge with label σ starts at n, create a new leaf n' and an edge (n, n') labeled with σ. Finally, delete every node from which none of the newly created leaves is reachable. For the resulting prefix tree we have $L(T_i) \cap \Sigma^r = \Sigma^r \setminus S[i, \dots, i + r - 1]$. It is readily seen that for every $s \in T_i$, all of its prefixes s' with $|s'| < |s|$ occur in S at position i.

To enable a classification in time $O(\ell)$, we construct an automaton by inserting failure links between the prefix trees of adjacent levels, similar to the well-known algorithm of Knuth, Morris and Pratt [18]. Briefly, the idea of our failure link method is as follows: If a mismatch occurs in a prefix tree T_i at a position k, then we need not restart from the root of tree T_{i+1}, but can go directly to the node in T_{i+1} that corresponds to the last $k-1$ symbols read. By inserting the failure links from right to left, turning the prefix trees into a prefix $\textsc{dag}$, we can inductively ensure that either such a node exists or there is no match at all.

We start by letting D be the disjoint union of $T_1, \dots, T_{\ell-r+1}$. Then we process the levels from $i = \ell - r$ down to 1 iteratively as follows: Consider every node n from T_i and every symbol $\sigma \in \Sigma$ where T_i has no outgoing edge with label σ. Let s be the string on the path from the root of T_i to n. Let $s' = s\sigma$ and let n' be the end node of the path from the root of T_{i+1} that is labeled by $s'[2 \dots |s'|]$. If this n' exists, we insert an edge from n to n' with label σ. By induction one can show that after every iteration i we have $L(T_i) \cap \Sigma^{l-i+1} = \textsc{chunk-nonself}(S[i, \dots, \ell], r)$. Finally, we

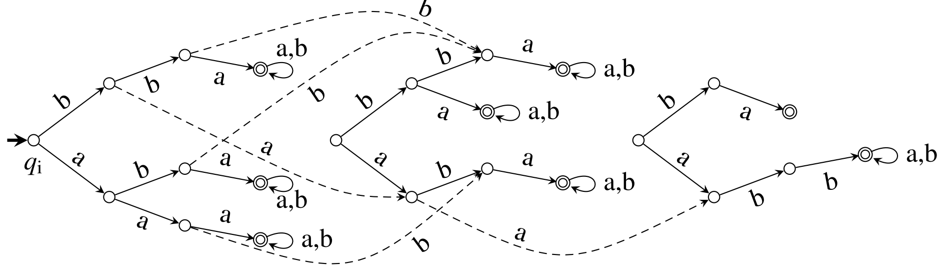

Figure 3: The constructed automaton M with $L(M) \cap \{a, b\}^5 = \textsc{chunk-nonself}(S, 3)$ where S is the self-set from Figure 2. The solid lines are from the prefix trees T_i, the dashed lines are failure links.

turn D into a finite automaton with the claimed property by making all leaves accepting states with self-loops for all $\sigma \in \Sigma$ and setting the initial state to the root of T_1. An example of this construction is shown in Figure 3.

Each prefix tree T_i can be constructed in time $O(|S|r|\Sigma|)$. The failure links between each pair of adjacent levels i and $i+1$ can be inserted in time $|S|r|\Sigma|$ by a simultaneous recursive traversal of T_i and T_{i+1}. Since the number of levels is $\ell - r + 1$, we obtain the claimed runtime. □

4. Negative Selection with Contiguous Detectors

In this section, we show how to efficiently construct automata for the languages $\textsc{cont}(S, r)$ and $\textsc{cont-nonself}(S, r)$, respectively. We first discuss the construction of an automaton for $\textsc{cont}(S, r)$, which will prove the following theorem:

Theorem 4.1. *There exists an algorithm that, given any $S \in \Sigma^\ell$ and $r \in \{1, \ldots, \ell\}$, constructs a finite automaton M with $L(M) \cap \Sigma^\ell = \textsc{cont}(S, r)$ in time $O(|S|\ell r|\Sigma|)$.*

The construction in this section is more complex than the one in the previous section since, in order to accept $\textsc{cont}(S, r)$, it does not suffice to determine non-occuring length-r substrings for the levels independently. Instead, we need to determine non-occuring substrings that can be extended by non-occuring substrings from other levels to form strings of length ℓ – the r-contiguous detectors.

