The Pigeonhole Principle
&
Functions

Problem Set Two due
in the box up front.

The **pigeonhole principle** is the following:

If m objects are placed into n bins,
where $m > n$, then some bin contains
at least two objects.

(We sketched a proof in Lecture #02)

Why This Matters

- The pigeonhole principle can be used to show results must be true because they are "too big to fail."
- Given a large enough number of objects with a bounded number of properties, eventually at least two of them will share a property.
- Can be used to prove some surprising results.

Using the Pigeonhole Principle

- To use the pigeonhole principle:
  - Find the m objects to distribute.
  - Find the $n < m$ buckets into which to distribute them.
  - Conclude by the pigeonhole principle that there must be two objects in some bucket.
- The details of how to proceeds from there are specific to the particular proof you're doing.

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

Thought: There are two colors here, so if we start picking points, we'll be dropping them into one of two buckets (red or blue).

How many points do we need to pick to guarantee that we get two of the same color?

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

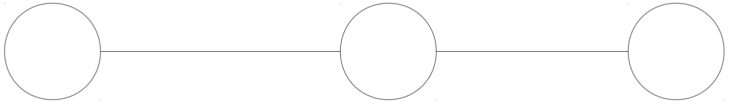

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

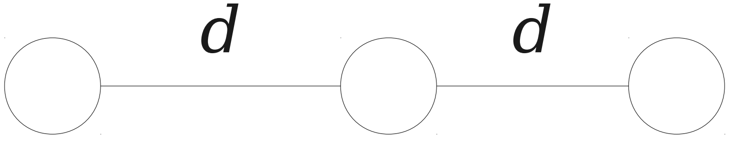

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

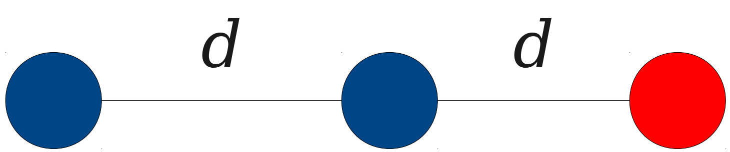

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

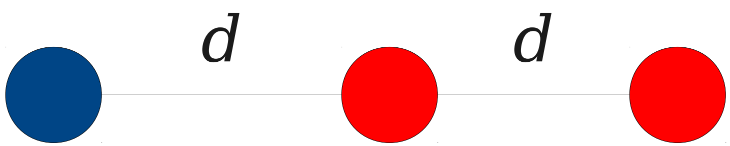

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

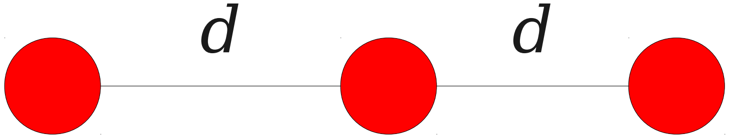

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

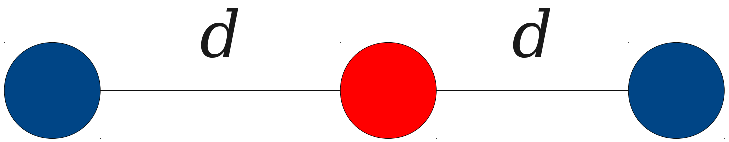

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

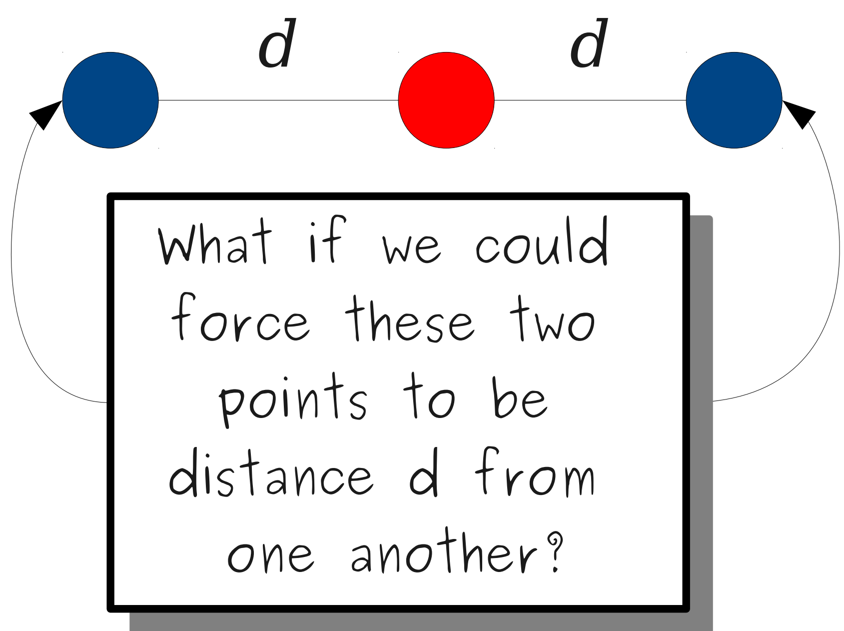

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

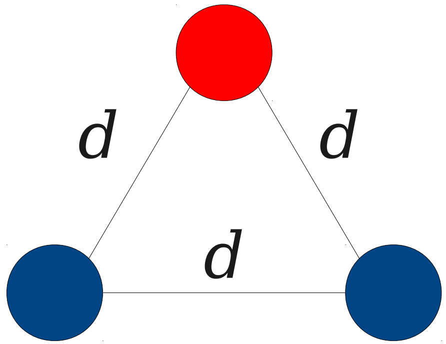

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

Any pair of these points is at distance d from one another. Since two must be the same color, there is a pair of points of the same color at distance d!

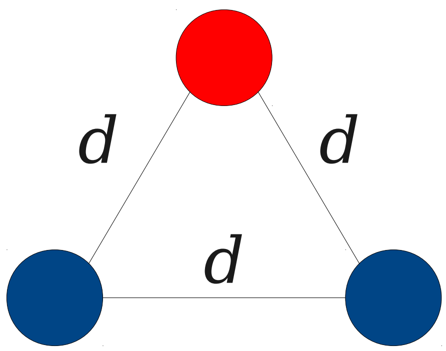

A Surprising Application

Theorem: Suppose that every point in the real plane is colored either red or blue. Then for any distance $d > 0$, there are two points exactly distance d from one another that are the same color.

Proof: Consider any equilateral triangle whose side lengths are d. Put this triangle anywhere in the plane. By the pigeonhole principle, because there are three vertices, two of the vertices must have the same color. These vertices are at distance d from each other, as required. ■

The Hadwiger-Nelson Problem

- No matter how you color the points of the plane, there will always be two points at distance 1 that are the same color.
- Relation to graph coloring:
  - Every point in the real plane is a node.
  - There's an edge between two points that are at distance exactly one.
- Question: What is the chromatic number of this graph? (That is, how many colors do you need to ensure no points at distance 1 are the same color?)
- This is the **Hadwiger-Nelson** problem. It's known that the number is between 4 and 7, but no one knows for sure!

