Theorem 12. U_{CFG} *is undecidable.*

4.2   Equivalence Problem

Since we now know that the class of context-free languages is not closed under either intersection or complement, it seems hopeless to believe that the equivalence problem for context-free languages is decidable. Indeed, since we can test the equivalence of two languages by testing the emptiness of their symmetric difference, and since the symmetric difference is defined in terms of intersection and complement, we cannot use the same approach here as we used for the equivalence problem for regular languages.

Like before, this observation itself doesn't actually prove that EQ_{CFG} is undecidable; it only tells us that we can't use the same technique we used previously to prove the decidability of EQ_{DFA}. We will return to this result later so that we can prove it formally.

Theorem 13. EQ_{CFG} *is undecidable.*

5   Undecidable Problems for Turing Machines

We now move on to considering decision problems for Turing machines. But wait... for each of the previous two models, we started by considering decidable problems. Here, we jumped directly to undecidable problems. Did we accidentally skip over one section? No: as it turns out, Turing machines are simply so powerful that no common decision problem is decidable for this model.

But wait, since Turing machines are so powerful, doesn't that mean they can solve all kinds of problems? Indeed, that's true, but they can't *decide* many problems. Remember, in order to decide a problem, the Turing machine must always halt and either accept or reject the input word corresponding to the problem instance. The main issue encountered by the Turing machine is in the "halt" step, since there's no guarantee that the machine will halt on every input word it receives.

Indeed, even the most basic decision problems are rendered undecidable on a Turing machine, simply because of the fundamental limitation that the machine may get caught in an infinite loop during its computation.

5.1   Membership Problem

We've seen that the membership problems for regular languages and context-free languages are both decidable, by virtue of the fact that the models recognizing such classes of languages always halt and either accept and reject their input words. For Turing machines, however, we lose this valuable decidability property.

Before we continue, let's review the notions of decidability and semidecidability from the previous lecture. Suppose that $\mathcal{M}$ is a Turing machine recognizing a language L. Recall that:

- L is *decidable* if (i) whenever $w \in L$, $\mathcal{M}$ accepts w; and (ii) whenever $w \notin L$, $\mathcal{M}$ rejects w; and

- L is *semidecidable* if (i) whenever $w \in L$, $\mathcal{M}$ accepts w; and (ii) whenever $w \notin L$, $\mathcal{M}$ either rejects w or enters a loop.

It's quite straightforward to show that the membership problem for Turing machines, A_{TM}, is semidecidable.

Theorem 14. A_{TM} *is semidecidable.*

Proof. Construct a Turing machine $\mathcal{M}_{\mathrm{ATM}}$ that takes as input $\langle \mathcal{M}, w \rangle$, where $\mathcal{M}$ is a Turing machine and w is a word, and performs the following steps:

1. Simulate $\mathcal{M}$ on input w.

2. (a) If $\mathcal{M}$ ever enters its accepting state q_{accept}, then accept.

   (b) If $\mathcal{M}$ ever enters its rejecting state q_{reject}, then reject. $\square$

The machine $\mathcal{M}_{\text{ATM}}$ that we constructed in the proof of Theorem 14 looks quite similar to the earlier machine $\mathcal{M}_{\text{ADFA}}$ we constructed to decide A_{DFA}. However, $\mathcal{M}_{\text{ATM}}$ only semidecides A_{TM}, since the machine has no way of rejecting its input word if it gets caught in an infinite loop.

Note also that $\mathcal{M}_{\text{ATM}}$ is the canonical example of a universal Turing machine, since we can give it an encoding of any Turing machine and any input word, and it will simulate the computation of that Turing machine on that input word.

Now, we will prove that A_{TM} is undecidable; that is, it is semidecidable but *not decidable*. The technique we will use is essentially a proof by contradiction: we will assume that we have some machine capable of deciding A_{TM}, and then show that the existence of such a machine leads to some logical absurdity when we give a particular input word to that machine.

Theorem 15. *A_{TM} is undecidable.*

Proof. Assume by way of contradiction that A_{TM} is decidable, and suppose that $\mathcal{H}$ is a Turing machine that decides A_{TM}. Then, given an encoding of a Turing machine $\mathcal{M}$ and an input word w,

- $\mathcal{H}$ accepts $\langle \mathcal{M}, w \rangle$ if $\mathcal{M}$ accepts w; and
- $\mathcal{H}$ rejects $\langle \mathcal{M}, w \rangle$ if $\mathcal{M}$ rejects w.

