Problem Set 9

What problems are beyond our capacity to solve? Why are they so hard? And why is anything that we've discussed this quarter at all practically relevant? In this problem set – the last one of the quarter! – you'll explore the absolute limits of computing power.

Before attempting any of the problems on this problem set, we strongly recommend reading over the *Guide to Self-Reference* and *Guide to the Lava Diagram* that are available on the course website, which provide a ton of extra background that you might find useful here.

As always, please feel free to drop by office hours or ask questions on Piazza if you have any questions. We'd be happy to help out.

Good luck, and have fun!

Due Friday, March 16th at 2:30PM

Because this problem set is due on the last day of class, no late days may be used and no late submissions will be accepted. Sorry about that! On the plus side, we'll release solutions as soon as the problem set comes due.

Problem One: Isn't Everything Undecidable?

(We recommend reading the Guide to Self-Reference on the course website before attempting this problem.)

In lecture, we proved that A_{TM} and the halting problem are undecidable – that, in some sense, they're beyond the reach of algorithmic problem-solving. The proofs we used involved the nuanced technique of self-reference, which can seem pretty jarring and weird the first time you run into it. The good news is that with practice, you'll get the hang of the technique pretty quickly!

One of the most common questions we get about self-reference proofs is why you can't just use a self-reference argument to prove that *every* language is undecidable. As is often the case in Theoryland, the best way to answer this question is to try looking at some of the ways you might use self-reference to prove that *every* language is undecidable, then see where those arguments break down.

To begin with, consider this proof:

>   ***Theorem:*** All languages are undecidable.

>   ***Proof:*** Suppose for the sake of contradiction that there is a decidable language *L*. This means there's a decider for *L*; call it `inL`.

>   Now, consider the following program, which we'll call *P*:

```
int main() {
    string input = getInput();

    /* Do the opposite of what's expected. */
    if (inL(input)) {
        reject();
    } else {
        accept();
    }
}
```

>   Now, given any input *w*, either $w \in L$ or $w \notin L$. If $w \in L$, then the call to `inL(input)` will return true, at which point *P* rejects *w*, a contradiction! Otherwise, if $w \notin L$, then the call to `inL(input)` will return false, at which point *P* accepts *w*, a contradiction!

>   In both cases we reach a contradiction, so our assumption must have been wrong. Therefore, no languages are decidable. ∎

This proof has to be wrong because we know of many decidable languages.

  i.   What's wrong with this proof? Be as specific as possible.

Go one sentence at a time and check that each claim is correct. Something is fishy here.

Here's another incorrect proof that all languages are undecidable:

Theorem: All languages are undecidable.

Proof: Suppose for the sake of contradiction that there is a decidable language *L*. This means that there is some decider *D* for the language *L*, which we can represent in software as a method `willAccept`. Then we can build the following self-referential program, which we'll call *P*:

```
int main() {
    string me = mySource();
    string input = getInput();

    /* See whether we'll accept, then do the opposite. */
    if (willAccept(me, input)) {
        reject();
    } else {
        accept();
    }
}
```

Now, given any input *w*, program *P* either accepts *w* or it does not accept *w*. If *P* accepts *w*, then the call to `willAccept(me, input)` will return true, at which point *P* rejects *w*, a contradiction! Otherwise, we know that *P* does not accept *w*, so the call to `willAccept(me, input)` will return false, at which point *P* accepts *w*, a contradiction!

In both cases we reach a contradiction, so our assumption must have been wrong. Therefore, no languages are decidable. ∎

It's a nice read, but this proof isn't correct.

ii. What's wrong with this proof? Be as specific as possible.

Many of the examples we've seen of undecidable languages involve checking for properties of Turing machines or computer programs, which might give you the sense that *every* question you might want to ask about TMs or programs is undecidable. That isn't the case, though, and this question explores why.