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

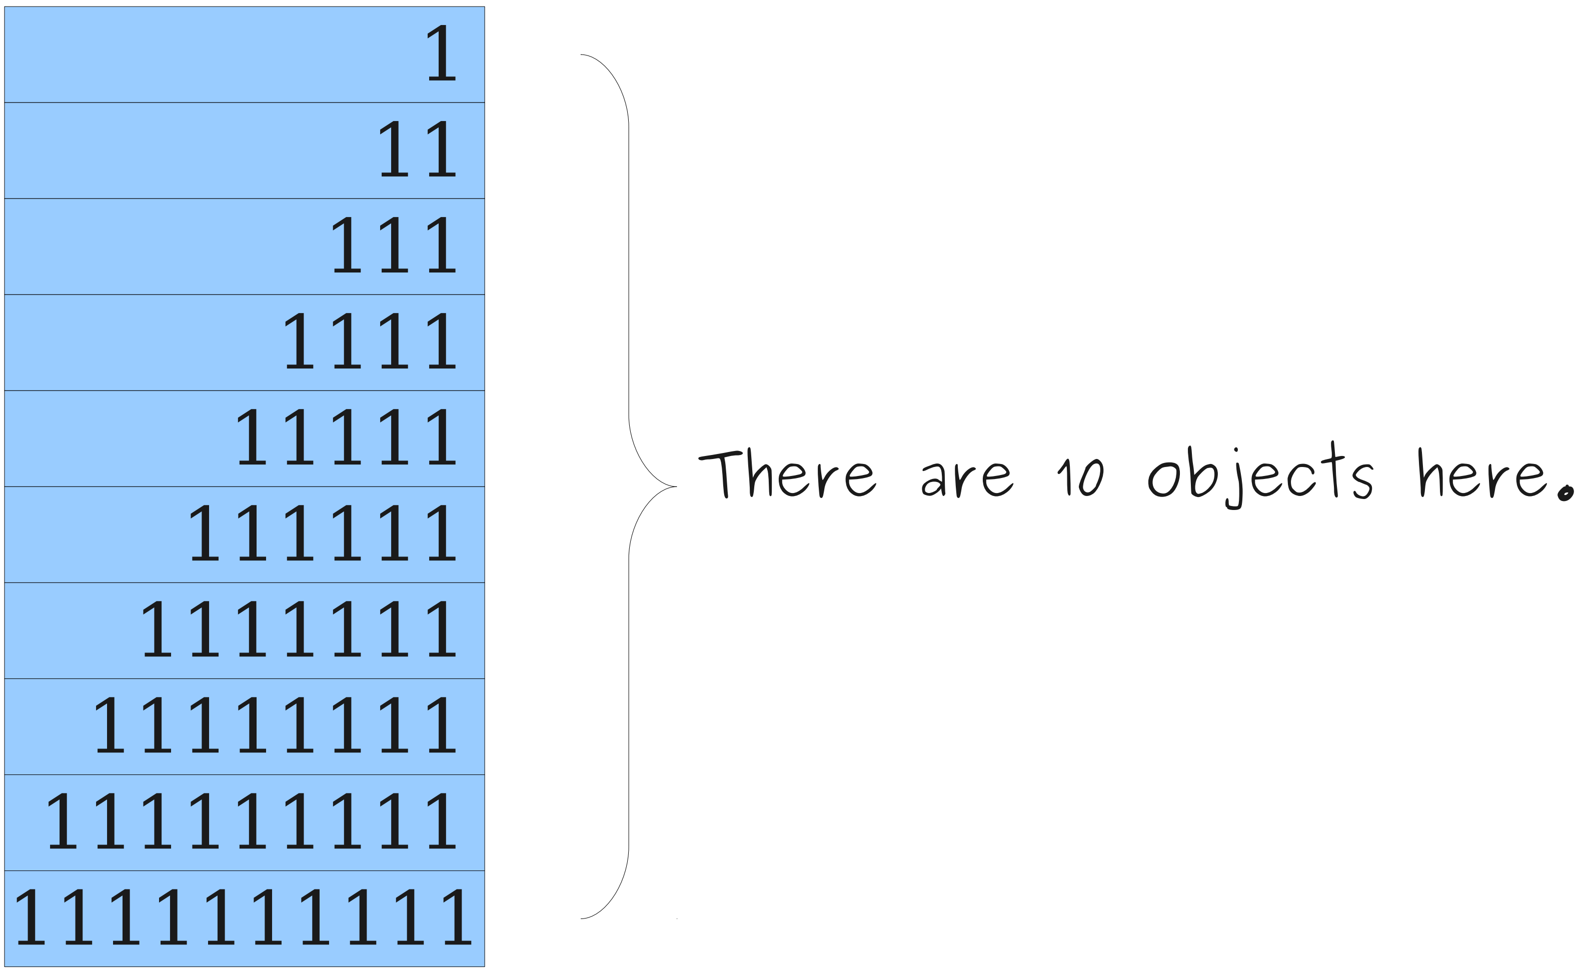

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

1
11
111
1111
11111
111111
1111111
11111111
111111111
1111111111

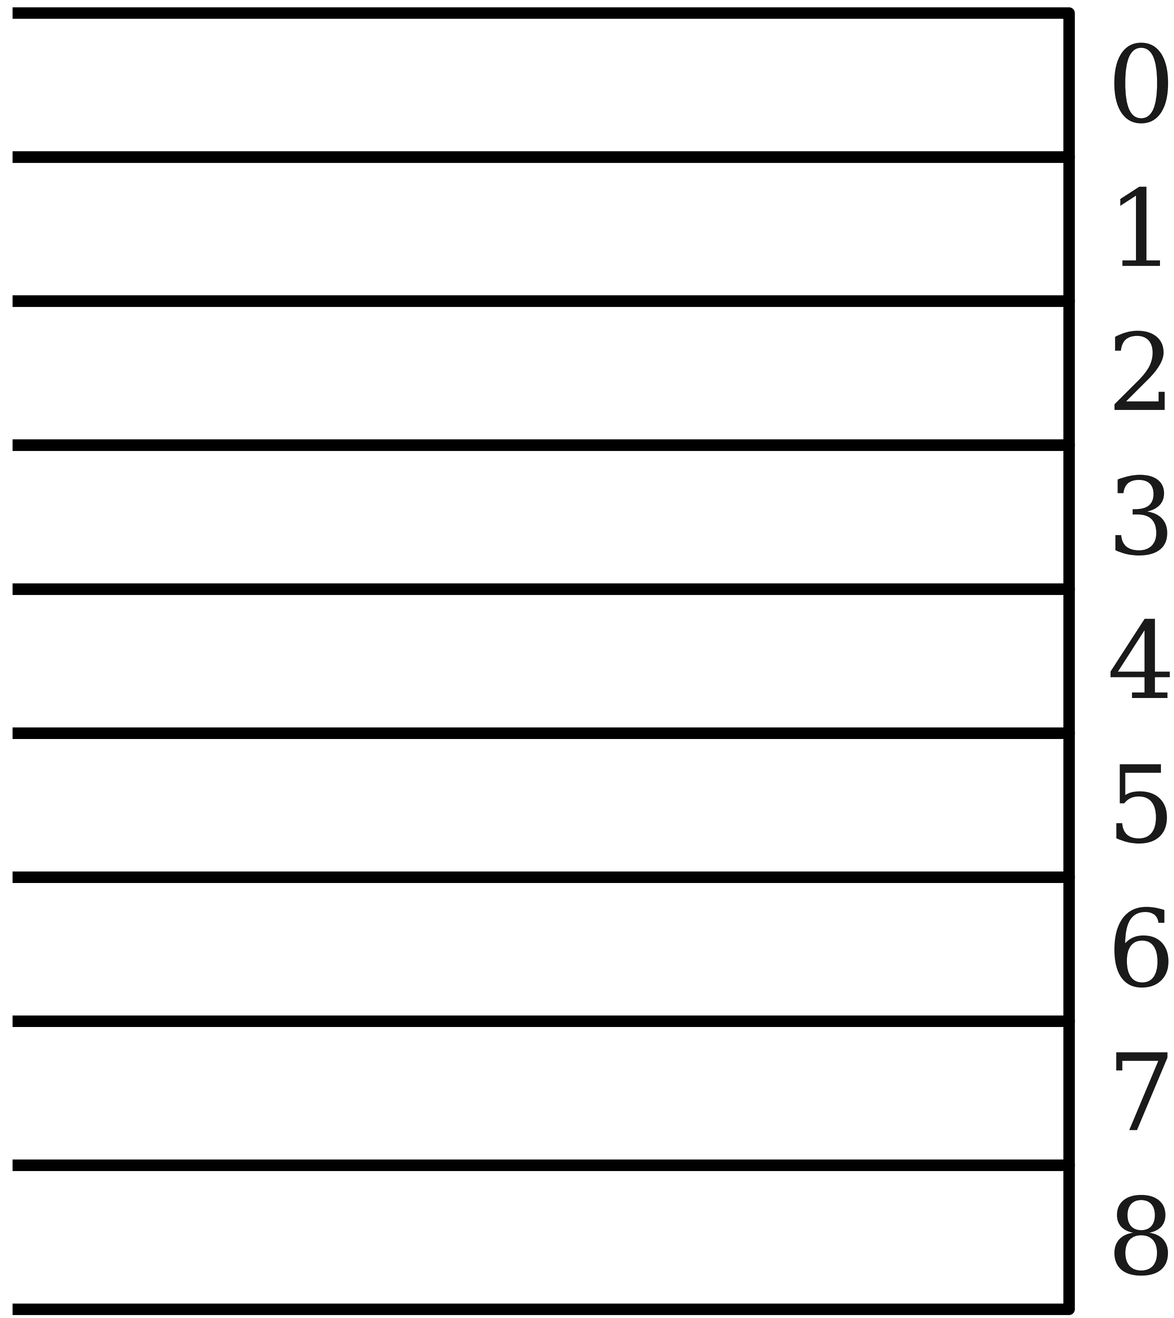

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

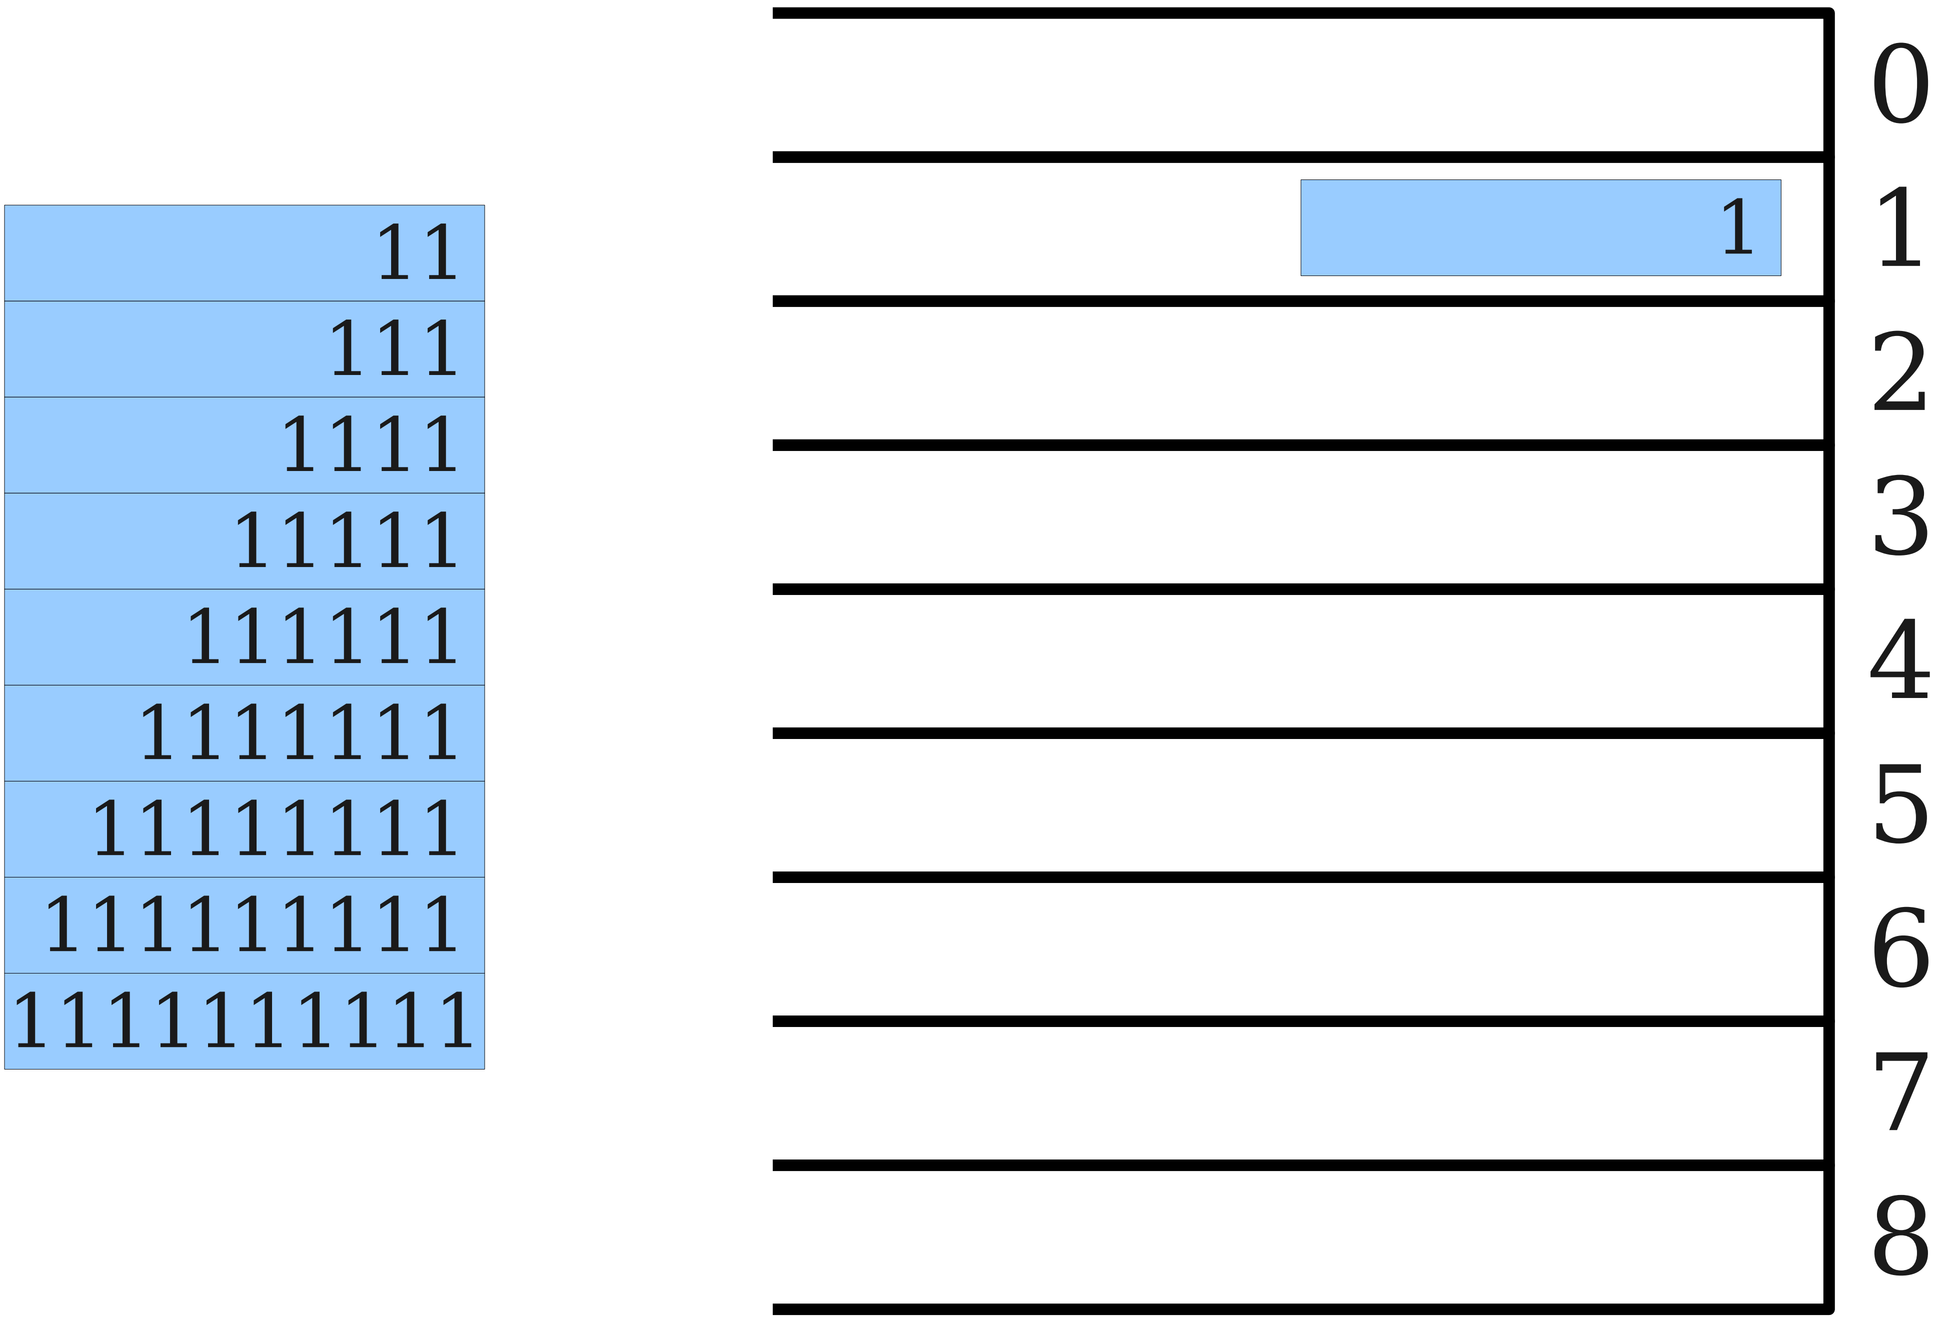

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

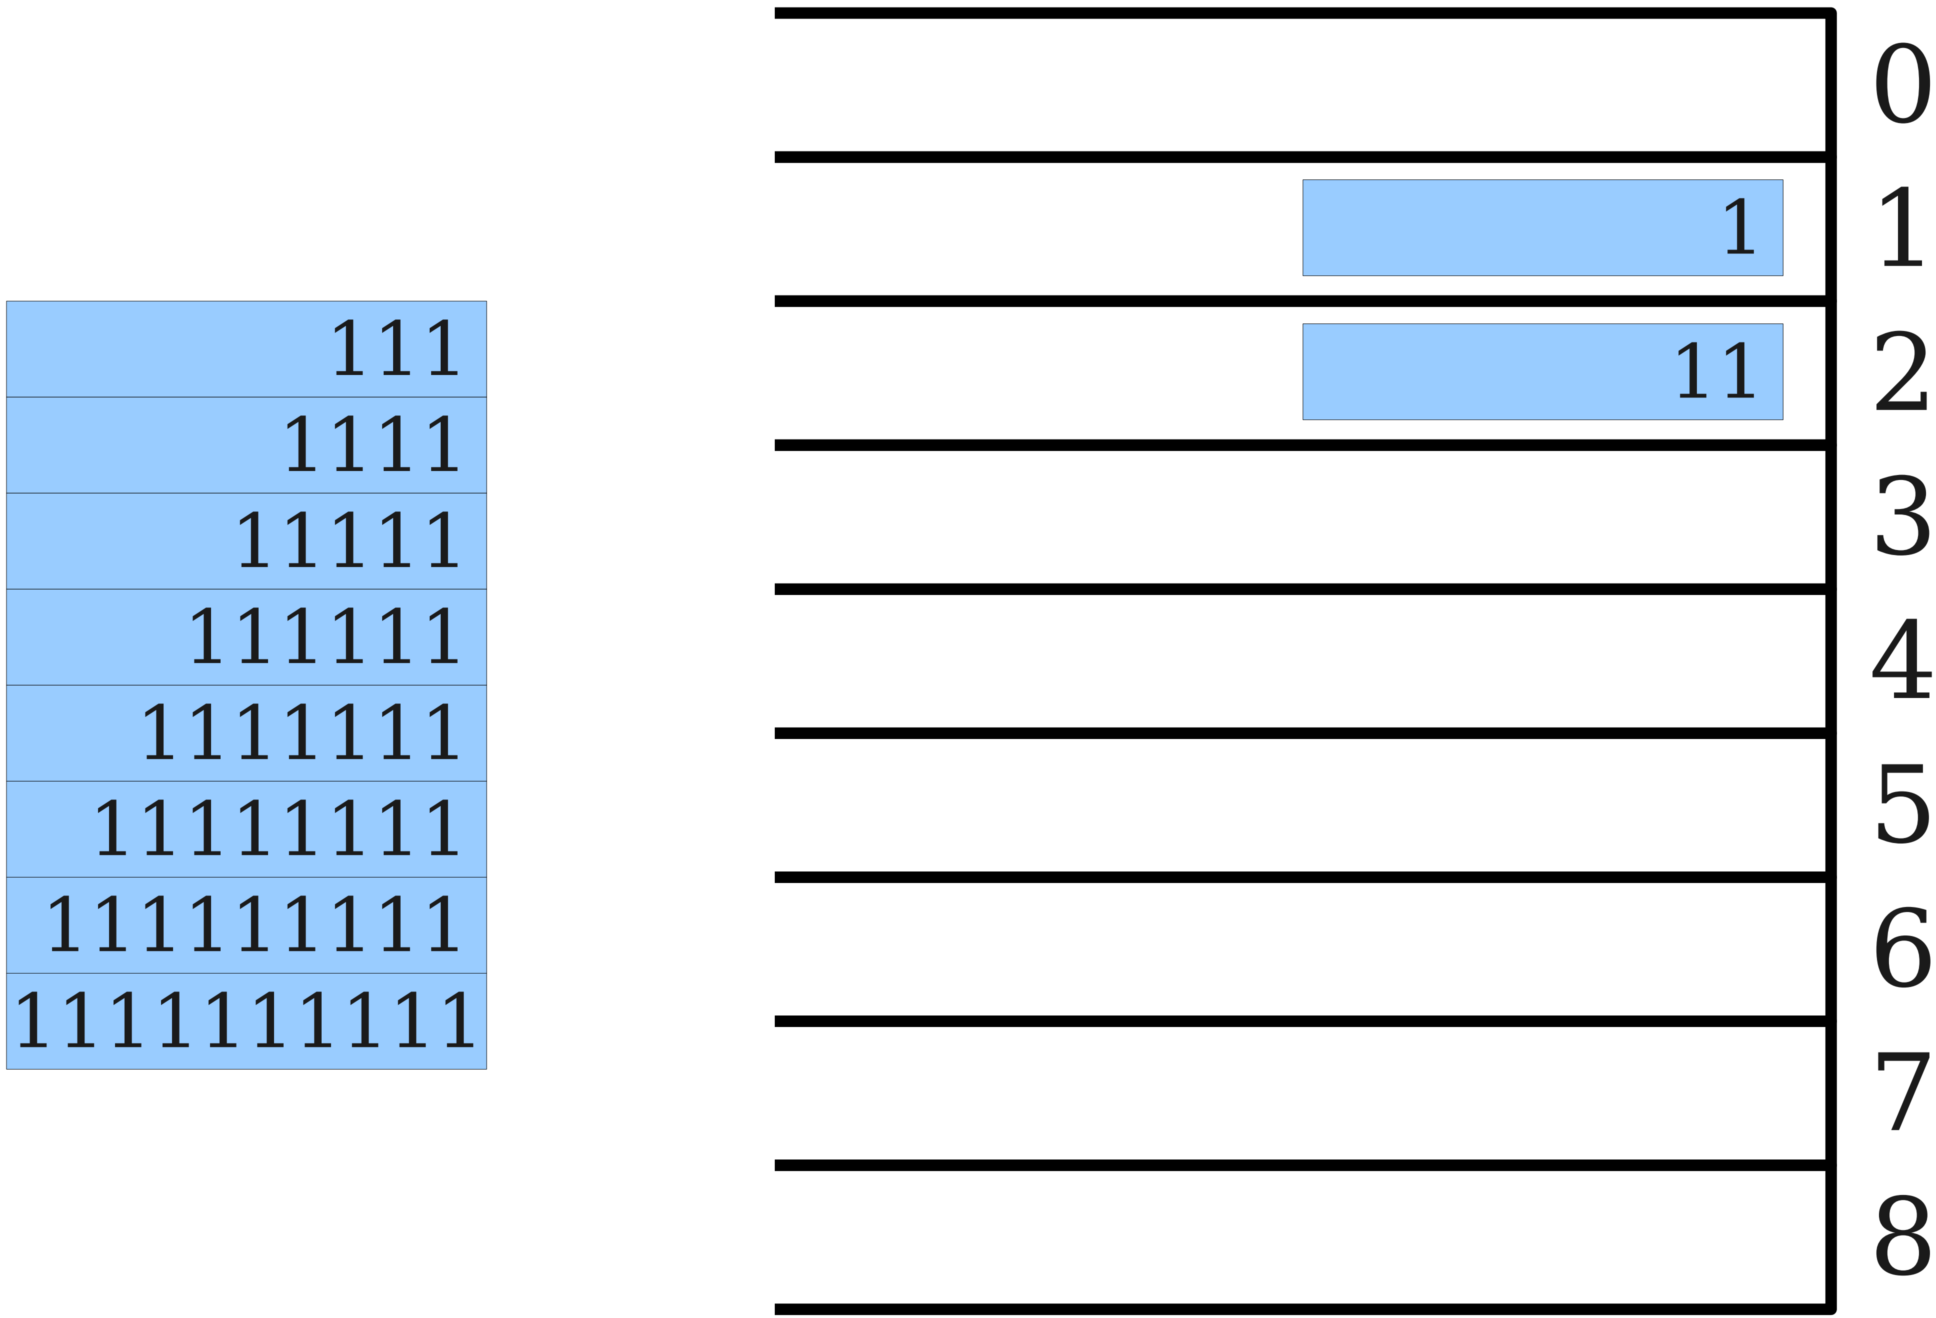

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

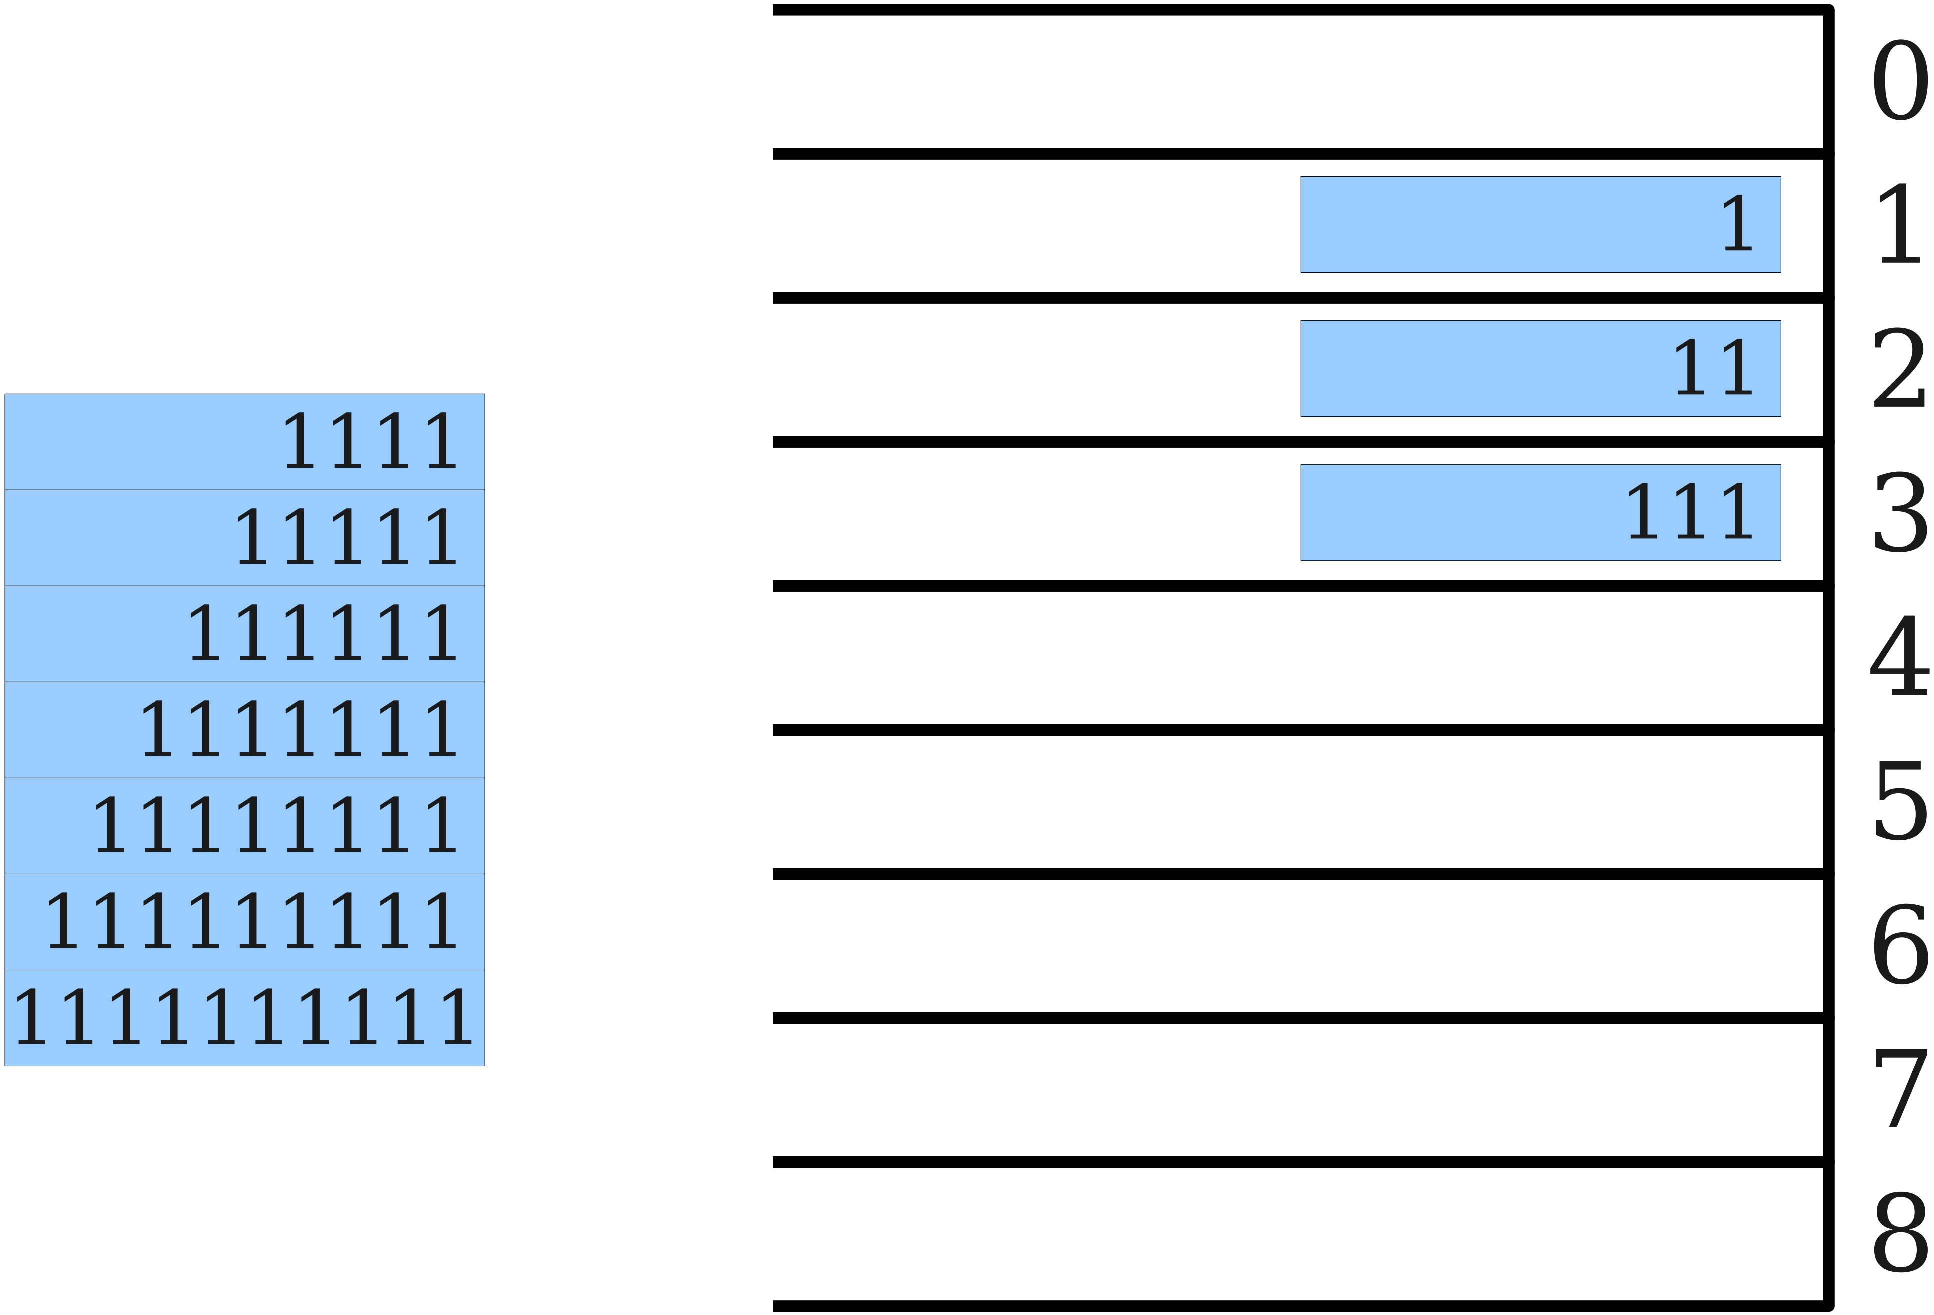

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

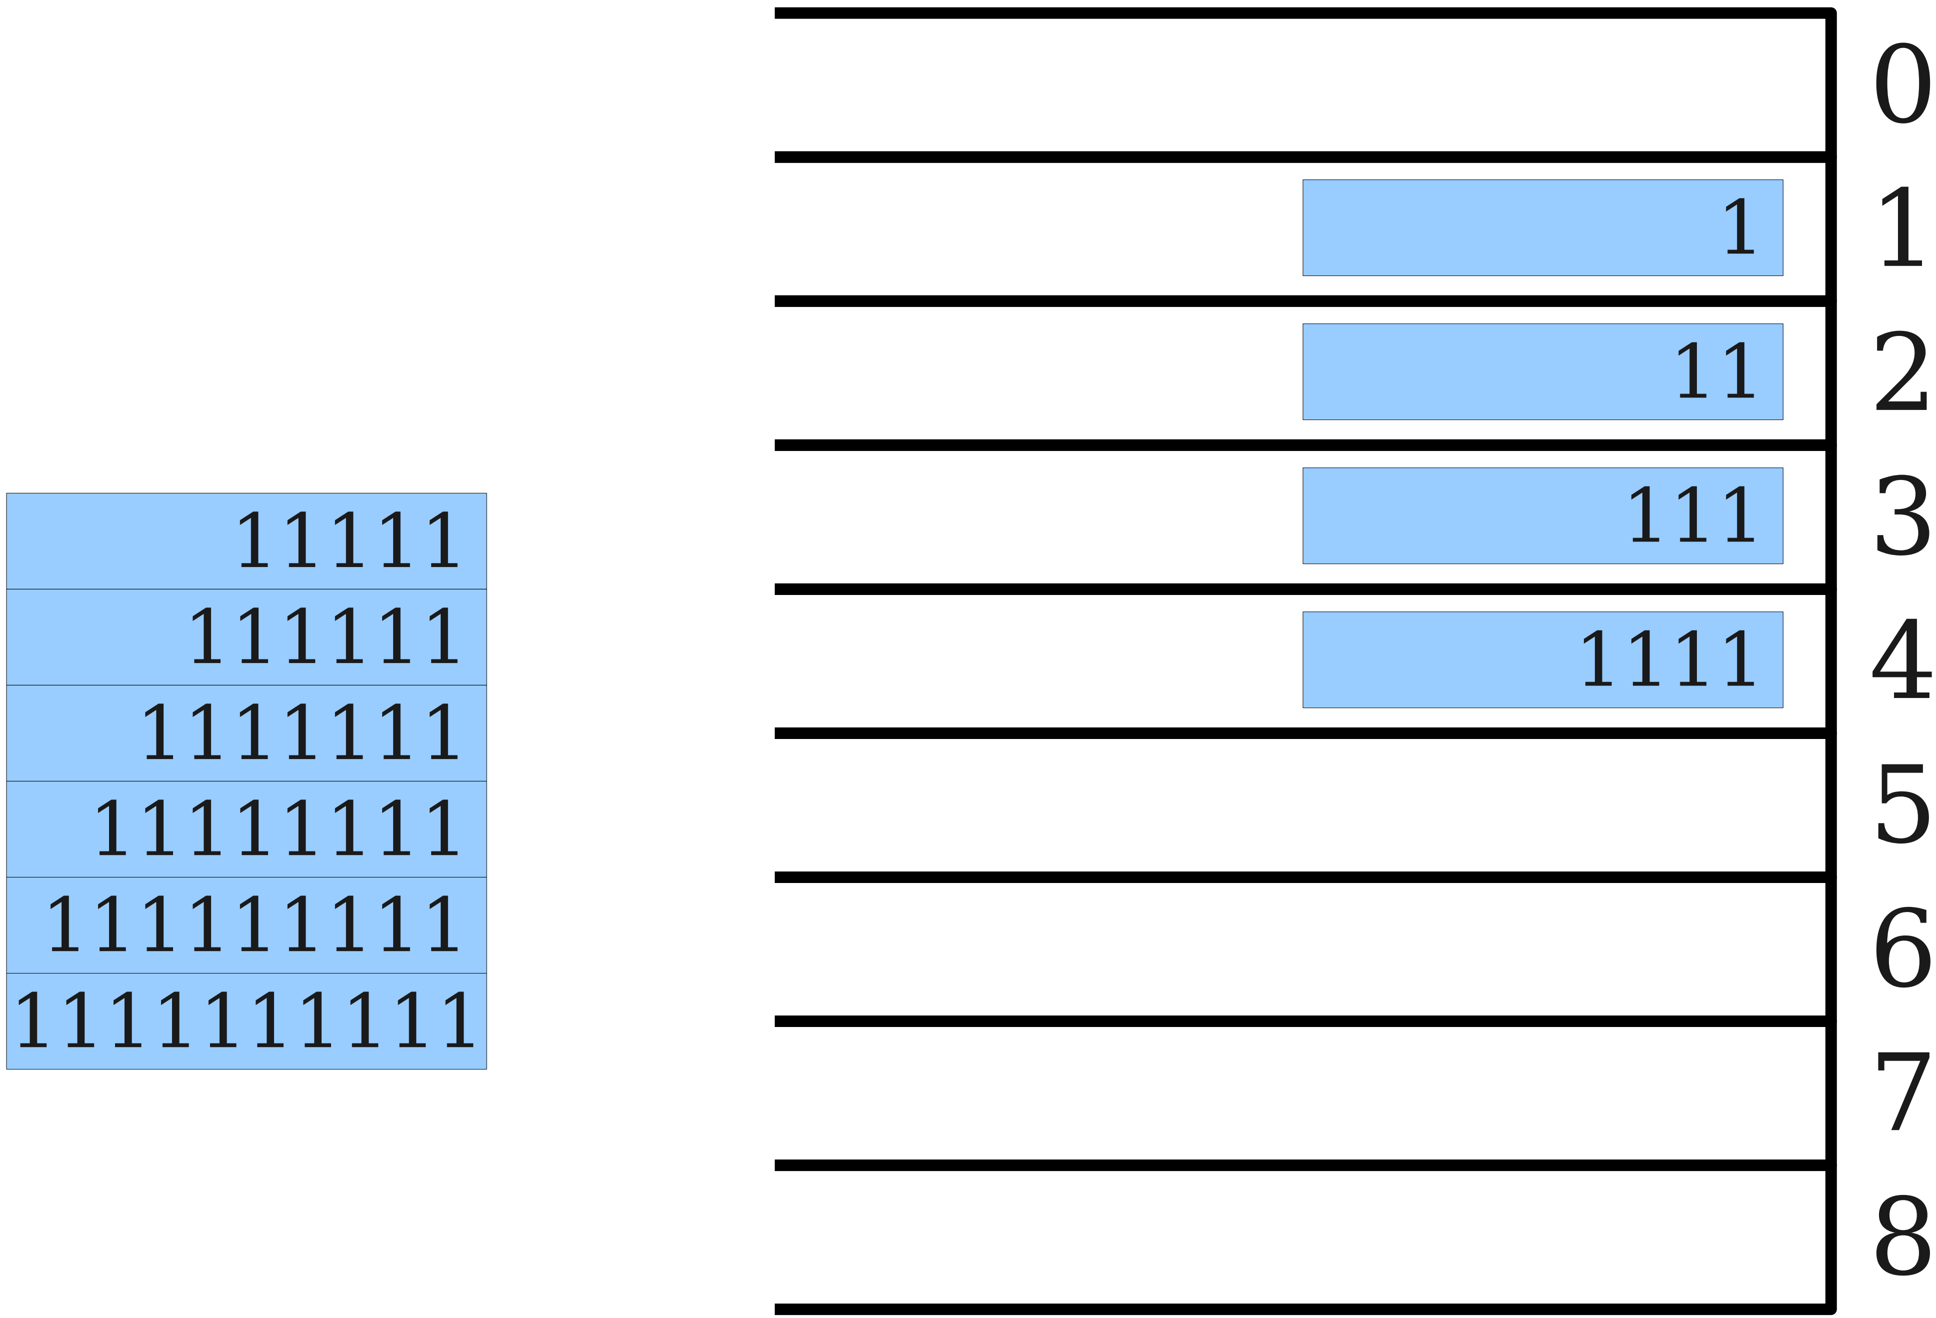

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

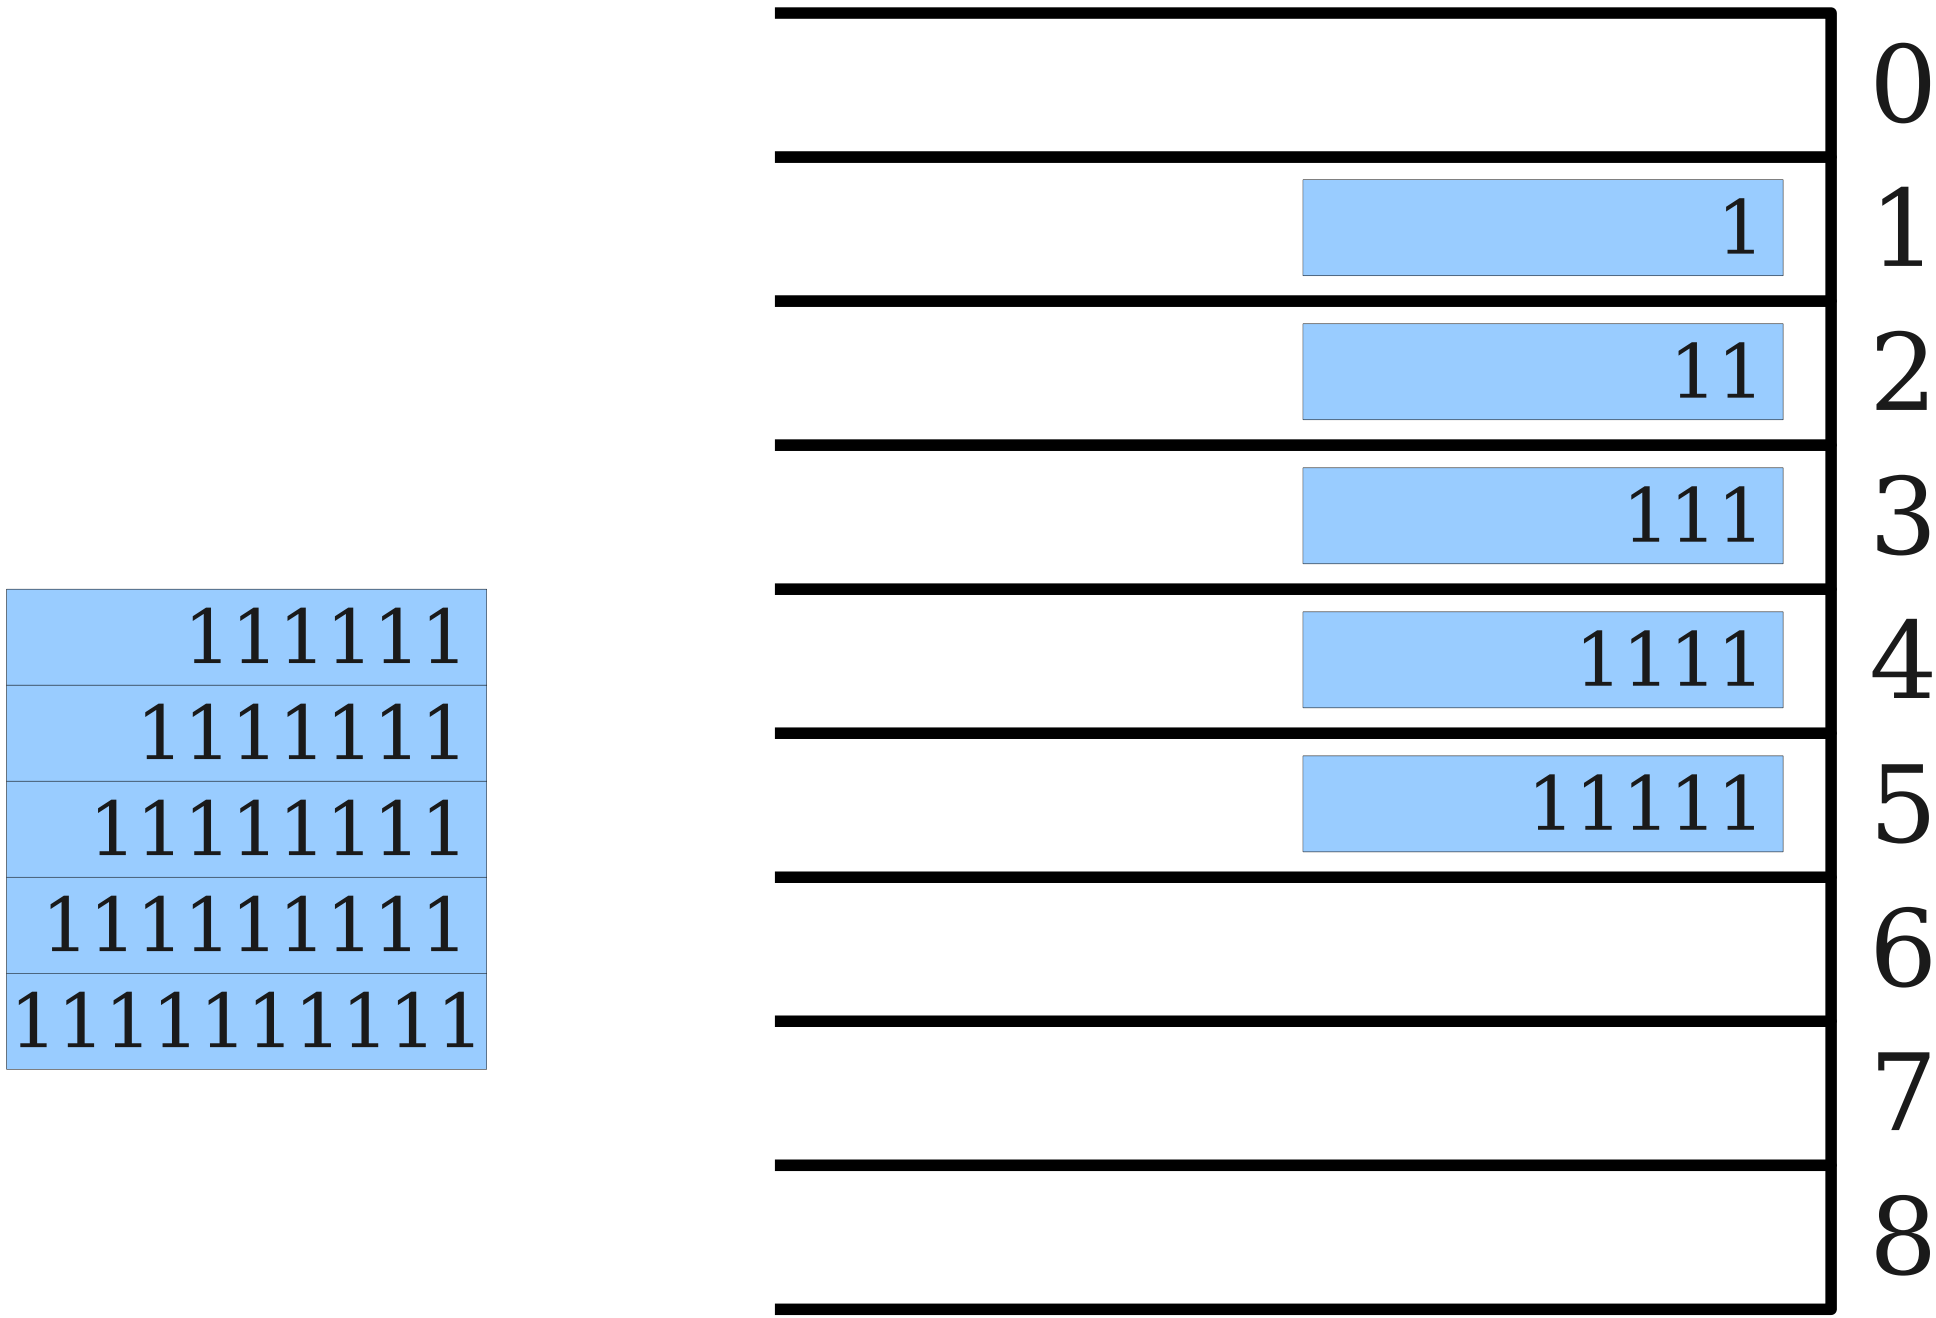

