

St. Francis Xavier University
Department of Computer Science
CSCI 541: Theory of Computing
Lecture 4: Decidability and Undecidability
Fall 2021

1 Decision Problems

One important aspect of studying various models of computation is determining precisely what kinds of questions about that model can be answered algorithmically. For example, does there exist an algorithm that can tell us whether or not a deterministic finite automaton accepts a given input word?

In our previous lecture, we introduced the notions of decidable and semidecidable languages. Using the terminology we've developed for defining and reasoning about languages, we can talk more generally not just about decidable languages, but about *decidable problems* for our various models of computation.

A *decision problem* is a problem for which each input instance corresponds to either a “yes” or a “no” answer. At its core, a decision problem is a language, and the elements of the language are instances with “yes” answers. Each element of the language is an encoding of whatever model of computation we're considering, in some cases with an input word given to that model: in the case of regular languages, we use encodings of finite automata, and so on.

If there exists a *decision algorithm* that produces the correct “yes” or “no” answer for any instance of a given decision problem, and the algorithm halts on all inputs, then we say that the decision problem is *decidable*. If no such decision algorithm exists, then we say that the decision problem is *undecidable*.¹

There are a number of common decision problems that we can ask about a model of computation X . The main decision problems we will focus on here are as follows:

- the *membership problem*: $A_X = \{\langle \mathcal{B}, w \rangle \mid \mathcal{B} \text{ is an } X \text{ that accepts input word } w\}$;
- the *emptiness problem*: $E_X = \{\langle \mathcal{B} \rangle \mid \mathcal{B} \text{ is an } X \text{ and } L(\mathcal{B}) = \emptyset\}$;
- the *universality problem*: $U_X = \{\langle \mathcal{B} \rangle \mid \mathcal{B} \text{ is an } X \text{ and } L(\mathcal{B}) = \Sigma^*\}$; and
- the *equivalence problem*: $EQ_X = \{\langle \mathcal{B}, \mathcal{C} \rangle \mid \mathcal{B} \text{ and } \mathcal{C} \text{ are both } X \text{ and } L(\mathcal{B}) = L(\mathcal{C})\}$.

In this lecture, then, we will get our first taste of “programming” a Turing machine to solve these decision problems, without having to construct the machine explicitly. We no longer need to specify *exactly* how the Turing machine is constructed, since the Church–Turing thesis tells us that any function that can be computed by an algorithm can also be computed by a Turing machine.

1.1 Decidable Problems for Regular Languages

We begin by considering our common decision problems applied to models of computation that recognize the class of regular languages. Regular languages (and the associated models that recognize such languages) are very useful for practical applications since, as we will see, each of the common decision problems are decidable for this class. The downside, of course, is that the class of regular languages is much smaller than the other language classes we know, which in turn limits our expressive power.

¹Note that “undecidable” doesn't mean the decision problem is impossible; “undecidable” only means “not decidable”, as in “no decision algorithm exists that *always* halts and gives a yes/no answer”. In this sense, “undecidable” has a meaning closer to our notion of semidecidability, where an algorithm may exist that gives a yes/no answer but doesn't halt on all inputs.

Membership Problem

For our first result, we will focus on the membership problem, and we will show that there exists an algorithm that allows us to determine whether or not some deterministic finite automaton \mathcal{B} accepts an input word w .

The technique we will apply in the proof of this result (and others) involves the use of a Turing machine to simulate the computation of the finite automaton. Then, whether the finite automaton accepts or doesn't accept the input word, the Turing machine returns the same result.

Theorem 1. A_{DFA} is decidable.

Proof. Construct a Turing machine \mathcal{M}_{ADFA} that takes as input $\langle \mathcal{B}, w \rangle$, where \mathcal{B} is a deterministic finite automaton and w is the input word to \mathcal{B} , and performs the following steps:

1. Simulate \mathcal{B} on input w .
2. If the simulation ends in an accepting state of \mathcal{B} , then accept. Otherwise, reject. \square

Since we know also that we can convert from nondeterministic finite automata to deterministic finite automata, and from regular expressions to deterministic finite automata, we get similar positive decidability results for these models.

Corollary 2. A_{NFA} is decidable.

Proof. Given a nondeterministic finite automaton \mathcal{B} , convert it to an equivalent deterministic finite automaton \mathcal{B}' and run \mathcal{M}_{ADFA} from the proof of Theorem 1 on input $\langle \mathcal{B}', w \rangle$. \square

Corollary 3. A_{RE} is decidable.

Proof. Given a regular expression R , convert it to an equivalent deterministic finite automaton \mathcal{S} and run \mathcal{M}_{ADFA} from the proof of Theorem 1 on input $\langle \mathcal{S}, w \rangle$. \square

Since a decision problem being decidable for the class DFA implies that it is also decidable for the classes NFA and RE, we will only focus on proofs for the class DFA going forward.