We now construct a new Turing machine $\mathcal{D}$ that relies on using $\mathcal{H}$ as an integral part of its computation. The machine $\mathcal{D}$ takes as input $\langle \mathcal{M} \rangle$, where $\mathcal{M}$ is a Turing machine, and performs the following steps:

1. Run $\mathcal{H}$ on input $\langle \mathcal{M}, \langle \mathcal{M} \rangle \rangle$.

2. (a) If $\mathcal{H}$ accepts, then reject.

    (b) If $\mathcal{H}$ rejects, then accept.

In other words, $\mathcal{D}$ gives as input to $\mathcal{H}$ an encoding of the Turing machine $\mathcal{M}$ and an encoded *description* of the Turing machine $\mathcal{M}$ as the input word to $\mathcal{M}$. Then, $\mathcal{D}$ returns the opposite of whatever $\mathcal{H}$ returns. That is,

- $\mathcal{D}$ accepts $\langle \mathcal{M} \rangle$ if $\mathcal{M}$ rejects $\langle \mathcal{M} \rangle$; and
- $\mathcal{D}$ rejects $\langle \mathcal{M} \rangle$ if $\mathcal{M}$ accepts $\langle \mathcal{M} \rangle$.

Diagrammatically, $\mathcal{D}$ functions in the following way. Observe that $\mathcal{H}$ is a "black box"; we don't know exactly how it performs its computation, we only assume that it can do so.
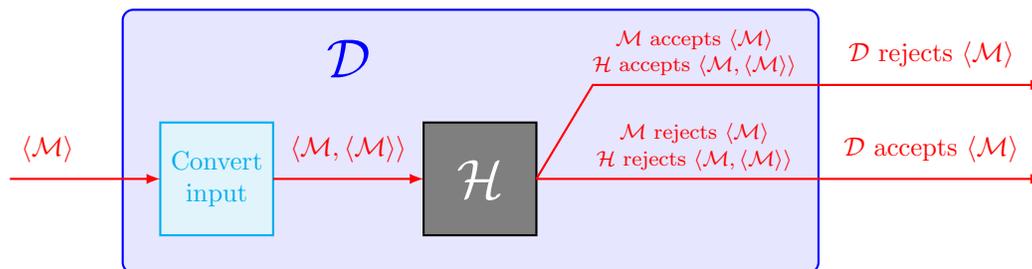
Now, consider what happens when we give $\mathcal{D}$ an encoding of *itself* as input. Suppose we give $\mathcal{D}$ the input $\langle \mathcal{D} \rangle$, so that $\mathcal{D}$ runs $\mathcal{H}$ on input $\langle \mathcal{D}, \langle \mathcal{D} \rangle \rangle$. Then, we get that

- $\mathcal{D}$ accepts $\langle \mathcal{D} \rangle$ if $\mathcal{D}$ rejects $\langle \mathcal{D} \rangle$; and
- $\mathcal{D}$ rejects $\langle \mathcal{D} \rangle$ if $\mathcal{D}$ accepts $\langle \mathcal{D} \rangle$.

In either case, $\mathcal{D}$ on input $\langle \mathcal{D} \rangle$ must return the opposite answer as the simulation of $\mathcal{D}$ on input $\langle \mathcal{D} \rangle$, which is of course impossible! As a result, we obtain a contradiction, and so no such machine $\mathcal{H}$ can exist to decide A_{TM}. $\qquad \square$

The underlying concept that makes the proof of Theorem 15 work is *Cantor's diagonal argument*. In 1891, the German mathematician Georg Cantor showed that there exist uncountably infinite sets, or sets whose elements cannot be placed into a bijection with the (countably infinite) set of natural numbers.

The crux of the diagonal argument is that, if we take a set T of all binary words and assume that it is countably infinite, then we can enumerate all the elements of T. We can then construct a new element s that doesn't belong to T by setting the ith symbol of s to be the inverse of the ith symbol in the ith word; that is, if the ith symbol of the ith word is 0, then we take the ith symbol of s to be 1, and vice versa. In this way, s differs from every other word in T in at least one position. Since s is not an element of T, s could not have been accounted for in the enumeration of T, and so T must not be countably infinite.

