

Lecture 2: Computability

MIEFis - Quantum Computation, L. S. Barbosa, 2018-19

Summary

- (1) The quest for a formal definition of algorithm.
 - (2) Turing machines. Universal Turing machines.
 - (3) Recursive functions. The Church-Turing thesis. The uncomputable.
-

1 The question

- What does it mean for a function to be computable?
- Are there any noncomputable functions?
- How does computational power depend on programming constructs?

Models	Grammars (Chomsky hierarchy)
finite memory: finite automata finite memory with stack: pushdown automata linear bounded automata unrestricted memory: Turing machines (Alan Turing) Post systems (Emil Post) μ -recursive functions (K. Gödel, J. Herbrand) λ -calculus (A. Church, S. Kleene) Combinatory logic (M. Schönfinkel, Haskell Curry)	right-linear grammars context-free grammars context-sensitive grammars unrestricted grammars

For those wondering on the use of grammars to specify computational models, notice that symbol manipulation or parsing a sentence in a language bears a strong resemblance to computation.

The quest for formalising the concept of effective computability started around the beginning of the twentieth century with the development of the formalist school of mathematics (Hilbert's programme to find a complete and consistent set of axioms for all mathematics) with the prospect of reducing all of mathematics to the formal manipulation of symbols. The formalist program was eventually shattered by Kurt Gödel's incompleteness theorem, which states that no matter how strong a deductive system for number theory you take, it will always be possible to construct simple statements that are true but unprovable. This theorem is essentially a statement

about computability. Actually, Gödel's first (there is no consistent system of axioms whose theorems can be listed by an effective procedure, i.e., an algorithm, able to prove all truths about the arithmetic of the natural numbers) and second (no formal system can prove its own consistency) incompleteness theorems, were the first of a series of results on the of formal systems, culminating on Turing's theorem that there is no algorithm to solve the halting problem.

Church-Turing thesis

All the formalisms above capture precisely the same intuition about what it means to be *effectively computable*.

2 Turing machines

$$M = (Q, \Sigma, \Gamma, \delta, \prec, \lambda, s, t, r)$$

A Turing machine consists of a finite set of states Q , a semi-infinite tape that is delimited on the left end by an endmarker \prec and is infinite to the right, and a head that can move left and right over the tape, reading and writing symbols.

The input string is of finite length and is initially written on the tape in contiguous tape cells snug up against the left endmarker \prec . The infinitely many cells to the right of the input all contain a special blank symbol λ . The machine starts in the start state s with its head scanning the left endmarker. In each step it reads the symbol on the tape under its head. Depending on that symbol and the current state, it writes a new symbol on that tape cell, moves its head either left or right one cell, and enters a new state. The action it takes in each situation is determined by a transition function

$$\delta : Q \times \Gamma \longrightarrow Q \times \Gamma \times \{L, R\}$$

The meaning of $\delta(p, a) = (q, b, d)$ is *when in state p scanning symbol a , write b on that tape cell, move the head in direction d , and enter state q .*

Restrictions:

$\forall_p \exists_q. \delta(p, \prec) = (q, \prec, R)$ $\delta(t, -) = (t, -, -)$ and $\delta(r, -) = (r, -, -)$	never moves off to the left of \prec never leave the accept (reject) state
----------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------

The Turing machine accepts its input by entering a special accept state t and rejects by entering a special reject state r . On some inputs it may run infinitely (loop on) without ever accepting or rejecting.

A configuration is a tuple in $Q \times \{w \lambda^\omega \mid w \in \Sigma^*\} \times \mathbb{N}$ and denotes a global state of the machine. The configuration $\alpha = (p, z, n)$ specifies a current state p of the finite control, current tape contents z , and current position of the read/write head ($n \geq 0$). Ex. the initial configuration on input $x \in \Sigma^*$: $(s, \prec x \lambda^\omega, 0)$.