Emptiness Problem

Let's now turn to the emptiness problem. What does it mean for the language of a finite automaton to be empty? If there exists no path from the initial state of the finite automaton to a final state, then the finite automaton can't accept any words. In this case, then, its language will be empty.

We will use this idea in the algorithm to decide the emptiness problem for deterministic finite automata. Since every finite automaton has a finite set of states, we can traverse the transitions of the finite automaton starting from the initial state and mark each state as we encounter it. If, by the end of this traversal, we never mark a final state, then this implies there exists no path from the initial state to any final state.

Theorem 4. E_{DFA} is decidable.

Proof. Construct a Turing machine \mathcal{M}_{EDFA} that takes as input $\langle \mathcal{B} \rangle$, where \mathcal{B} is a deterministic finite automaton, and performs the following steps:

1. Mark the initial state of \mathcal{B} .
2. Repeat until no new states are marked:
 - (a) Mark all states that have an incoming transition from any previously-marked state.
3. If no final state of \mathcal{B} is marked, then accept. Otherwise, reject. \square

Universality Problem

We now move on to considering the universality problem, which is closely related to the emptiness problem. In fact, the two problems are complementary: if $L = \Sigma^*$, then $\overline{L} = \emptyset$, where \overline{L} denotes the *complement* of the language L . Therefore, all we need to decide the universality problem is a way of complementing the language of a deterministic finite automaton. Then, we can simply reuse the Turing machine $\mathcal{M}_{\text{EDFA}}$ that we specified earlier.

Fortunately, it is quite straightforward to show that the class of languages recognized by deterministic finite automata is closed under complement.

Lemma 5. *The class DFA is closed under the operation of complement.*

Proof. Suppose we are given a deterministic finite automaton $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$. The language of this automaton, $L(\mathcal{A})$, consists of all words that take us from the initial state q_0 to a final state in the subset F . Therefore, the complement of $L(\mathcal{A})$ consists of all words that *do not* take us to a final state.

We construct a finite automaton $\mathcal{A}' = (Q, \Sigma, \delta, q_0, F')$ recognizing the language $\overline{L(\mathcal{A})}$ by taking $F' = Q \setminus F$; that is, all non-final states in \mathcal{A} are final states in \mathcal{A}' , and vice versa. \square

With our procedure to complement the language of a deterministic finite automaton, we can now decide the universality problem using the same algorithm that we constructed to decide the emptiness problem. In order to decide whether $L(\mathcal{B}) = \Sigma^*$, we simply decide whether $\overline{L(\mathcal{B})} = \emptyset$.

Theorem 6. *U_{DFA} is decidable.*

Proof. Construct a Turing machine $\mathcal{M}_{\text{UDFA}}$ that takes as input $\langle \mathcal{B} \rangle$, where \mathcal{B} is a deterministic finite automaton, and performs the following steps:

1. Convert \mathcal{B} to a deterministic finite automaton \mathcal{B}' recognizing the language $\overline{L(\mathcal{B})}$ using the construction from the proof of Lemma 5.
2. Run $\mathcal{M}_{\text{EDFA}}$ from the proof of Theorem 4 on input $\langle \mathcal{B}' \rangle$.
3. If $\mathcal{M}_{\text{EDFA}}$ accepts, then accept. Otherwise, reject. \square

Equivalence Problem


Finally, we consider the equivalence problem. Given two deterministic finite automata \mathcal{B} and \mathcal{C} , how can we test whether $L(\mathcal{B}) = L(\mathcal{C})$? In theory, we could take all the words in $L(\mathcal{B})$ and give them as input to \mathcal{C} , and vice versa. However, this won't work out very well if either of $L(\mathcal{B})$ or $L(\mathcal{C})$ is infinite.

Instead, we'll make the following observation: if $L(\mathcal{B}) = L(\mathcal{C})$, then every word in both languages will belong to the intersection of the two languages. That is, no word will belong only to one of the two languages. Thus, in order to test whether two languages are equivalent, we just need to test whether the non-intersecting parts of each language are empty!

The “non-intersecting parts” of each language are more properly referred to as the *symmetric difference* of the languages. Given two languages $L(\mathcal{B})$ and $L(\mathcal{C})$, their symmetric difference is the language

$$L(\mathcal{B}) \triangle L(\mathcal{C}) = (L(\mathcal{B}) \cap \overline{L(\mathcal{C})}) \cup (\overline{L(\mathcal{B})} \cap L(\mathcal{C})).$$

We can illustrate the symmetric difference of $L(\mathcal{B})$ and $L(\mathcal{C})$ using a Venn diagram as follows:



Before we continue, note that the symmetric difference is defined using three operations: union, complement, and intersection. We know that the class of regular languages is closed under union and complement, but what about intersection? Fortunately, proving closure under intersection is again straightforward, so we'll do that here.