Consider the following language *L*:

$$L = \{ \langle P \rangle \mid P \text{ is a syntactically valid C++ program } \}$$

Below is a purported proof that *L* is undecidable:

> ***Theorem:*** The language *L* is undecidable.
>
> ***Proof:*** Suppose for the sake of contradiction that *L* is decidable. That means that there's some decider *D* for *L*, which we can represent in software as a function `isSyntactically Valid` that takes as input a program and then returns whether that program has correct syntax. Given this function, consider the following program *P*:
>
> ```
> int main() {
>     string me = mySource();
>
>     /* Execute a line based on whether our syntax is right. */
>     if (isSyntacticallyValid(me)) {
>         oops, this line of code isn't valid C++!
>     } else {
>         int num = 137; // Perfectly valid syntax!
>     }
> }
> ```
>
> Now, either this program *P* is syntactically valid or it is not. If *P* has valid syntax, then when *P* is run on any input, it will get its own source code, determine that it is syntactically valid, then execute a syntactically invalid line of code – a contradiction! Otherwise, if *P* is not syntactically valid, then when *P* is run on any input, it will get its own source code, determine that it is not syntactically valid, at which point it executes a perfectly valid line of C++ code – a contradiction!
>
> In either case we reach a contradiction, so our assumption must have been incorrect. Therefore, *L* is undecidable. ∎

This proof, unfortunately, is incorrect.

iii. What's wrong with this proof? Be as specific as possible.

Problem Two: Password Checking

(We recommend reading the Guide to Self-Reference on the course website before attempting this problem.)

If you're an undergraduate here, you've probably noticed that the dorm staff have master keys they can use to unlock any of the doors in the residences. That way, if you ever lock yourself out of your room, you can, sheepishly, ask for help back in. (Not that I've ever done that or anything.) Compare this to a password system. When you log onto a website with a password, you have the presumption that your password is the only possible password that will log you in. There shouldn't be a "master key" password that can unlock any account, since that would be a huge security vulnerability. But how could you tell? If you had the source code to the password checking system, could you figure out whether your password was the only password that would grant you access to the system?

Let's frame this question in terms of Turing machines. If we wanted to build a TM password checker, "entering your password" would correspond to starting up the TM on some string, and "gaining access" would mean that the TM accepts your string. Let's suppose that your password is the string `iheartquokkas`. A TM that would work as a valid password checker would be a TM M where $\mathscr{L}(M) = \{\texttt{iheartquokkas}\}$: the TM accepts your string, and it doesn't accept anything else. Given a TM, is there some way you could tell whether the TM was a valid password checker?

Consider the following language L:

$$L = \{ \langle M \rangle \mid M \text{ is a TM and } \mathscr{L}(M) = \{\texttt{iheartquokkas}\} \}$$

Your task in this problem is to prove that L is undecidable (that is, $L \notin \mathbf{R}$). This means that there's no algorithm that can mechanically check whether a TM is suitable as a password checker. Rather than dropping you headfirst into this problem, we've split this problem apart into a few smaller pieces.

Let's suppose for the sake of contradiction that $L \in \mathbf{R}$. That means that there is some function

<div align="center">

`bool isPasswordChecker(string program)`

</div>

with the following properties:

- If `program` is the source of a program that accepts just the string `iheartquokkas`, then calling `isPasswordChecker(program)` will return `true`.

- If `program` is not the source of a program that accepts just the string `iheartquokkas`, then calling `isPasswordChecker(program)` will return `false`.

We can try to build a self-referential program that uses the `isPasswordChecker` function to obtain a contradiction. Here's a first try:

```
int main() {
    string me = mySource();
    string input = getInput();

    if (isPasswordChecker(me)) {
        reject();
    } else {
        accept();
    }
}
```

This code is, essentially, a (minimally) modified version of the self-referential program we used to get a contradiction for the language A_{TM}.

(Continued on the next page.)

    i.   Prove that the above program *P* is not a valid password checker.

    ii.   Suppose that this program is ***not*** a valid password checker. Briefly explain why no contradiction arises in this case – no formal justification is necessary.