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

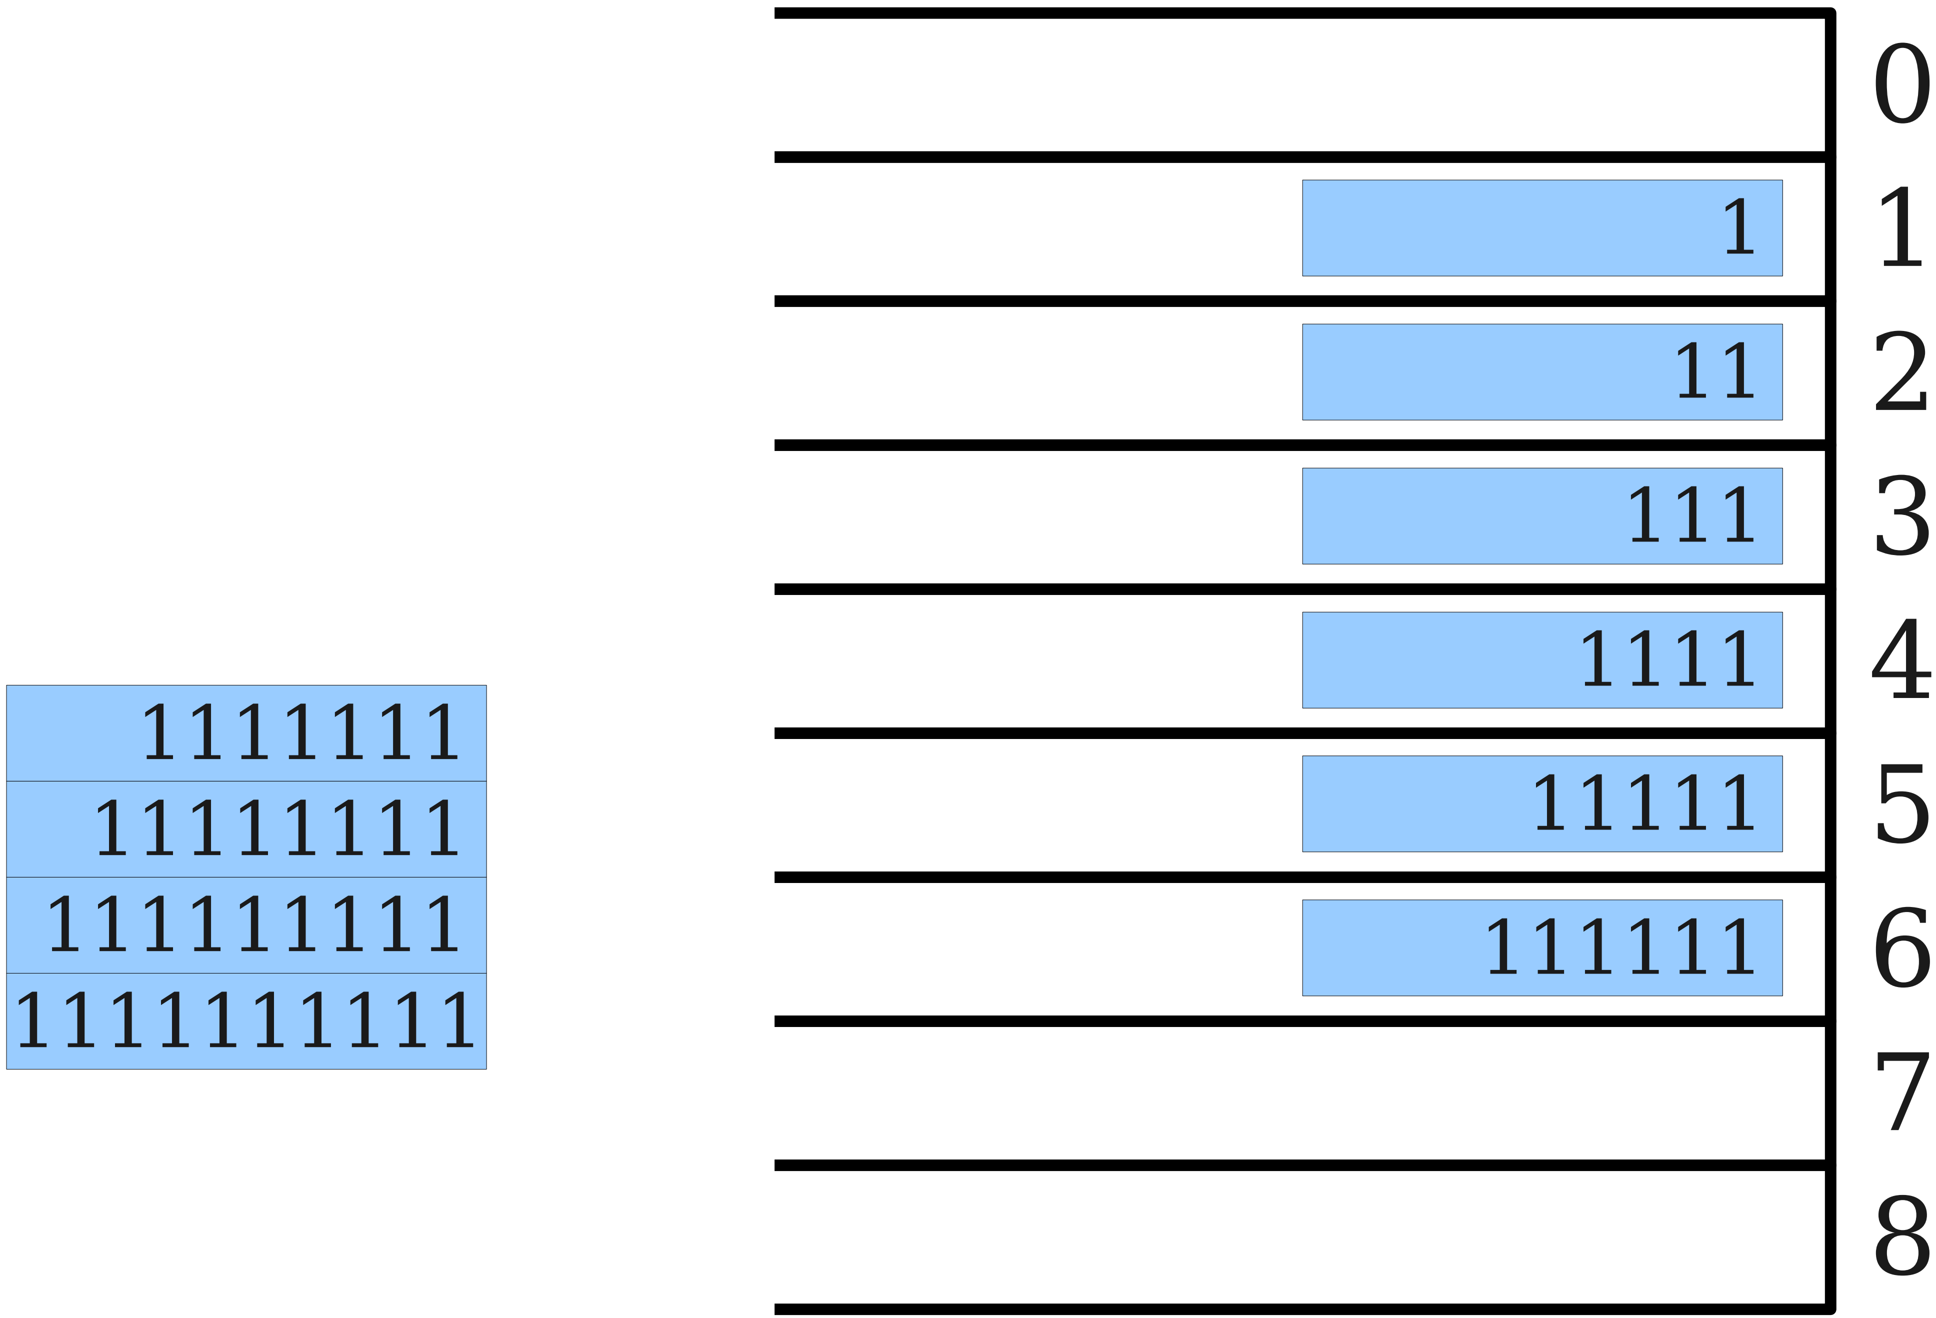

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

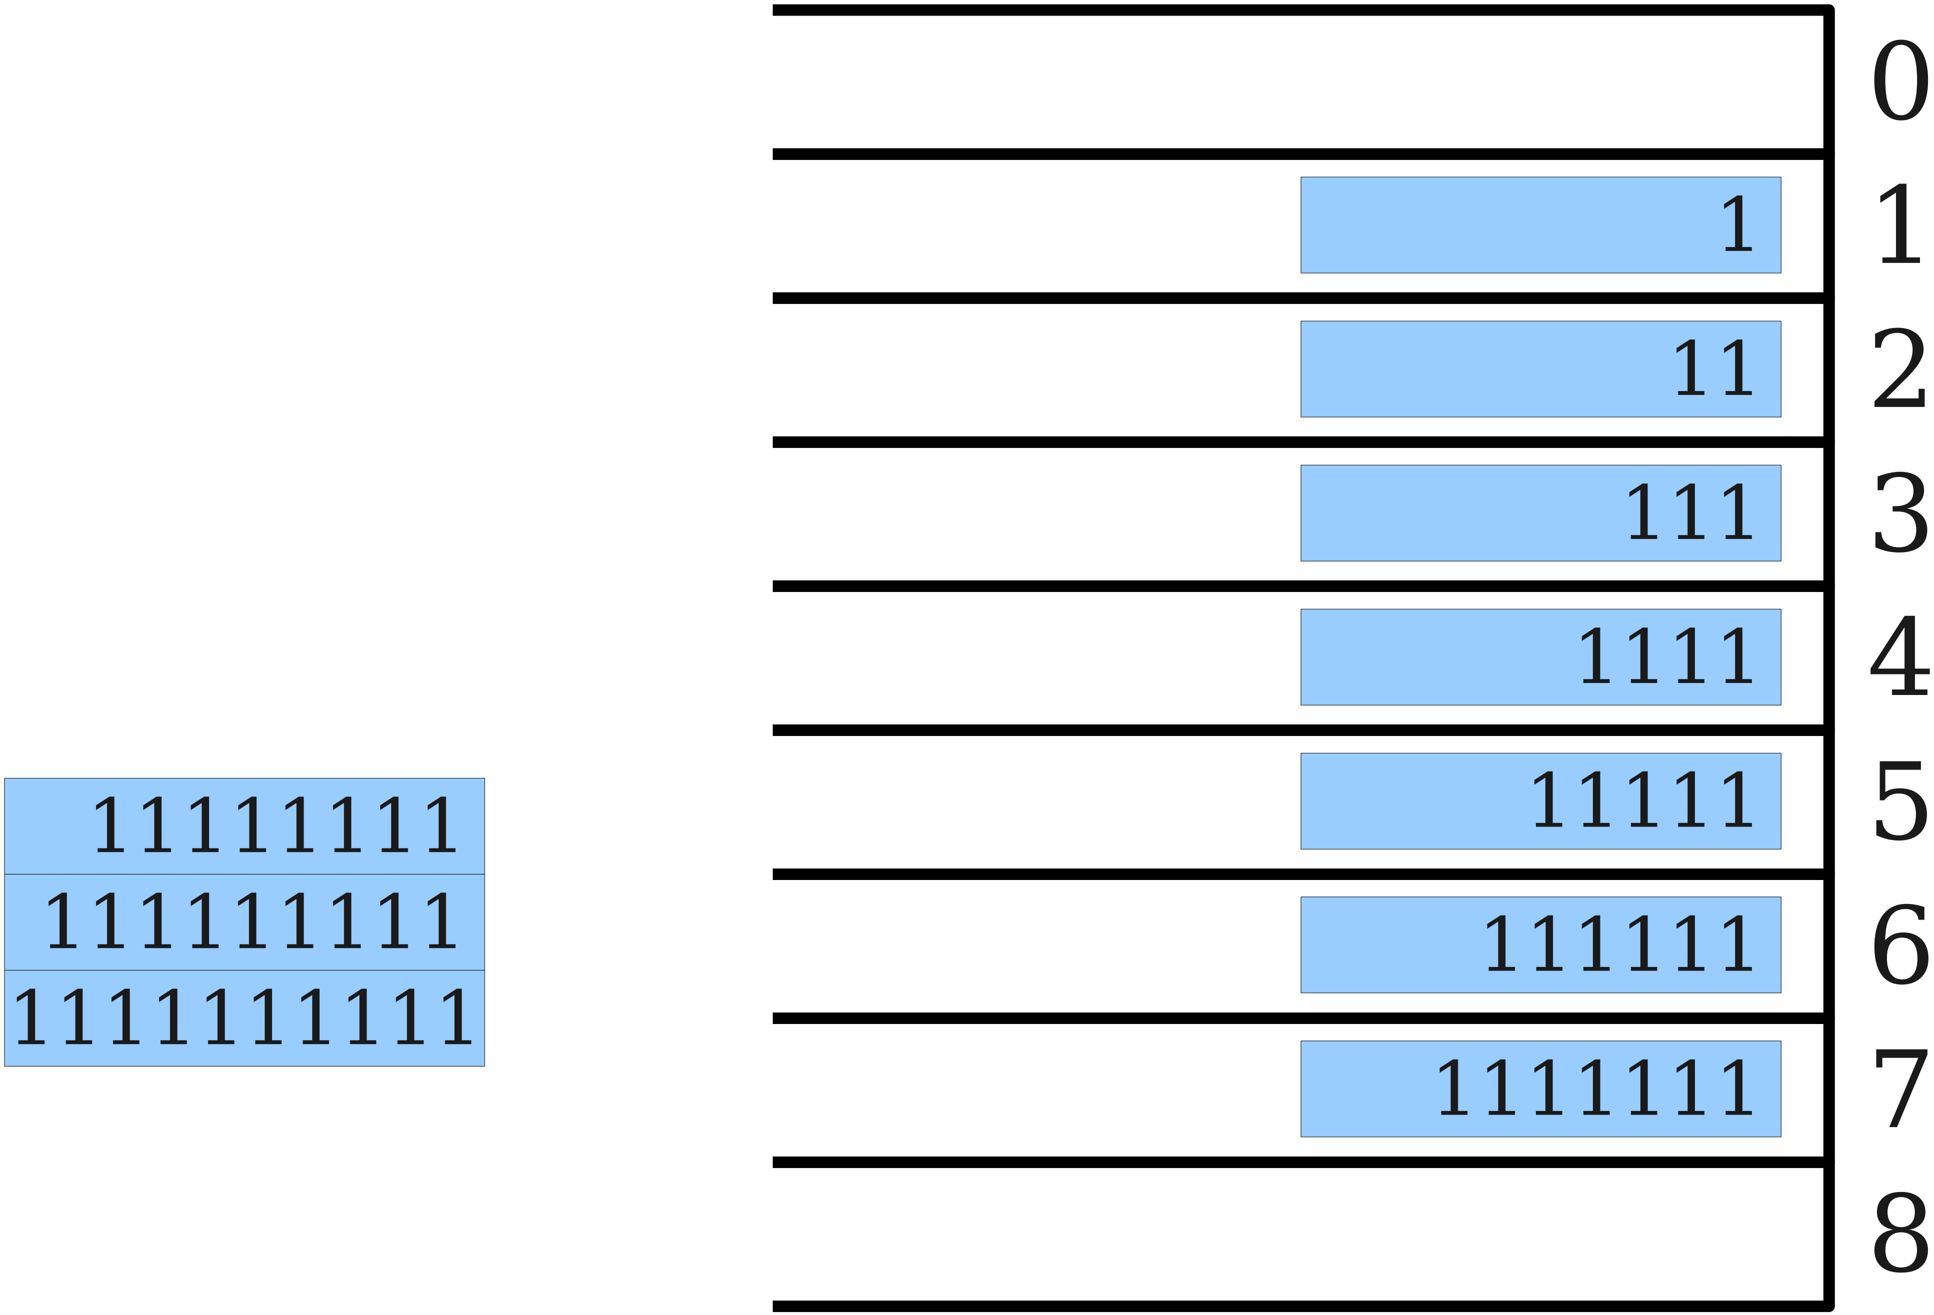

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

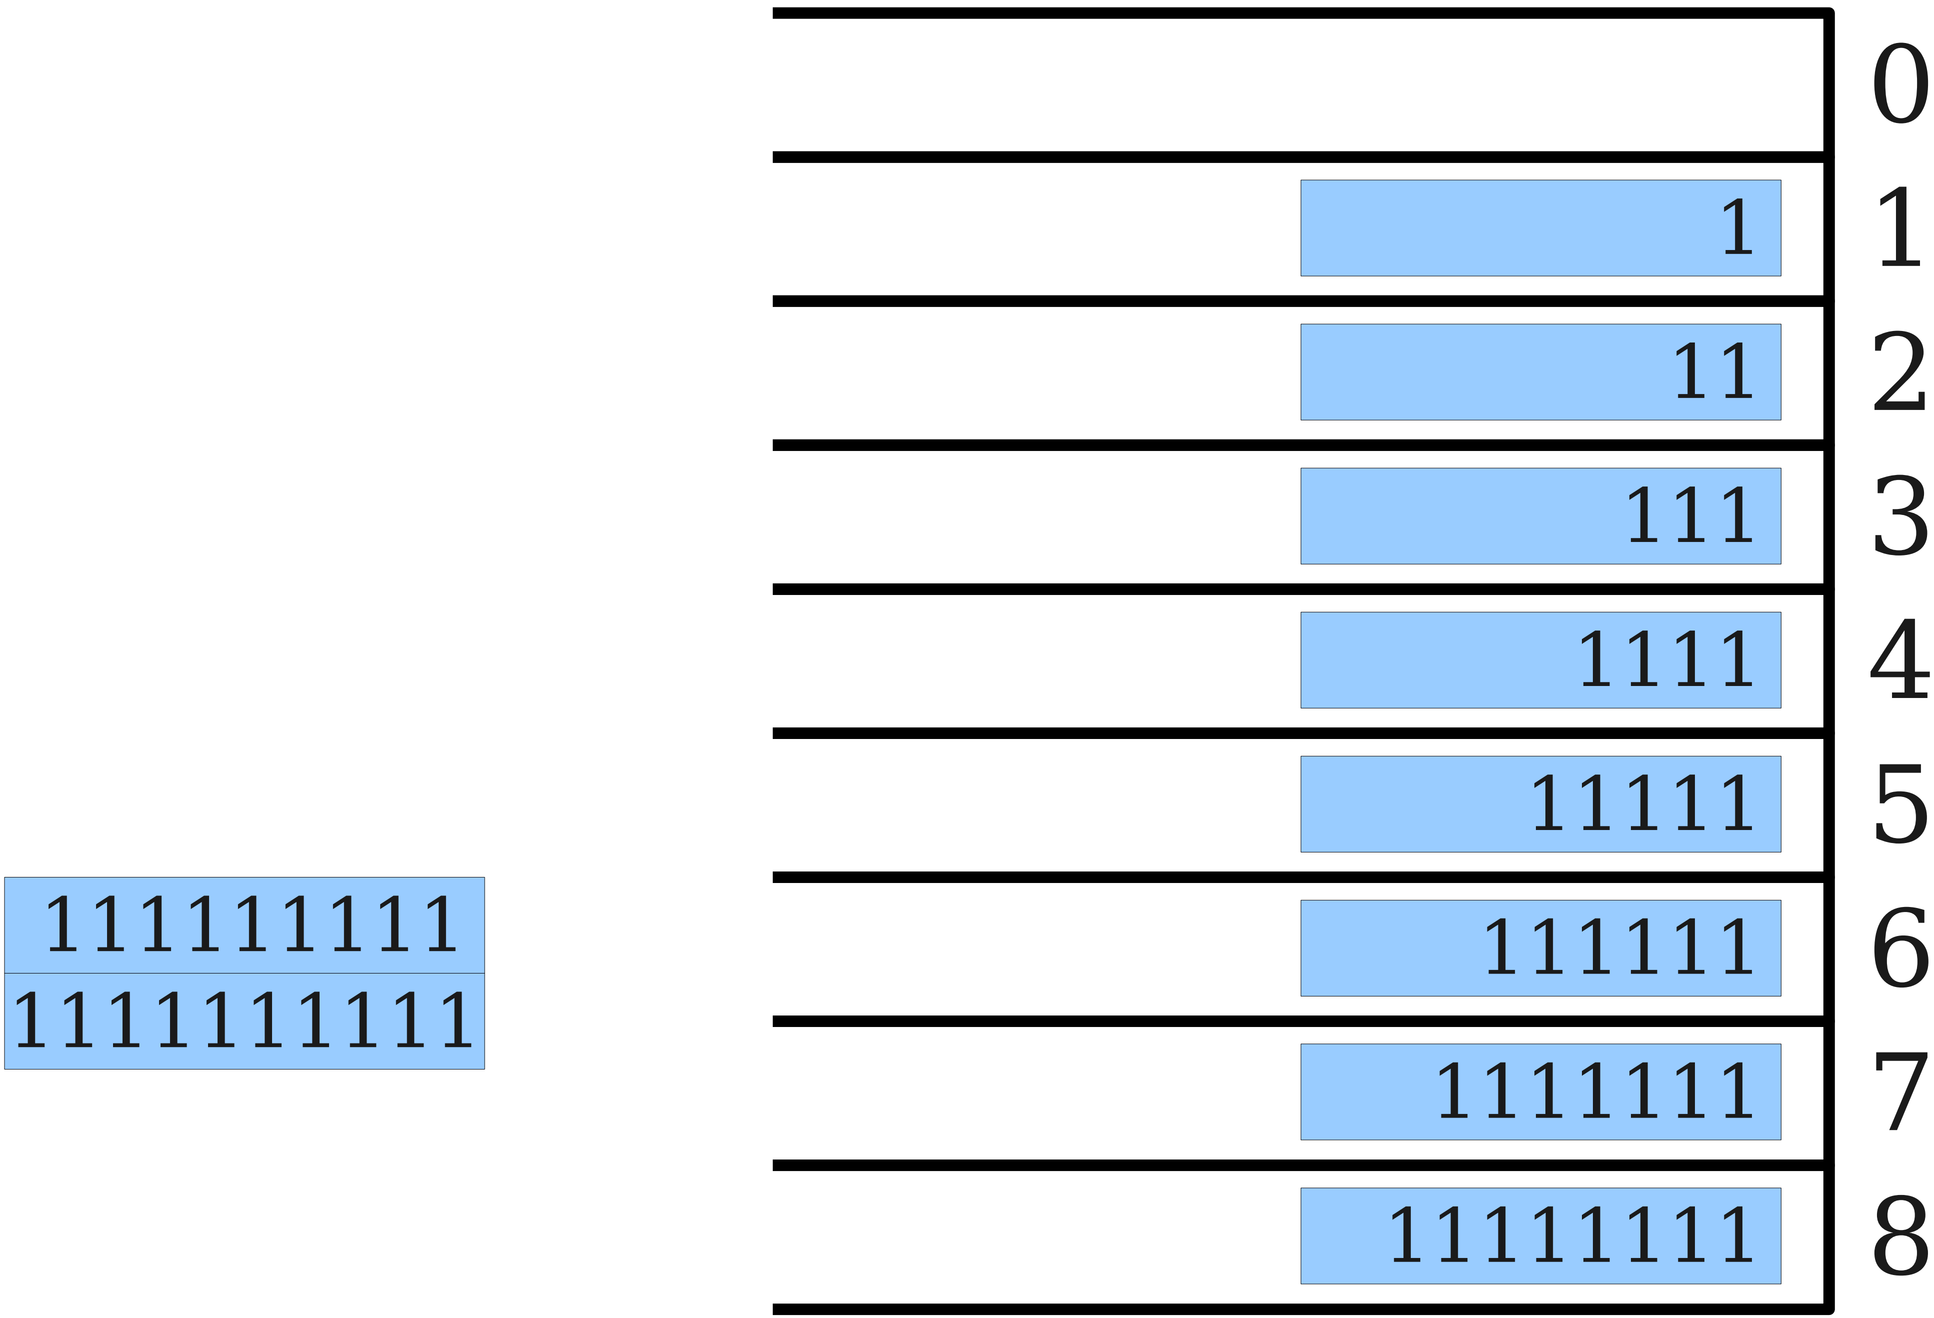

Theorem: For any nonzero natural number *n*, there is a nonzero multiple of *n* whose digits are all 0s and 1s.

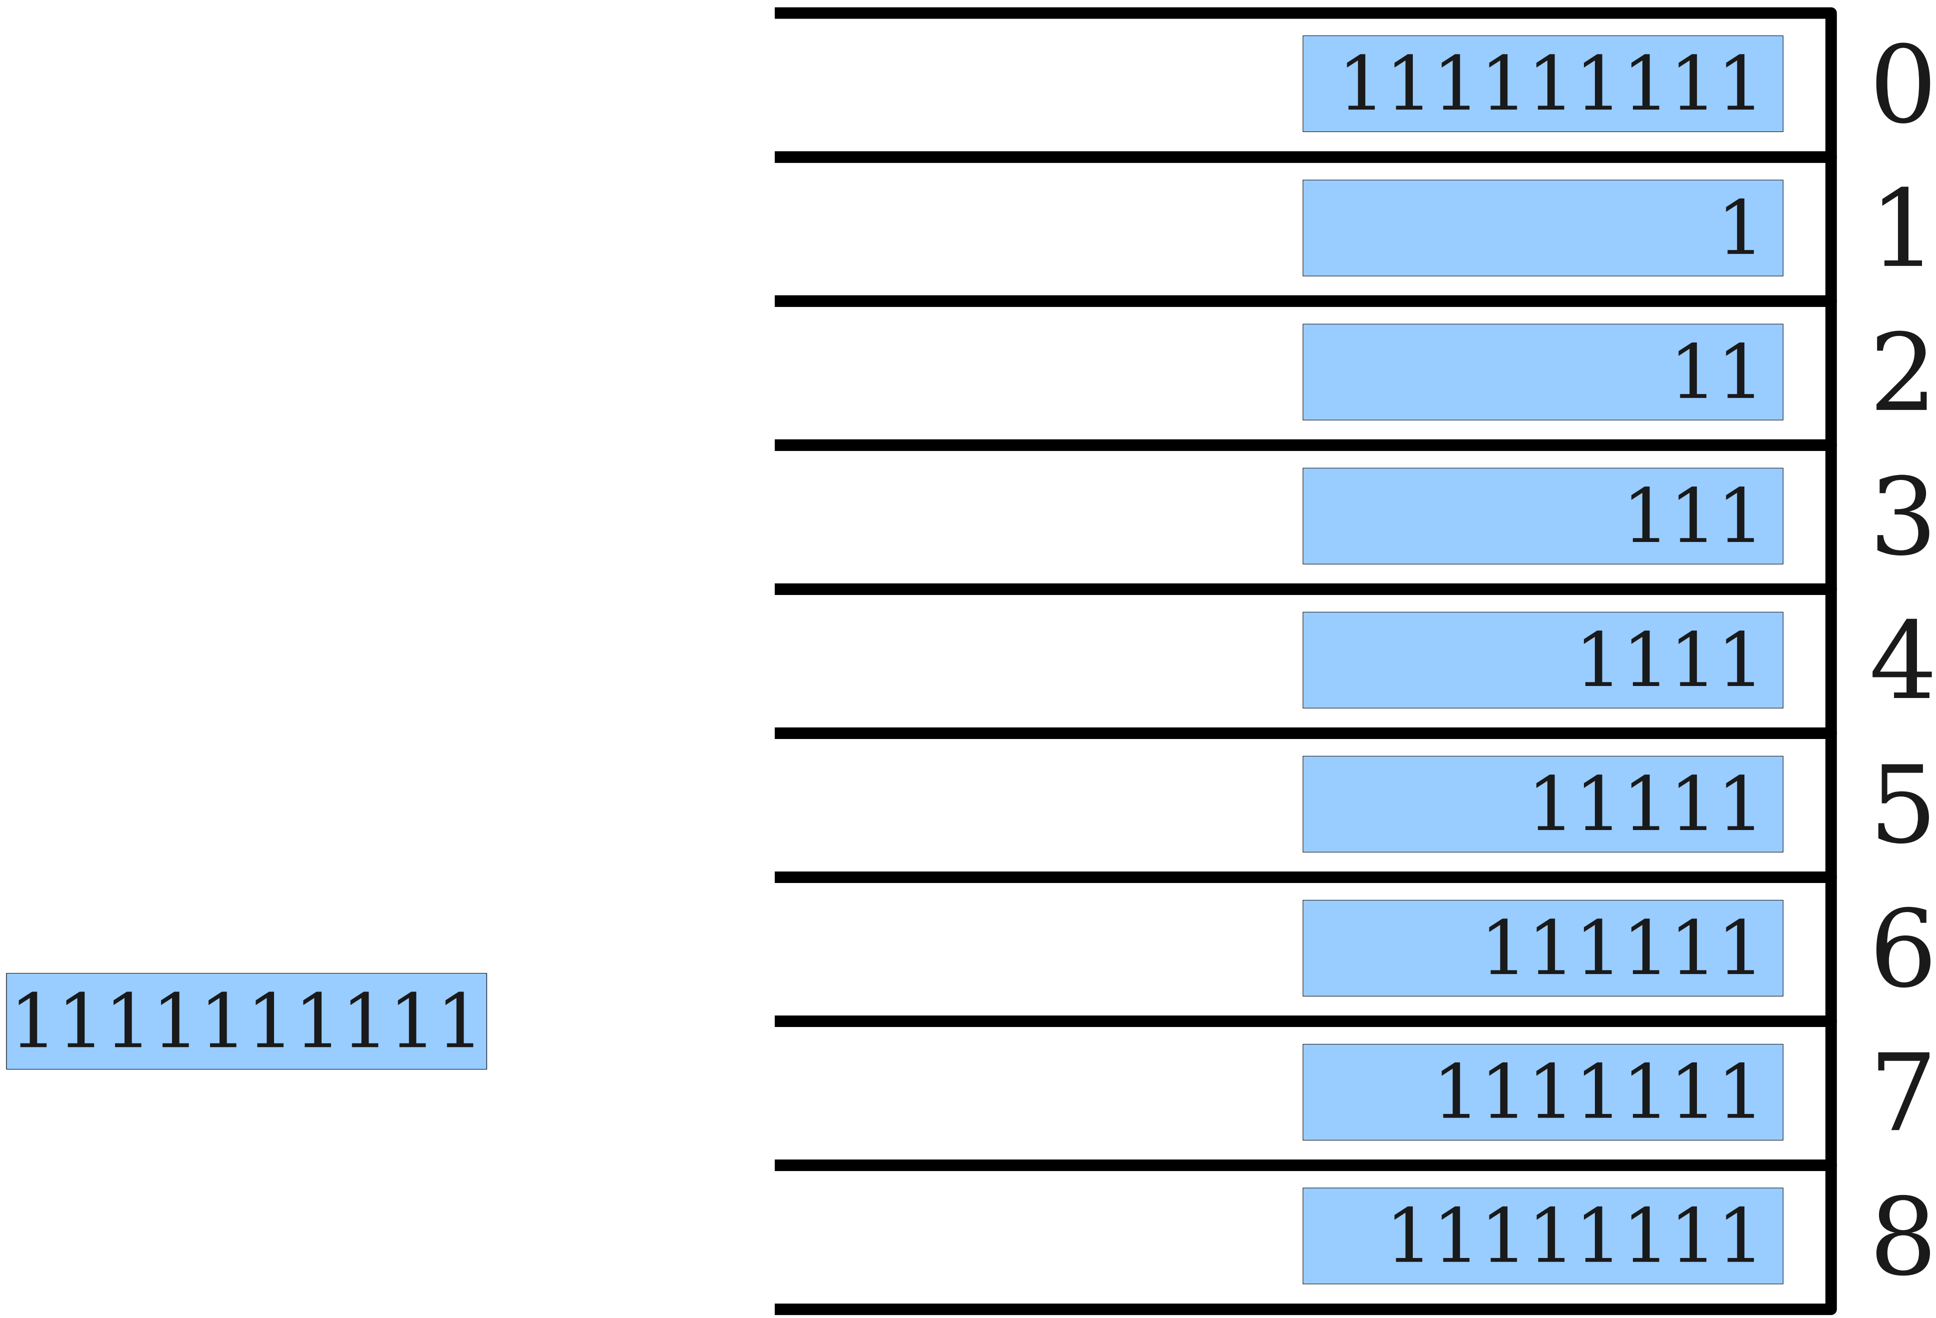

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

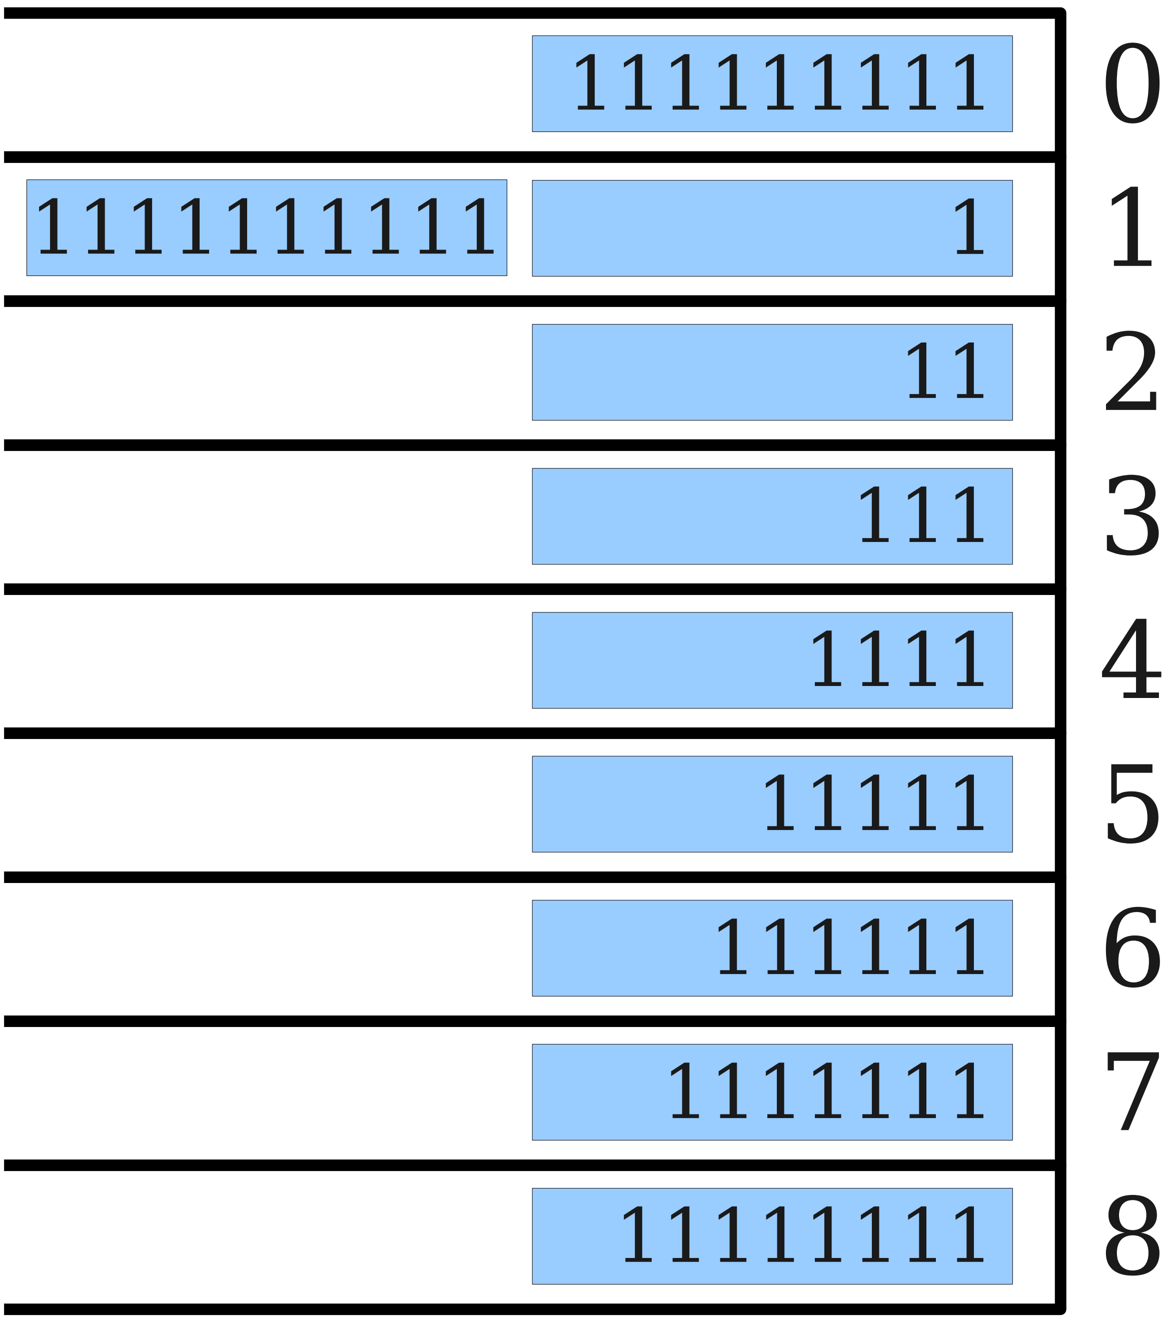

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

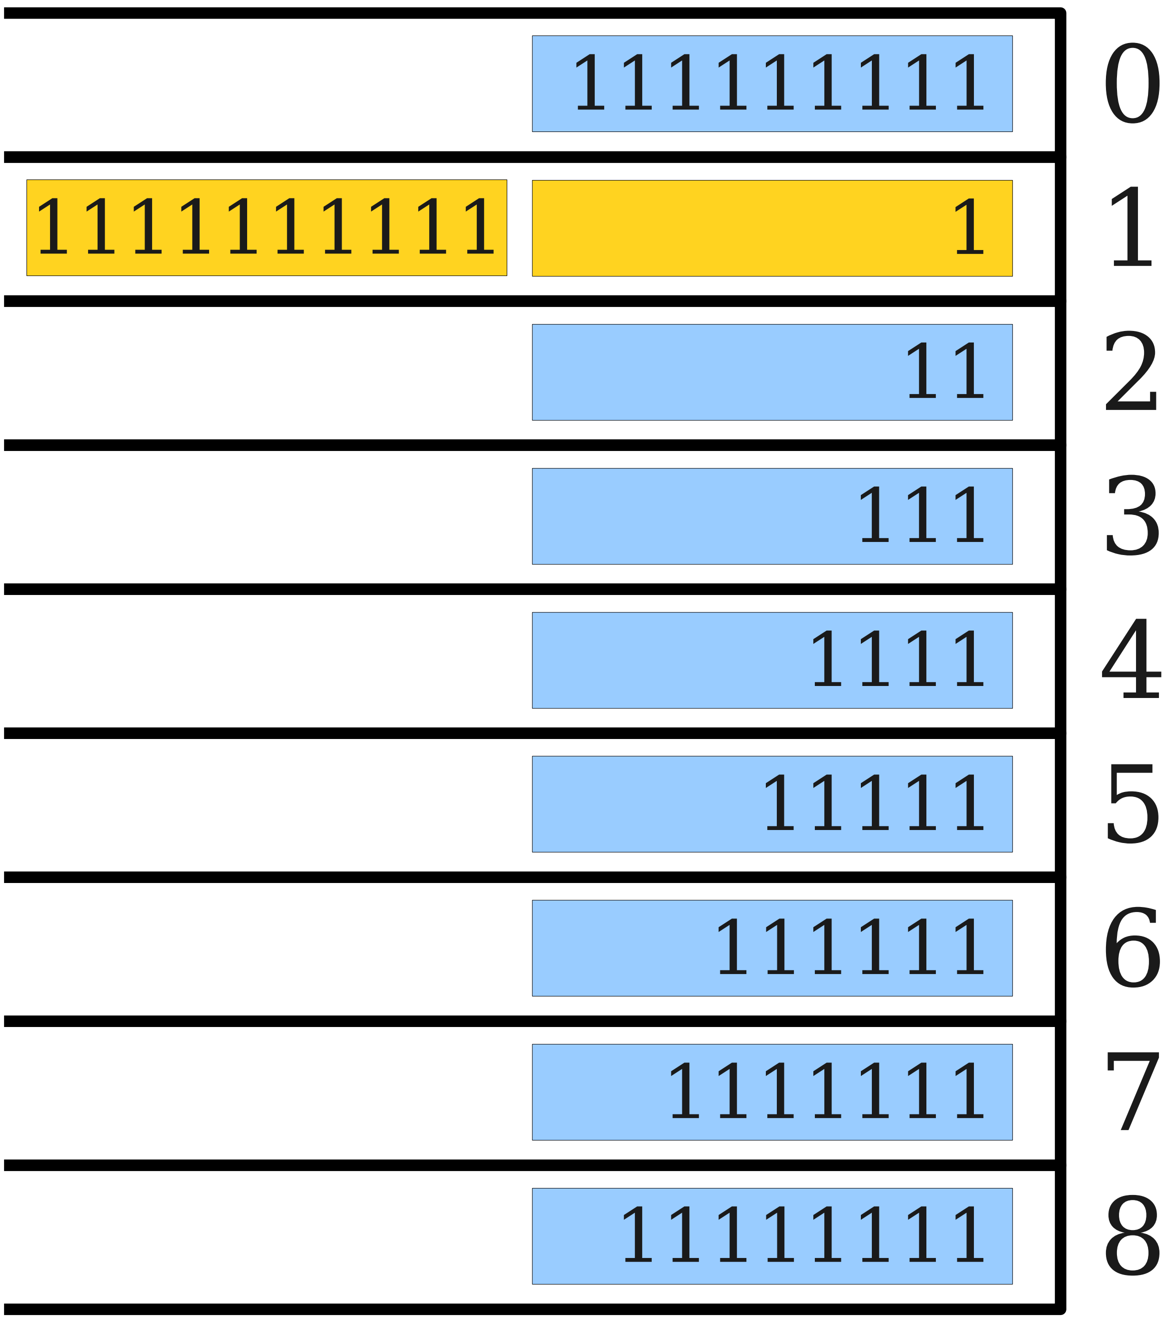

Theorem: For any nonzero natural number n, there is a nonzero multiple of n whose digits are all 0s and 1s.

