
Notes on State Minimization

These notes present a technique to prove a lower bound on the number of states of any DFA that recognizes a given language. The technique can also be used to prove that a language is not regular. (By showing that *for every* k one needs at least k states to recognize the language.) We also present an efficient algorithm to convert a given DFA into a DFA for the same language and with a minimum number of states.

1 Distinguishable and Indistinguishable States

It will be helpful to keep in mind the following two languages over the alphabet $\Sigma = \{0, 1\}$ as examples: the language $EQ = \{0^n 1^n \mid n \geq 1\}$ of strings containing a sequence of zeroes followed by an equally long sequence of ones, and the language $A = (0 \cup 1)^* \cdot 1 \cdot (0 \cup 1)$ of strings containing a 1 in the second-to-last position.

We start with the following basic notion.

Definition 1 (Distinguishable Strings) *Let L be a language over an alphabet Σ . We say that two strings x and y are **distinguishable** with respect to L if there is a string z such that $xz \in L$ and $yz \notin L$, or vice versa.*

For example the strings $x = 0$ and $y = 00$ are distinguishable with respect to EQ , because if we take $z = 1$ we have $xz = 01 \in EQ$ and $yz = 001 \notin EQ$. Also, the strings $x = 00$ and $y = 01$ are distinguishable with respect to A as can be seen by taking $z = 0$.

On the other hand, the strings $x = 0110$ and $y = 10$ are *not* distinguishable with respect to EQ because for every z we have $xz \notin EQ$ and $yz \notin EQ$.

Exercise 2 *Find two strings that are not distinguishable with respect to A .*

The intuition behind Definition 1 is captured by the following simple fact.

Lemma 3 *Let L be a language, M be a DFA that decides L , and x and y be distinguishable strings with respect to L . Then the state reached by M on input x is different from the state reached by M on input y .*

PROOF: Suppose by contradiction that M reaches the same state q on input x and on input y . Let z be the string such that $xz \in L$ and $yz \notin L$ (or vice versa). Let us call q' the state reached by M on input xz . Note that q' is the state reached by M starting from q and given the string z . But also, on input yz , M must reach the same state q' , because M reaches state q given y , and then goes from q to q' given z . This means that M either accepts both xz and yz , or it rejects both. In either case, M is incorrect and we reach a contradiction. \square

Consider now the following generalization of the notion of distinguishability.

Definition 4 (Distinguishable Set of Strings) *Let L be a language. A set of strings $\{x_1, \dots, x_k\}$ is distinguishable if for every two distinct strings x_i, x_j we have that x_i is distinguishable from x_j .*

For example one can verify that $\{0, 00, 000\}$ are distinguishable with respect to EQ and that $\{00, 01, 10, 11\}$ are distinguishable with respect to A .

We now prove the main result of this section.

Lemma 5 *Let L be a language, and suppose there is a set of k distinguishable strings with respect to L . Then every DFA for L has at least k states.*

PROOF: If L is not regular, then there is no DFA for L , much less a DFA with less than k states. If L is regular, let M be a DFA for L , let x_1, \dots, x_k be the distinguishable strings, and let q_i be the state reached by M after reading x_i . For every $i \neq j$, we have that x_i and x_j are distinguishable, and so $q_i \neq q_j$ because of Lemma 3. So we have k different states q_1, \dots, q_k in M , and so M has at least k states. \square

Using Lemma 5 and the fact that the strings $\{00, 01, 10, 11\}$ are distinguishable with respect to A we conclude that every DFA for A has at least 4 states.

For every $k \geq 1$, consider the set $\{0, 00, \dots, 0^k\}$ of strings made of k or fewer zeroes. It is easy to see that this is a set of distinguishable strings with respect to EQ . This means that there cannot be a DFA for EQ , because, if there were one, it would have to have at least k states for every k , which is clearly impossible.

2 State Minimization

Let L be a language over an alphabet Σ . We have seen in the previous section the definition of distinguishable strings with respect to L . We say that two strings x and y are **indistinguishable**, and we write it $x \approx_L y$ if they are not distinguishable. That is, $x \approx_L y$ means that, for every string z , the string xz belongs to L if and only if the string yz does. By definition, $x \approx_L y$ if and only if $y \approx_L x$, and we always have $x \approx_L x$. It is also easy to see that if $x \approx_L y$ and $y \approx_L w$ then we must have $x \approx_L w$. In other words:

Fact 6 *The relation \approx_L is an equivalence relation over the strings in Σ^* .*

As you may remember from such classes as Math55 or CS70, when you define an equivalence relation over a set you also define a way to partition the set into a collection of subsets, called *equivalence class*. An equivalence class in Σ^* with respect to \approx_L is a set of strings that are all indistinguishable from one another, and that are all distinguishable from all the others not in the set. We denote by $[x]$ the equivalence class that contains the string x .