Ultimately, the goal of building a self-referential program here is to have the program cause a contradiction regardless of whether or not it's a password checker. As you've seen in part (ii), this particular program does not cause a contradiction if it isn't a password checker. Consequently, if we want to prove that $L \notin \mathbf{R}$, we need to modify it so that it leads to a contradiction in the case where it is not a password checker.

    iii.   Modify the above code so that it causes a contradiction regardless of whether it's a password checker. Then, briefly explain why your modified program is correct. (No formal proof is necessary here; you're going to do that in the next step.)

    iv.   Formalize your argument in part (iii) by proving that $L \notin \mathbf{R}$. Use the proof that A$_{TM} \notin \mathbf{R}$ as a template for your proof.

Problem Three: L_D, Cantor's Theorem, and Diagonalization

Here's another perspective of the proof that $L_D \notin \mathbf{RE}$. Suppose we let *TM* be the set of all encodings of Turing machines. That is,

$$TM = \{ \langle M \rangle \mid M \text{ is a TM} \}$$

We can then define a function $\hat{L} : TM \to \wp(TM)$ as follows:

$$\hat{L}(\langle M \rangle) = \mathscr{L}(M) \cap TM$$

This question explores some properties of this function.

    i.   Briefly describe, in plain English, what $\hat{L}(\langle M \rangle)$ represents.

    ii.   Trace through the proof of Cantor's theorem from the Guide to Cantor's Theorem, assuming that the choice of the function f in the proof is the function $\hat{L}$. What is the set D that is produced in the course of the proof? Why?

Problem Four: Double Verification

This problem explores the following beautiful and fundamental theorem about the relationship between the **R** and **RE** languages:

$$\text{If } L \text{ is a language, then } L \in \mathbf{R} \text{ if and only if } L \in \mathbf{RE} \text{ and } \overline{L} \in \mathbf{RE}$$

This theorem has a beautiful intuition: it says that a language L is decidable ($L \in \mathbf{R}$) precisely if for every string in the language, it's possible to prove it's in the language ($L \in \mathbf{RE}$) and, simultaneously, for every string not in the language, it's possible to prove that the string is not in the language ($\overline{L} \in \mathbf{RE}$). In this problem, we're going to ask you to prove one of the two directions of this theorem.

Let L be a language where $L \in \mathbf{RE}$ and $\overline{L} \in \mathbf{RE}$. This means that there's a verifier V_{yes} for L and a verifier V_{no} for $\overline{L}$. In software, you could imagine that V_{yes} and V_{no} correspond to methods with these signatures:

```
bool imConvincedIsInL(string w, string c)
bool imConvincedIsNotInL(string w, string c)
```

Prove that $L \in \mathbf{R}$ by writing pseudocode for a function

```
bool isInL(string w)
```

that accepts as input a string w, then returns true if $w \in L$ and returns false if $w \notin L$. Then, write a brief proof explaining why your pseudocode meets these requirements. You don't need to write much code here. If you find yourself writing ten or more lines of pseudocode, you're probably missing something.

The theorem you proved in this problem is extremely useful for building an intuition for what languages are decidable. You'll see this in the next problem.

What other constructions have we done on verifiers? How did they work?

Problem Five: The Lava Diagram

Below is a Venn diagram showing the overlap of different classes of languages we've studied so far. We have also provided you a list of twelve numbered languages. For each of those languages, draw where in the Venn diagram that language belongs. As an example, we've indicated where Language 1 and Language 2 should go. No proofs or justifications are necessary – the purpose of this problem is to help you build a better intuition for what makes a language regular, **R**, **RE**, or none of these.

We ***strongly*** recommend reading over the Guide to the Lava Diagram before starting this problem.

To submit your answers, edit the file `LavaDiagram.h` in the `src/` directory of the starter files for this problem set.