The *transition relation*:

$$(p, z, n) \rightarrow \begin{cases} (q, z[b/n], n - 1) \Leftarrow \delta(p, z_n) = (q, b, L) \\ (q, z[b/n], n + 1) \Leftarrow \delta(p, z_n) = (q, b, R) \end{cases}$$

Example: $\{a^n b^n c^n \mid n \geq 0\}$

- Start in state s and scans to the right over the input string to check that it is of the form $a^* b^* c^*$.
- Does not write anything on the way across (formally, it writes the same symbol it reads).
- When founding the first blank symbol λ , it overwrites it with a right endmarker \succ
- Then it scans left, erasing the first c it sees, then the first b it sees, then the first a it sees, until it comes to \prec .
- Then scans right, erasing one a , one b , and one c .
- It continues to sweep left and right over the input, erasing one occurrence of each letter in each pass. If on some pass it sees at least one occurrence of one of the letters and no occurrences of another, it rejects. Otherwise, it eventually erases all the letters and makes one pass between \prec and \succ seeing only blanks, at which point it accepts.

Example: $\{ww \mid w \in \{a, b\}^*\}$

- In a first phase, scans out the input to the first blank symbol, counting the number of symbols mod 2 to make sure the input is of even length and rejecting immediately if not.
- It lays down a right endmarker \succ , then repeatedly scans back and forth over the input.
- In each pass from right to left, it marks the first unmarked a or b it sees with an overline.
- In each pass from left to right, it marks the first unmarked a or b it sees with an underline.
- It continues this until all symbols are marked. The objective is to find the center of the input string.

\prec a a b b a a a b b a λ λ λ \dots
 \prec a a b b a a a b \bar{a} λ λ λ \dots
 \prec \underline{a} a b b a a a b \bar{a} λ λ λ \dots
 \prec \underline{a} a b b a a a b \bar{b} \bar{a} λ λ λ \dots
 \prec \underline{a} \underline{a} b b a a a b \bar{b} \bar{a} λ λ λ \dots
 \prec \underline{a} \underline{a} b b a a a \bar{b} \bar{b} \bar{a} λ λ λ \dots
 \dots

< a a b b a a a b b a λ λ λ ...

In a second phase, repeatedly scans left to right over the input.

- In each pass it erases the first symbol it sees marked with underline but remembers that symbol in its finite control.
- It then scans forward until it sees the first symbol marked with overline, checks that that symbol is the same, and erases it.
- If the two symbols are not the same, it rejects. Otherwise, when it has erased all the symbols, it accepts.

< a a b b a a a b b a λ λ λ ...
 < λ a b b a λ a b b a λ λ λ ...
 < λ λ b b a λ λ b b a λ λ λ ...
 ...
 < λ λ λ λ a λ λ λ λ a λ λ λ ...
 < λ λ λ λ λ λ λ λ λ λ λ λ λ λ ...

Exercise 1

Specify a total TM that accepts its input string if its length is prime.

Hint. Give an implementation of the Sieve of Eratosthenes. To check whether n is prime, start writing down all the numbers from 2 to n in order. Then repeat: find the smallest number in the list, declare it prime, then cross off all multiples of that number. Repeat until each number in the list has been either declared prime or crossed off as a multiple of a smaller prime.

The set $L(M) = \{w \in \Sigma^* \mid M \text{ accepts } w\}$ is called the set (or *language*) accepted by Turing machine M . A Turing Machine is *total* if it halts (either by accepting or rejecting) in all inputs.

A language is

- *recursively enumerable* if it is the language $L(M)$ recognised by some Turing Machine M , i.e.

$$\begin{cases} M \text{ halts} & \Leftrightarrow w \in L \\ M \text{ halts in the non-acceptance state or loops forever} & \Leftrightarrow w \notin L \end{cases}$$

- *co-recursively enumerable* if its complement is recursively enumerable;

- *recursive* if it is the language $L(M)$ recognised by some *total* Turing Machine M , i.e.

$$\begin{cases} M \text{ halts in the acceptance state} & \Leftrightarrow w \in L \\ M \text{ halts in the rejection state} & \Leftrightarrow w \notin L \end{cases}$$

A property ϕ (over strings) is

- *decidable* if the set of all strings exhibiting ϕ is *recursive*, i.e. if there is a total Turing machine that accepts all input strings that have property ϕ and rejects those that do not.
- *semidecidable* if the set of all strings exhibiting ϕ is *recursively enumerable*, i.e. if there is a Turing machine that accepts all input strings that have property ϕ and rejects or loops if not.

