

Lecture 10: DFA minimization


19 February 2009

In this lecture, we will see that every language has a unique minimal DFA. We will see this fact from two perspectives. First, we will see a practical algorithm for minimizing a DFA, and provide a theoretical analysis of the situation.

1 On the number of states of DFA


1.1 Starting a DFA from different states

Consider the DFA on the right. It has a particular defined start state. However, we could start it from any of its states. If the original DFA was named M , define M_q to be the DFA with its start state changed to state q . Then the language L_q , is the one accepted if you start at q .



For example, in this picture, L_3 is $(a+b)^*$, and L_6 is the same. Also, L_2 and L_5 are both $b^*a(a+b)^*$. Finally, L_7 is \emptyset .

Suppose that $L_q = L_r$, for two states q and r . Then once we get to q or r , the DFA is going to do the same thing from then on (i.e., its going to accept or reject *exactly* the same strings).



So these two states can be merged. In particular, in the above automata, we can merge 2 and 5 and the states 3 and 6. We can the new automata, depicted on the right.

1.2 Suffix Languages

Let Σ be some alphabet.

Definition 1.1 Let $L \subseteq \Sigma^*$ be any language.

The *suffix language* of L with respect to a word $x \in \Sigma^*$ is defined as

$$\llbracket L/x \rrbracket = \{y \mid xy \in L\}.$$

In words, $\llbracket L/x \rrbracket$ is the language made out of all the words, such that if we append x to them as a prefix, we get a word in L .

The *class of suffix languages* of L is

$$\mathcal{C}(L) = \left\{ \llbracket L/x \rrbracket \mid x \in \Sigma^* \right\}.$$


Example 1.2 For example, if $L = 0^*1^*$, then:

- $\llbracket L/\epsilon \rrbracket = 0^*1^* = L$
- $\llbracket L/0 \rrbracket = 0^*1^* = L$
- $\llbracket L/0^i \rrbracket = 0^*1^* = L$, for any $i \in \mathbb{N}$
- $\llbracket L/1 \rrbracket = 1^*$
- $\llbracket L/1^i \rrbracket = 1^*$, for any $i \geq 1$
- $\llbracket L/10 \rrbracket = \left\{ y \mid 10y \in L \right\} = \emptyset$.

Hence there are only three suffix languages for L : 0^*1^* , 1^* , \emptyset . So $\mathcal{C}(L) = \left\{ 0^*1^*, 1^*, \emptyset \right\}$.

As the above example demonstrates, if there is a word x , such that any word w that have x as a prefix is not in L , then $\llbracket L/x \rrbracket = \emptyset$, which implies that \emptyset is one of the suffix languages of L .

Example 1.3 The above suggests the following automata for the language of Example 1.2: $L = 0^*1^*$.



And clearly, this is the automata with the smallest number of states that accepts this language.

1.2.1 Regular languages have few suffix languages

Now, consider a DFA $M = (Q, \Sigma, \delta, q_0, F)$ accepting some language L . Let $x \in \Sigma^*$, and let M reach the state q on reading x . The suffix language $\llbracket L/x \rrbracket$ is precisely the set of strings w , such that xw is in L . But this is exactly the same as L_q . That is, $\llbracket L/x \rrbracket = L_q$, where q is the state reached by M on reading x . Hence the suffix languages of a regular language accepted by a DFA are precisely those languages L_q , where $q \in Q$.

Notice that the definition of suffix languages is more general, because it can also be applied to non-regular languages.

Lemma 1.4 *For a regular language L , the number of different suffix languages it has is bounded; that is $\mathcal{C}(L)$ is bounded by a constant (that depends on L).*

Proof: Consider the DFA $M = (Q, \Sigma, \delta, q_0, F)$ that accepts L . For any string x , the suffix language $\llbracket L/x \rrbracket$ is just the languages associated with L_q , where q is the state M is in after reading x .

Indeed, the suffix language $\llbracket L/x \rrbracket$ is the set of strings w such that $xw \in L$. Since the DFA reaches q on x , it is clear that the suffix language of x is precisely the language accepted by M starting from the state q , which is L_q . Hence, for every $x \in \Sigma^*$, $\llbracket L/x \rrbracket = L_{\delta(q_0, x)}$, where q is the state the automaton reaches on x .

As such, any suffix language of L is realizable as the language of a state of M . Since the number of states of M is some constant k , it follows that the number of suffix languages of L is bounded by k . ■

An immediate implication of the above lemma is the following.