A fancy way of stating Lemma 5 is to say that every DFA for L must have at least as many states as the number of equivalence class in Σ^* with respect to \approx_L . Perhaps surprisingly, the converse is also true: there is always a DFA that has *precisely* as many states as the number of equivalence classes.

Theorem 7 (Myhill-Nerode) *Let L be a language over Σ . If Σ^* has infinitely many equivalence classes with respect to \approx_L , then L is not regular. Otherwise, L can be decided by a DFA whose number of states is equal to the number of equivalence classes in Σ^* with respect to \approx_L .*

PROOF: If there are infinitely many equivalence classes, then it follows from Lemma 5 that no DFA can decide L , and so L is not regular.

Suppose then that there is a finite number of equivalence class. We define an automaton that has a state for each equivalence class. The start state is the class $[\epsilon]$ and every state of the form $[x]$ for $x \in L$ is a final state.

It remains to describe the transition function. From a state $[x]$, reading the character a , the automaton moves to state $[xa]$. We need to make sure that this definition makes sense. If $x \approx_L x'$, then the state $[x]$ and the state $[x']$ are the same, so we need to verify that the state $[xa]$ and the state $[x'a]$ are also the same. That is, we need to verify that, for every string z , $xaz \in L$ if and only if $x'az \in L$; this is clearly true because x and x' are indistinguishable and so appending the string az makes xaz an element of the language L if and only if it also makes $x'az$ an element of the language.

So the automaton is well defined. Let now $x = x_1x_2 \cdots x_n$ be an input string for the automaton. The automaton starts in $[\epsilon]$, then moves to state $[x_1]$, and so on, and at the end is in state $[x_1 \cdots x_n]$; this is an accepting state if and only if $x \in L$, and so the automaton works correctly on x . \square

The Myhill-Nerode theorem shows that one can use the distinguishability method to prove optimal lower bounds on the number of states of a DFA for a given language, but it does not give an efficient way to construct an optimal DFA.

We will now see a *polynomial time* algorithm that given a DFA finds an equivalent DFA with a minimum number of states. There is even an $O(n \log n)$ algorithm for this problem, where $n = |Q| \cdot |\Sigma|$ is the size of the input, but we will not see it.

Before describing the algorithm, consider the following definition.

Definition 8 (Equivalence of states) *Let $M = (Q, \Sigma, \delta, q_0, F)$ be a DFA. We say that two states $p, q \in Q$ are **equivalent**, and we write it $p \equiv q$, if for every string $x \in \Sigma^*$ the state that M reaches from p given x is accepting if and only if the state that M reaches from q given x is accepting.*

An equivalent way to think of the definition is that if we change M so that p is the start state the language that we recognize is the same as if we change M so that q is the start state.

We can verify that \equiv is an equivalence relation among the set of states Q , and so \equiv partitions Q into a set of equivalence classes.

The intuition behind this definition is that if $p \equiv q$, then it is redundant to have two different states p and q . We would then like to compute the set of equivalence classes of Q with respect to \equiv and then construct a new automaton that has only one state for each equivalence class. The problem here is that it is not clear how to compute the relation \equiv , since it looks like we need to test infinitely many cases (one for each string $x \in \Sigma^*$) in order to verify that $p \equiv q$. Fortunately, there is a shortcut. To discuss the shortcut, we need one more definition.

Definition 9 *Let $M = (Q, \Sigma, \delta, q_0, F)$ be a DFA. We say that two states $p, q \in Q$ are **equivalent for length n** , and we write it $p \equiv_n q$, if for every string $x \in \Sigma^*$ of length $\leq n$ the state that M reaches from p given x is accepting if and only if the state that M reaches from q given x is accepting.*

In other words, if we change M so that p becomes the start state, the set of strings of length at most n that we accept is precisely the same as if we had made q be the start state.

The equivalence relations \equiv_n can be computed recursively.

- First, note that $p \not\equiv_0 q$ if and only if p is accepting and q is not (or vice versa). That is, the equivalence classes for \equiv_0 are the set F and the set $Q - F$.

- Suppose we have computed \equiv_{n-1} . Now we have that $p \equiv_n q$ if and only if
 - $p \equiv_{n-1} q$ and
 - for every $a \in \Sigma$ we have $\delta(p, a) \equiv_{n-1} \delta(q, a)$.

The second part requires some justification, and we prove it as a lemma below.