Exercise 2

Given a decidable property, what is the corresponding recursive set? Conversely, which decidable property corresponds to a given recursive set?

Note: Properties (decidable) vs languages (recursive)

$$\begin{aligned} \phi \text{ is decidable} & \Leftrightarrow \{w \mid \phi(w)\} \text{ is recursive} \\ S \text{ is recursive} & \Leftrightarrow w \in S \text{ is decidable} \end{aligned}$$

Universal Turing Machines.

Turing machines are not restricted to do just a single computational task but can be *programmed* to do many different ones. Actually a Turing machine can simulate other Turing machines whose descriptions are presented as part of the input.

The key issue is to fix a reasonable encoding scheme for Turing machines over the alphabet $\{0, 1\}$, e.g.

$$0^n 10^m 10^k 10^s 10^t 10^u 10^v 1,$$

and then construct U such that

$$L(U) = \{M\#s \mid s \in L(M)\}$$

where symbol $\#$ is just a new symbol to separate M from its input. If the encodings of M and s are valid, U makes a step-by-step simulation of machine M .

This encoding scheme should be such that all the data associated with a machine M — the set of states, the transition function, the input and tape alphabets, the endmarker, the blank symbol,

and the start, accept, and reject states—can be determined easily by another machine reading the encoded description of M . For example, the expression above might indicate that the machine has n states represented by the numbers 0 to $n - 1$; it has m tape symbols represented by the numbers 0 to $m - 1$, of which the first k represent input symbols; the start, accept, and reject states are s , t and T , respectively; and the end marker and blank symbol are u and v , respectively. The remainder of the string can consist of a sequence of substrings specifying the transitions. For example, the substring

$$0^a 10^p 10^q 10^b 10$$

might indicate that δ contains the transition $((p, a), (q, b, L))$, where the direction to move the head encoded by the final digit.

We may then discuss the expressive power and limitations of Turing machines using no other instruments other than the machines themselves. It is precisely this self-referential property that Gödel exploited to embed statements about arithmetics in statements of arithmetics in his Incompleteness Theorem. The embedding in the Theorem is the same as the encoding of Turing machines into input forms acceptable for universal machines and is achieved by converting the finite description of a Turing machine into a unique non-negative integer. The conversion is possible as we are only dealing here with machines having a finite number of states, a finite number of symbols in its alphabet, and only a finite number of movements for their heads.

3 Undecidability

Our quest for a precise notion of an *algorithm* is concluded with the identification

$$\boxed{\text{Algorithm} \equiv \text{a Turing machine that halts on all inputs}}$$

This, however, opens the possibility of formally showing that there are some computational problems which cannot be solved by any algorithm.

Turing machines (like other formalisms such as finite or push-down automata, regular, context-free or unrestricted grammars) can be represented by sequences of symbols (cf, the discussion on universal Turing machines). As there is only a countable number of languages defined over an alphabet, the number of languages specified by Turing machines (cf, recursive and recursively enumerable languages) is also countable. Turing machines decide or semidecide on an infinitesimal fraction of all possible languages. According to the Church-Turing thesis, we have reached a fundamental limitation: *computational tasks that cannot be performed by a Turing machine are undecidable.*

$$\boxed{\text{The halting problem}}$$

Suppose one can write an algorithm $\text{dec}(p, x)$ which receives a program p and a value x as input and decides whether it terminates or not. Such a wonderful program could be used to write the following procedure:

Algorithm 1: DecideHalt(P).

```
1 if dec(P, P) then
  | go to 1;
else
  | halt;
end
```

Exercise 3

Explain why $\text{dec}(p, x)$ cannot exist and conclude that no algorithm to decide whether an arbitrary program would halt or loop does not exist. **Hint.** What is the result of $\text{DecideHalt}(\text{DecideHalt})$? It halts iff $\text{dec}(\text{DecideHalt}, \text{DecideHalt})$ returns false ...