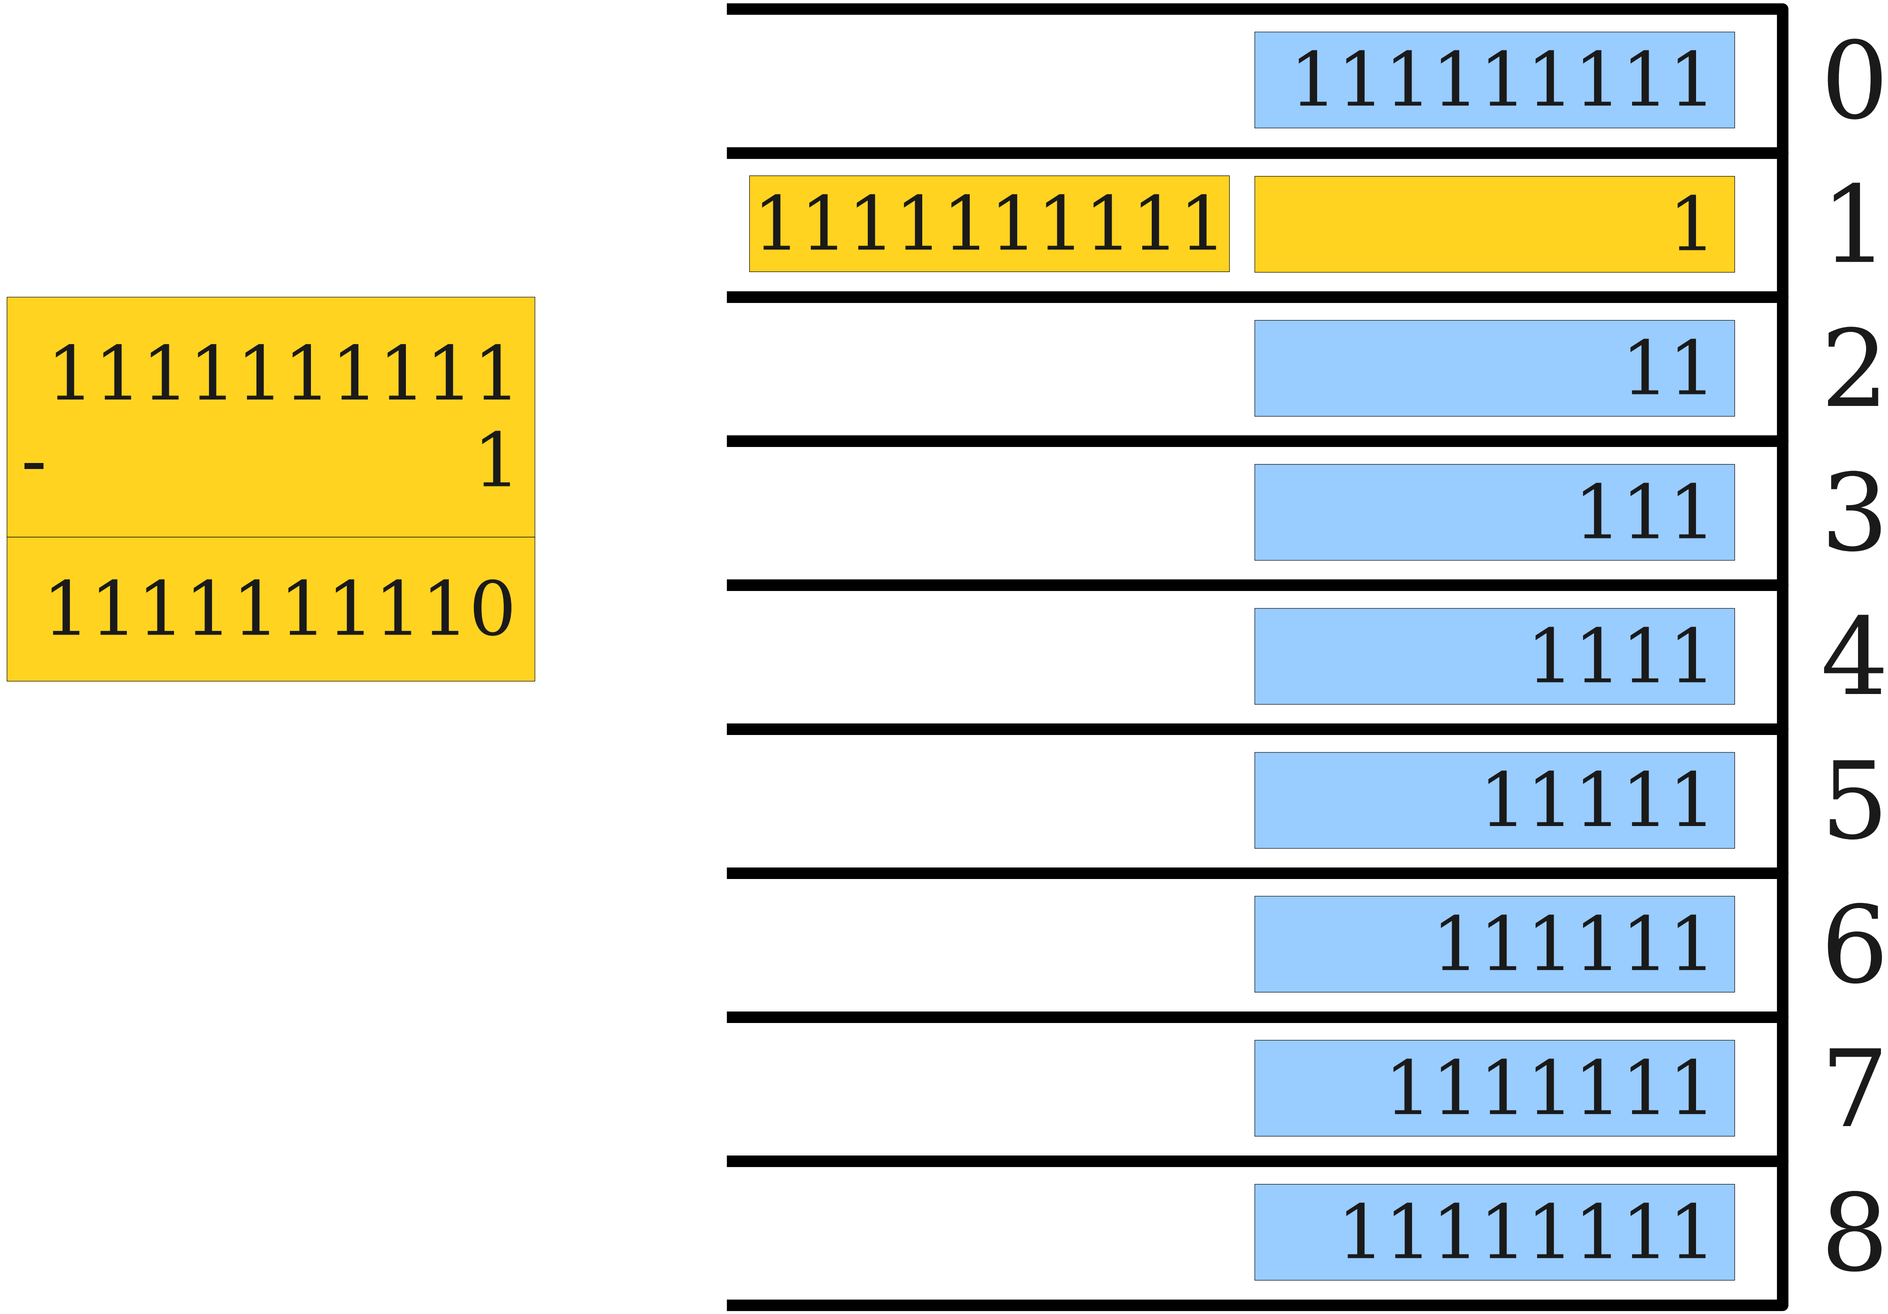

Proof Idea

- Generate the numbers 1, 11, 111, … until $n + 1$ numbers are generated.
- There are n possible remainders modulo n, so two of these numbers have the same remainder.
- Their difference is a multiple of n.
- Their difference consists of 1s and 0s.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \le k \le n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r)$$

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t$$

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \le k \le n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i .$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t = n(q_s - q_t).$$

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \le k \le n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t = n(q_s - q_t).$$

Therefore, $X_s - X_t$ is a multiple of n.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \le k \le n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t = n(q_s - q_t).$$

Therefore, $X_s - X_t$ is a multiple of n. Moreover, we have

$$X_s - X_t = \sum_{i=0}^{s} 10^i - \sum_{i=0}^{t} 10^i = \sum_{i=t+1}^{s} 10^i.$$

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \le k \le n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t = n(q_s - q_t).$$

Therefore, $X_s - X_t$ is a multiple of n. Moreover, we have

$$X_s - X_t = \sum_{i=0}^{s} 10^i - \sum_{i=0}^{t} 10^i = \sum_{i=t+1}^{s} 10^i.$$

So $X_s - X_t$ is a sum of distinct powers of ten, so its digits are 0s and 1s.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \le k \le n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t = n(q_s - q_t).$$

Therefore, $X_s - X_t$ is a multiple of n. Moreover, we have

$$X_s - X_t = \sum_{i=0}^{s} 10^i - \sum_{i=0}^{t} 10^i = \sum_{i=t+1}^{s} 10^i.$$

So $X_s - X_t$ is a sum of distinct powers of ten, so its digits are 0s and 1s. Therefore $X_s - X_t$ is a nonzero multiple of n whose digits are all 0s and 1s.

Theorem: Every positive natural number has a nonzero multiple whose digits are all 0s and 1s.

Proof: Let n be an arbitrary positive natural number. For every natural number k in the range $0 \leq k \leq n$, define X_k as

$$X_k = \sum_{i=0}^{k} 10^i.$$

Now, consider the remainders of the X_k's modulo n. Since there are $n + 1$ X_k's and n remainders modulo n, by the pigeonhole principle there must be at least two X_k's that leave the same remainder modulo n. Let X_s and X_t be two of these numbers and let r be that remainder. Without loss of generality, let $s > t$.

Since $X_s \equiv_n r$, there exists a $q_s \in \mathbb{Z}$ such that $X_s = nq_s + r$. Since $X_t \equiv_n r$, there exists a $q_t \in \mathbb{Z}$ such that $X_t = nq_t + r$. Therefore,

$$X_s - X_t = (nq_s + r) - (nq_t + r) = nq_s - nq_t = n(q_s - q_t).$$

Therefore, $X_s - X_t$ is a multiple of n. Moreover, we have

$$X_s - X_t = \sum_{i=0}^{s} 10^i - \sum_{i=0}^{t} 10^i = \sum_{i=t+1}^{s} 10^i.$$

So $X_s - X_t$ is a sum of distinct powers of ten, so its digits are 0s and 1s. Therefore $X_s - X_t$ is a nonzero multiple of n whose digits are all 0s and 1s. ■

Announcements!

Friday Four Square!
Today at 4:15PM, in front of Gates

Problem Set Three

- Problem Set Two due at the start of today's lecture, or Monday with a late period.
- Problem Set Three out.
  - Checkpoint due next Monday at the start of lecture.
  - Rest of the problem set due Friday.
  - Play around with graphs, relations, and the pigeonhole principle!

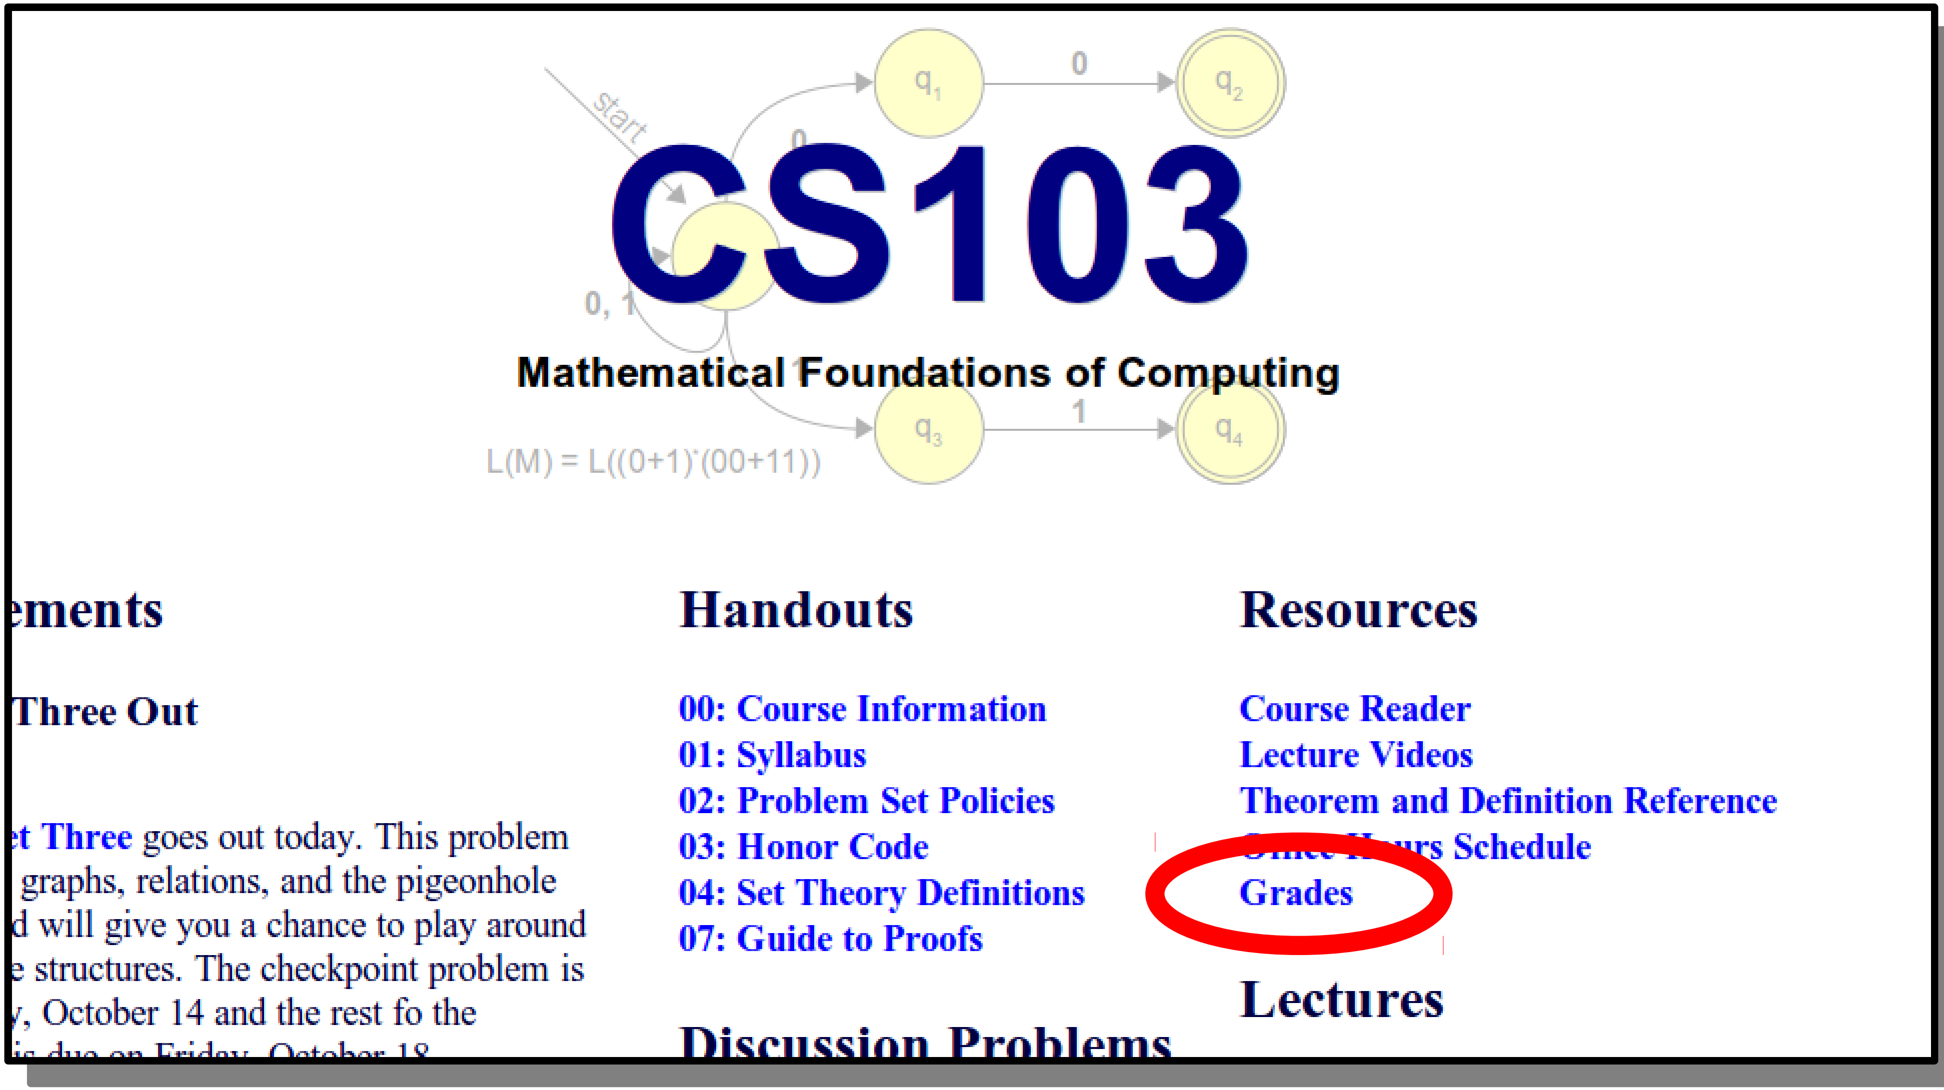
CS103
Mathematical Foundations of Computing
Handouts
00: Course Information
01: Syllabus
02: Problem Set Policies
03: Honor Code
04: Set Theory Definitions
07: Guide to Proofs
Resources
Course Reader
Lecture Videos
Theorem and Definition Reference
Grades
Lectures

Your Questions

“How do you decide whether a statement needs to be proved with a lemma or is counted as logical reasoning?”

"Can we email you or TAs questions we have about homework?"

Yes! Please!

Functions

A **function** is a means of associating each object in one set with an object in some other set.

