LING3701/PSYCH3371: Lecture Notes 2
Models of Thought

Contents

Language seems similar to thought, but we can distinguish them.

Let's start with thought.

2.1 Decision theory [von Neumann & Morgenstern, 1944]

We model thought as a **decision process**: we choose actions to maximize **average expected utility**.

(There's lots of other thought: enjoying, reminiscing, etc., but this is about what helps us survive.)

A decision process assumes a set of **plans** p, a **world model** m and a **reward** $\mathsf{R}(m)$.

For example:

- p may be a plan to walk to a hill,
- m may include our location and the knowledge that a step may take us closer to a goal,
- $R(m)$ may be the prestige we get from reaching the hill.

Average (over time t) expected reward for a plan p is a sum over **outcomes** e of **actions** a in p:

$$\mathsf{AEU}(t, m, p) = \overbrace{\mathsf{R}(m)}^{\text{reward}} \cdot \begin{cases} \text{if } \exists_a \overbrace{p \wedge m \rightarrow a}^{\text{next action } a \text{ of } p}: & \overbrace{\sum_e \mathsf{P}(e \mid m \wedge a)}^{\text{sum all outcome events } e \text{ of } a} \cdot \overbrace{\mathsf{AEU}(t{+}1, \underbrace{m \wedge a \wedge e}_{\text{new } m \text{ at next } t}, p)}^{\text{repeat with } m,\, a \text{ and } e \text{ as new } m} \\ \text{otherwise:} & \dfrac{1}{t} \end{cases}$$

(This assumes the plan p is perfectly specific: at most one action a for each possible world m.)

For example, if the goal hill is one step away, we get a reward in one step, so $\mathsf{AEU}(t, m, p) = 1$.

But if it's muddy and we slip half the time and don't go anywhere, then:

$$\mathsf{AEU}(t, m, p) = \begin{cases} .5 \text{ (no slip)} & \times \frac{1}{1} \text{ (arrive in 1 step)} \\ +.5 \text{ (slip)} & \times \begin{cases} .5 \text{ (no slip)} & \times \frac{1}{2} \text{ (arrive in 2 steps)} \\ +.5 \text{ (slip)} & \times \begin{cases} .5 \text{ (no slip)} & \times \frac{1}{3} \text{ (arrive in 3 steps)} \\ +.5 \text{ (slip)} & \times \dots \end{cases} \end{cases} \end{cases}$$
$$\approx .7$$

So if we have two plans (clear path and muddy path) and we know mud slows us, we can avoid it.

We speculate language provides us an advantage by letting us **share** plans and world knowledge.

So what are these plans and world knowledge?

2.2 Lambda calculus [Church, 1940]

We'll use **lambda calculus** to model complex ideas that make up plans and world knowledge:

Lambda calculus expressions are of the following types:

- **entity terms**: references to things that can be predicated over, like people and places;
- **propositions**: things that can be true or false, like the proposition that a place is muddy;
- **functions** from any type as input to any type as output (including other functions).

  (We can think of functions as sets of pairs, e.g. $\mathsf{Muddy} = \{\langle \mathsf{MyHill}, \mathsf{True}\rangle, \langle \mathsf{YourHill}, \mathsf{False}\rangle\}$.)

Lambda calculus expressions consist of:

- **applications** of functions f to arguments x to get outputs $y = f\ x$, like $\overbrace{\underbrace{\mathsf{Muddy}}_{\text{function}}\ \underbrace{\mathsf{MyHill}}_{\text{entity}}}^{\text{truth value (proposition)}}$

- **abstractions** over argument variables x to get functions $f = \lambda_x \dots x \dots$, like $\overbrace{\lambda_x\ \underbrace{\mathsf{Muddy}\ x}_{\text{truth value (proposition)}}}^{\text{function from } x \text{ to truth value}}$

  (Functions from entities to truth values also define **sets** of entities x that return true: $\{\mathsf{MyHill}\}$.)