We can now give a mathematically rigorous version of this paradox (somehow paradigmatic within last century's scientific culture). Note that we have now a formalised notion of an algorithm and a sort of universal programming language – the Turing machine. In this setting we can define a language which is not recursive and prove that indeed such is the case.

$$Y = \{ "Mw" \mid \text{Turing machine } M \text{ halts on input } x \}$$

Clearly, Y is recursively enumerable: it is the language recognised by an universal Turing machine. Such a machine halts exactly when its input is in Y .

Suppose that Y is decidable by some machine N . Then, given a particular Turing machine M semideciding language $L(M)$, one could design another machine that actually *decides* $L(M)$ by writing " M " w " in its input tape and then simulating N on this input. If such is the case, every recursively enumerable language would be recursive.

However, Y is not recursive. Actually, if it were recursive, language

$$Z = \{ "M" \mid \text{Turing machine } M \text{ halts on input "M"} \}$$

would also be recursive (cf. put " M " M " on the tape of a new machine and hand control to N). As it can be proved that recursive languages are closed for complements, it suffices to show that the complement of Z is not recursive.

$$\bar{Z} = \{ w \mid \Psi(w) \}$$

where $\Psi(w) = w$ is not encoding of a Turing machine, or it is the encoding of a Turing machine M that does not halt on " M ".

\bar{Z} , finally, is not recursively enumerable, let alone recursive. Suppose that there exists a machine K semideciding \bar{Z} . Is " K " in Z ?

- " K " $\in \bar{Z}$ iff K does not accept input " K ";

- but, K is supposed to semidecide Y ; so " K " $\in \bar{Z}$ iff K accepts input " K ".

Thus, we get an example of a non recursive language and proved an important theorem: the set of recursive languages is a strict subset of the set of recursive enumerable functions.

The Halting problem: *there is no algorithm that decides, for an arbitrary Turing machine M and input w , whether or not M accepts w .*

Other undecidable problems for Turing machines:

- does M halt on the empty tape?
- is there any string at all on which M halts?
- does M halts on every input?
- given two Turing machines, do they halt on the same input?
- does M fails to halt on input w ?

Other problems can be reduced to one of these (e.g. the tiling problem).

Exercise 4

Is it decidable whether a given Turing machine

- has at least 30 states?
- ever moves its head more than 30 tape cells away from the left?
- accepts the null string ϵ ?
- accepts every string?
- accepts a finite set?

Hint. To show that a problem is decidable amounts to give a total Turing machine that accepts exactly the positive instances and rejects the others. On the other hand, undecidable problems are identified by showing that a decision procedure for it could be used to construct a decision procedure for the halting problem, which does not exist.

- Build a Turing machine that, given the encoding of M written on its input tape, counts the number of states of M and accepts or rejects depending on whether the number is at least 30.
- Suppose it is possible to decide whether a given machine accepts ϵ . Thus the halting problem could be decided as follows: Suppose we are given a Turing machine M and string x to determine whether M halts on x . Construct a new machine N that does the following on input y : i) erases its input y ; ii) writes x on its tape (x is hard-wired in N finite control); iii) runs M on input x (M is hard-wired in N finite control); (iv) accepts if M halts on x . Note if M halts on x , then N accepts

its input y ; and if M does not halt on x , then N does not halt on y , therefore does not accept y . Moreover, this is true for every y . Thus,

$$L(N) = \begin{cases} \Sigma^* & \text{if } M \text{ halts on } x \\ \emptyset & \text{if } M \text{ does not halt on } x \end{cases}$$

Now if we know whether a given machine accepts ϵ , we could apply this decision procedure to the N , and this would tell whether M halts on x , thus obtaining a decision procedure for halting.

Exercise 5

Turing completeness: *A computational system (e.g. a programming language) is Turing complete if it can compute every Turing-computable function, i.e. if it can simulate a universal Turing machine. Is HASKELL Turing complete?*

Hint. Apply the following rule-of-thumb on Turing completeness: i) can it express a finite state machine?; ii) Does it have an unbounded amount of memory (scratch space) available to it?; iii) Is the use of memory unrestricted? Or does it impose an order — e.g. LIFO or FIFO — to restrict access?