Lemma 1.5 *If a language L has infinite number of suffix languages, then L is not regular.*

1.2.2 The suffix languages of a non-regular language

Consider the language $L = \{a^n b^n \mid n \in \mathbb{N}\}$. The suffix language of L for a^i is

$$\llbracket L/a^i \rrbracket = \{a^{n-i} b^n \mid n \in \mathbb{N}\}.$$

Note, that $b^i \in \llbracket L/a^i \rrbracket$, but this is the only string made out of only bs that is in this language. As such, for any i, j , where i and j are different, the suffix language of L with respect to a^i is different from that of L with respect to a^j (i.e. $\llbracket L/a^i \rrbracket \neq \llbracket L/a^j \rrbracket$). Hence L has infinitely many suffix languages, and hence is not regular, by Lemma 1.5.

Let us summarize what we had seen so far:

- Any state of a DFA of a language L is associated with a suffix language of L .
- If two states are associated with the same suffix language, that we can merge them into a single state.
- At least one non-regular language $\{a^n b^n \mid n \in \mathbb{N}\}$ has an infinite number of suffix languages.

It is thus natural to conjecture that the number of suffix languages of a language, is a good indicator of how many states an automata for this language would require. And this is indeed true, as the following section testifies.

2 Regular Languages and Suffix Languages

2.1 A few easy observations

Lemma 2.1 *If $\epsilon \in \llbracket L/x \rrbracket$ if and only if $x \in L$.*

Proof: By definition, if $\epsilon \in \llbracket L/x \rrbracket$ then $x = x\epsilon \in L$. Similarly, if $x \in L$, then $x\epsilon \in L$, which implies that $\epsilon \in \llbracket L/x \rrbracket$. ■

Lemma 2.2 *Let L be a language over alphabet Σ . For all $x, y \in \Sigma^*$ we have that if $\llbracket L/x \rrbracket = \llbracket L/y \rrbracket$ then for all $\mathbf{a} \in \Sigma$ we have $\llbracket L/x\mathbf{a} \rrbracket = \llbracket L/y\mathbf{a} \rrbracket$.*

Proof: If $w \in \llbracket L/x\mathbf{a} \rrbracket$, then (by definition) $xaw \in L$. But then, $\mathbf{a}w \in \llbracket L/x \rrbracket$. Since $\llbracket L/x \rrbracket = \llbracket L/y \rrbracket$, this implies that $\mathbf{a}w \in \llbracket L/y \rrbracket$, which implies that $yaw \in L$, which implies that $w \in \llbracket L/y\mathbf{a} \rrbracket$. This implies that $\llbracket L/x\mathbf{a} \rrbracket \subseteq \llbracket L/y\mathbf{a} \rrbracket$, a symmetric argument implies that $\llbracket L/y\mathbf{a} \rrbracket \subseteq \llbracket L/x\mathbf{a} \rrbracket$. We conclude that $\llbracket L/x\mathbf{a} \rrbracket = \llbracket L/y\mathbf{a} \rrbracket$. ■

2.2 Regular languages and suffix languages

We can now state a characterization of regular languages in term of suffix languages.

Theorem 2.3 (Myhill-Nerode theorem.) *A language $L \subseteq \Sigma^*$ is regular if and only if the number of suffix languages of L is finite (i.e. $\mathcal{C}(L)$ is finite).*

Moreover, if $\mathcal{C}(L)$ contains exactly k languages, we can build a DFA for L that has k states; also, any DFA accepting L must have k states.

Proof: If L is regular, then $\mathcal{C}(L)$ is a finite set by Lemma 1.4.

Second, let us show that if $\mathcal{C}(L)$ is finite, then L is regular. Let the suffix languages of L be

$$\mathcal{C}(L) = \{ \llbracket L/x_1 \rrbracket, \llbracket L/x_2 \rrbracket, \dots, \llbracket L/x_k \rrbracket \}. \quad (1)$$

Note that for any $y \in \Sigma^*$, $\llbracket L/y \rrbracket = \llbracket L/x_j \rrbracket$, for some $j \in \{1, \dots, k\}$.

We will construct a DFA whose states are the various suffix languages of L ; hence we will have k states in the DFA. Moreover, the DFA will be designed such that after reading y , the DFA will end up in the state $\llbracket L/y \rrbracket$.

The DFA is $M = (Q, \Sigma, q_0, \delta, F)$ where

- $Q = \{ \llbracket L/x_1 \rrbracket, \llbracket L/x_2 \rrbracket, \dots, \llbracket L/x_k \rrbracket \}$
- $q_0 = \llbracket L/\epsilon \rrbracket$,
- $F = \{ \llbracket L/x \rrbracket \mid \epsilon \in \llbracket L/x \rrbracket \}$. Note, that by Lemma 2.1, if $\epsilon \in \llbracket L/x \rrbracket$ then $x \in L$.
- $\delta(\llbracket L/x \rrbracket, \mathbf{a}) = \llbracket L/x\mathbf{a} \rrbracket$ for every $\mathbf{a} \in \Sigma$.