Let $S \subseteq \Sigma^\ell$, $d \in \Sigma^\ell$, $r \in \{1, \ldots, \ell\}$, and $(d', i) \in \Sigma^{\leq r} \times \{1, \ldots, \ell - r + 1\}$. The string d is an *(S, r)-avoiding right-completion of* (d', i) if (1) d' occurs in d at position i and S avoids (d', i), and (2) for all $j \in \{i + 1, \ldots, \ell - r + 1\}$, there is a string $d'' \in \Sigma^{\leq r}$ such that d'' occurs in d at position j and S avoids (d'', j). If property (2) is phrased with j ranging from 1 to $i - 1$, then d is an *(S, r)-avoiding left-completion of* (d', i). With this definition we have $d \in \textsc{cont}(S, r)$ iff there exists $(d', i) \in \Sigma^{\leq r} \times \{1, \ldots, \ell - r + 1\}$ such that d is both an (S, r)-avoiding left-completion and an (S, r)-avoiding right-completion of (d', i).

To prove Theorem 4.1, we first prove the following Lemma:

Lemma 4.2. *There exists an algorithm that, given any $S \subseteq \Sigma^\ell$ and $r \in \{1, \ldots, \ell\}$, constructs a prefix* DAG *D with roots $\rho_1, \ldots, \rho_{\ell-r+1}$ such that $L(D, \rho_i) \cap \Sigma^{\ell-i+1} = \textsc{cont}(S, r)[i \ldots \ell]$ for every $i \in \{1, \ldots, \ell - r + 1\}$ in time $O(|S|\ell r|\Sigma|)$.*

Proof. The construction of D is done in four phases, presented and discussed in the next four paragraphs. While the following proof text explains the basic ideas and their correctness, the detailed computational steps are shown by the pseudocode in Figure 4, which also provides an implementation blueprint.

For all $i \in \{1, \ldots, \ell - r + 1\}$, let T_i be prefix trees with $L(T_i) \cap \Sigma^r = \Sigma^r \setminus S[i \ldots i + r - 1]$ from the proof of Theorem 3.1; by definition we know that every r-contiguous detector contains a string at position i that occurs in T_i. However, there are still strings in T_i that do not occur in any r-contiguous detector at position i. Those are precisely the strings that have no (S, r)-avoiding left-completion or no (S, r)-avoiding right-completion.

We trim the trees $T_1, \ldots, T_{\ell-r+1}$ to obtain new trees $T_1^\mathrm{R}, \ldots, T_{\ell-r+1}^\mathrm{R}$ where every T_i^R contains exactly the strings from T_i that have (S, r)-avoiding right-completions. This holds directly for all strings from the rightmost level, so $T_{\ell-r+1}^\mathrm{R} = T_{\ell-r+1}$. We trim the other trees in a right-to-left pass from $i = \ell - r$ down to 1. Each time we initialize T_i^R to

Procedure CONSTRUCT-DETECTOR-DAG(S, r)

```
1   for i = 1 to ℓ − r + 1 do              | construct prefix trees
2       T_i ← prefix tree with L(T_i) ∩ Σ^r = Σ^r \ S[i ... i + r − 1]
3   T^R_{ℓ−r+1} ← T_{ℓ−r+1}                 | trim the trees in a right-to-left pass
4   for i = ℓ − r down to 1 do
5       T^R_i ← empty prefix tree
6       for each string s ∈ T_i do
7           if there exists s′ ∈ T^R_{i+1} such that s[2 ... |s|] is a prefix of s′ then insert s into T^R_i
8   T^L_1 ← T^R_1                           | trim the trees in a left-to-right pass
9   for i = 2 to ℓ − r + 1 do
10      T^L_i ← empty prefix tree
11      for each string s ∈ T^R_i do
12          if there exists s′ ∈ T^L_{i−1} such that s′[2 ... |s′|] is a prefix of s then insert s into T^L_i
13  D_{ℓ−r+1} ← T^L_{ℓ−r+1}                 | weave the trees together into a prefix DAG
14  for i = ℓ − r down to 1 do
15      D_i ← disjoint union of D_{i+1} and T^L_i;  ρ_i ← root of T^L_i
16      for each string s ∈ T^L_i do
17          (n, n′, σ) ← last edge on the s-path from ρ_i in T^L_i, and its label
18          n″ ← end node of the s[2 ... |s|]-path from ρ_{i+1} in D_{i+1}
19          delete edge (n, n′, σ) from D_i and insert edge (n, n″, σ)
20  output D ← D_1                          | final prefix DAG with roots ρ_1,...,ρ_{ℓ−r+1}
```