The same diagonal argument can be applied to Turing machines in order to prove Theorem 15. If we could enumerate all Turing machines, then we could construct a table where each row corresponds to a Turing machine $\mathcal{M}_i$ and each column corresponds to an encoding of a Turing machine $\langle \mathcal{M}_j \rangle$. Then, the entries of the table could be taken to be the result of running $\mathcal{H}$ on input $\langle \mathcal{M}_i, \langle \mathcal{M}_j \rangle \rangle$.

	$\langle \mathcal{M}_1 \rangle$	$\langle \mathcal{M}_2 \rangle$	$\langle \mathcal{M}_3 \rangle$	$\langle \mathcal{M}_4 \rangle$	$\cdots$
$\mathcal{M}_1$	accept	accept	reject	reject	
$\mathcal{M}_2$	accept	reject	reject	accept	
$\mathcal{M}_3$	reject	reject	accept	accept	
$\mathcal{M}_4$	accept	accept	accept	reject	
$\vdots$					$\ddots$

By our assumption that $\mathcal{H}$ existed, and by the fact that we constructed $\mathcal{D}$ using $\mathcal{H}$, we know that $\mathcal{D}$ must appear in this table. We also know each entry in the row corresponding to $\mathcal{D}$, since the entry in column i is the opposite of the entry at index (i, i) of the table. However, this presents a problem: what is the entry corresponding to $\langle \mathcal{D}, \langle \mathcal{D} \rangle \rangle$? It somehow needs to be the opposite of itself, and we therefore arrive at the same contradiction as before.

	$\langle \mathcal{M}_1 \rangle$	$\langle \mathcal{M}_2 \rangle$	$\langle \mathcal{M}_3 \rangle$	$\langle \mathcal{M}_4 \rangle$	$\cdots$	$\langle \mathcal{D} \rangle$	$\cdots$
$\mathcal{M}_1$	accept	accept	reject	reject		accept	
$\mathcal{M}_2$	accept	reject	reject	accept		reject	
$\mathcal{M}_3$	reject	reject	accept	accept		reject	
$\mathcal{M}_4$	accept	accept	accept	reject		accept	
$\vdots$					$\ddots$		
$\mathcal{D}$	reject	accept	reject	accept		?	
$\vdots$							$\ddots$

5.2   Non-membership Problem

Now that we know the membership problem A_{TM} is semidecidable but not decidable, we can be fairly sure that obtaining a positive decidability result for other Turing machine decision problems is hopeless. But what can we establish about variants of the Turing machine membership problem—like, say, the complement of the membership problem?

Let us consider the Turing machine "non-membership problem", defined as follows:

$$\overline{A_{\mathsf{TM}}} = \{\langle \mathcal{M}, w \rangle \mid \mathcal{M} \text{ is a Turing machine that does not accept input word } w\}.$$

Instead of proving the decidability status of $\overline{A_{\mathsf{TM}}}$ directly, as we did in proving A_{TM} was undecidable, we will use a different approach with $\overline{A_{\mathsf{TM}}}$ that focuses on semidecidability.

Consider $\overline{L}$, the complement of a language L. If the original language L is semidecidable, then we say that the complement language $\overline{L}$ is *co-semidecidable*. Saying that $\overline{L}$ is co-semidecidable is another way of saying that we can *semidecide* the *co*mplement of $\overline{L}$; in other words, we can semidecide $\overline{\overline{L}} = L$.

Since the original language L and the language of all words *not* in the complement language $\overline{L}$ are one and the same, we can use any Turing machine $\mathcal{M}$ semideciding L also to co-semidecide $\overline{L}$.

Note that $\overline{L}$ is co-semidecidable whenever L is semidecidable, but this does not necessarily imply that $\overline{L}$ is itself semidecidable. If $\overline{L}$ is semidecidable, then of course L is co-semidecidable by our definition. In the situation where L is both semidecidable and co-semidecidable, we can "upgrade" the class to which L belongs.