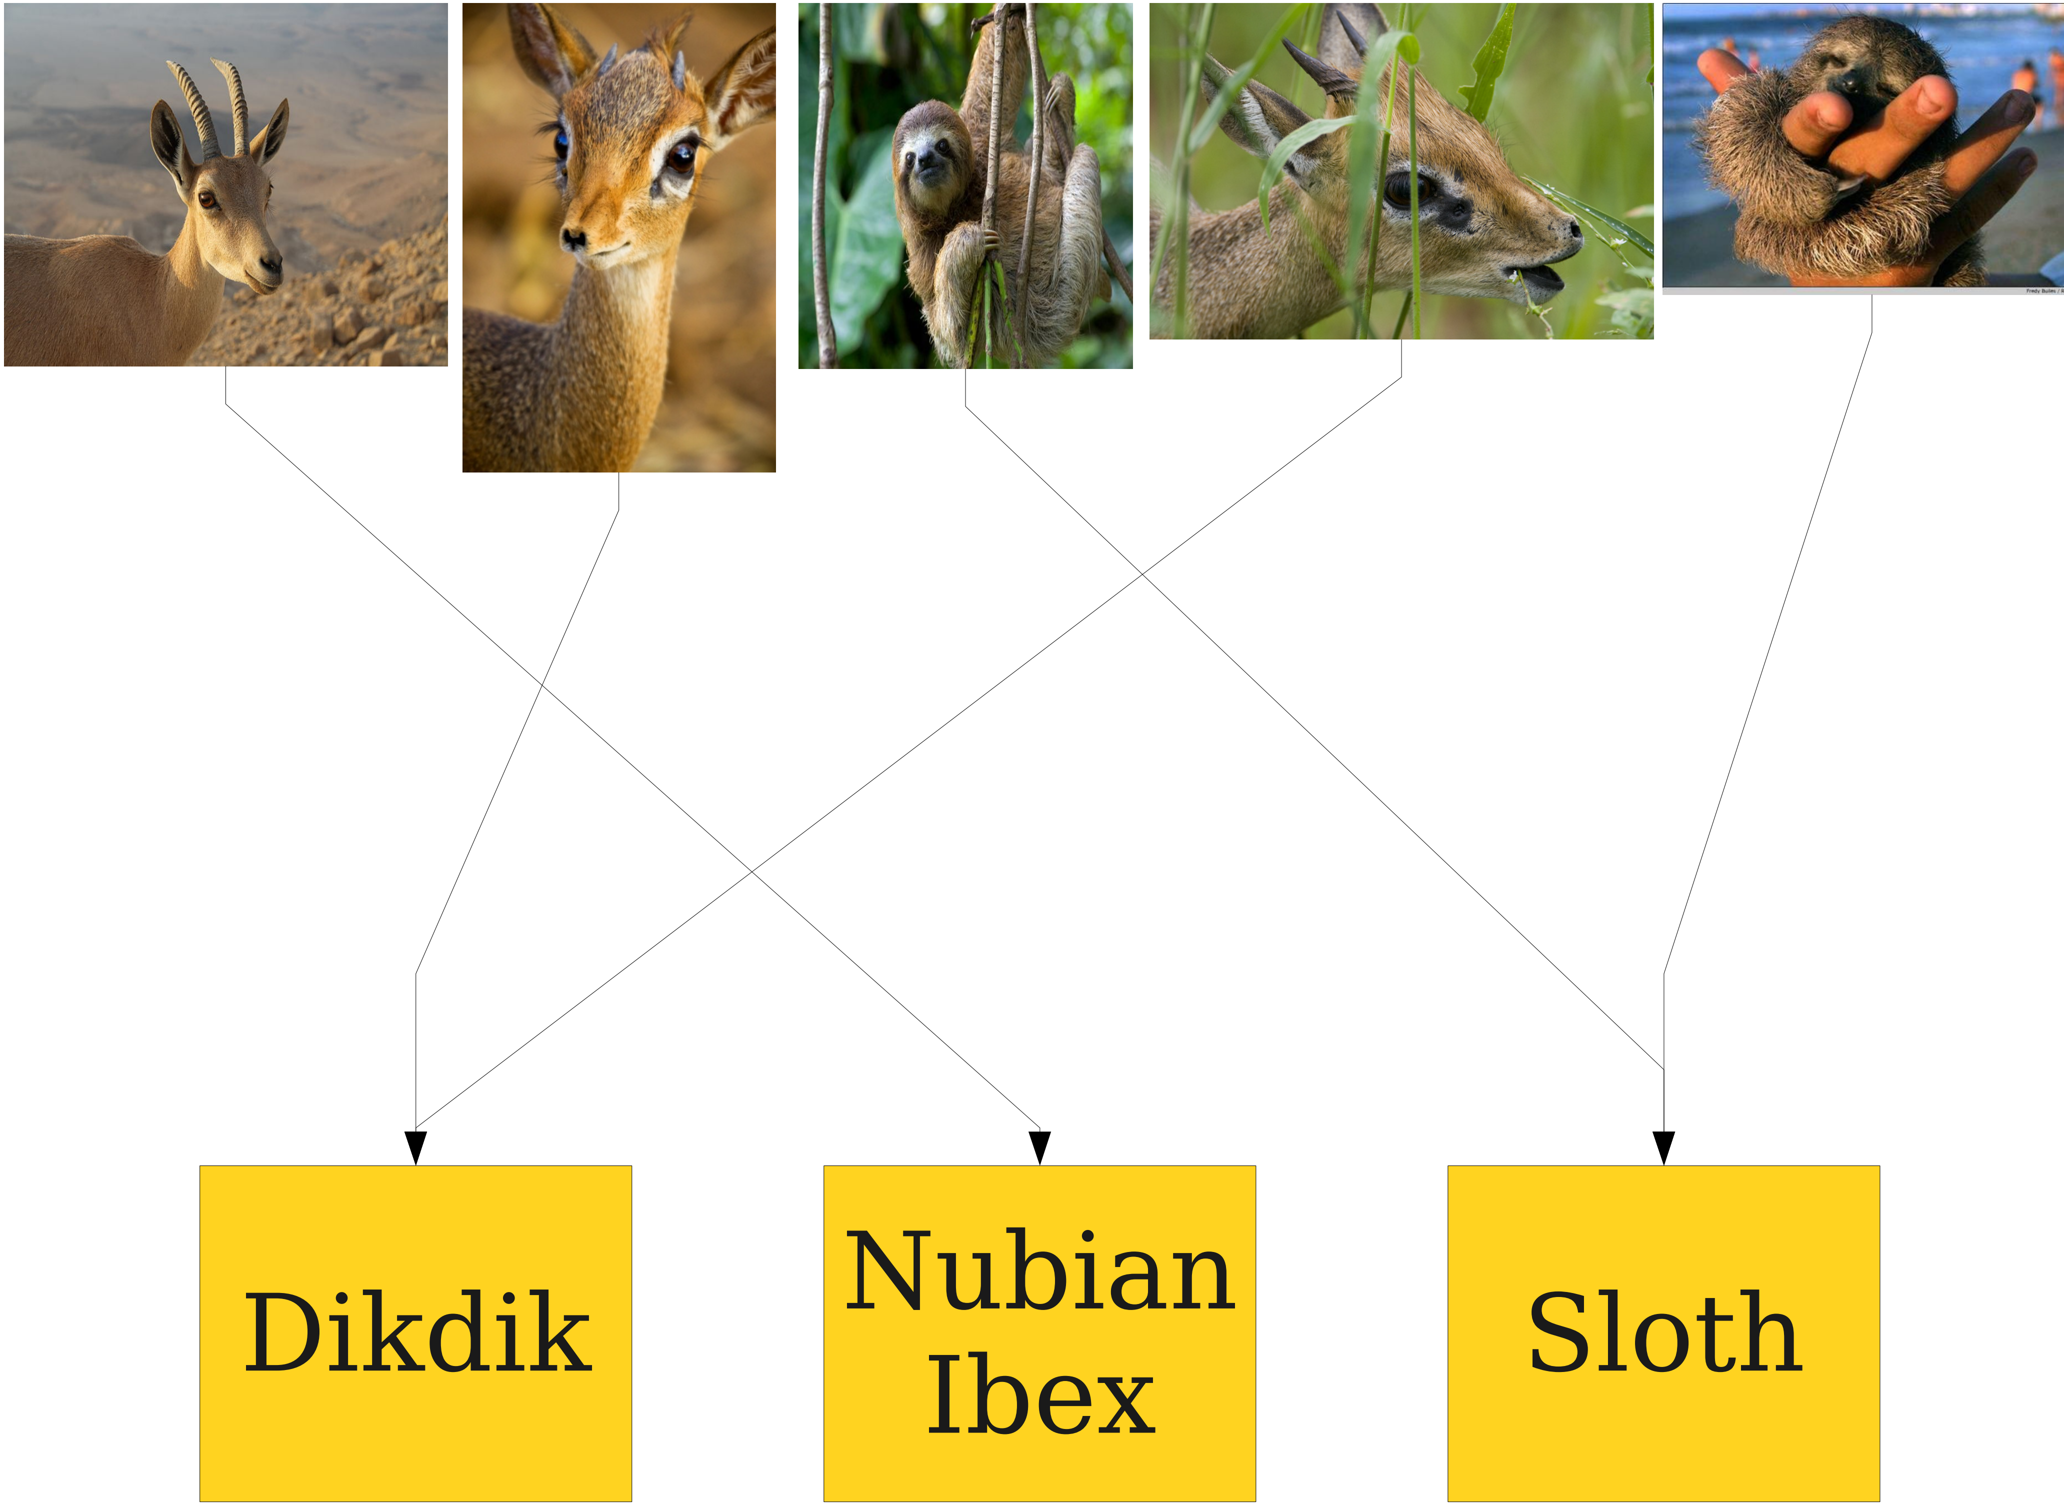
Dikdik
Nubian Ibex
Sloth

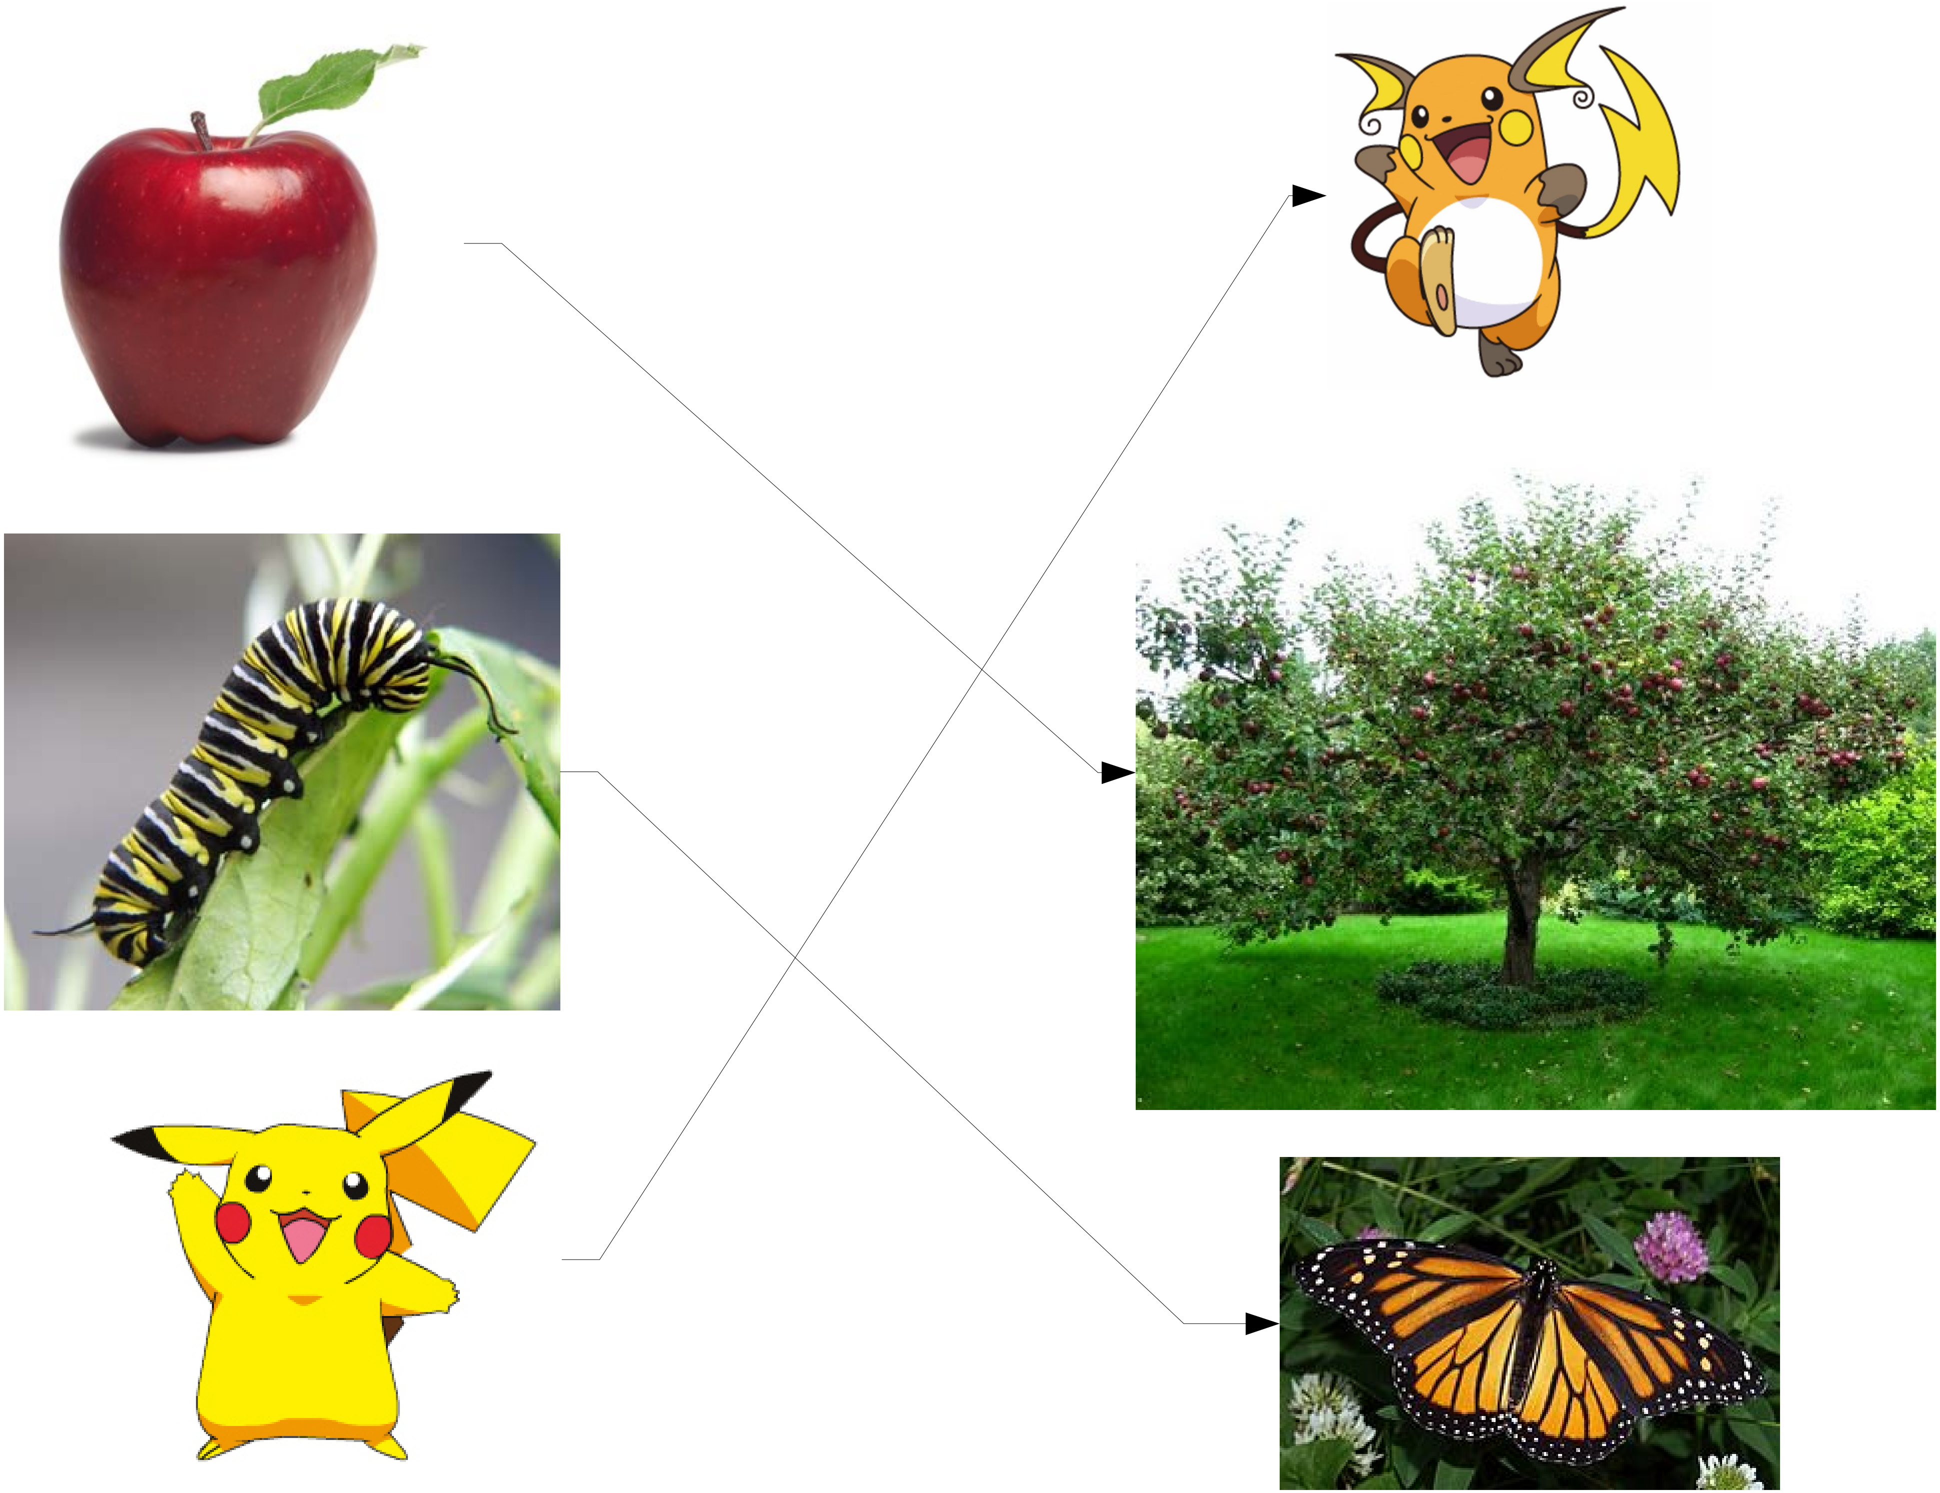

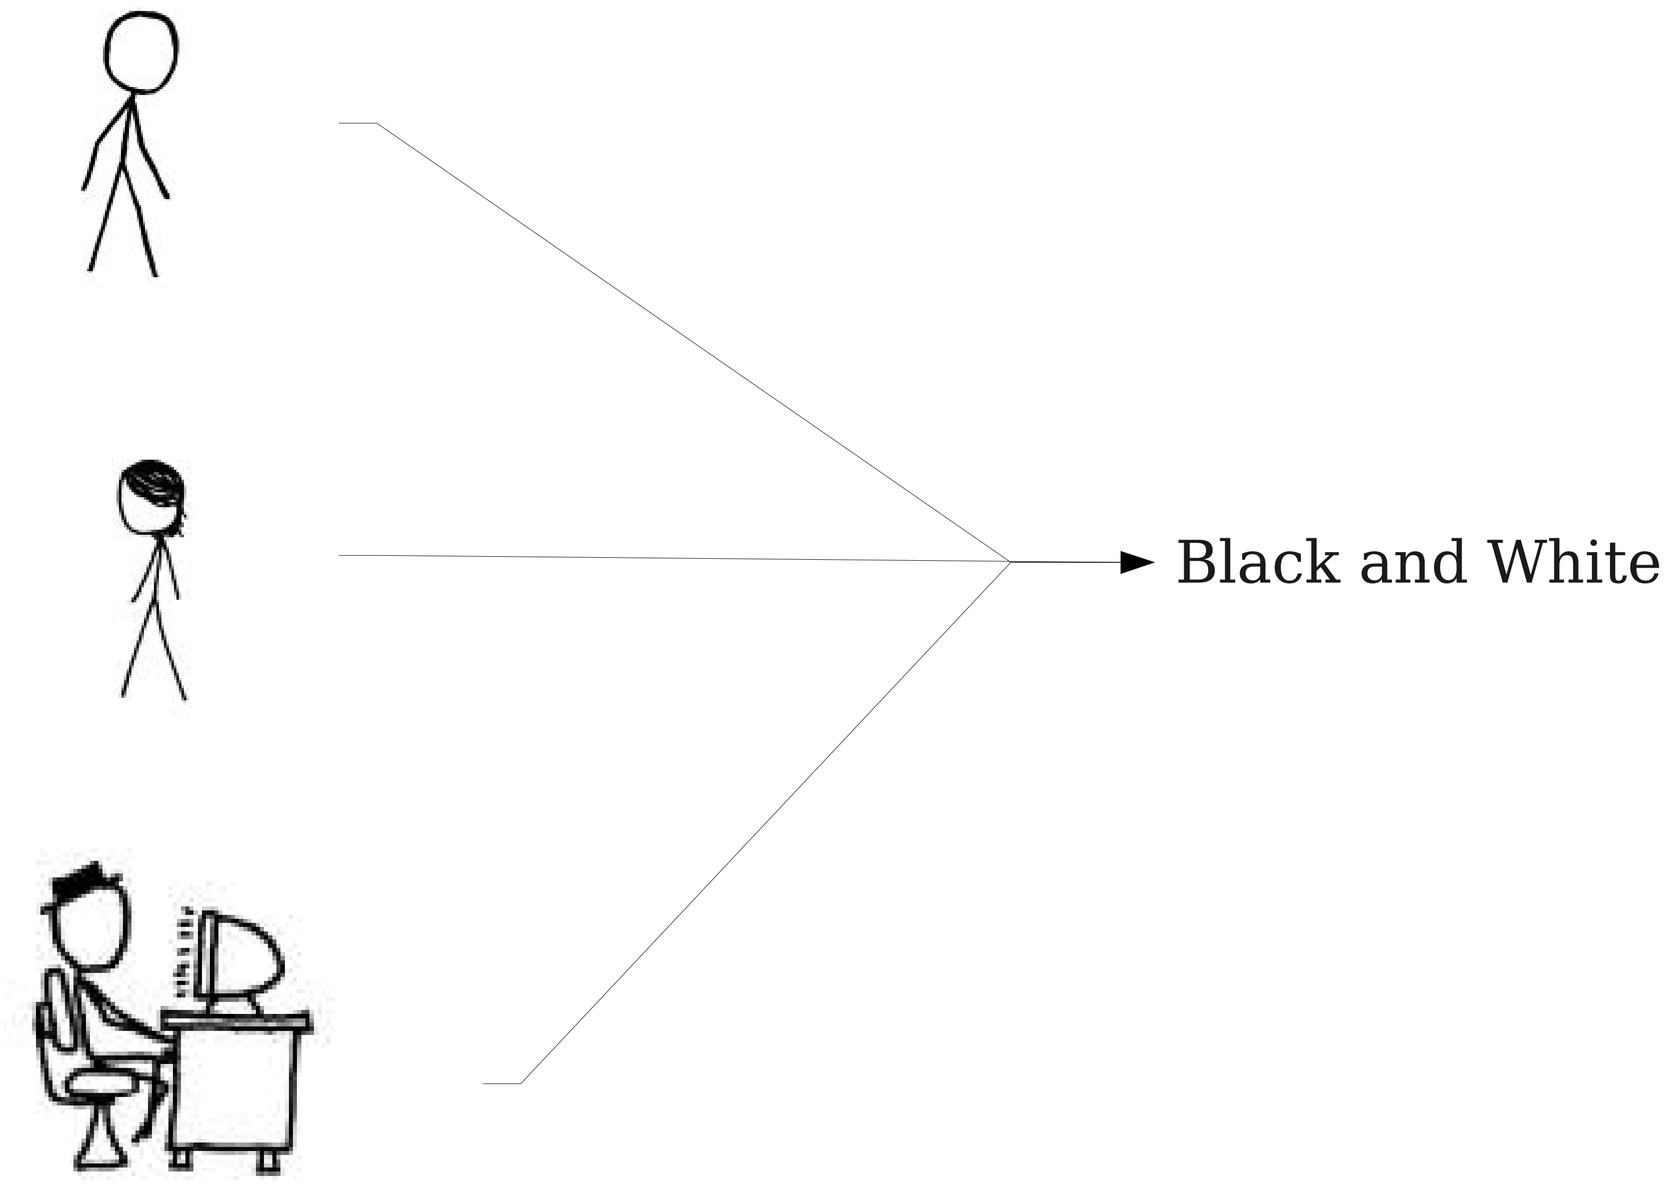
Black and White

Terminology

- A **function** f is a mapping such that every element of A is associated with a single element of B.
  - For each $a \in A$, there is some $b \in B$ with $f(a) = b$.
  - If $f(a) = b_0$ and $f(a) = b_1$, then $b_0 = b_1$.
- If f is a function from A to B, we say that f is a **mapping** from A to B.
  - We call A the **domain** of f.
  - We call B the **codomain** of f.
- We denote that f is a function from A to B by writing

$$\boldsymbol{f : A \to B}$$

Defining Functions

- Typically, we specify a function by describing a rule that maps every element of the domain to some element of the codomain.
- Examples:
  - $f(n) = n + 1$, where $f : \mathbb{Z} \to \mathbb{Z}$
  - $f(x) = \sin x$, where $f : \mathbb{R} \to \mathbb{R}$
  - $f(x) = \lceil x \rceil$, where $f : \mathbb{R} \to \mathbb{Z}$
- Notice that we're giving both a rule and the domain/codomain.

Defining Functions

Typically, we specify a function by describing a rule that maps every element of the domain to some
codomain.

Examples:

$f(n) = n + 1$, where $f : \mathbb{Z}$

$f(x) = \sin x$, where $f : \mathbb{R} \to \mathbb{R}$

- $f(x) = \lceil x \rceil$, where $f : \mathbb{R} \to \mathbb{Z}$

This is the ceiling function – the smallest integer greater than or equal to x. For example, $\lceil 1 \rceil = 1$, $\lceil 1.37 \rceil = 2$, and $\lceil \pi \rceil = 4$.

Notice that we're giving both a rule and the domain/codomain.

Defining Functions

- Typically, we specify a function by describing a rule that maps every element of the domain to some element of the codomain.
- Examples:
  - $f(n) = n + 1$, where $f : \mathbb{Z} \to \mathbb{Z}$
  - $f(x) = \sin x$, where $f : \mathbb{R} \to \mathbb{R}$
  - $f(x) = \lceil x \rceil$, where $f : \mathbb{R} \to \mathbb{Z}$
- Notice that we're giving both a rule and the domain/codomain.

Is this a function from A to B?

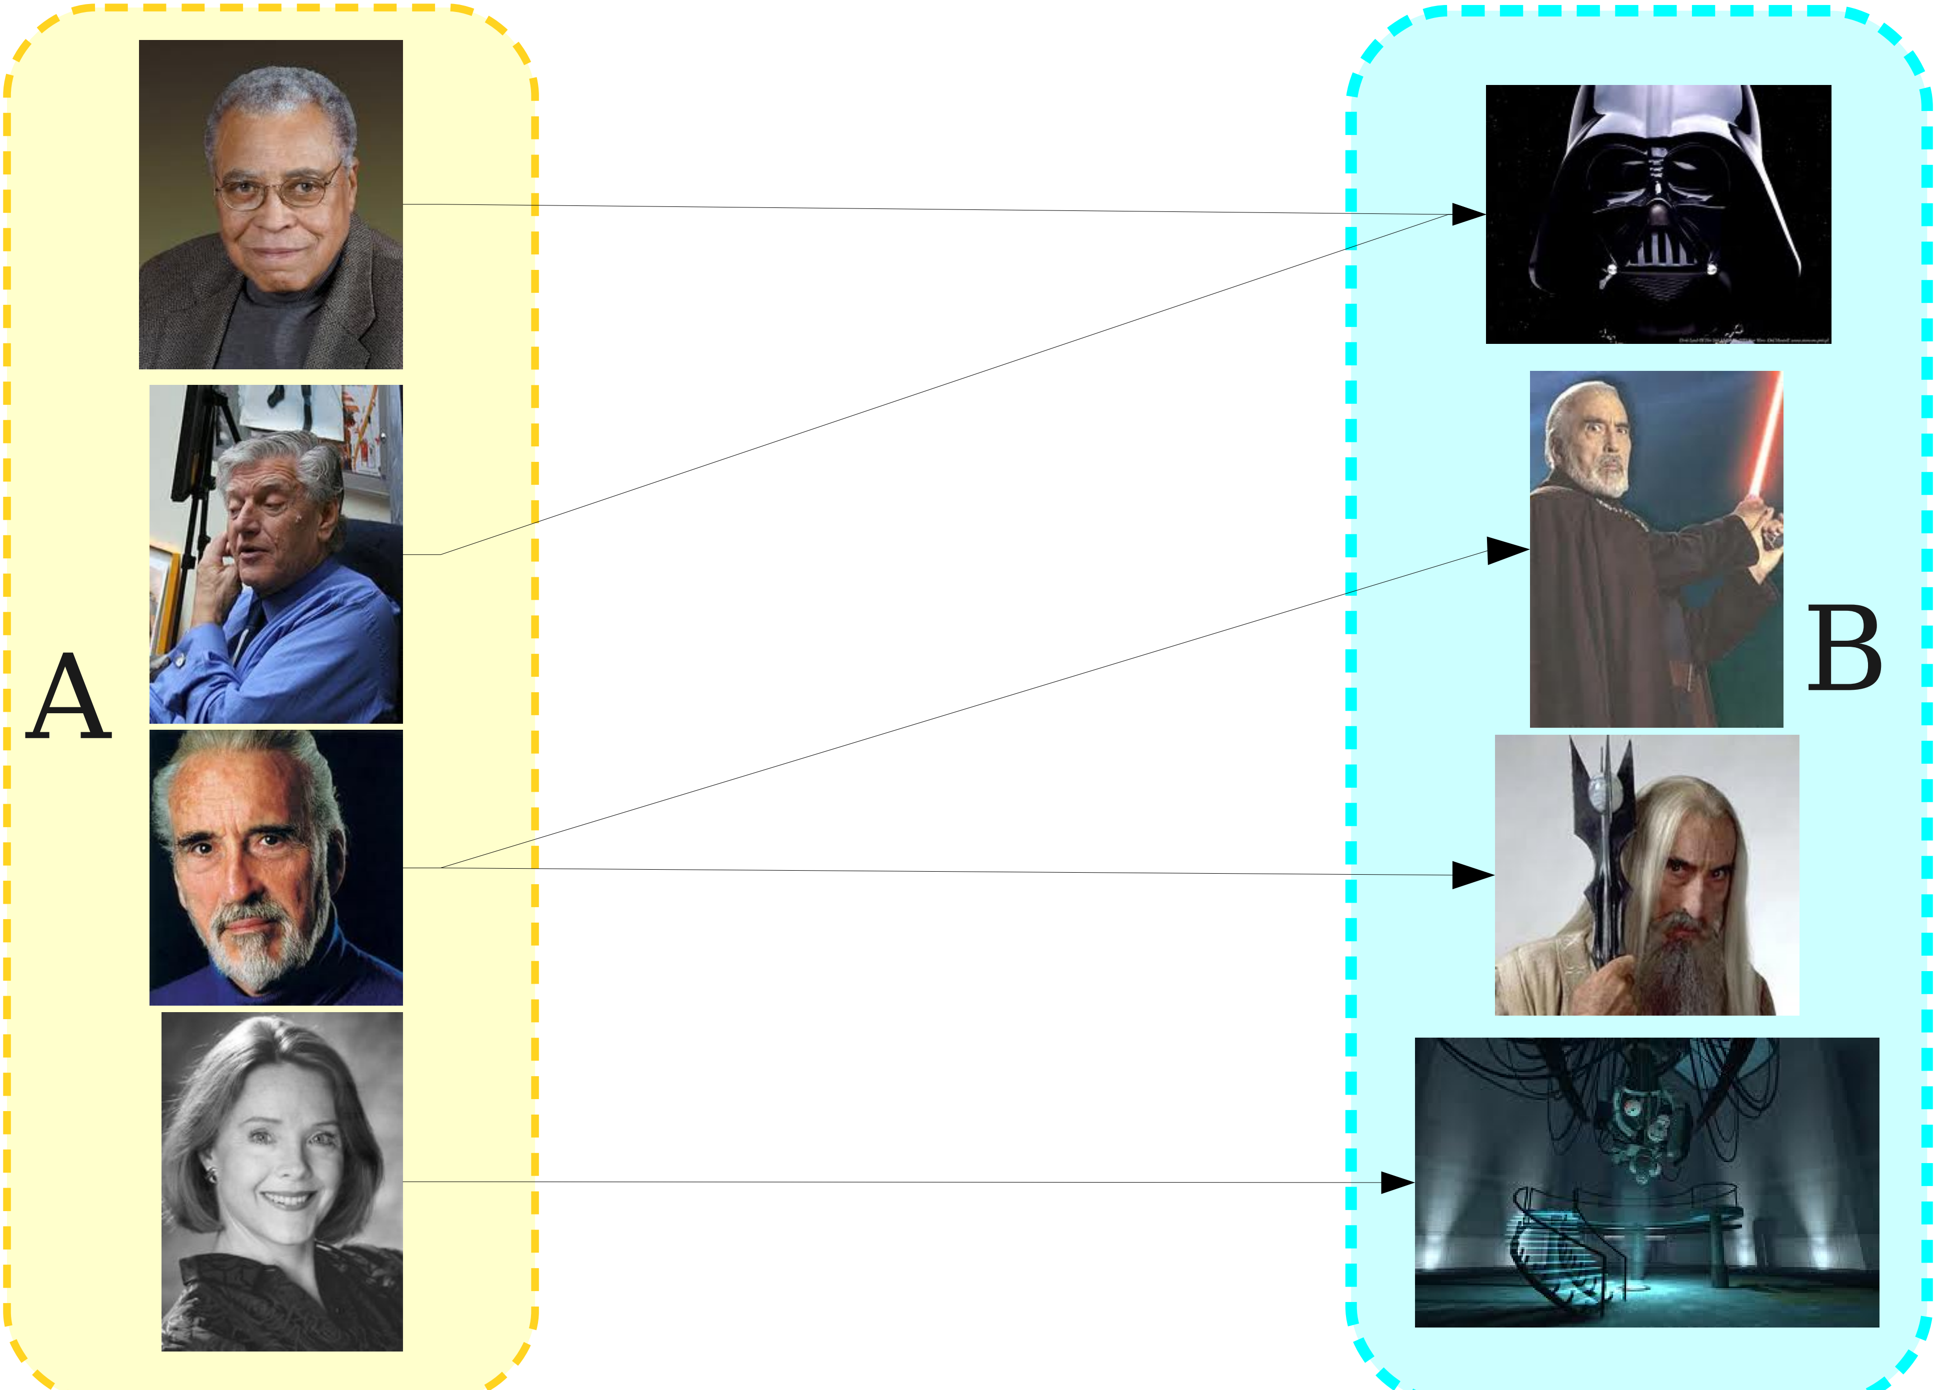

Is this a function from A to B?

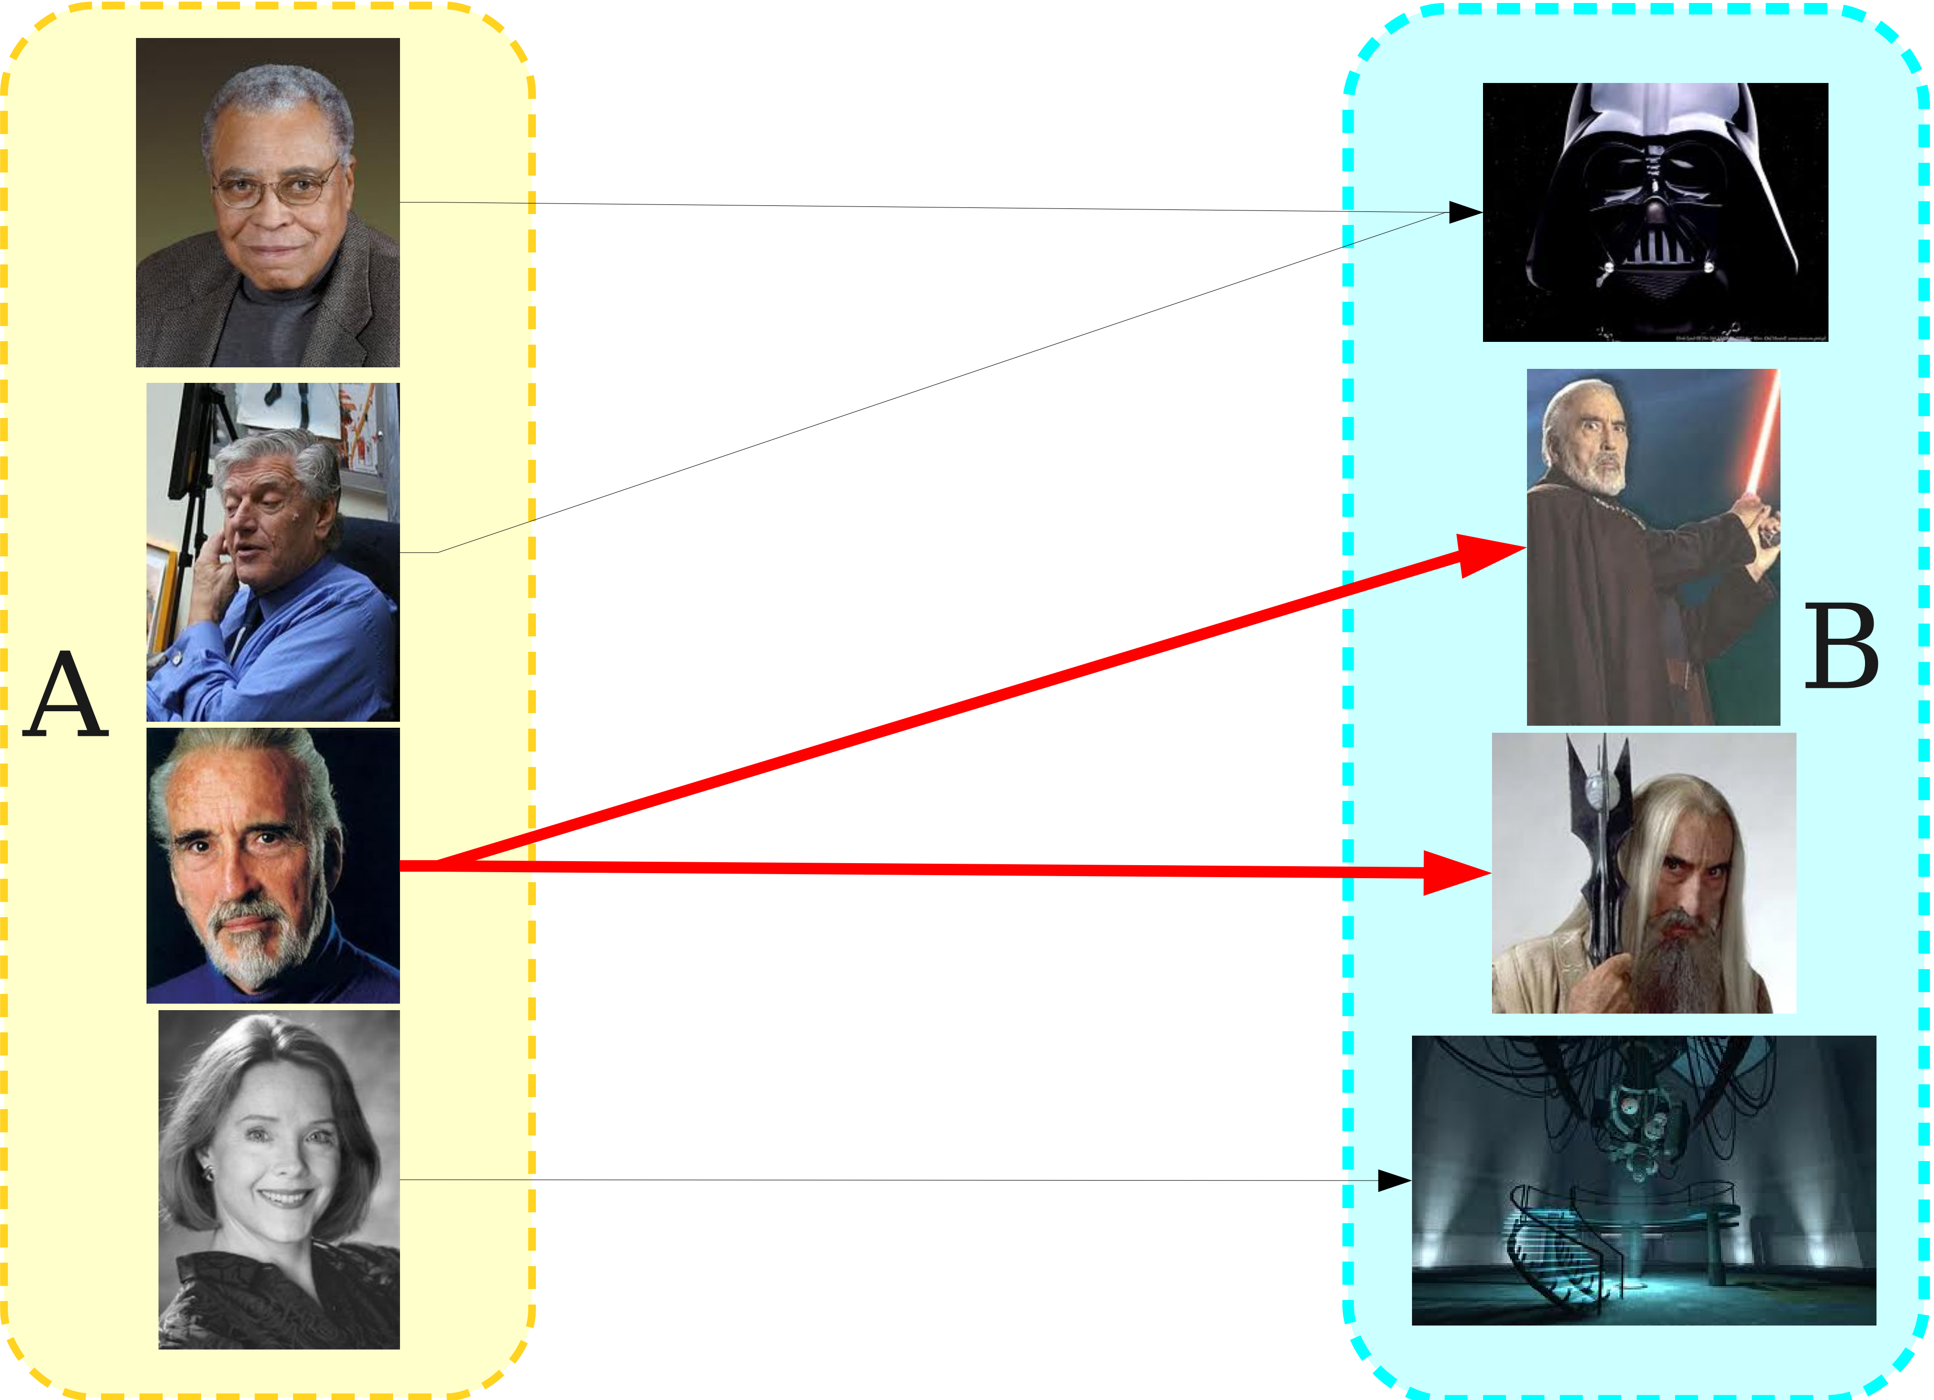

Is this a function from A to B?