Figure 4: For a given self-set $S \subseteq \Sigma^\ell$ and number $r \in \{1, \ldots, \ell\}$, this procedure constructs a prefix DAG D with roots $\rho_1, \ldots, \rho_{l-r+1}$ such that $L(D, \rho_i) \cap \Sigma^{\ell-i+1} = \text{CONT}(S, r)[i \ldots \ell]$ for every $i \in \{1, \ldots, \ell - r + 1\}$ in time $O(|S|\ell r|\Sigma|)$. Thus, in particular we have $L(D, \rho_1) \cap \Sigma^\ell = \text{CONT}(S, r)$.

be the empty tree. Then we consider every $s \in T_i$ and insert it into T^R_i if $s[2 \ldots |s|]$ is a prefix of some $s' \in T^R_{i+1}$. There are two potential reasons for a string $s \in T_i$ not to be contained in T^R_i: (1) It may be the case that already no string from T_{i+1} starts with $s[2 \ldots |s|]$, which, in turn, implies that a proper prefix of $s[2 \ldots |s|]$ lies in T_{i+1}. Since $s \in T_i$, all of its proper prefixes occur in S at position i and, thus, all proper prefixes of $s[2 \ldots |s|]$ occur in S at position $i + 1$. This is a contradiction and, thus, this case can never occur. (2) The second possibility is that there is a string that starts with $s[2 \ldots |s|]$ in T_{i+1}, but not in T^R_{i+1}. By induction, one can prove that this is due to the fact that $(s[2 \ldots |s|], i + 1)$ has no (S, r)-avoiding right-completion and, therefore, also (s, i) has none.

Next, we construct a set of trees $T^L_1 \ldots T^L_{\ell-r+1}$ containing only the strings that have both left- and right-completions by an analogous left-to-right pass. Thus, $L(T^L_i) \cap \Sigma^r = \text{CONT}(S, r)[i, \ldots, i + r - 1]$ holds.

Finally, we weave the trees together into a prefix DAG as follows: For the rightmost level $i = \ell - r + 1$, we set $D_{\ell-r+1} = T^L_{\ell-r+1}$, since by construction $L(T^L_{\ell-r+1}) \cap \Sigma^r = \text{CONT}(S, r)[\ell - r + 1 \ldots \ell]$. Now we prove the lemma by decreasing induction on i going from $i = \ell - r$ down to 1. For the induction step, suppose we have a prefix DAG D_{i+1} with $L(D_{i+1}) \cap \Sigma^{\ell-i} = \text{CONT}(S, r)[i + 1 \ldots \ell]$. For all $s \in T^L_i$, let n denote the corresponding leaf in T^L_i. Let n' denote the end node on the path from the root of T^L_{i+1} with label s, which exists by induction assumption because $s[2 \ldots |s|]$ is a prefix of some $d \in \text{CONT}(S, r)[i + 1 \ldots \ell]$. Create a new edge from the parent of n to n' and delete the leaf n along with all nodes and edges from which only n can be reached. After all leaves have been iterated through, let D_i be the resulting graph. Let $d \in \text{CONT}(S, r)[i \ldots \ell]$. Then d starts with a prefix from T^L_i and, thus, $d[2 \ldots |d|] \in L(D_{i+1})$. Hence, $d \in L(D_i)$ by construction. Conversely, let $d \in L(D_i)$ with $|d| = \ell - i + 1$. Then d starts with a nonempty prefix that has both an (S, r)-avoiding right-completion and an (S, r)-avoiding left-completion. Furthermore, $d[2 \ldots |d|] \in L(D_{i+1})$. Hence $d \in \text{CONT}(S, r)[i \ldots \ell]$. Now by setting $D = D_1$ we obtain a DAG with the properties claimed by the Lemma.