The transition function δ is well-defined because of Lemma 2.2.

We can now prove, by induction on the length of x , that after reading x , the DFA reaches the state $\llbracket L/x \rrbracket$. If $x \in L$, then $\epsilon \in \llbracket L/x \rrbracket$, which implies that $\delta(q_0, x) = \llbracket L/x \rrbracket \in F$. Thus,

$x \in L(M)$. Similarly, if $x \in L(M)$, then $\llbracket L/x \rrbracket \in F$, which implies that $\epsilon \in \llbracket L/x \rrbracket$, and by Lemma 2.1 this implies that $x \in L$. As such, $L(M) = L$.

We had shown that the DFA M accepts L , which implies that L is regular, furthermore M has k states.

We next prove that *any* DFA for L must have at least k states. So, let $N = (Q', \Sigma, \delta_N q_{\text{init}}, F)$ any DFA accepting L . The language L has k suffix languages, generated by the strings x_1, x_2, \dots, x_k , see Eq. (1).

For any $i \neq j$, we have that $\llbracket L/x_i \rrbracket \neq \llbracket L/x_j \rrbracket$. As such, there must exist a word w such that $w \in \llbracket L/x_j \rrbracket$ and $w \notin \llbracket L/x_i \rrbracket$ (the symmetric case where $w \in \llbracket L/x_i \rrbracket \setminus \llbracket L/x_j \rrbracket$ is handled in a similar fashion. But then, $x_i w \in L$ and $x_j w \notin L$. Namely, $N(q_{\text{init}}, x) \neq N(q_{\text{init}}, y)$, and the two states that N reaches for x_i and x_j respectively, are distinguishable. Formally, let $q_i = \delta(q_{\text{init}}, x_i)$, for $i = 1, \dots, k$. All these states are pairwise distinguishable, which implies that N must have at least k states. ■

Remark 2.4 The full Myhill-Nerode theorem also shows that all minimal DFAs for L are isomorphic, i.e. have identical transitions as well as the same number of states, but we will not show that part.

This is done by arguing that any DFA for L that has k states must be *identical* to the DFA we created above. This is a bit more involved notationally, and is proved by showing a 1 – 1 correspondence between the two DFAs and arguing they must be connected the same way. We omit this part of the theorem and proof.

2.3 Examples

Let us explain the theorem we just proved using examples.

2.3.1 The suffix languages of a non-regular language

Consider the language $L \subseteq \{a, b\}^*$:

$$L = \left\{ w \mid w \text{ has an odd number of } a\text{'s} \right\}.$$

The suffix language of $x \in \Sigma^*$, where x has an even number of a 's is:

$$\llbracket L/x \rrbracket = \left\{ w \mid w \text{ has an odd number of } a\text{'s} \right\} = L.$$

The suffix language of $x \in \Sigma^*$, where x has an odd number of a 's is:

$$\llbracket L/x \rrbracket = \left\{ w \mid w \text{ has an even number of } a\text{'s} \right\}.$$

Hence there are only two distinct suffix languages for L . By the theorem, we know L must be regular and the minimal DFA for L has two states. Going with the construction of the DFA mentioned in the proof of the theorem, we see that we have two states, $q_0 = \llbracket L/\epsilon \rrbracket$ and $q_1 = \llbracket L/a \rrbracket$. The transitions are as follows:

- From $q_0 = \llbracket L/\epsilon \rrbracket$, on a we go to $\llbracket L/a \rrbracket$, which is the state q_1 .

- From $q_0 = \llbracket L/\epsilon \rrbracket$, on b we go to $\llbracket L/b \rrbracket$, which is same as $\llbracket L/\epsilon \rrbracket$, i.e. the state q_0 .
- From $q_1 = \llbracket L/a \rrbracket$, on a we go to $\llbracket L/aa \rrbracket$, which is same as $\llbracket L/\epsilon \rrbracket$, i.e. the state q_0 .
- From $q_1 = \llbracket L/a \rrbracket$, on b we go to $\llbracket L/ab \rrbracket$, which is same as $\llbracket L/a \rrbracket$, i.e. the state q_1 .

The initial state is $\llbracket L/\epsilon \rrbracket$ which is the state q_0 , and the final states are those states $\llbracket L/x \rrbracket$ that have ϵ in them, which is the set $\{q_1\}$.