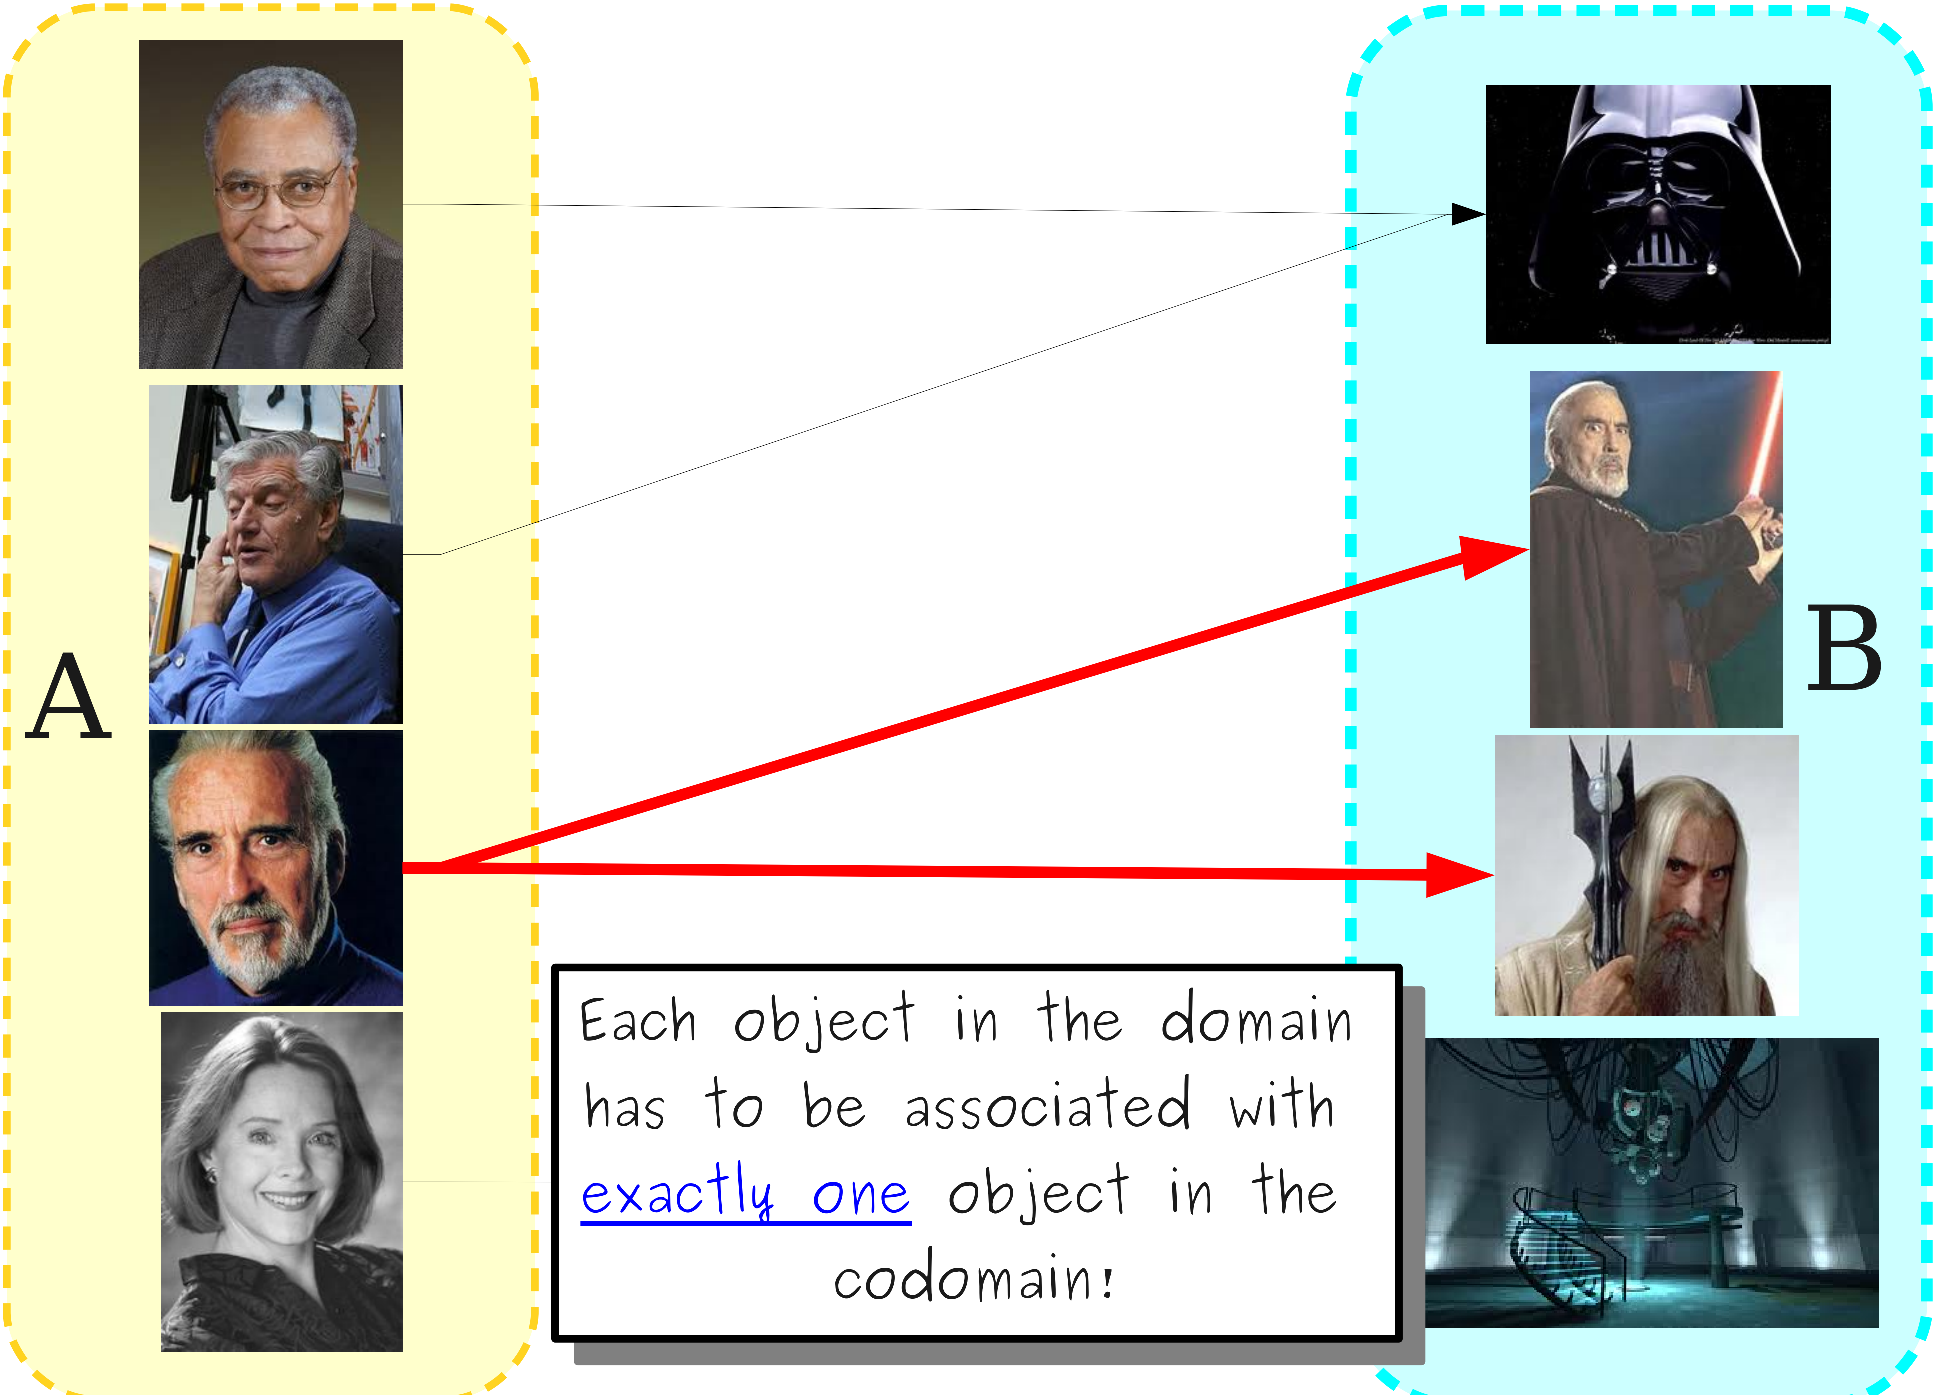

Is this a function from A to B?

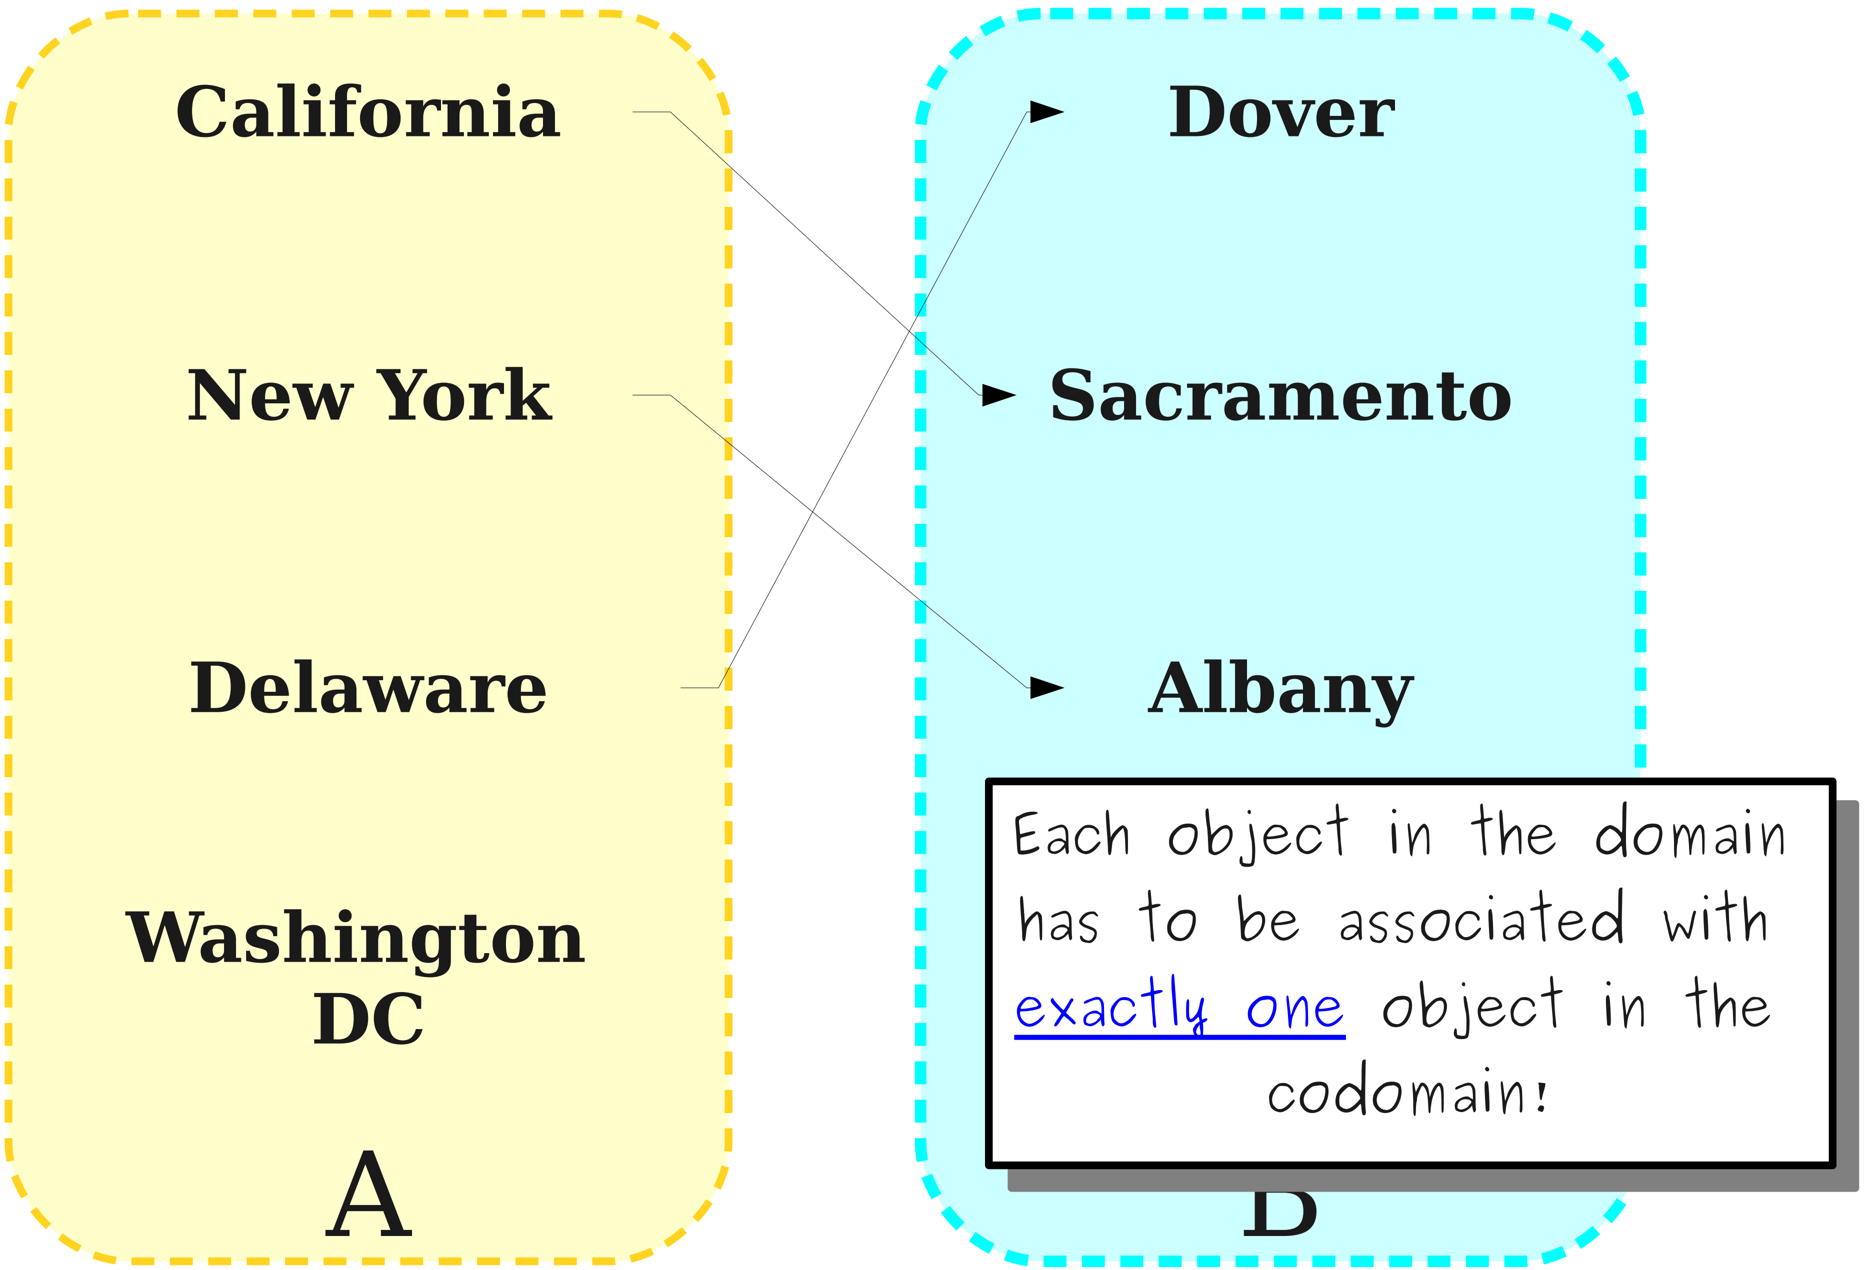

Is this a function from A to B?

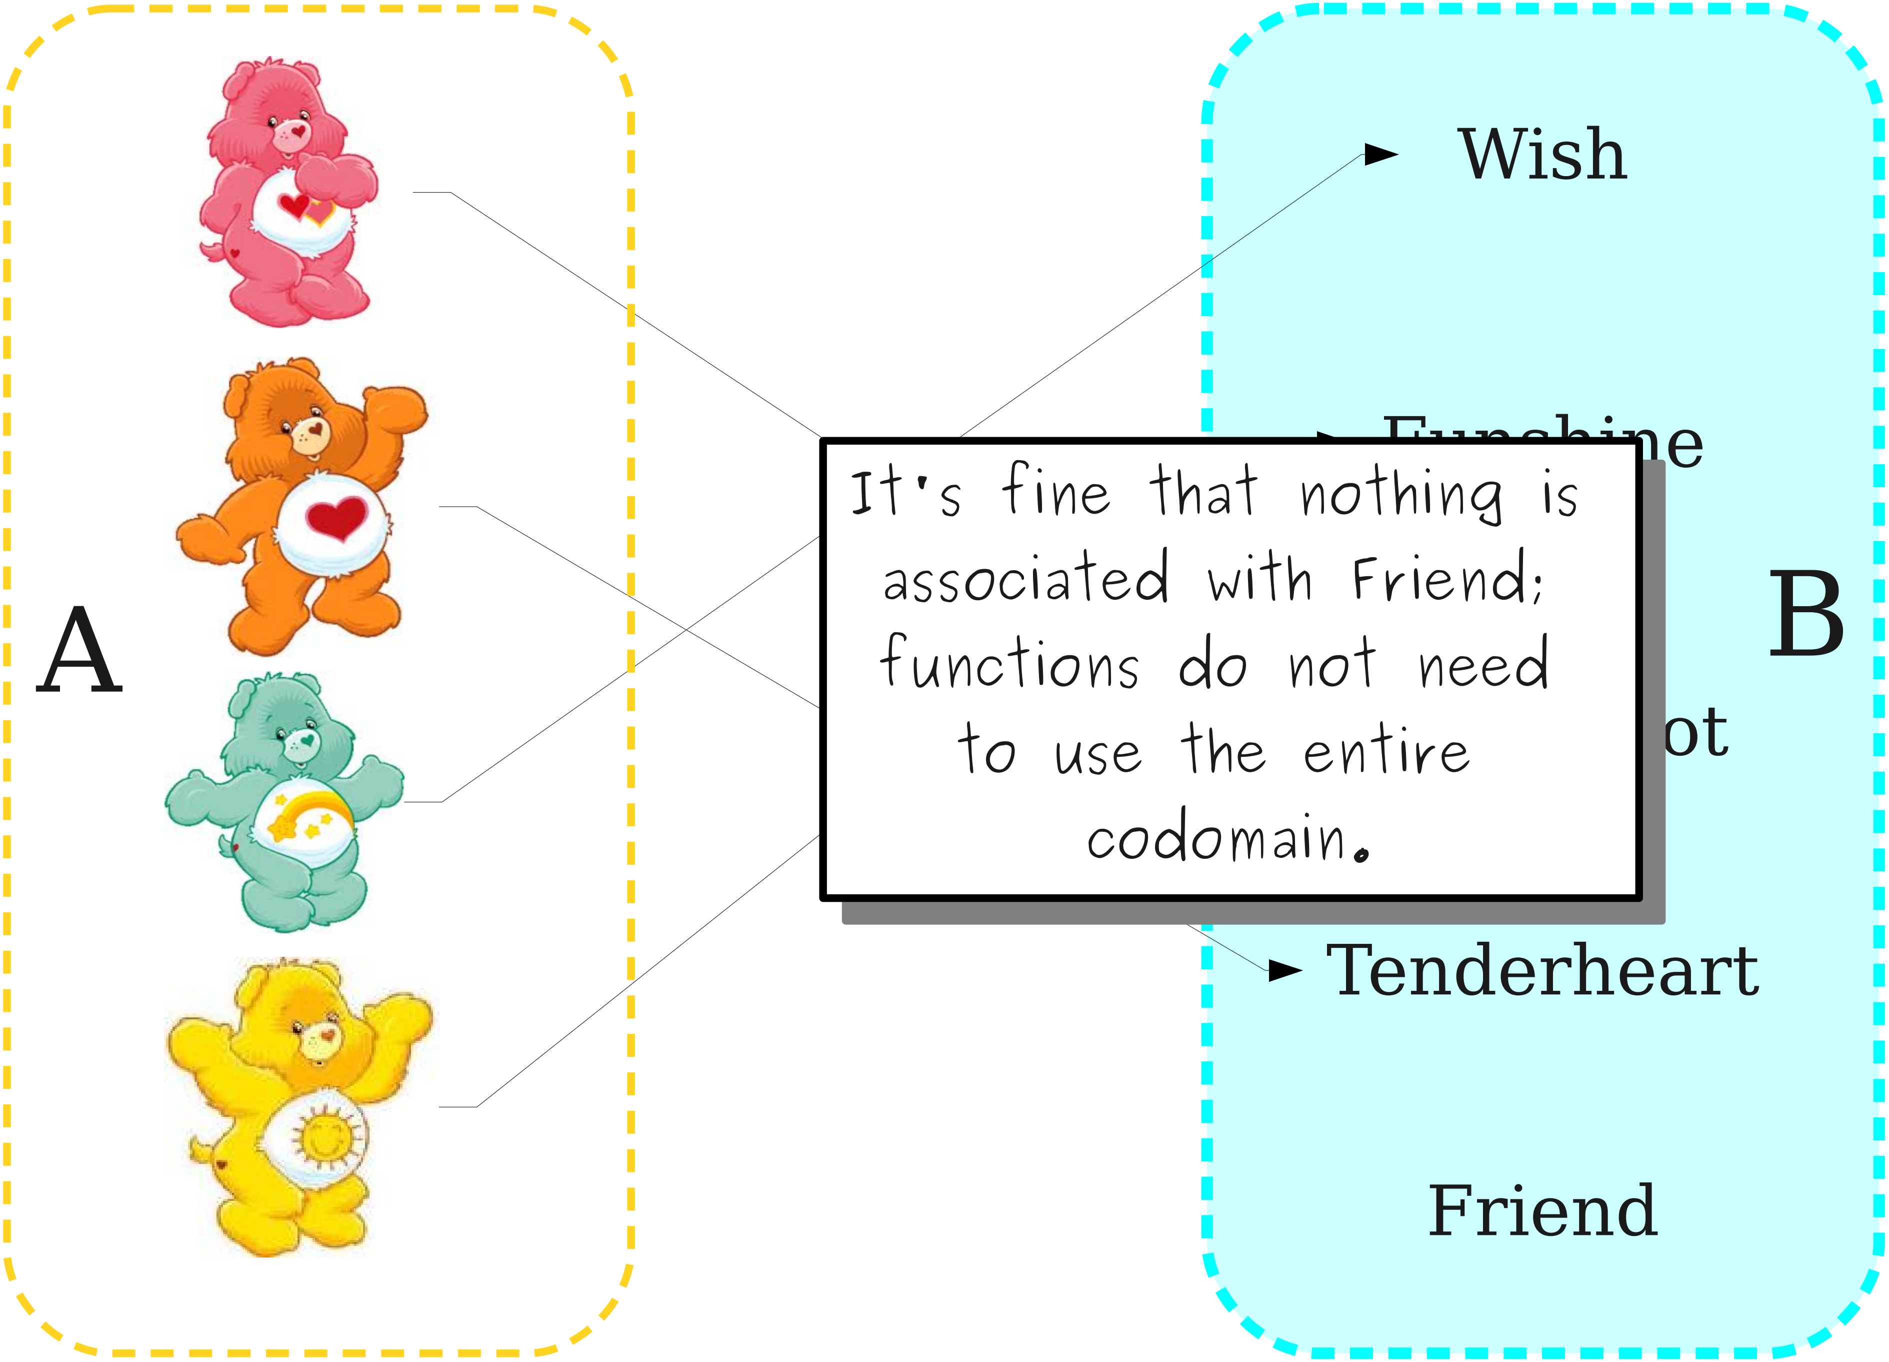

Piecewise Functions

- Functions may be specified **piecewise**, with different rules applying to different elements.
- Example:

$$f(n) = \begin{cases} -n/2 & \text{if } n \text{ is even} \\ (n+1)/2 & \text{otherwise} \end{cases}$$

- When defining a function piecewise, it's up to you to confirm that it defines a legal function!