The runtime of the construction can be easily determined from the pseudocode given in Figure 4. As stated in Theorem 3.1, constructing the prefix trees in lines 1 and 2 takes time $O(|S|\ell r|\Sigma|)$. The inner loops in the right-to-left passes in lines 3–7 and 13–19 as well as in the left-to-right pass in lines 8–12 can be implemented by a simultaneous recursion through the trees on adjacent levels. This yields a worst-case runtime of $O(|S|\ell r|\Sigma|)$ for each of the passes and, hence, of the overall algorithm. □

Proof (Theorem 4.1). Let D with roots $\rho_1,\dots,\rho_{\ell-r+1}$ be the prefix dag from Lemma 4.2. We transform D into an automaton $M = (Q, q_\mathrm{i}, Q_\mathrm{a}, \Sigma, \Delta)$ with $L(M) \cap \Sigma^\ell = \textsc{cont}(S, r)$: For every leaf n of D and $\sigma \in \Sigma$ we append a self-loop with label σ to n. Then Q and Δ are the set of nodes and set of labeled edges, respectively, Q_a contains all former leafs, and $q_\mathrm{a} = \rho_1$. Figure 5 shows an example of such an automaton. □

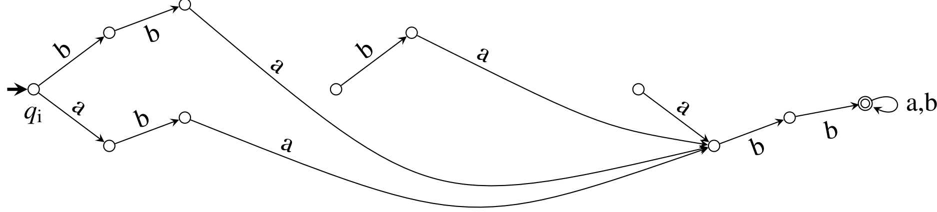

Figure 5: The constructed automaton M with $L(M) \cap \{a, b\}^5 = \textsc{cont}(S, 3)$ where S is the self-set from Figure 2.

In addition to describing the language $\textsc{cont}(S, r)$, the prefix dag D can already be used to classify a string $m \in \Sigma^\ell$ in time $O(\ell r)$: Consider every position $i \in \{1, \dots, \ell - r + 1\}$ and test whether $m[i \dots i + r - 1] \in L(D, \rho_i)$. If there exists a position where this is true, then m is "non-self" and "self", otherwise. At the end of this section we will speed up the classification to time $O(\ell)$. But first let us show how to use the prefix dag D for counting the number of detectors.

Corollary 4.3. *There exists an algorithm that, given $S \subseteq \Sigma^\ell$ and $r \in \{1, \dots, \ell\}$, outputs $|\textsc{cont}(S, r)|$ in time $O(|S|\ell r|\Sigma|)$.*

Proof. Our task is simply to count the number of strings of length ℓ in $L(D, \rho_1)$, where D is the prefix dag constructed in Lemma 4.2. First, for each node $n \in D$, compute the number of different paths leading from ρ_1 to n. Denote this quantity by $P[n]$, and let $\delta(\rho_1, n)$ denote the distance between ρ_1 and n in D (note that by construction, all paths leading from ρ_1 to n in D have the same length). Then $|\textsc{cont}(S, r)| = \sum_{n \text{ is a leaf of } D} P[n] \cdot |\Sigma|^{\ell - \delta(\rho_1, n)}$. Since D is acyclic, computing $P[n]$ can be done by a dynamic program that traverses D in breadth-first order from ρ_1. For the desired time bound note that the number of nodes and edges in D is bounded by $O(|S|\ell r|\Sigma|)$. □

Finally, let us discuss how to classify a single string in time $O(\ell)$. Again the solution is akin to failure links, however this time it is less simple than in Section 3 since the set of potential partial matches is no longer described as a set of prefix trees.

Our approach is to augment the automaton constructed by Theorem 4.1 with edge outputs. The outputs will be numbers and their partial sums will equal the lengths of maximal partial matches to r-contiguous detectors. Formally, we use Mealy automata that output numbers and define a proper language based on these outputs.

Definition 4.4. A *Mealy automaton* is a tuple $M = (Q, q_\mathrm{i}, Q_\mathrm{a}, \Sigma, \Delta, \Omega, \omega)$ where $(Q, q_\mathrm{i}, Q_\mathrm{a}, \Sigma, \Delta)$ is a finite automaton, Ω is the *output alphabet*, and $\omega : \Delta \to \Omega$ is the *output function*. Let $m \in \Sigma^*$ and $t_1, \dots, t_{|m|} \in \Delta$ be the sequence of transitions made by M for input m, then the *output of M on input m* is the string $\omega(M, m) = \omega(t_1) \dots \omega(t_{|m|}) \in \Omega^*$. If Ω is a set of numbers, we define the *r-threshold language $L(M, r)$* to be the set of strings $m \in \Sigma^*$ where there exists an $i \le |m|$ with $\sum_{j=1}^{i} \omega(m)[j] \ge r$.