1. Σ^*

2. L_D

3. $\{ a^n \mid n \in \mathbb{N} \}$

4. $\{ a^n \mid n \in \mathbb{N}$ and is a multiple of 137 $\}$

5. $\{ 1^n + 1^{m^2} = 1^{n+m} \mid m, n \in \mathbb{N} \}$

6. $\{ \langle M \rangle \mid M$ is a Turing machine and $\mathscr{L}(M) \neq \varnothing \}$

7. $\{ \langle M \rangle \mid M$ is a Turing machine and $\mathscr{L}(M) = \varnothing \}$

8. $\{ \langle M \rangle \mid M$ is a Turing machine and $\mathscr{L}(M) = L_D \}$

9. $\{ \langle M, n \rangle \mid M$ is a TM, $n \in \mathbb{N}$, and M ***accepts*** all strings in its input alphabet of length at most $n \}$

10. $\{ \langle M, n \rangle \mid M$ is a TM, $n \in \mathbb{N}$, and M ***rejects*** all strings in its input alphabet of length at most $n \}$

11. $\{ \langle M, n \rangle \mid M$ is a TM, $n \in \mathbb{N}$, and M ***loops*** on all strings in its input alphabet of length at most $n \}$

12. $\{ \langle M_1, M_2, M_3, w \rangle \mid M_1, M_2,$ and M_3 are TMs, w is a string, and at least two of $M_1, M_2,$ and M_3 accept w. $\}$

Problem Six: The Big Picture

We have covered a *lot* of ground in this course throughout our whirlwind tour of computability and complexity theory. This last question surveys what we have covered so far by asking you to see how everything we have covered relates.

Take a minute to review the hierarchy of languages we explored:

$$\textbf{REG} \subsetneq \textbf{CFL} \subsetneq \textbf{P} \overset{?}{=} \textbf{NP} \subsetneq \textbf{R} \subsetneq \textbf{RE} \subsetneq \textbf{ALL}$$

The following questions ask you to provide examples of languages at different spots within this hierarchy.

    i.   Give an example of a regular language. How might you prove that it is regular? You don't need to actually prove that it's regular – just tell us what proof technique you'd use.

    ii. Give an example of a context-free language is not regular. How might you prove that it is context-free? How might you prove that it is not regular? Just tell us what techniques you'd use for these proofs; no formal proof is actually necessary here.

    iii. Give an example of a language in **P**.

    iv. Give an example of an **NP**-complete language. *(We'll talk about this on Wednesday.)*

    v.   Give an example of a language in **RE** not contained in **R**. How might you prove that it is **RE**? How might you prove that it is not contained in **R**? Again, we just need to proof techniques you'd use, not actual formal proofs.

    vi. Give an example of a language that is not in **RE**. How might you prove it is not contained in **RE**? As before, we're just looking for proof techniques, not actual proofs.

Problem Seven: Class Participation Opt-Out

You have the option to choose to shift the 5% weight of your grade that's normally associated with your class participation onto the final exam. If you do this, you will receive no credit based on participation, and your final exam will be worth 40% of your grade rather than the regular 35%.

If you'd like to opt out of your class participation grade and shift the weight to the final, follow **this link** and fill out the form. ***Do not submit your answer through GradeScope; Keith and Cynthia won't see it***.

If you'd like to keep 5% of your grade for participation, you don't need to do anything. You're all set!

Optional Fun Problem: Quine Relays (1 Point Extra Credit)

Write four different C++ programs with the following properties:

- Running the first program prints the complete source code of the second program.

- Running the second program prints the complete source code of the third program.

- Running the third program prints the complete source code of the fourth program.

- Running the fourth program prints the complete source code of the first program.

- None of the programs perform any kind of file reading.

In other words, we'd like a collection of four different programs, each of which prints the complete source of the next one in the sequence, wrapping back around at the end. You can download starter files for this assignment from the course website and should submit your files through GradeScope.

This is actually a really fun problem to try. Once you figure out the trick, it's not that hard to code it up.

Grand Challenge Problem: P $\overset{?}{=}$ NP (Worth an A+, $1,000,000, and a Ph.D)

Prove or disprove: **P = NP**.

Take fifteen minutes and try this. Seriously. And if you can't crack this problem, feel free to submit your best effort, or the silliest answer you can think of.