Theorem 16. *A language L is decidable if and only if L is both semidecidable and co-semidecidable.*

Proof. ($\Rightarrow$): Suppose L is decidable. Since all decidable languages are also semidecidable, we have that L is semidecidable. If we take the Turing machine deciding L and exchange the accepting and rejecting outputs, then we get a machine that decides $\overline{L}$. Thus, $\overline{L}$ is semidecidable, and so L is co-semidecidable.

($\Leftarrow$): Suppose L is both semidecidable and co-semidecidable. Then there exists a Turing machine $\mathcal{M}_{\text{SD}}$ semideciding L and a Turing machine $\mathcal{M}_{\text{coSD}}$ co-semideciding L. Using these two Turing machines, we can construct a Turing machine $\mathcal{M}_{\text{D}}$ that takes as input a word w and performs the following steps:

1. Run both $\mathcal{M}_{\text{SD}}$ and $\mathcal{M}_{\text{coSD}}$ on the input word w in parallel.

2. (a) If $\mathcal{M}_{\text{SD}}$ accepts, then accept.

   (b) If $\mathcal{M}_{\text{coSD}}$ accepts, then reject.

Diagrammatically, $\mathcal{M}_{\text{D}}$ functions in the following way.
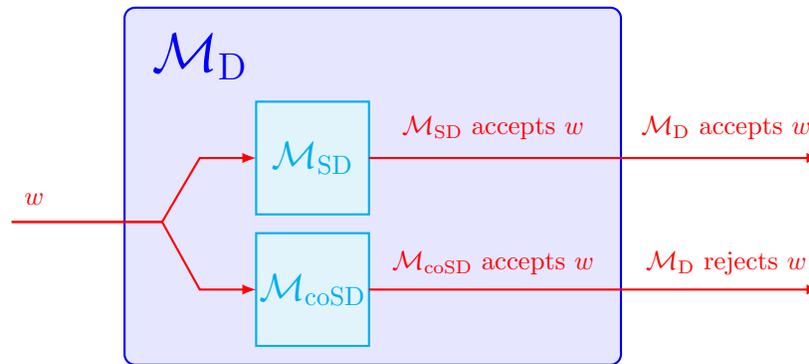
Since every word w must belong to either L or $\overline{L}$, either $\mathcal{M}_{\text{SD}}$ or $\mathcal{M}_{\text{coSD}}$ must accept w. Moreover, since $\mathcal{M}_{\text{D}}$ halts whenever either $\mathcal{M}_{\text{SD}}$ or $\mathcal{M}_{\text{coSD}}$ accepts, we have that $\mathcal{M}_{\text{D}}$ decides L, and so L is decidable. $\square$

Since Theorem 14 tells us that A_{TM} is semidecidable, we know that $\overline{A_{\text{TM}}}$ is co-semidecidable. However, if $\overline{A_{\text{TM}}}$ were not just co-semidecidable but also semidecidable, then we would have that A_{TM} is co-semidecidable. This presents a problem, since this would ultimately suggest that A_{TM} is decidable! Therefore, we can conclude the following.

Corollary 17. $\overline{A_{\text{TM}}}$ *is not semidecidable.*

Proof. Assume by way of contradiction that $\overline{A_{\text{TM}}}$ is semidecidable. Then A_{TM} would be co-semidecidable, and by Theorem 16, A_{TM} would be decidable, which contradicts Theorem 15. $\square$

Corollary 17 reveals a remarkable fact: $\overline{A_{\text{TM}}}$ is not only undecidable, but it also isn't even semidecidable, and so it lies entirely outside of our language hierarchy! In the next lecture, we'll see a number of other examples of languages that are neither decidable nor semidecidable; in some cases, these languages are co-semidecidable, while in other cases, we can't even establish co-semidecidability.

Summary

Speaking of our language hierarchy, we now have enough information to complete the "big picture" of the relationships between languages. We know by Corollary 10 that all context-free languages are also decidable. We also know by Theorem 16 that a language is decidable if and only if it is both semidecidable and co-semidecidable, so the class of decidable languages forms a subset positioned at the intersection between the classes of semidecidable and co-semidecidable languages.

Adding these new language classes to our hierarchy, we get the following diagram. Each of the decision problems we discussed in this lecture have been added to the appropriate class in this diagram, along with a number of other decision problems we will investigate further in the next lecture.