We hence have a DFA for L , and in fact this is the minimal automaton accepting L .

3 Minimization algorithm

The above discussion leaves us with a way to decide what is the minimum number of states of a DFA that accepts a language, but it is not clear how to turn this into an algorithm (in particular, we do not have an efficient way to compute suffix languages of a language).

The idea is to work directly on a given DFA and compute a minimum DFA from it. So consider the DFA of Figure 1. It is more complex than it needs to be.

The DFA minimization algorithm first removes any states which are not reachable from the start state, because they obviously aren't contributing anything to what the DFA accepts. (D in this example.) It then marks which of the remaining states are distinct. States not marked as distinct can then be merged, to create a simpler DFA.

3.1 Idea of algorithm

Suppose the given DFA is $M = (Q, \Sigma, \delta, q_0, F)$.

We know by the above discussion that two states p and q are distinct if their two languages are different. Namely, there is some word w that belongs to L_p but w is not in L_q (or vice versa). It is not clear however how to detect when two states have different suffix languages. The idea is to start with the “easy” case, and then propagate the information.

As such, p and q are *distinct* if there exists w , such that $\delta(p, w)$ is an accept state and $\delta(q, w)$ is not an accept state. In particular, for $w = \epsilon$, we have that p and q are distinct if $p \in F$ and $q \notin F$, or vice versa.

for $w = c_1c_2 \dots c_m$, we have that p and q are *distinct* if

$$\delta(\delta(p, c_1), c_2c_3 \dots c_m) \in F \text{ and } \delta(\delta(q, c_1), c_2c_3 \dots c_m) \notin F,$$

or vice versa.

In particular, this implies that if p and q are distinct because of word w of length m , then $\delta(p, c_1)$ and $\delta(q, c_1)$ are distinct because of a word $w' = c_2 \dots c_m$ of length $m - 1$.

Thus, its easy to compute the pairs of states distinct because of empty words, and if we computed all the states distinct because of words of length $m - 1$, we can “propagate” this information for pairs of states distinct by states of length m .

3.2 The algorithm

The algorithm for marking distinct states follows the above (recursive) definition. Create a table **Distinct** with an entry for each pair of states. Table cells are initially blank.

- (1) For every pair of states (p, q)

If p is final and q is not, or vice versa,

Set $\text{Distinct}(p, q)$ to be ϵ .

- (2) Loop until there is no change in the table contents

For each pair of states (p, q) and each character a in the alphabet:

if $\text{Distinct}(p, q)$ is empty and $\text{Distinct}(\delta(p, a), \delta(q, a))$ is not empty

Set $\text{Distinct}(p, q)$ to be a .

- (3) Two states p and q are distinct iff $\text{Distinct}(p, q)$ is not empty.

3.2.1 Example of how the algorithm works




Figure 1: (a) Original automata, (b) minimized automata.

The following is the execution of the algorithm on the DFA of Figure 1.
After step (1):

q_0								
q_1								
q_2								
q_3								
q_4								
q_5	ϵ	ϵ	ϵ	ϵ	ϵ			
q_6	ϵ	ϵ	ϵ	ϵ	ϵ			
q_7						ϵ	ϵ	
	q_0	q_1	q_2	q_3	q_4	q_5	q_6	q_7

(Note, that for a pair of states (q_i, q_j) we need only a single entry since (q_j, q_i) is equivalent, and we do not need to consider pair on the diagonal of the form (q_i, q_i) .)

q_0								
q_1								
q_2	1	1						
q_3	1	1						
q_4	0	0	0	0				
q_5	ϵ	ϵ	ϵ	ϵ	ϵ			
q_6	ϵ	ϵ	ϵ	ϵ	ϵ			
q_7			1	1	0	ϵ	ϵ	
	q_0	q_1	q_2	q_3	q_4	q_5	q_6	q_7

After one iteration of step (2)

\Rightarrow

q_0								
q_1								
q_2	1	1						
q_3	1	1						
q_4	0	0	0	0				
q_5	ϵ	ϵ	ϵ	ϵ	ϵ			
q_6	ϵ	ϵ	ϵ	ϵ	ϵ			
q_7	1	1	1	1	0	ϵ	ϵ	
	q_0	q_1	q_2	q_3	q_4	q_5	q_6	q_7

After the second iteration of step (2)

Third iteration of step (2) makes no changes to the table, so we halt. The cells (q_0, q_1) , (q_2, q_3) and (q_5, q_6) are still empty, so these pairs of states are not distinct. Merging them produces the following simpler DFA recognizing the same language.