The most common functions we need are:

- **predicates**: e.g. $\overbrace{\mathsf{Person}\ x}^{\text{proposition}}$ or $\overbrace{\mathsf{At}\ x\ y}^{\text{proposition}}$, which map entities ($x$, y) to truth values (propositions)

- **generalized quantifiers**: e.g. $\overbrace{\mathsf{All}\ \overbrace{(\lambda_x\ \overbrace{\mathsf{Person}\ x}^{\text{proposition}})}^{\text{'restriction' set}}\ \overbrace{(\lambda_x\ \overbrace{\mathsf{At}\ x\ \mathsf{MyHill}}^{\text{proposition}})}^{\text{'nuclear scope' set}}}^{\text{proposition}}$, map sets to truth values

2.3 Practice: notation

Using the predicates Dog x, which means x is a dog, and Mammal x, which means x is a mammal, write a lambda calculus expression stating that all dogs are mammals.

2.4 Generalized quantifiers [Barwise & Cooper, 1981]

Generalized quantifiers compare sizes of sets, specifically the restriction R and intersection $R \cap S$.

Here are some common generalized quantifiers over restriction R and nuclear scope S sets:

- None R S $\Leftrightarrow$ $|R \cap S| = 0$ — true if none of the R's are S's:

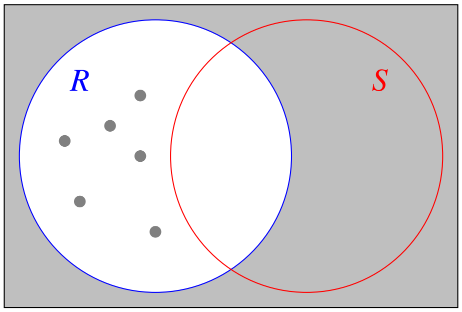

- Some R S $\Leftrightarrow$ $|R \cap S| > 0$ — true if some of the R's are S's:

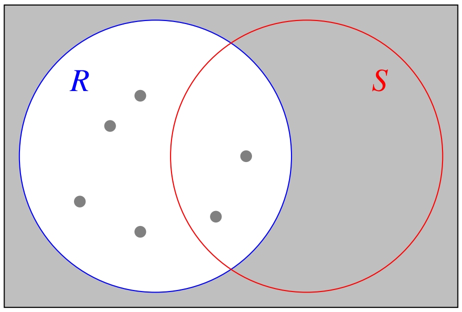

- Half R S $\Leftrightarrow$ $\frac{|R \cap S|}{|R|} = 0.5$ — true if half of the R's are S's:

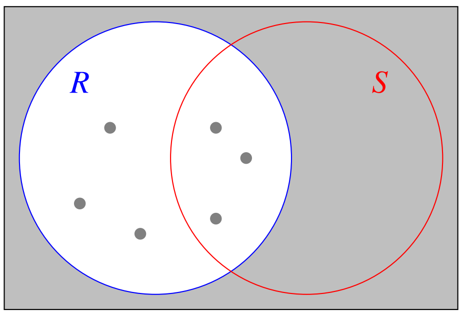

- Most $R\ S \iff \frac{|R \cap S|}{|R|} > 0.5$ — true if most of the R's are S's:

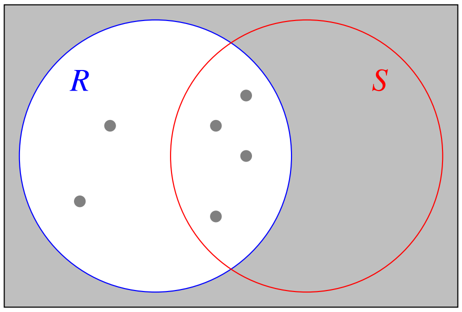

- All $R\ S \iff \frac{|R \cap S|}{|R|} = 1.0$ — true if all of the R's are S's:

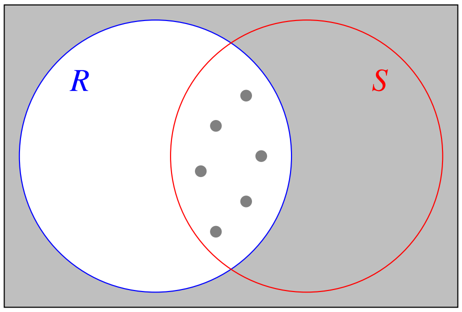

Note the similarity to diagrams of conditional probabilities from the previous lecture notes.