Similar to finite automaton, a Mealy automaton can be represented by a graph where every edge label represents both the symbol that triggers the corresponding transition and the output of the transition. For example, for the Mealy automaton $M =$ 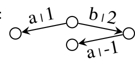

we have $ba \in L(M, 2)$ and $a \in L(M, 1)$, but $a \notin L(M, 2)$.

Theorem 4.5. *There exists an algorithm that, given any $S \in \Sigma^\ell$ and $r \in \{1, \dots, \ell\}$, constructs a Mealy automaton M with output alphabet $\Omega = \{-r, \dots, r\}$ such that $L(M, r) \cap \Sigma^\ell = \textsc{cont-nonself}(S, r)$ in time $O(|S|\ell r|\Sigma|)$.*

Proof. Let M be the finite automaton constructed in the proof of Theorem 4.1 and let $\rho_1, \dots, \rho_{\ell-r+1}$ be the roots of its underlying graph. We turn M into a Mealy automaton with output alphabet $\Omega = \{-r, \dots, r\}$ such that $L(M, r) \cap \Sigma^\ell =$

Procedure CONSTRUCT-NONSELF-MEALY-AUTOMATON(S, r)

```
1   M ← Finite automaton from Theorem 4.1 with output 1 for all transitions
2   ρ_1, ..., ρ_{ℓ−r+1} ← root nodes of M's graph
3   for i = ℓ − r down to 1 do     | insert failure links with outputs in right-to-left pass
4       for each node n reachable from ρ_i but not from ρ_{i+1} do
5           for each σ ∈ Σ where n has no outgoing σ-edge do
6               p ← path from ρ_i to n ; s ← string on p ; s' ← sσ
7               if there exists a path p' for s'[2 ... |s'|] from ρ_{i+1} then
8                   w ← sum of outputs on p ; w' ← sum of outputs on p' ; n' ← end node of p'
9                   create a transition (n, n', σ) with output w' − w
10  output M
```

Figure 6: The procedure sketched in the proof of Theorem 4.5, which transforms the finite automaton M constructed by Theorem 4.1 into a Mealy automaton with $L(M, r) \cap \Sigma^\ell$ = CONT-NONSELF(S, r). Note that the language $L(M, r)$, formalized in Definition 4.4, depends solely on the output of M, regardless its accepting states.

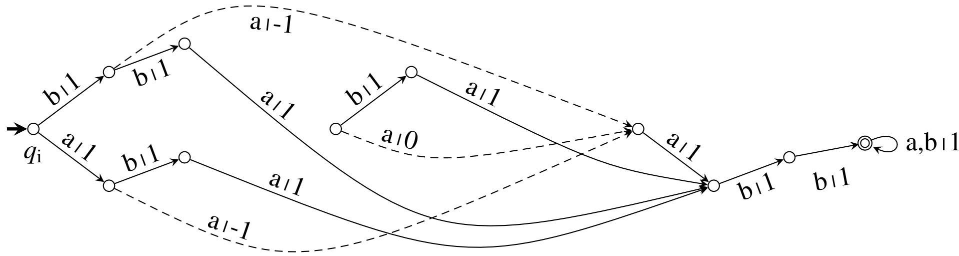

Figure 7: The Mealy automaton M with $L(M, 3)$ = CONT$(S, 3)$ where S is the self-set from Figure 2. The solid straight edges are the ones that remain from the initial prefix trees. The dashed lines are failure links, inserted to admit a linear time classification of a given string. Every edge is labeled with both the symbol that triggers the corresponding transition, and the number output of the transition.

CONT-NONSELF(S, r) holds. We describe the main ideas of the construction and discuss its correctness. For a presentation of the detailed computation steps, we refer to the pseudocode in Figure 6. An example of the constructed automaton is shown in Figure 7.