☿

♀

♂

♃

♄

♅

♆

☿
♀
♂
♃
♄
♅
♆

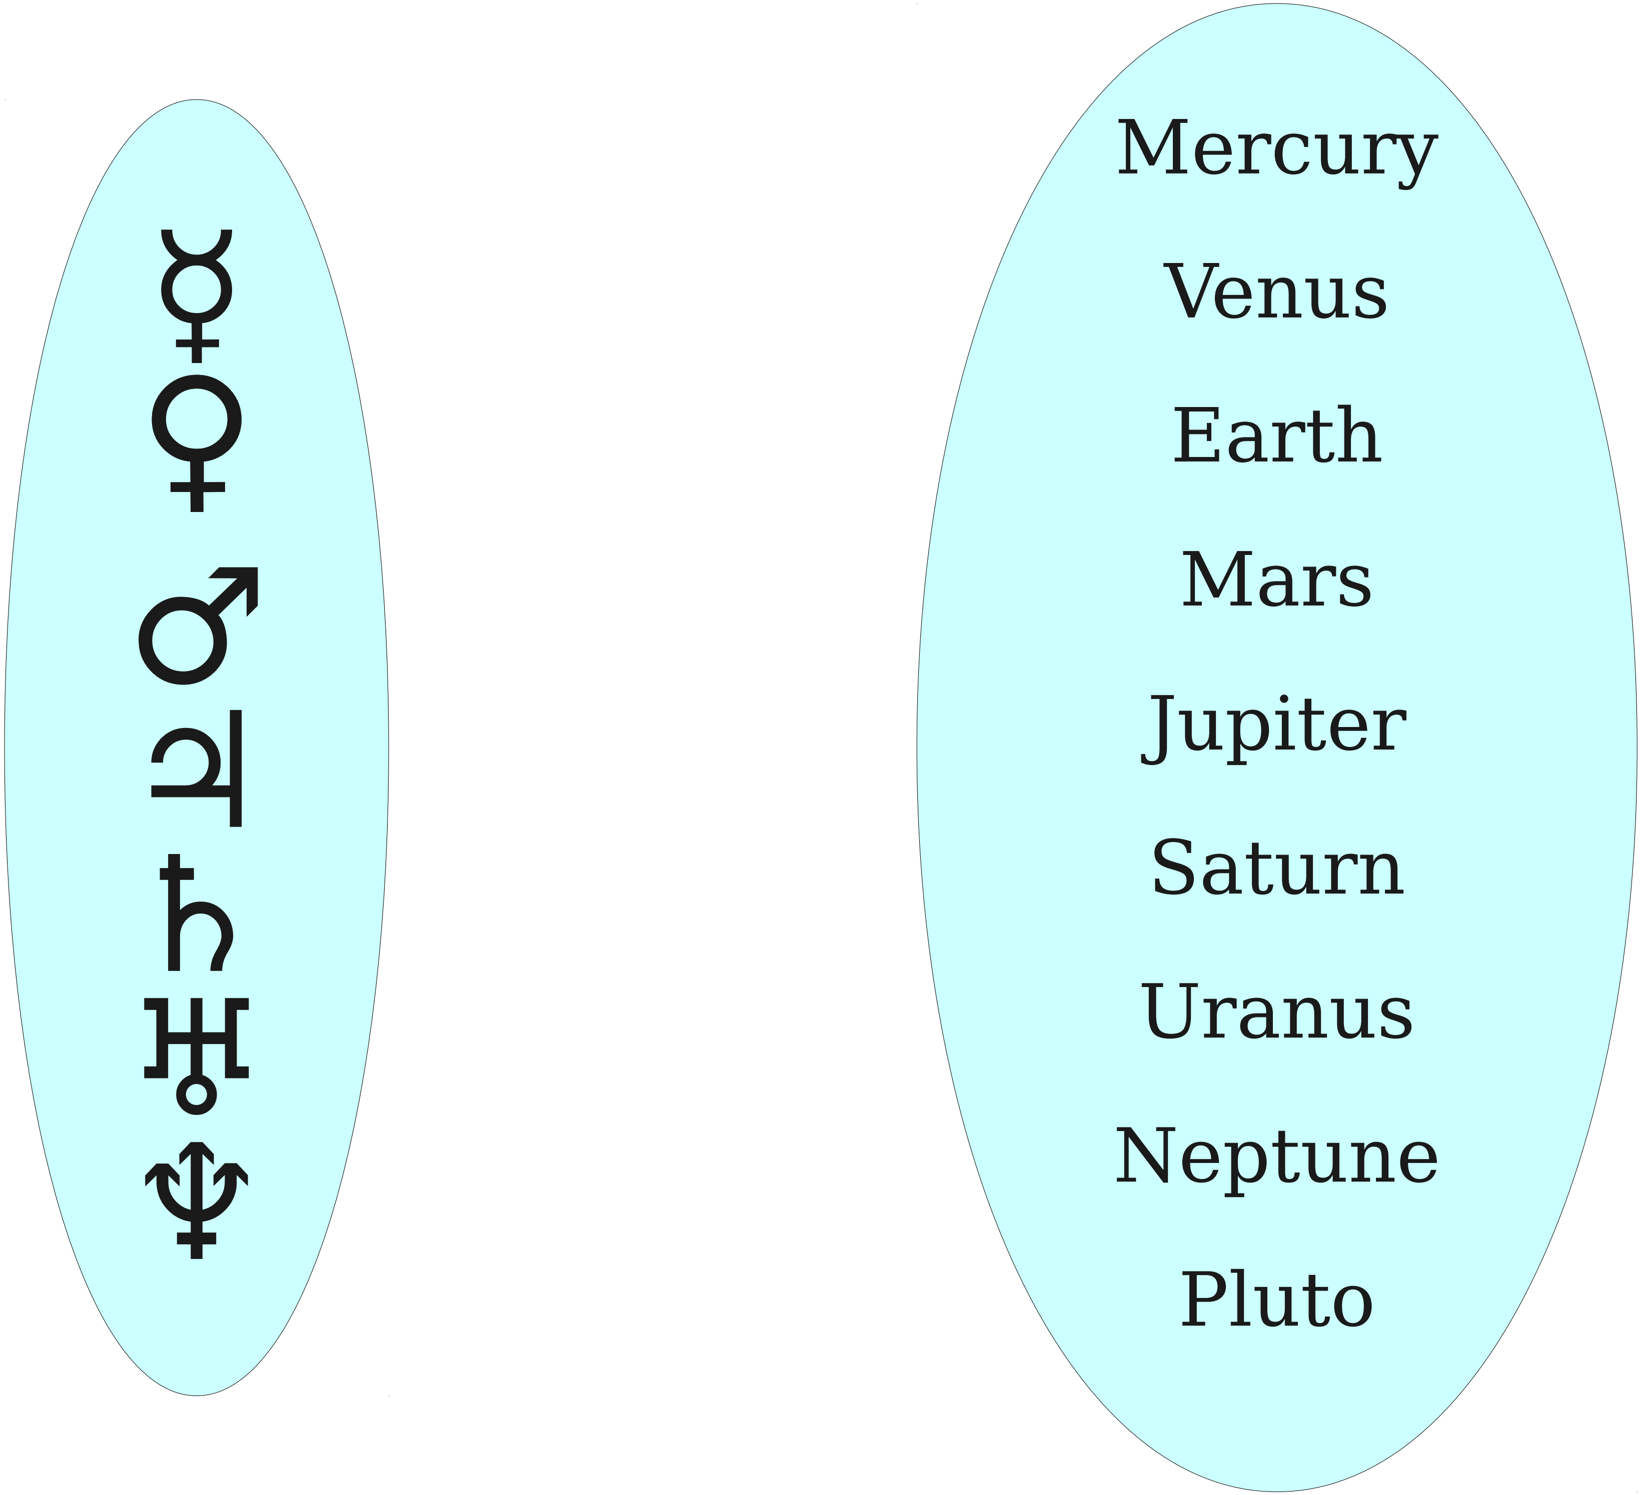
☿
♀
♂
♃
♄
♅
♆
Mercury
Venus
Earth
Mars
Jupiter
Saturn
Uranus
Neptune
Pluto

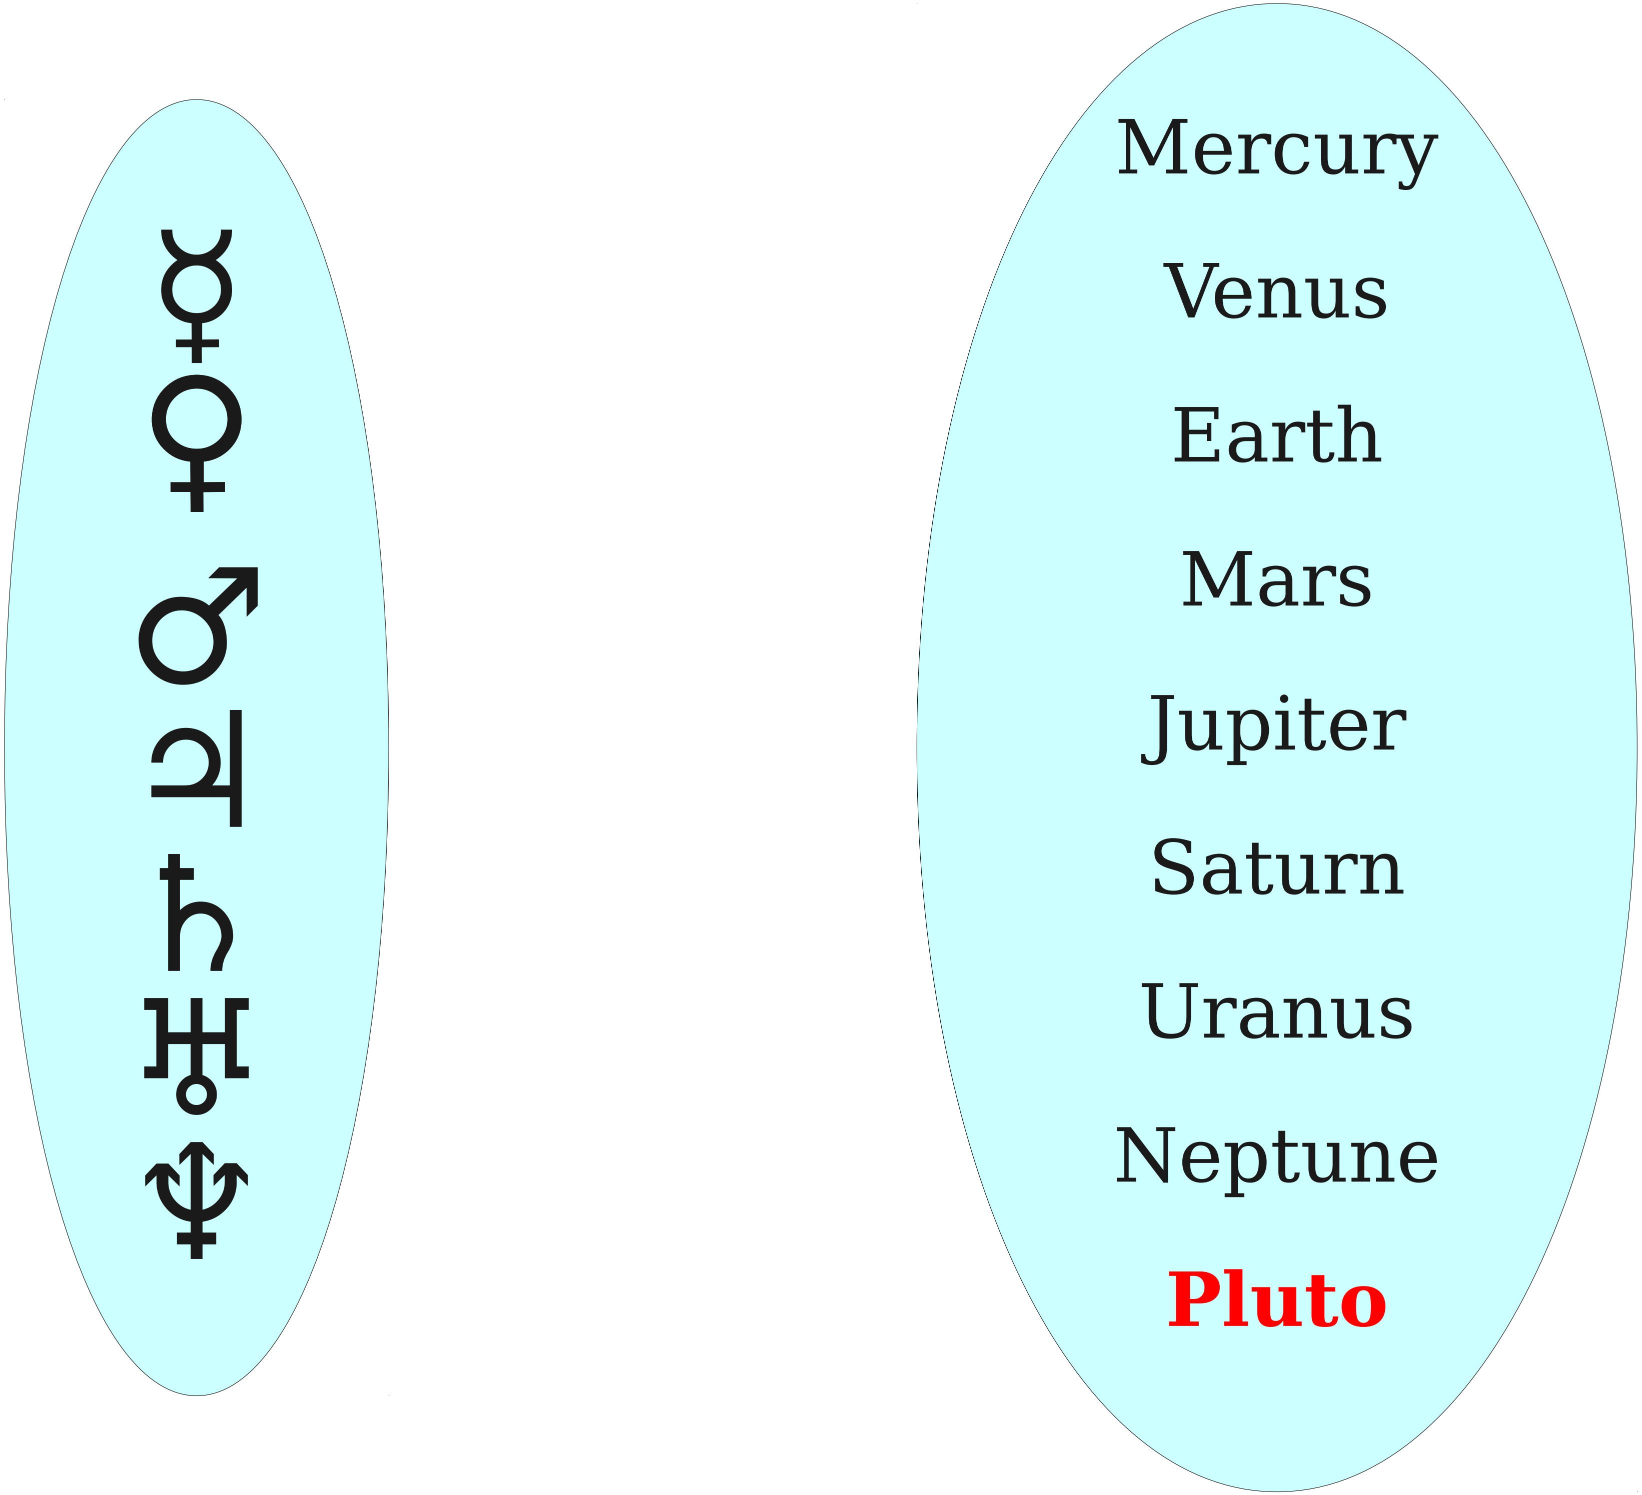
☿
♀
♂
♃
♄
♅
♆
Mercury
Venus
Earth
Mars
Jupiter
Saturn
Uranus
Neptune
Pluto

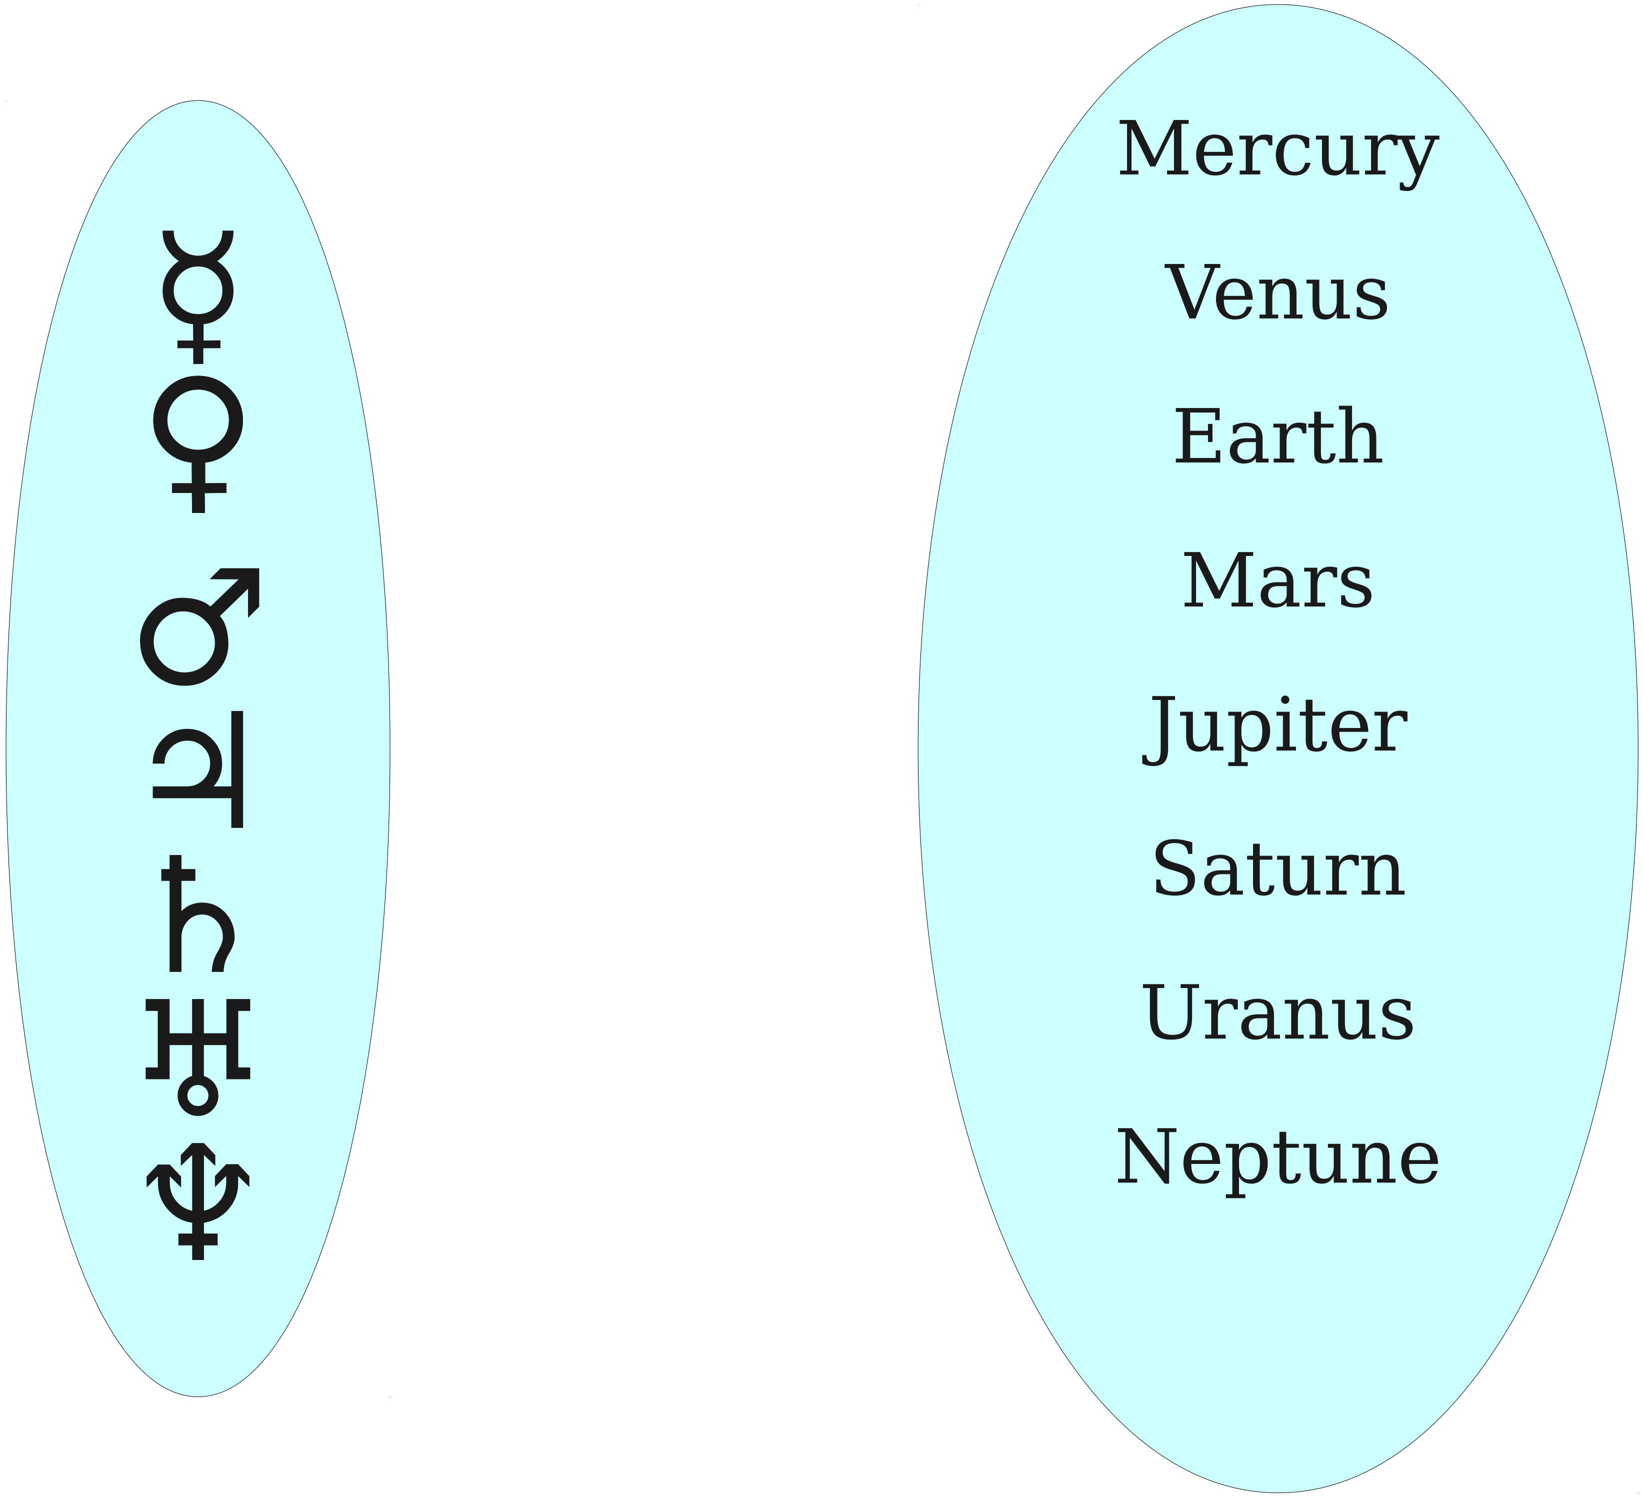
☿
♀
♂
♃
♄
♅
♆
Mercury
Venus
Earth
Mars
Jupiter
Saturn
Uranus
Neptune

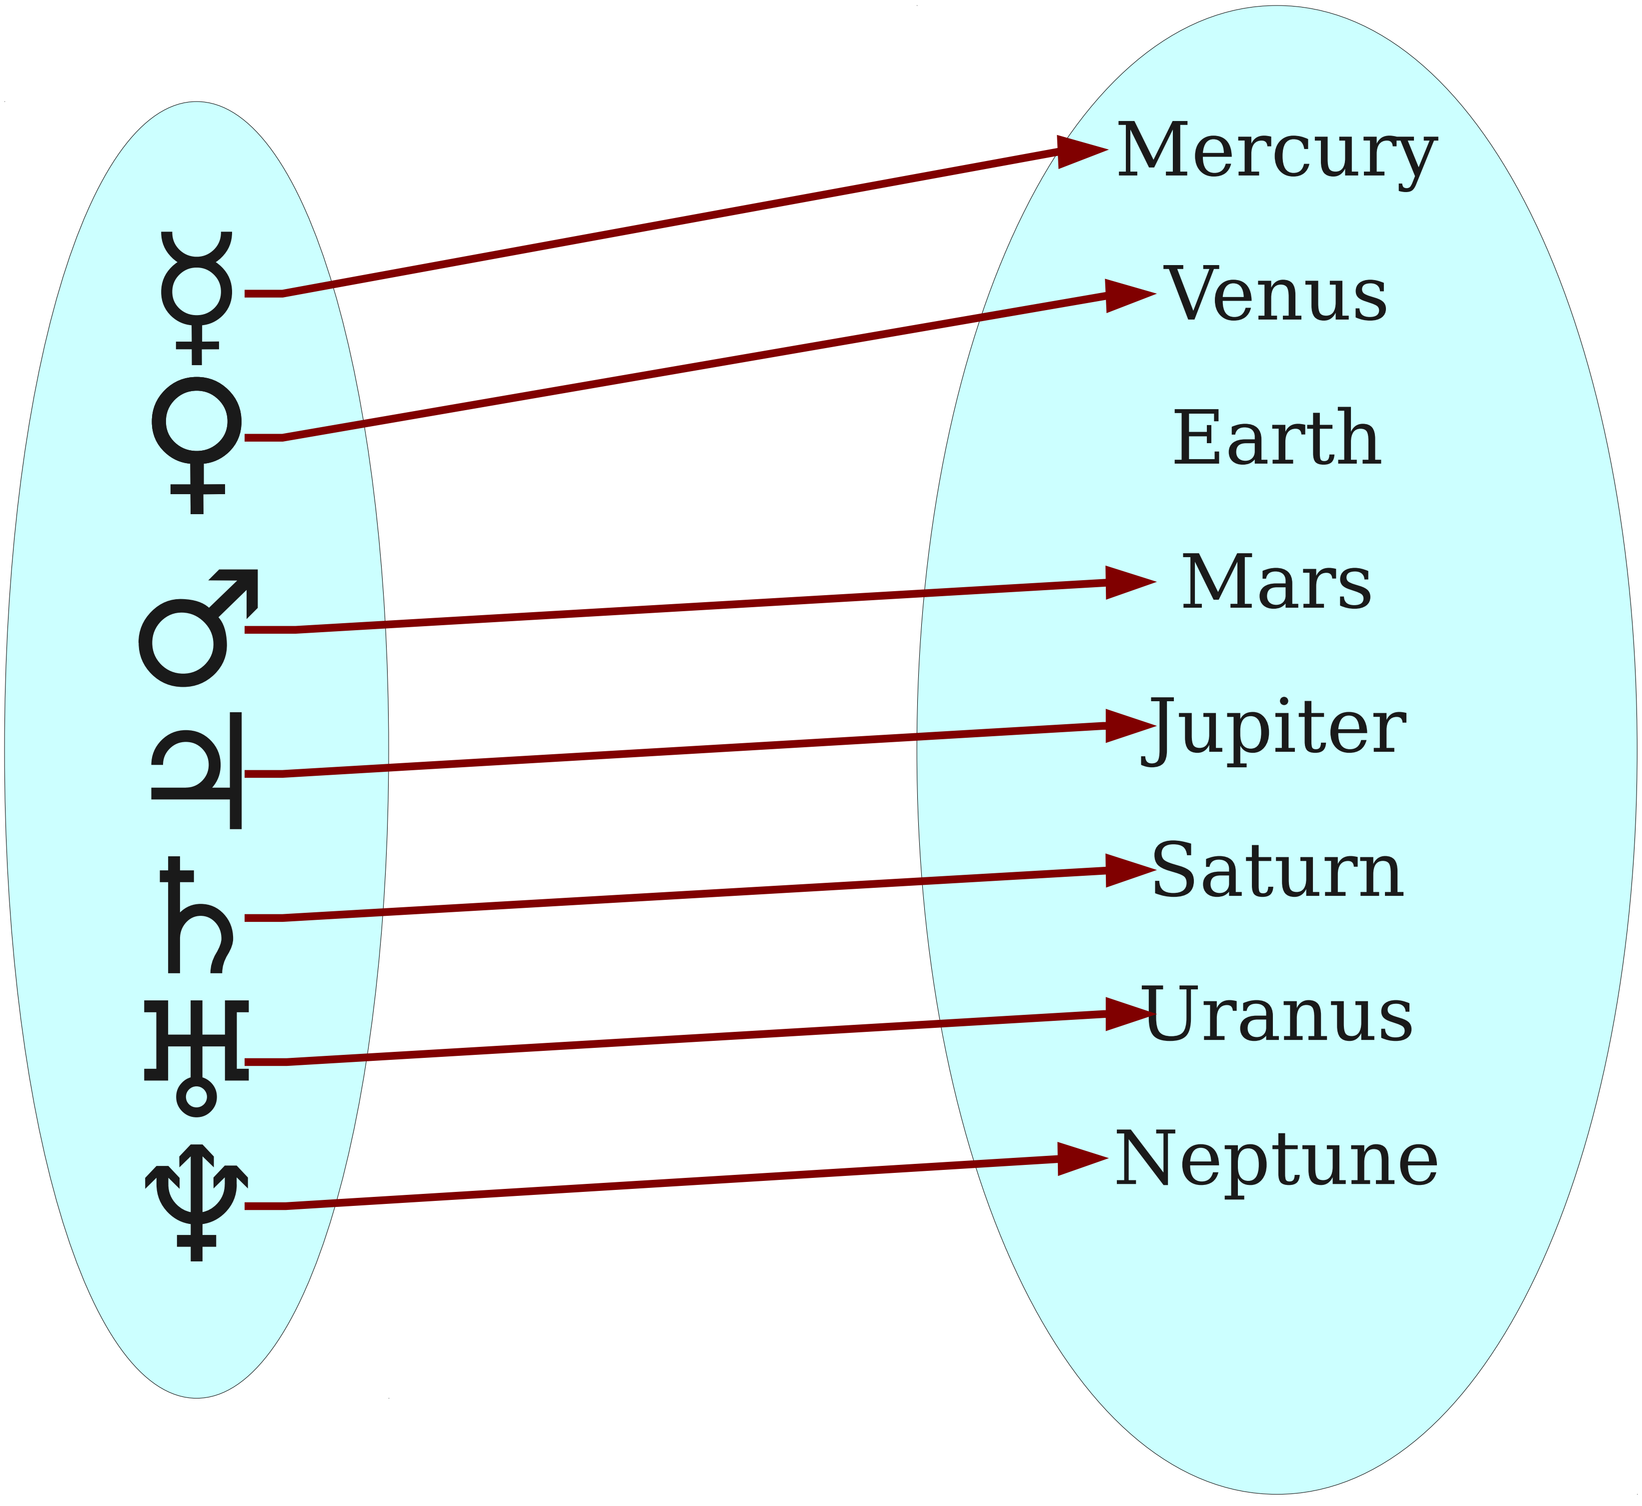
Mercury
Venus
Earth
Mars
Jupiter
Saturn
Uranus
Neptune

Injective Functions

- A function $f : A \to B$ is called **injective** (or **one-to-one**) if each element of the codomain has at most one element of the domain that maps to it.
  - A function with this property is called an **injection**.
- Formally, $f : A \to B$ is an injection iff

$$\textbf{For any } x_0, x_1 \in A\textbf{:}$$
$$\textbf{if } f(x_0) = f(x_1)\textbf{, then } x_0 = x_1$$

- An intuition: injective functions label the objects from A using names from B.