Generalized quantifiers represent conditional probabilities $\mathsf{P}(S \mid R)$, so we use them for reasoning!

(Specifically, we use them to represent probabilistic outcome events e in our decision processes.)

2.5 Practice: meaning

Given this set of Shape entities (where Red and Square have their usual meanings):

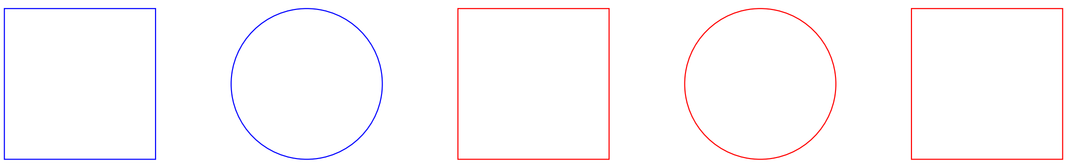

what is the value of the following lambda calculus expression?

$$\mathsf{Most}\ (\lambda_x\ \mathsf{Shape}\ x\ \wedge \mathsf{Red}\ x)\ (\lambda_x\ \mathsf{Square}\ x)$$

2.6 Practice: another meaning

Given the same set of shapes above, what is the value of the following lambda calculus expression?

$$\mathsf{Most}\ (\lambda_x\ \mathsf{Shape}\ x)\ (\lambda_x\ \mathsf{Square}\ x\ \wedge \mathsf{Red}\ x)$$

2.7 Quantifiers over substances

Generalized quantifiers model **substances** as sets of infinitesimal (arbitrarily small) ‘minimal parts’:

$$\mathsf{Most}\ (\lambda_v\ \mathsf{VolumeOf}\ v\ \mathsf{MilkyWayGalaxy})\ (\lambda_v\ \mathsf{EmptySpace}\ v)$$

Proportions over infinitesimals can be well defined using random sampling:

$$\mathsf{Most}\ R\ S\ \Leftrightarrow\ \lim_{K\to\infty} \mathsf{E}_{D_K\sim\pi} \frac{|D_K \cap R \cap S|}{|D_K \cap R|} > 0.5$$

(Here $\mathsf{E}_{D_K\sim\pi}$ ... is expected value of ... for K-element set D_K randomly drawn from distribution π.)

(Think of this as setting K high enough to be reliable and ensure a non-zero denominator.)

This lets us use the same quantifier functions (e.g. Most) for objects and substances.

2.8 Complex (nested) propositions

Recall that quantifiers are propositions that can contain propositions:

$$\overbrace{\mathsf{All}\ (\lambda_x\ \overbrace{\mathsf{Person}\ x}^{\text{proposition}})\ (\lambda_x\ \overbrace{\mathsf{At}\ x\ \mathsf{MyHill}}^{\text{proposition}})}^{\text{proposition}}$$

This means we can stuff them inside each other like Russian dolls or turduckens:

$$\overbrace{\mathsf{All}\ (\lambda_x\ \overbrace{\mathsf{Person}\ x}^{\text{proposition}})\ (\lambda_x\ \overbrace{\mathsf{Some}\ (\lambda_y\ \overbrace{\mathsf{Place}\ y}^{\text{proposition}})\ (\lambda_y\ \overbrace{\mathsf{At}\ x\ y}^{\text{proposition}})}^{\text{proposition}})}^{\text{proposition}}$$

This is called **nesting**.