Lemma 7. *The class DFA is closed under the operation of intersection.*

Proof. Suppose we are given two deterministic finite automata \mathcal{A} and \mathcal{B} recognizing languages $L(\mathcal{A})$ and $L(\mathcal{B})$, respectively. By De Morgan's laws, we know that

$$L(\mathcal{A}) \cap L(\mathcal{B}) = \overline{\overline{L(\mathcal{A})} \cup \overline{L(\mathcal{B})}}.$$

Since $L(\mathcal{A})$ and $L(\mathcal{B})$ are regular, we know that $\overline{L(\mathcal{A})}$ and $\overline{L(\mathcal{B})}$ are regular by closure under complement. We also know that $\overline{L(\mathcal{A})} \cup \overline{L(\mathcal{B})}$ is regular by closure under union. Therefore, $\overline{\overline{L(\mathcal{A})} \cup \overline{L(\mathcal{B})}}$ is regular again by closure under complement, and so $L(\mathcal{A}) \cap L(\mathcal{B})$ is regular. \square

Since we now know that the class of regular languages is closed under each of union, complement, and intersection, we can construct a deterministic finite automaton \mathcal{D} whose language is $L(\mathcal{D}) = L(\mathcal{B}) \Delta L(\mathcal{C})$. We will use this finite automaton \mathcal{D} in our decision algorithm for the equivalence problem.

The idea behind our decision algorithm for the equivalence problem, as we mentioned before, is to test equivalence by testing the emptiness of the symmetric difference language. If $L(\mathcal{D})$ is empty, then we know that all words belong to the intersection of $L(\mathcal{B})$ and $L(\mathcal{C})$, and therefore $L(\mathcal{B}) = L(\mathcal{C})$.

Theorem 8. *EQ_{DFA} is decidable.*

Proof. Construct a Turing machine $\mathcal{M}_{EQ_{DFA}}$ that takes as input $\langle \mathcal{B}, \mathcal{C} \rangle$, where \mathcal{B} and \mathcal{C} are deterministic finite automata, and performs the following steps:

1. Construct a deterministic finite automaton \mathcal{D} recognizing the language $L(\mathcal{D}) = L(\mathcal{B}) \Delta L(\mathcal{C})$.
2. Run \mathcal{M}_{EDFA} from the proof of Theorem 4 on input $\langle \mathcal{D} \rangle$.
3. If \mathcal{M}_{EDFA} accepts, then accept. Otherwise, reject. \square

1.2 Decidable Problems for Context-Free Languages

Moving on to the class of context-free languages, we will again consider each of the common decision problems, but this time applied to the context-free grammar model. We could alternatively consider each decision problem applied to the pushdown automaton model, but if we were to do that, we would need to manage the stack of the machine as we simulate its computation. By comparison, we can simulate a derivation using a context-free grammar simply by applying the rules of the grammar appropriately. Since we know that context-free grammars and pushdown automata are equivalent in terms of recognition power, we will make our lives easier and work with grammars here.

Membership Problem

As before, we will begin by considering the membership problem. If we're given an input $\langle G, w \rangle$, where G is a context-free grammar and w is a word over the grammar's terminal alphabet Σ , the membership problem asks whether G is capable of generating the word w .

The naïve approach to determine whether G generates w checks every possible derivation using the rules of G . However, this isn't a good approach, since we may have to check infinitely-many derivations. In fact, if we tried this approach and G actually couldn't generate w , then our decision algorithm would never halt! As a result, this approach *semidecides* the problem, but it doesn't *decide* the problem.

We must therefore ensure that we only check some finite number of derivations using the rules of G . Thinking back to our discussion of context-free grammars in Chomsky normal form, we made an important observation that will help us greatly in this regard: if a grammar in Chomsky normal form generates a word w , then the derivation of w will take $2|w| - 1$ steps. This lets us place an upper bound on the length of the derivation.

Our decision algorithm, then, will convert its grammar to Chomsky normal form and check only derivations requiring $2|w| - 1$ steps. If w can indeed be generated by G , then its derivation will be found by the algorithm.

Theorem 9. A_{CFG} is decidable.