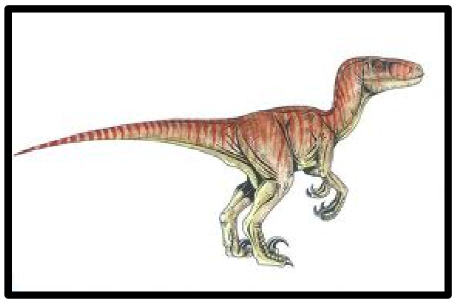


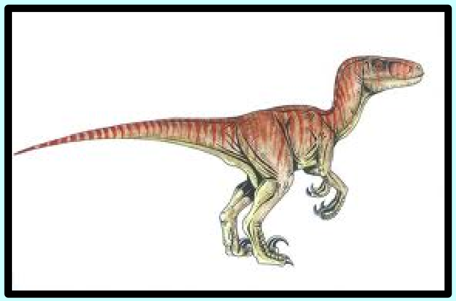

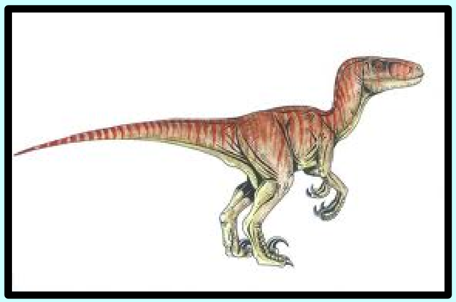

Front Door

Balcony Window

Bedroom Window

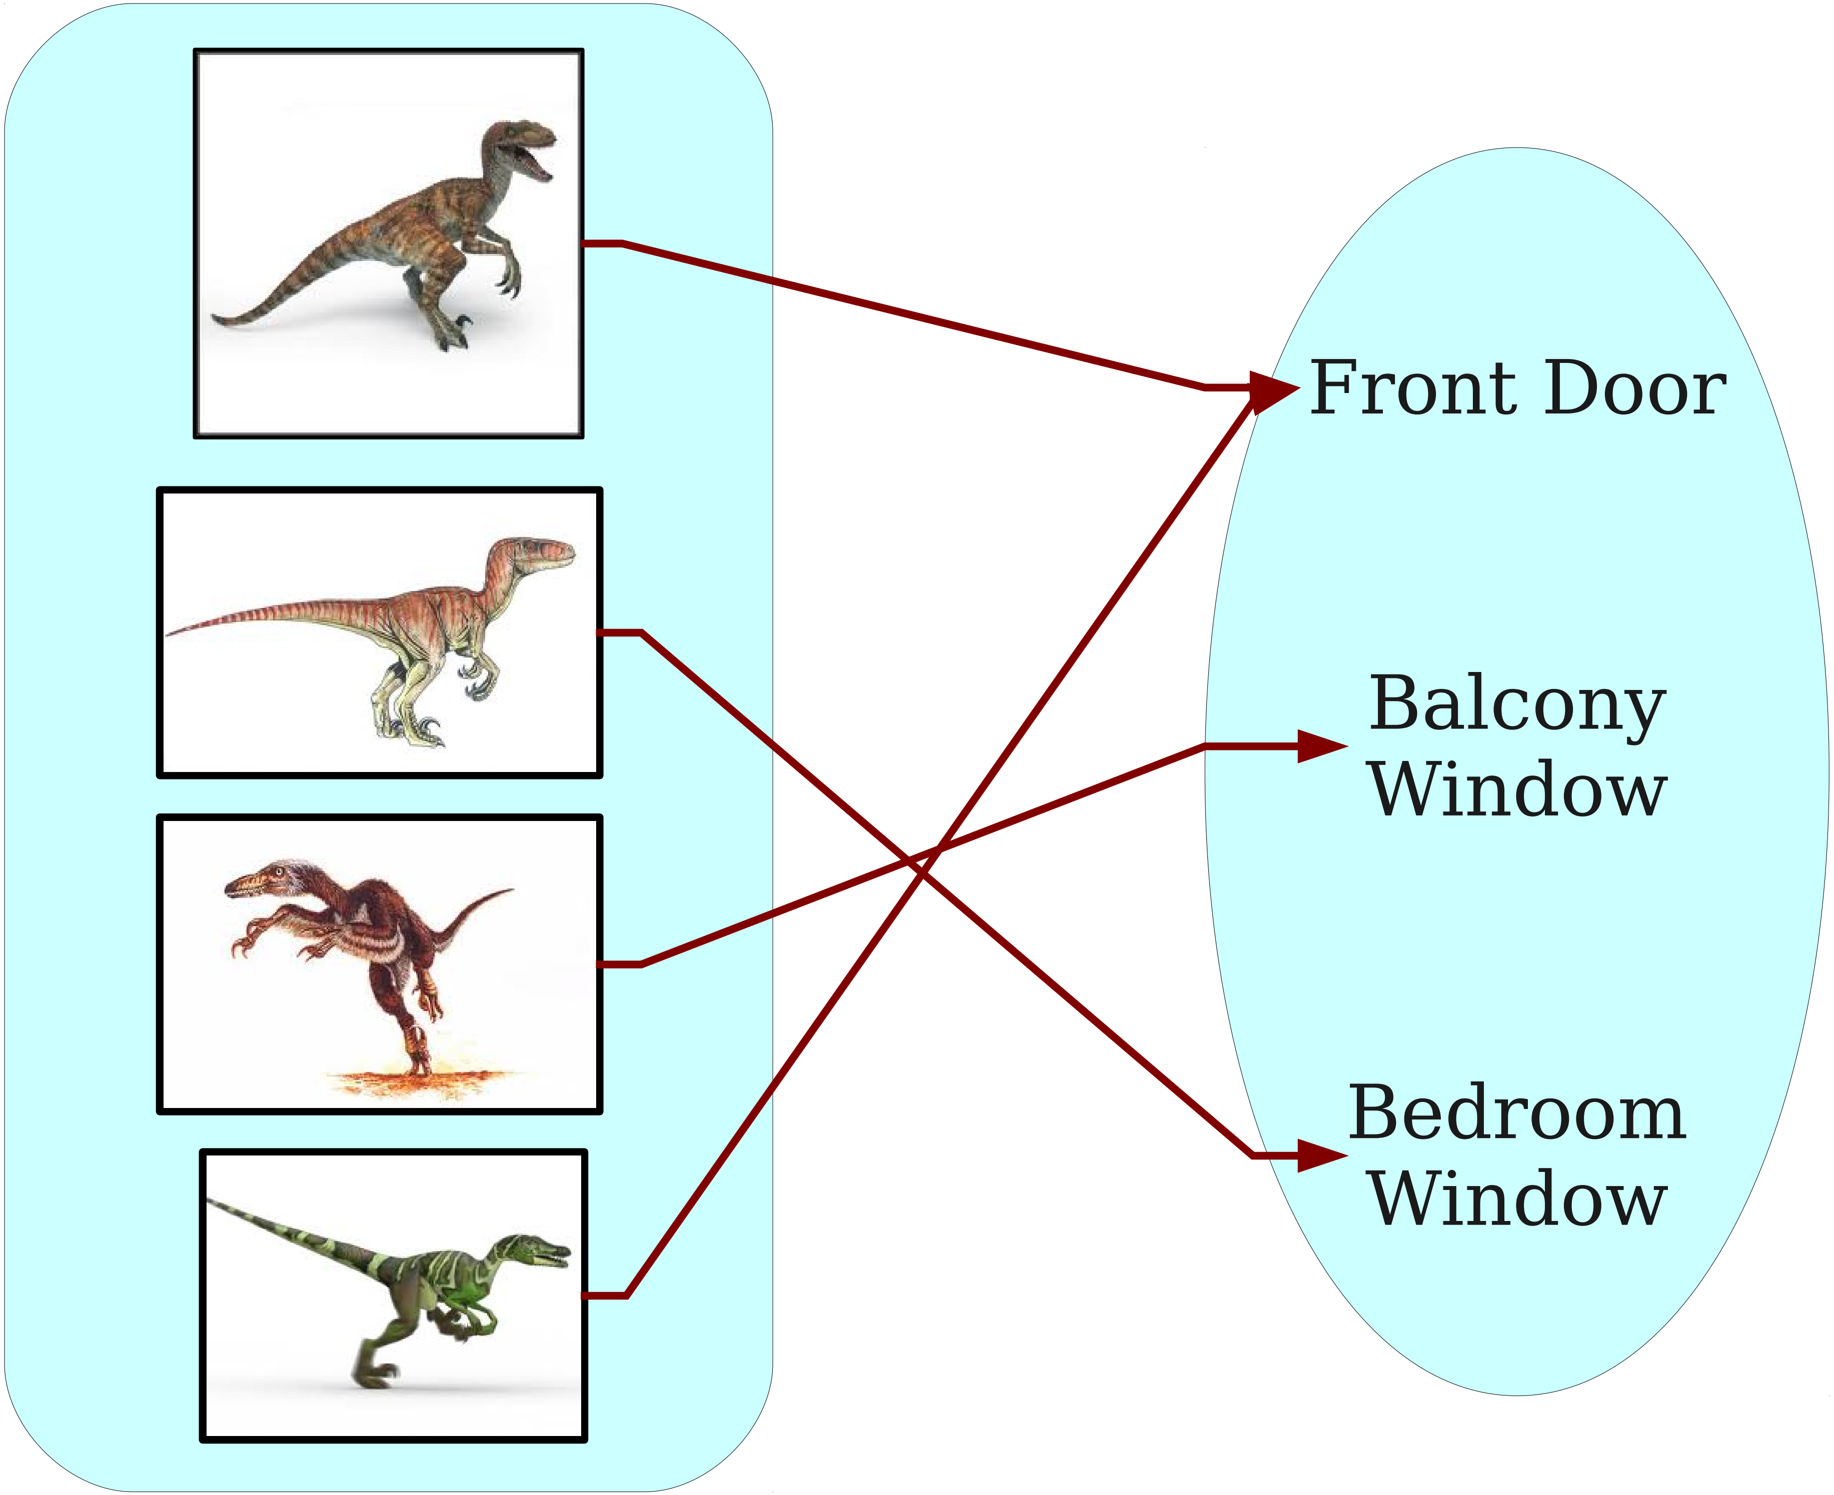
Front Door
Balcony Window
Bedroom Window

Surjective Functions

- A function $f : A \to B$ is called **surjective** (or **onto**) if each element of the codomain has at least one element of the domain that maps to it.
  - A function with this property is called a **surjection**.
- Formally, $f : A \to B$ is a surjection iff

For every $b \in B$, there exists at least one $a \in A$ such that $f(a) = b$.

- Intuition: surjective functions cover every element of B with at least one element of A.

Injections and Surjections

- An injective function associates ***at most*** one element of the domain with each element of the codomain.
- A surjective function associates ***at least*** one element of the domain with each element of the codomain.
- What about functions that associate ***exactly one*** element of the domain with each element of the codomain?

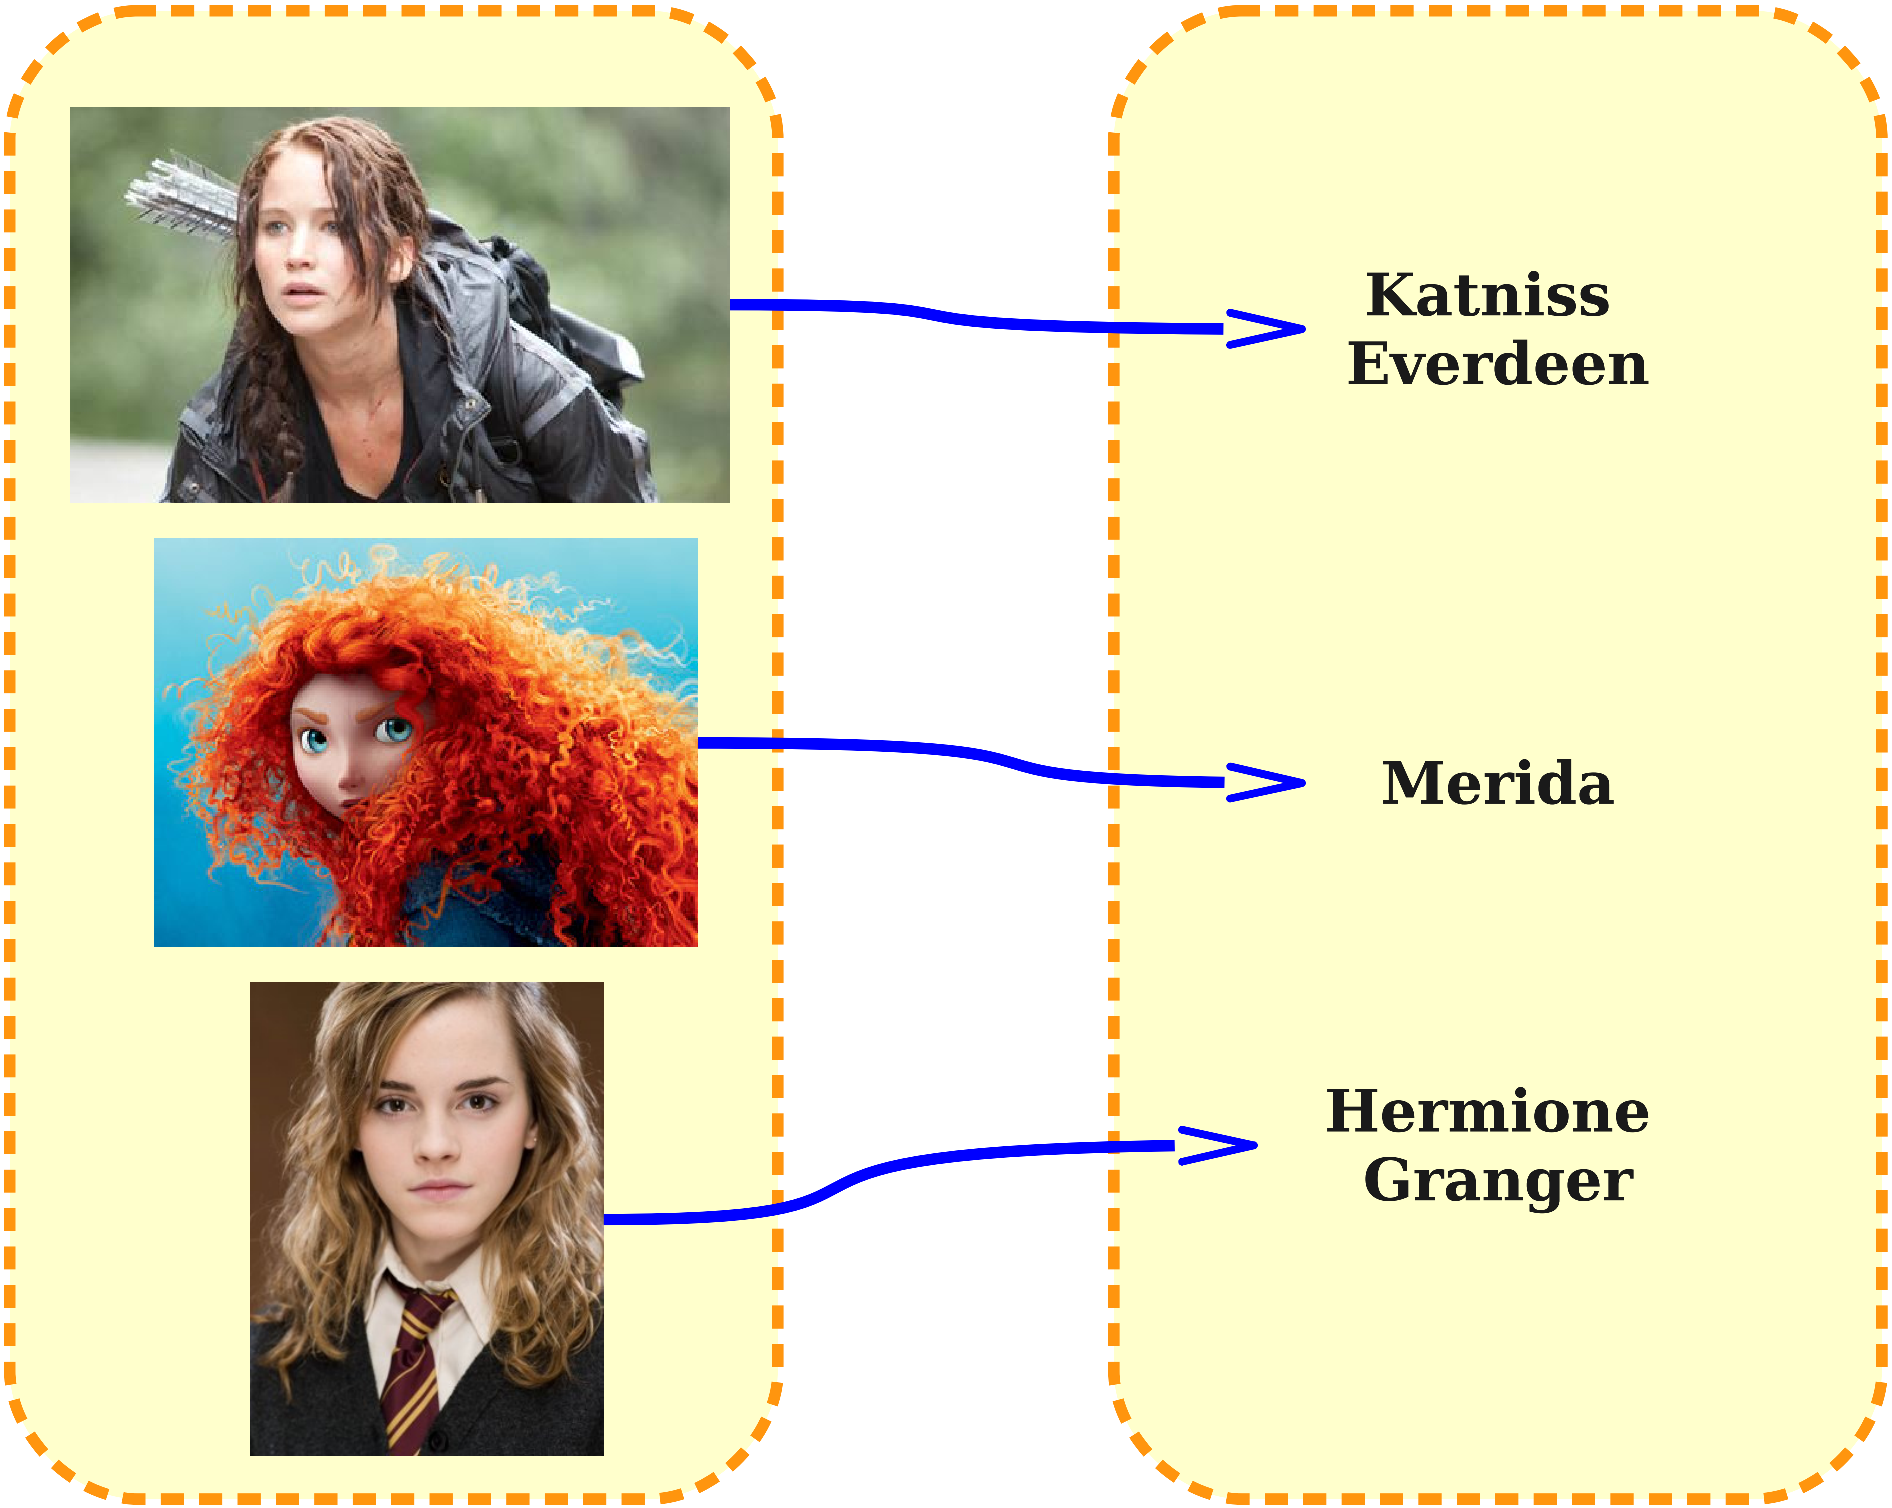
Katniss
Everdeen
Merida
Hermione
Granger

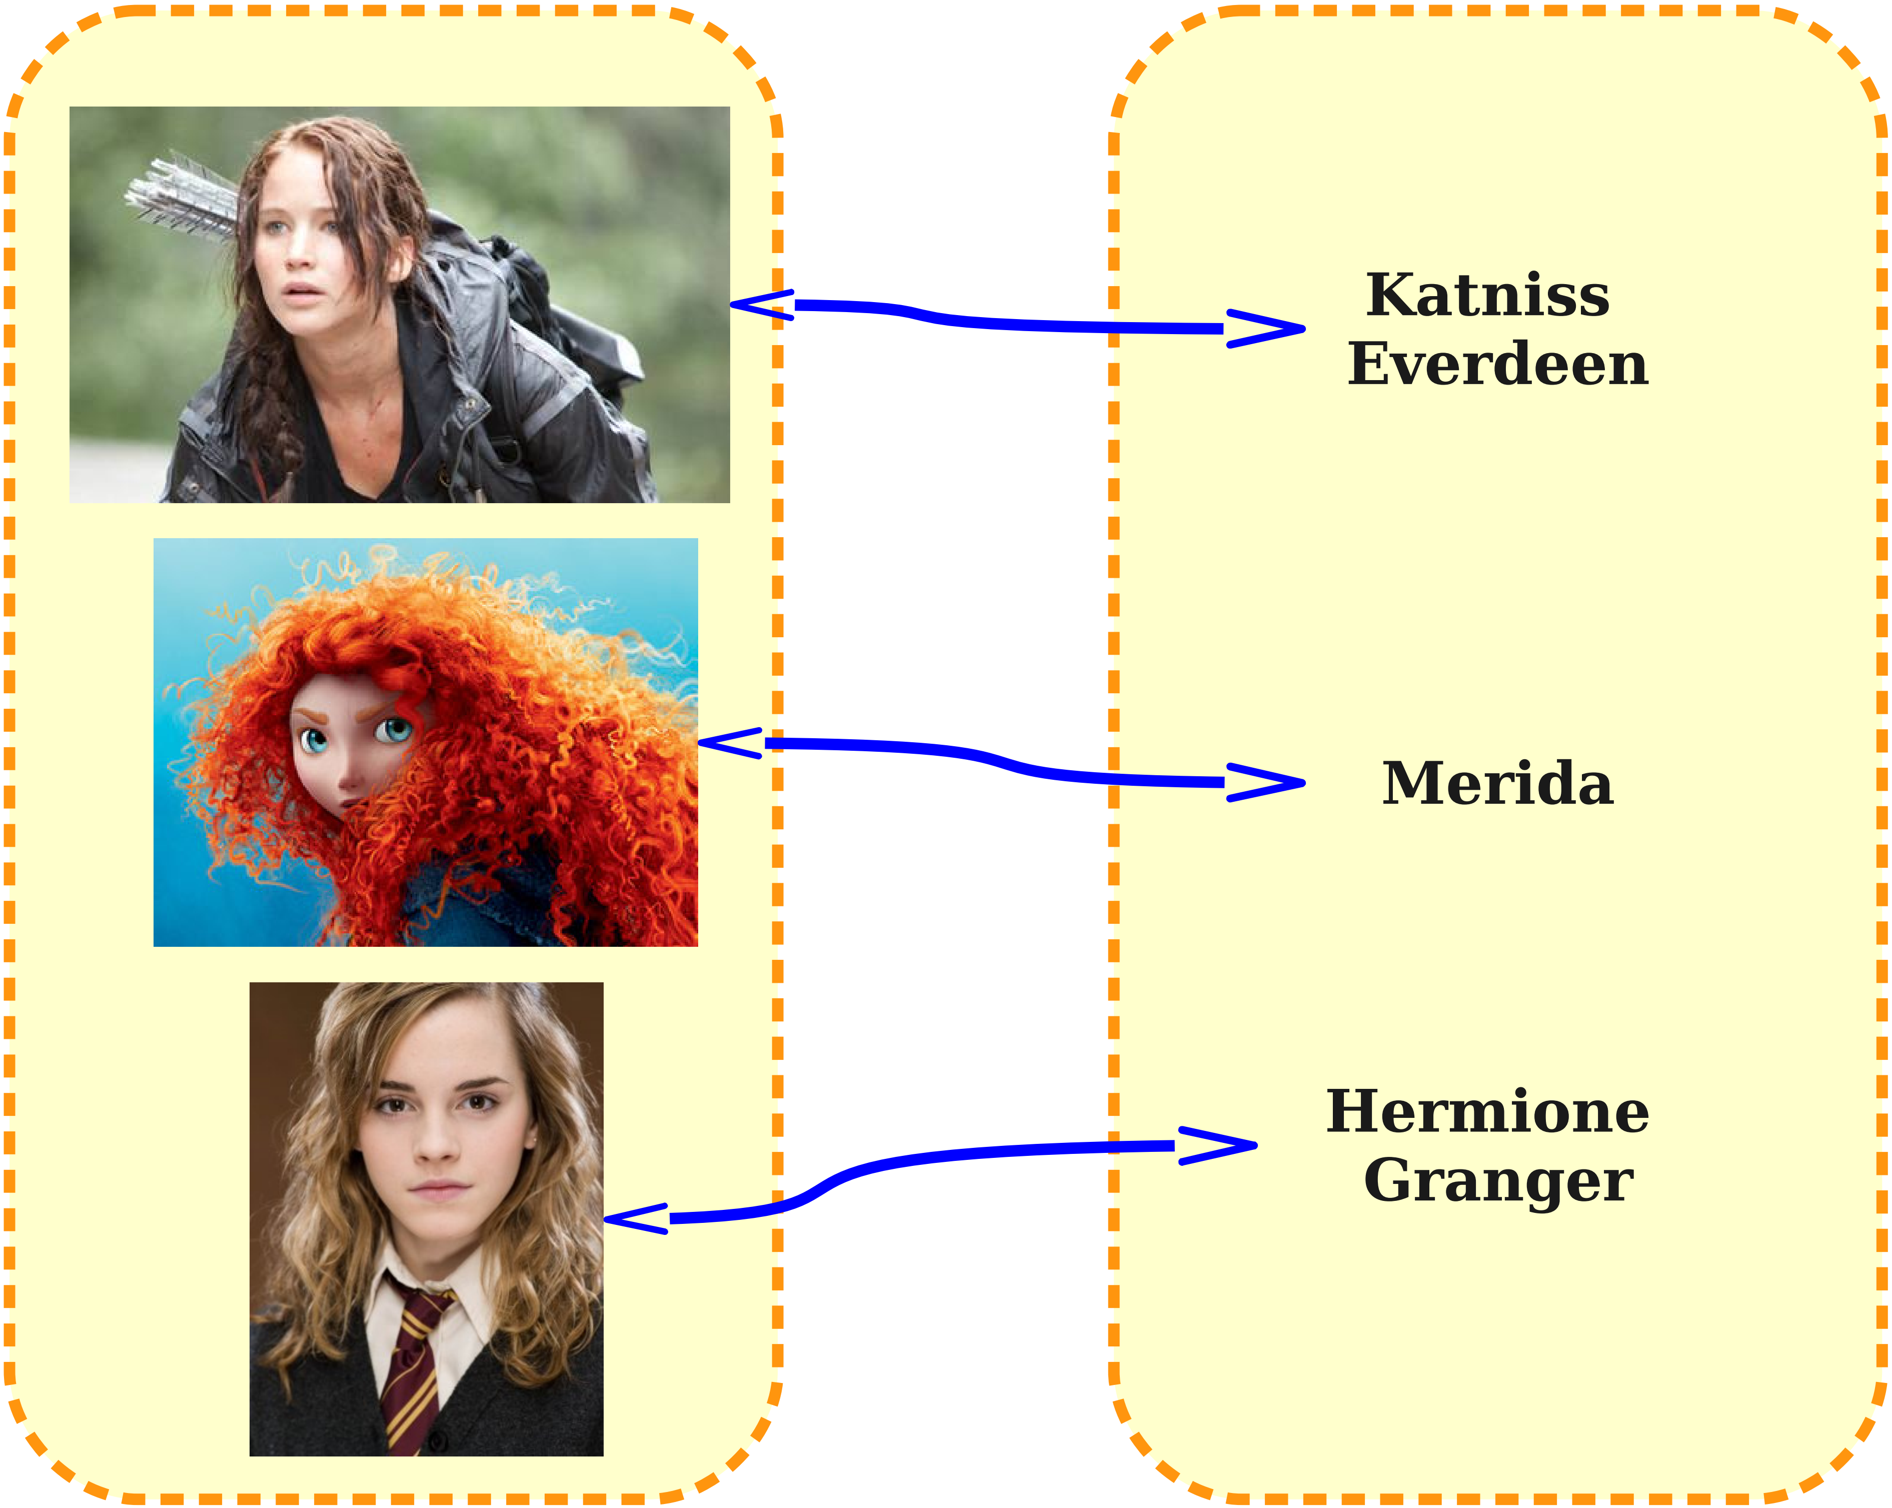
Katniss Everdeen
Merida
Hermione Granger

Bijections

- A function that associates each element of the codomain with a unique element of the domain is called **bijective**.
  - Such a function is a **bijection**.
- Formally, a bijection is a function that is both **injective** and **surjective**.
- Bijections are sometimes called **one-to-one correspondences**.
  - Not to be confused with "one-to-one functions."

Compositions

www.microsoft.com

www.apple.com

www.google.com

www.microsoft.com | **www.apple.com** | **www.google.com**

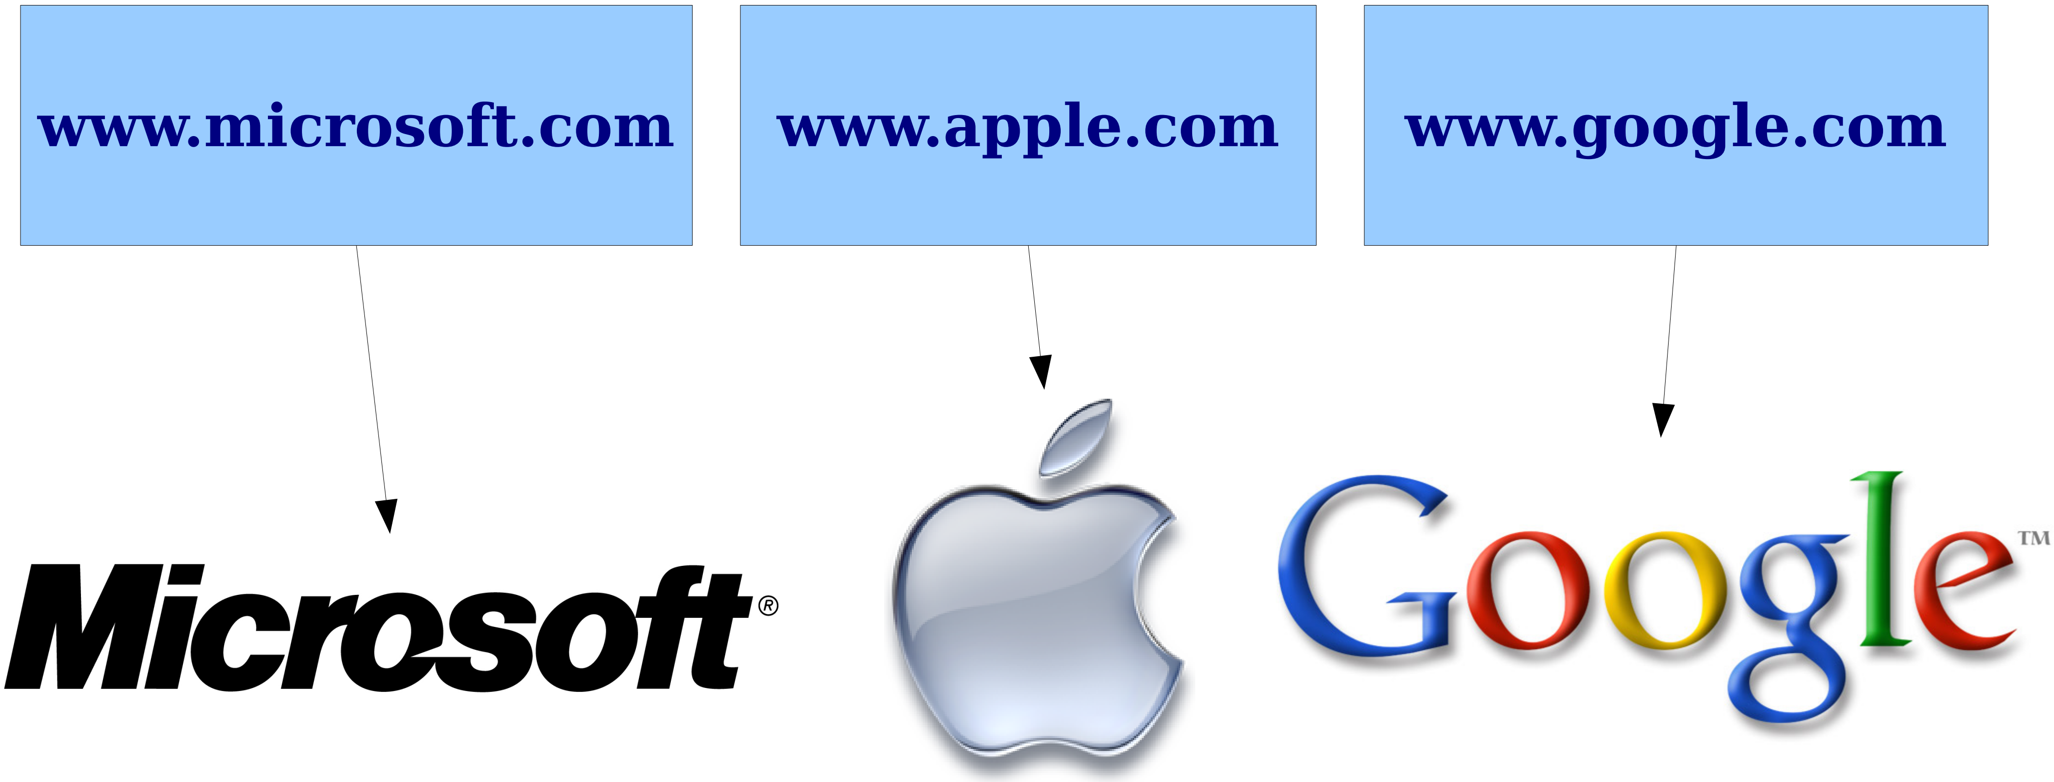
www.microsoft.com
www.apple.com
www.google.com
Microsoft
Google


www.microsoft.com
www.apple.com
www.google.com
Microsoft®
Google™
zune


www.microsoft.com
www.apple.com
www.google.com
Microsoft®
Google™
zune


www.microsoft.com
www.apple.com
www.google.com
Microsoft®
Google™
zune
cingular
9:41 AM
Text
Calendar
Photos
Camera
Calculator
Stocks
Maps
Weather
Notes
Clock
Settings
2:30
Google
Gmail
Maps
YouTube
Market

www.microsoft.com

www.apple.com

www.google.com

Function Composition

- Let $f : A \to B$ and $g : B \to C$.
- The **composition of f and g** (denoted **$g \circ f$**) is the function $g \circ f : A \to C$ defined as

$$(g \circ f)(x) = g(f(x))$$

- Note that f is applied first, but f is on the right side!
- Function composition is **associative**:

$$h \circ (g \circ f) = (h \circ g) \circ f$$

Function Composition

- Suppose $f : A \to A$ and $g : A \to A$.
- Then both $g \circ f$ and $f \circ g$ are defined.
- Does $g \circ f$ always equal $f \circ g$?
- **In general, no:**
  - Let $f(x) = 2x$
  - Let $g(x) = x + 1$
  - $(g \circ f)(x) = g(f(x)) = g(2x) = 2x + 1$
  - $(f \circ g)(x) = f(g(x)) = f(x + 1) = 2x + 2$

Next Time

- **Cardinality**
  - Formalizing infinite cardinalities
- **Diagonalization**
  - $|\mathbb{N}| \stackrel{?}{=} |\mathbb{R}|$
  - Formalizing Cantor's Theorem