We start by assigning to all existing transitions of M the output 1. Our aim is to transform M in a right-to-left pass that inductively ensures the following property: Let $m \in \Sigma^\ell$ and let $1 \le i \le j \le \ell$. Let $k \ge 0$ denote the length of the longest suffix of $m[i \dots j]$ that is also a suffix of some $d' \in$ CONT$(S, r)[i \dots j]$. If $k \ge r - \ell + j$, then there exists a path from ρ_i for $m[i \dots j]$, and the sum of outputs on this path is equal to k. Otherwise there is no such path. Hence, if such a path exists and we have $k \ge r$, then $m \in$ CONT-NONSELF(S, r); otherwise, k is the length of the longest partial match between $m[i \dots j]$ and some $d \in$ CONT$(S, r)[i \dots j]$ that can still be extended to length $\ge r$.

The property already holds for $i = \ell - r + 1$. For i decreasing from $\ell - r$ to 1, we iteratively transform the graph of M as follows: For every node n in M that is reachable from ρ_i, but not from ρ_{i+1}, consider all $\sigma \in \Sigma$ where n has no outgoing edge labeled with σ. Let s be the string on the path p from ρ_i to n, w be the total weight on p, and $s' = s\sigma$. If there exists a path p' labeled with $s'[2 \dots |s'|]$ from ρ_{i+1}, let w' denote the sum of weights on this path. Create an edge from n to the last node of p' and label it with $w' - w$. Now there is a path from ρ_i labeled with s' with weight w', fulfilling the required property. Again, the correctness of this procedure is easily proved by induction, and we obtain a Mealy automaton with the desired property. Similarly as in Lemma 4.2, the described transformation can be implemented in time $O(|S|\ell r|\Sigma|)$ by simultaneous recursion from ρ_i and ρ_{i+1}. □

Assuming that we can compute the sum of integers in unit time, we can compute the membership test for the r-threshold language $L(M, r)$ in time $O(\ell)$ and thus obtain a negative selection algorithm with time $O(|S|\ell r|\Sigma|)$ for the training phase and time $O(\ell)$ for classifying one string. However, it is possible to get rid of the unit cost assumption

by using a finite automaton whose states store the value of the partial sums. For this construction, we would need to invest an additional runtime factor r in the training phase.

Corollary 4.6. *There exists an algorithm that, given any $S \in \Sigma^\ell$ and $r \in \{1, \ldots, \ell\}$, constructs a finite automaton M with $L(M) \cap \Sigma^\ell =$ CONT-NONSELF(S, r) in time $O(|S|\ell r^2|\Sigma|)$.*

5. Conclusions

We have shown how to construct automata that simulate the classification results of negative selection algorithms with r-contiguous and r-chunk detectors. The constructions take time $O(|S|\ell r|\Sigma|)$ and enable subsequent classification of each string in linear time $O(\ell)$. Table 1 in the introduction compares the runtimes of previously published algorithms with those from the present paper. As a corollary, our result establishes that the question if any r-contiguous detectors can be generated for a given self-set [28] can be answered in polynomial time. We leave it as an open problem whether the asymptotic time and space complexities of our constructions are optimal. It is conceivable, for example, that by applying techniques similar to those used in on-line construction of suffix trees [29] the runtime of the construction could be further improved.

References

[1] Uwe Aickelin. Special issue on artificial immune systems – editorial. *Evolutionary Intelligence*, 1(2):83–84, 2008.

[2] Modupe Ayara, Jon Timmis, Rogério de Lemos, Leandro N. de Castro, and Ross Duncan. Negative selection: How to generate detectors. In Jon Timmis and Peter J. Bentley, editors, *1st International Conference on Artificial Immune Systems*, pages 89–98, University of Kent at Canterbury, September 2002. Unversity of Kent at Canterbury Printing Unit.

[3] Justin Balthrop, Fernando Esponda, Stephanie Forrest, and Matthew Glickman. Coverage and generalization in an artificial immune system. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002)*, pages 3–10, 2002.

[4] Christopher M. Bishop. Novelty detection and neural network validation. *IEE Proceedings on Vision and Image Signal Processing*, 141:217–222, 1994.

[5] William Cavnar and John M. Trenkle. N-gram-based text categorization. In *Proceedings of SDAIR-94, 3rd Annual Symposium on Document Analysis and Information Retrieval*, pages 161–175, 1994.

[6] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58, 2009.

[7] Maxime Crochemore, Christophe Hancart, and Thierry Lecroq. *Algorithms on Strings*. Cambridge University Press, 1 edition, April 2007.