Proof. Construct a Turing machine \mathcal{M}_{ACFG} that takes as input $\langle G, w \rangle$, where G is a context-free grammar and w is a word, and performs the following steps:

1. Convert G to an equivalent grammar G' in Chomsky normal form.
2. (a) If $|w| = 0$, then list all derivations using G' that take a single step.
(b) If $|w| \geq 1$, then list all derivations using G' that take $2|w| - 1$ steps.
3. If any of these derivations generate w , then accept. Otherwise, reject. □

With our Turing machine \mathcal{M}_{ACFG} , we can obtain an important corollary connecting the class of context-free languages to the class of decidable languages recognized by Turing machines. Since \mathcal{M}_{ACFG} is a Turing machine simulating the computation of a context-free grammar, we can "recognize" the context-free language generated by that grammar using a Turing machine.

Here, it becomes more evident why we chose context-free grammars as our model in this section instead of pushdown automata. If we converted a pushdown automaton to a nondeterministic Turing machine, then we could convert it to an equivalent deterministic Turing machine with no problems. However, we can't tell in advance whether some branch of the pushdown automaton's computation tree goes on forever without accepting, and as a result, we can't guarantee that our Turing machine decides its language. Thus, instead of worrying about this issue, we'll simply use the machine we already have to decide membership for context-free grammars, since we know that this machine is guaranteed to halt.

Corollary 10. Every context-free language is also a decidable language.

Proof. Construct a Turing machine \mathcal{M}_G that takes as input $\langle G, w \rangle$, where G is a context-free grammar and w is a word, and performs the following steps:

1. Run \mathcal{M}_{ACFG} from the proof of Theorem 9 on input $\langle G, w \rangle$.
2. If \mathcal{M}_{ACFG} accepts, then accept. Otherwise, reject. □

Emptiness Problem

For the emptiness problem, we again have a naïve algorithm to check whether $L(G) = \emptyset$ for some context-free grammar G : just try to generate all words w over the terminal alphabet Σ . This is, of course, not a good approach for the same reason as last time: we have infinitely-many words w that we must try to generate, so if the language truly were empty, our decision algorithm would never halt.

Instead, we will take a “reverse” approach to testing the emptiness of the grammar’s language. We will check, for each nonterminal in the grammar, whether that nonterminal is able to generate some sequence of terminal symbols. If this is the case, then whenever that nonterminal appears in a derivation, we know that some sequence of terminal symbols will appear later in the derivation.

Our decision algorithm for E_{CFG} works a lot like our decision algorithm for E_{DFA} , except backwards: while the algorithm for E_{DFA} marked states starting from the initial state and leading to a final state, our algorithm for E_{CFG} will mark symbols starting from the terminal symbols and returning to the start nonterminal.

Theorem 11. E_{CFG} is decidable.

Proof. Construct a Turing machine $\mathcal{M}_{E_{CFG}}$ that takes as input $\langle G \rangle$, where G is a context-free grammar, and performs the following steps:

1. Mark all terminal symbols in G .
2. Repeat until no new symbols are marked:
 - (a) If G has a rule of the form $A \rightarrow \alpha_1 \dots \alpha_k$ and each symbol $\alpha_1, \dots, \alpha_k$ has already been marked, then mark the symbol A .
3. If the start nonterminal of G has not been marked, then accept. Otherwise, reject. □

1.3 Undecidable Problems for Context-Free Languages

Unlike the situation with the class of regular languages, the class of context-free languages has *just* enough expressive power to render certain decision problems undecidable. In order to prove that a decision problem is undecidable for a certain model, we typically need to use a special technique known as a *reduction*, which we will study in greater depth later. For now, we will simply present some undecidable problems for context-free languages and informally intuit why these problems can’t be decided.

Universality Problem

When we established the decidability of the universality problem for regular languages, we relied on the fact that the class of regular languages was closed under complement, and this allowed us to use the decidability of E_{DFA} to decide U_{DFA} .

For context-free languages, we know that E_{CFG} is decidable, so we might be tempted to take a similar approach to show that U_{CFG} is decidable. However, we run into one big problem: the class of context-free languages is *not* closed under complement!

We saw earlier that, if a class of languages is closed under both union and intersection, then it must be closed under complement by De Morgan’s laws. While it is true that context-free languages are closed under union, it is not true that they’re also closed under intersection. We can reason about why this is the case with the following example:

- $L_1 = \{a^n b^n c^m \mid m, n \geq 0\}$ is context-free; and
- $L_2 = \{a^m b^n c^n \mid m, n \geq 0\}$ is also context-free; but
- $L_1 \cap L_2 = \{a^n b^n c^n \mid n \geq 0\}$ is not context-free.

Thus, the class of context-free languages is not closed under intersection, and it is therefore not closed under complement. This explains why we can’t use the same technique as we did for regular languages to show that U_{CFG} is decidable. However, it doesn’t formally prove that U_{CFG} is undecidable; we will leave that proof for later.

Theorem 12. U_{CFG} is undecidable.