2.9 Example: grid-world navigation

Now we can make a plan and world model that moves a person toward a hill in a grid world:

- a **plan** p to move Me to MyHill (assuming the following predicates:
  - NewTime t is true for the most recent time point t in an AEU evaluation;
  - PrecedeOrEqual s t is true if time s precedes or is equal to t;
  - TryMoveToward t a x is true if a tries to move toward x at time t;

  where ‘All’s iterate over entities, ‘None’s give conditions):

$$\begin{aligned}&\mathsf{All}\ (\lambda_t\ \mathsf{NewTime}\ t \wedge\\ &\qquad\quad \mathsf{None}\ (\lambda_s\ \mathsf{PrecedeOrEqual}\ 0\ s\ \wedge \mathsf{PrecedeOrEqual}\ s\ t)\\ &\qquad\qquad\quad (\lambda_s\ \mathsf{At}\ s\ \mathsf{Me}\ \mathsf{MyHill}))\\ &\quad (\lambda_t\ \mathsf{TryMoveToward}\ t\ \mathsf{Me}\ \mathsf{MyHill})\end{aligned}$$

(Here time '0' is when the plan is created – I have to reach the goal *after* that to succeed.)

When used in an AEU, this gives us our actions a – in this case: TryMoveToward predicates.

- **world knowledge** m that moving through mud may fail (assuming the following predicates:
  - Adjacent x y is true if grid squares x and y are adjacent;
  - Aligned x y z is true if a grid square y lies on a line from x to z;
  - Muddy y and Clear y are true if grid square y is muddy or clear, respectively;
  - Consecutive t u is true if time t immediately precedes time u;

  where 'Half''s give probability cost):

$$
\begin{aligned}
&\mathsf{All}\ (\lambda_{t,a,z}\ \mathsf{TryMoveToward}\ t\ a\ z)\\
&\quad (\lambda_{t,a,z}\ \mathsf{All}\ (\lambda_x\ \mathsf{At}\ t\ a\ x)\\
&\qquad\qquad (\lambda_x\ \mathsf{All}\ (\lambda_y\ \mathsf{Adjacent}\ x\ y\ \wedge \mathsf{Aligned}\ x\ y\ z\ \wedge \mathsf{Muddy}\ y)\\
&\qquad\qquad\quad (\lambda_y\ \mathsf{Half}\ (\lambda_u\ \mathsf{Consecutive}\ t\ u)\ (\lambda_u\ \mathsf{At}\ u\ a\ y)\ \wedge\\
&\qquad\qquad\qquad \mathsf{Half}\ (\lambda_u\ \mathsf{Consecutive}\ t\ u)\ (\lambda_u\ \mathsf{At}\ u\ a\ x)))) \wedge
\end{aligned}
$$

  and that moving through clear terrain always succeeds:

$$
\begin{aligned}
&\mathsf{All}\ (\lambda_{t,a,z}\ \mathsf{TryMoveToward}\ t\ a\ z)\\
&\quad (\lambda_{t,a,z}\ \mathsf{All}\ (\lambda_x\ \mathsf{At}\ t\ a\ x)\\
&\qquad\qquad (\lambda_x\ \mathsf{All}\ (\lambda_y\ \mathsf{Adjacent}\ x\ y\ \wedge \mathsf{Aligned}\ x\ y\ z\ \wedge \mathsf{Clear}\ y)\\
&\qquad\qquad\quad (\lambda_y\ \mathsf{All}\ (\lambda_u\ \mathsf{Consecutive}\ t\ u)\ (\lambda_u\ \mathsf{At}\ u\ a\ y))))
\end{aligned}
$$

  When used in an AEU, this gives us our outcome events e – in this case: At predicates.

Dropping these plans and world models into the AEU function defines a rational decision process.

Since they are *simple* and *work*, they are in some sense a 'natural' representation of complex ideas.

We'll therefore use these expressions as the complex ideas that get communicated using language.

But note these look very different from how we might represent these ideas in natural language.

References

[Barwise & Cooper, 1981] Barwise, J. & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4.

[Church, 1940] Church, A. (1940). A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5(2), 56–68.

[von Neumann & Morgenstern, 1944] von Neumann, J. & Morgenstern, O. (1944). Theory of games and economic behavior. *Science and Society*, 9(4), 366–369.