[8] Maxime Crochemore, Filippo Mignosi, and Antonio Restivo. Automata and forbidden words. *Information Processing Letters*, 67:111–117, 1998.

[9] Patrick D'haeseleer. An immunological approach to change detection: Theoretical results. In *Proceedings of the 9th IEEE Computer Security Foundations*, pages 18–26. IEEE Computer Society, 1996.

[10] Patrick D'haeseleer, Stephanie Forrest, and Paul Helman. An immunological approach to change detection: Algorithms, analysis, and implications. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 110–119. IEEE Computer Society, 1996.

[11] Ted Dunning. Statistical identification of language. Technical report, New Mexico State University, 1994.

[12] Michael Elberfeld and Johannes Textor. Efficient algorithms for string-based negative selection. In *Proceedings of the 8th International Conference on Artificial Immune Systems (ICARIS 2009)*, volume 5666 of *Lecture Notes in Computer Science*, pages 109–121. Springer, 2009.

[13] Stephanie Forrest, Steven A. Hofmeyr, Anil Somayaji, and Thomas A. Longstaff. A sense of self for unix processes. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 120–128, Washington, DC, USA, 1996. IEEE Computer Society.

[14] Stephanie Forrest, Alan S. Perelson, Lawrence Allen, and Rajesh Cherukuri. Self-nonself discrimination in a computer. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, pages 202–212. IEEE Computer Society Press, 1994.

[15] Charles Janeway, Paul Travers, Mark Walport, and Mark Shlomchick. *Immunobiology*. Garland Science, 2005.

[16] Zhou Ji and Dipankar Dasgupta. Revisiting negative selection algorithms. *Evolutionary Computation*, 15(2):223–251, 2007.

[17] Jeffrey O. Kephart. A biologically inspired immune system for computers. In *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 130–139. MIT Press, 1994.

[18] Donald E. Knuth, Jr. Morris, and Vaughan R. Pratt. Fast pattern matching in strings. *SIAM Journal on Computing*, 6(2):323–350, 1977.

[19] Maciej Liśkiewicz and Johannes Textor. Negative selection algorithms without generating detectors. In *Proceedings of Genetic and Evolutionary Computation Conference (GECCO'10)*. ACM, 2010.

[20] Emanuel Parzen. On the estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1065–1076, 1962.

[21] Jerome K. Percus, Ora E. Percus, and Alan S. Perelson. Predicting the size of the T-cell receptor and antibody combining region from consideration of efficient self-nonself discrimination. *Proceedings of the National Academy of Sciences of the United States of America*, 90(5):1691–1695, March 1993.

[22] Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.

[23] Thomas Stibor. *On the Appropriateness of Negative Selection for Anomaly Detection and Network Intrusion Detection*. PhD thesis, Darmstadt University of Technology, 2006.

[24] Thomas Stibor. Phase transition and the computational complexity of generating r-contiguous detectors. In *Proceedings of the 6th International Conference on Artificial Immune Systems (ICARIS 2007)*, volume 4628 of *Lecture Notes in Computer Science*, pages 142–155. Springer, 2007.

[25] Thomas Stibor. Foundations of r-contiguous matching in negative selection for anomaly detection. *Natural Computing*, 8:613–641, 2009.

[26] Thomas Stibor, Kpatscha M. Bayarou, and Claudia Eckert. An investigation of r-chunk detector generation on higher alphabets. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2004)*, volume 3102 of *Lecture Notes in Computer Science*, pages 299–307. Springer, 2004.

[27] Thomas Stibor, Philipp Mohr, Jonathan Timmis, and Claudia Eckert. Is negative selection appropriate for anomaly detection? In *Proceedings of the Genetic And Evolutionary Computation Conference (GECCO 2005)*, pages 321–328, 2005.

[28] Jonathan Timmis, Andrew Hone, Thomas Stibor, and Edward Clark. Theoretical advances in artificial immune systems. *Theoretical Computer Science*, 403:11–32, 2008.

[29] Esko Ukkonen. On-line construction of suffix-trees. *Algorithmica*, 14(3):249–260, 1995.

[30] Slawomir T. Wierzchoń. Generating optimal repertoire of antibody strings in an artificial immune system. In *Intelligent Information Systems*, Advances in Soft Computing, pages 119–133. Physica-Verlag, 2000.