Lemma 10 *Let $M = (Q, \Sigma, \delta, q_0, F)$ be a DFA. For any two states p, q and integer $n \geq 1$, we have that $p \not\equiv_n q$ if and only if $p \not\equiv_{n-1} q$ or there is an $a \in \Sigma$ such that $\delta(p, a) \not\equiv_{n-1} \delta(q, a)$.*

PROOF: If $p \not\equiv_n q$, then there is a string x such that M accepts when starting from p and given x , but it rejects when starting from q and given x . If the length of x is $\leq n-1$, then we have $p \not\equiv_{n-1} q$. Otherwise, let us write $x = ax'$, where a is the first character of x , and let us call $p' = \delta(p, a)$ and $q' = \delta(q, a)$. Then the string x' has length $n-1$ and it shows that $p' \not\equiv_{n-1} q'$.

For the other direction, if $p \not\equiv_{n-1} q$ then clearly also $p \not\equiv_n q$. Otherwise, if there is an a such that $\delta(p, a) \not\equiv_{n-1} \delta(q, a)$, then let x' be the string of length $\leq n-1$ that shows that $\delta(p, a) \not\equiv_{n-1} \delta(q, a)$. Then the string ax' of length $\leq n$ shows that $p \not\equiv_n q$. \square

This gives an $O(|\Sigma| \cdot |Q|)$ time algorithm to compute \equiv_n given \equiv_{n-1} , and so an $O(n \cdot |\Sigma| \cdot |Q|)$ time algorithm to compute \equiv_n from scratch.

Now we come to an important observation: If, for some k , we have that the relation \equiv_k is the same as the relation \equiv_{k+1} , then also \equiv_k is the same as \equiv_n for all $n > k$. To see that it is true, note that the process that we use to go from \equiv_k to \equiv_{k+1} is independent of k , and so if it does not change the relation when applied once, it will not change the relation if applied an arbitrary number of times.

This shows that the algorithm converges to a fixed partition in $\leq |Q| - 1$ steps, because it starts with two equivalence classes, it cannot create more than $|Q|$ equivalence classes, and at each step it must either increase the number of equivalence classes or reach the final partition.

In particular, if we let $k = |Q| - 1$, then we have that \equiv_k is the same as \equiv_n for every $n \geq k$, and it can be computed in time at most $O(|Q|^2 \cdot |\Sigma|)$. We can also see that \equiv_k is the same relation as \equiv . This is true because if $p \equiv q$ then certainly $p \equiv_k q$. But also, if $p \not\equiv q$, then there is a string x that shows this is the case, and if we call n the length of x we have that $p \not\equiv_n q$. If $n \leq k$, then certainly this also means that $p \not\equiv_k q$; if $n > k$ we can still say that $p \not\equiv_k q$ because we have observed above that \equiv_n and \equiv_k are the same for $n > k$.

We can now finally describe our state minimization algorithm. Given $M = (Q, \Sigma, \delta, q_0, F)$.

- Let $k = |Q| - 1$ and compute the equivalence classes of Q with respect to \equiv_k . Define a new automaton $M' = (Q', \Sigma, \delta', q'_0, F')$ as follows. There is a state in Q' for each equivalence class. The initial state q'_0 is the equivalence class $[q_0]$. The set F' contain all the equivalence classes that contain final states in F .¹ Define $\delta'([q], a) = [\delta(q, a)]$.
- Remove from Q' all the states that are not reachable from q'_0 . Such states can be easily found doing a depth-first search of the graph of the automaton starting from q'_0 . The removal of these states does not change the language accepted by automaton because they never occur in a computation. Let $M'' = (Q'', \Sigma, \delta'', q''_0, F'')$ be the new automaton.

¹Note that an equivalence class contains either only states in F or only states not in F .

The reader should verify that our definition of δ' makes sense and that the automaton M' decides the same language as M .

To see that the algorithm constructs an optimal automaton, let $t = |Q''|$ be number of states of the automaton. Every state $[q]$ of the automaton is reachable from $[q_0]$, so for every state $[q]$ there is at least a string $x_{[q]}$ such that M' reaches $[q]$ when given in input $x_{[q]}$.

Consider now two different states $[p] \neq [q]$. We have $p \neq q$ in M , and so there must be some string y such that starting from p in M we accept when given y , but starting from q we reject (or vice versa). Also, in M'' , starting from $[p]$ we accept when given y but starting from $[q]$ we reject. This means that y shows that the strings $x_{[q]}$ and $x_{[p]}$ are distinguishable, and so we have a set of t distinguishable strings, which implies that it is impossible to decide the same language with fewer than t states.