Exercise 6

Try some examples in a Turing machine simulator on the web. Look for such simulators written in HASKELL

4 Recursive functions

Alternative approach to computability, focussed on *what is computed* rather than on an explicit *model of computation*.

- k -ary *zero* functions: $z(n_1, \dots, n_k) = 0$
- k -ary, j -*projection* functions: $\text{id}_k, j(n_1, \dots, n_k) = n_j$
- the *successor* function: $s(n) = n + 1$
- function *composition* of a k -ary function g with k m -ary functions h_i :

$$f(n_1, \dots, n_m) = g(h_1(n_1, \dots, n_m), h_2(n_1, \dots, n_m), \dots, h_k(n_1, \dots, n_m))$$

- recursive definition by a k -ary g and a $k + 2$ -ary h function:

$$\begin{aligned} f(n_1, \dots, n_k, 0) &= g(n_1, \dots, n_k) \\ f(n_1, \dots, n_k, m + 1) &= h(n_1, \dots, n_k, m, f(n_1, \dots, n_k, m)) \end{aligned}$$

Exercise 7

Show that if $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ is primitive recursive, so is function

$$\text{sum}_f(n, m) = \sum_{i \in \{0, 1, \dots, m\}} f(n, i)$$

Exercise 8

Show that the set of primitive recursive functions is countable. Use this fact to prove that most functions in \mathbb{N} are not primitive recursive, and therefore, according to Church thesis, uncomputable.

Exercise 9

Show that the factorial function, the function that computes the greatest common divisor and the prime predicate (which returns 1 if its argument is a prime number) are all primitive recursive functions.

μ -recursive functions

Let p be a $k + 1$ -ary function and define the following *unbounded iteration* scheme:

$$\mu_m[p(n_1, \dots, n_k, m) = 1] = \begin{cases} \text{the least } m \text{ such that } p(n_1, \dots, n_k, m) = 1 & \Leftarrow \text{ such a } m \text{ exists} \\ 0 & \Leftarrow \text{ otherwise} \end{cases}$$

The obvious way to compute this function is through unbounded iteration, i.e. through a `while` loop:

```

m := 0;
while p(n1, ..., nk, m) ≠ 1 do
  | m := m + 1
end
return m

```

This, however, may fail to terminate. We call function p *minimalizable* if the minimization scheme above terminates for every input.

Functions defined by the basic functions above, composition, recursion and minimization are called μ -*recursive*.

Example

Consider function $\log(b, n)$ standing for the logarithm of $n + 1$ over base $b + 2$. Formally,

$$\log(b, n) = \mu_m[\text{geq}(b + 2)^m, n + 1]$$

Exercise 10

Why are we using $b + 2$ and $n + 1$ above? Show that the minimization algorithm terminates, i.e. function

$$f(b, n, m) = \text{geq}((b + 2)^m, n + 1)$$

is minimizable.

Theorem

A function $f : \mathbb{N}^n \rightarrow \mathbb{N}$ is μ -recursive iff it is recursive, i.e. computable by a Turing machine.

Exercise 11

Suppose that g is a μ -recursive bijection over \mathbb{N} . Show that its inverse g^{-1} is μ -recursive as well.

References.

My preference on complexity theory is Papadimitriou's wonderful book [?]; reference [?] provides an interesting alternative. S. Arora and B. Barak book [?] is a more recent textbook covering recent achievements in complexity theory (including challenges from quantum computation) and putting them in the context of the classical results. Worth reading.

References.

Both the textbook of H. Lewis and D. Papadimitriou [?] or the lecture notes by D. Kozen [?] provide excellent introductions to computability and the (classical) theory of computation. A quite interesting book by N. Yanofsky [?] may help to build up the correct intuitions which often is as important as mastering the technicalities. As a side reading on Gödel's incompleteness theorem and connections to computability I suggest reference [?]. A. Hodges biography of Alan Turing [?] makes a most pleasant weekend reading (see also the book website at www.turing.org.uk).