HS Game-Theoretical Semantics and Pragmatics (Summer 2007, Buschbom/Schäfer)

Script: Introduction I

Roland Schäfer
Seminar für Englische Philologie Göttingen

8. März 2009

1 Benz/Jäger/van Rooij, „An Introduction to Game Theory for Linguists“

- badly written
- best there is
- worst principle: math first, then examples! - so: Read ahead when things get complicated.
- relation to linguistics sometimes lost
- bad structuring: p. 3, par. 4 Äs an illustration...", but the illustration doesn't immediately follow!
- undefined symbols: J on p. 2

2 Today: pages 1-18. What are they about?

- A simple theory of informativity and information learning
  - People have knowledge about how probable it is that certain facts are true/false in the world. This is all there is to knowledge.
  - When people learn knew facts, they just change the probability they assign to facts they previously deemed at least not absolute improbable.
- Decision theory
  - Ultimately, information is used to make decisions. (What else is?)
  - Learning a fact thus influences our decisions (maybe even crucially) by making certain outcomes for actions more or less probable.
  - Important: No action can have a guaranteed effect, only a more or less probable one. Knowledge only favors certrain outcomes for certain actions.
- Measures of relevance and utility of learned facts

- The influence of learned facts on decisions we make can be used to define how releveant an information is.
- A simple measure of relevance is based just on how much a learned fact influences our beliefs (i.e. our knowledge after learning the fact).
- A usually more useful definition of relevance is based on how much an information influences our decisions.
- A piece of information which leads to a crucially different behavior of a human agent could thus be taken to be more informative than one which has no or little influence on our behavior.

- Games and their formalization
  - Games are models of decision procedures, just like the model of decision problems described in 2 and 3 above.
  - However, they take into account behavior of other agents or players.
  - Players' preferences are modelled as receiving payoffs in outcomes of the game, the payoffs depending both on their own decision and on the decision of the other players.
  - The normal form displays the payoffs in relation to the actions chosen by all players (usually 2 players) in a table.
  - Especially in sequential games, when players can make move after move, the extended form is suitable, which displays decisions and outcomes in a kind of tree diagram.

3 Knowledge and Bayesian Learning

The *rain* example:

1. Kermit thinks there is a 50% chance (a probability of 0.5) that it's cloudy.
2. Write $P(cloudy) = 0.5$ ('The probability that it is cloudy is 0.5'.)
3. Notice: Probabilities are assigned to **events** which can be expressed as propositions like 'It's going to rain', abbreviated usually as uppercase letters (like R for the aforementioned event).
4. Notice 2: Probabilities are always measured between 0 (impossibility) and 1 (certainty). No negative probabilities!
5. Notice 3: Actually, we would have to distinguish between *probability* and *likelihood*.
6. If it were cloudy, then Kermit thinks there is a 66% chance that it would also be raining.
7. Write $P(raining|cloudy) = 0.\overline{6}$ ('The probability that it is raining if it is cloudy is $0.\overline{6}$.)
8. This is all that Kermit 'knows' (i.e., assumes to be the case with the assigned probabilities) – at least as far as cloudiness and rain are concerned.
9. The probability that **both** is the case (in absence of further information) is written as $P(raining \cap cloudy)$ and calculated by simple multiplication of the probabilities:

   $$P(raining \cap cloudy) = P(cloudy) \times P(raining|cloudy) = \frac{1}{2} \times \frac{2}{3} = \frac{2}{6} = \frac{1}{3} = 0.\overline{3}$$

10. Note: Actually it is the case that:
$P(H|E) = P(H) \times P(E|H) = P(E) \times P(H|E)$

11. Assume now that Kermit learns from Piggy that it is cloudy. How is that going to change his beliefs about whether it is raining?

12. In probability parlance, the **prior** belief will be used to calculate a **posterior** belief.

13. This is where the Bayesian update comes into play:

$$P(raining|cloudy) = \frac{P(rain \cap cloudy)}{P(cloudy)} = \frac{\frac{1}{3}}{\frac{1}{2}} = \frac{2}{3} = 0.\overline{6}$$

14. Now you see (what you don't see in the introductory text) that there is a connection between the formula calculating the joint probability and the update formula. Once you learn that something is the case (cloudy), and you know that when that is the case, something else (rain) is the case with some probability, then this probability becomes the absolute probability of the second event (rain). You just 'eliminate' (informally speaking) the propability of the pre-condition.

15. The prior belief/probability 'that H' is written $P(H)$, the posterior one (after learning E) $P^{+} = P(H|E)$.

16. There is a mathematical tweak to the above calculations: $P(E)$ must never equal 0 (else division by 0 for the update). This can be interpreted as saying that we never learn (= come to believe) anything that we deemed completely improbable before.

17. **Important: If the two events H and E are independent (their probabilities are not connected), then learning E or H will not change the probability of the other.**

About the **complement** thing: $\overline{H}$

1. The complement of some event (like H) is its negation, and it's written $\overline{H}$.

2. If H represents 'it will rain', then $\overline{H}$ represents 'it won't rain'.

3. Trivially, since for every event H either itself or its complement $\overline{H}$ happens, we have: $P(H) + P(\overline{H}) = 1$.

4. If the update influences H, it will to the same degree also influence $\overline{H}$.

5. For the *rain* example (now with R for *raining* and C for *cloudy*):

$$\text{prior: } P(C) = \frac{1}{2} \text{ and } P(\overline{R}|C) = \frac{1}{3}$$
$$\text{so: } P(\overline{R} \cap C) = P(\overline{C}) \times P(\overline{R}|C) = \frac{1}{2} \times \frac{1}{3} = \frac{1}{6}$$
$$\text{update, learn } C\text{: } P(\overline{R}|C) = \frac{P(\overline{R} \cap C)}{P(C)} = \frac{\frac{1}{6}}{\frac{1}{2}} = \frac{2}{6} = \frac{1}{3} = 0.\overline{3}$$

4 Relevance as argumentative force (Merin)

- What was referred to as an event above can be expressed as a proposition, as statement that the event has happened/is happening/will happen.
- Hearing such a statement (and believing, of course) is represented as Bayesian learning.
- We have the intuition, however, that learning that the albedo of Jupiter is 0.52 will influence our beliefs about what counts as a subjacency violation in English to a neglectable degree.
- The impact of a learned fact becomes visible in no change or smaller/larger changes from prior to posterior belief.
- This change can be best read off from the ratio of the change from prior to posterior for the event and the change from the prior to the posterior of its complement. This is what the Good/Merin measure of relevance calculates.

A derivation of Merin's relevance measure:

1. Bayes Theorem tells us how to get from $P(H|E)$ to $P(E|H)$:

   $$P(E|H) = P(H|E) \times \frac{P(E)}{P(H)}$$

2. By way of example, this gives us the probability that it is cloudy when it is raining as the product of the probability that it rains when it is cloudy and the fraction of the probability that it is cloudy and that it is raining. To calculate this, we actually have to have a prior probability value for $P(R|\overline{C})$, i.e. that it rains when it is not cloudy, see footnote. Using the values from above:[1]

   $$P(C|R) = P(R|C) \times \frac{P(C)}{P(R)} = \frac{2}{3} \times \frac{\frac{1}{2}}{\sim \frac{1}{3}} = \frac{2}{3} \times \sim \frac{3}{2} = \sim 1$$

3. The Bayesian Theorem can be transformed to give us a new definition of $P(H|E)$ which shows its interdependence with $P(E|H)$:

   $$P(E|H) = P(H|E) \times \frac{P(E)}{P(H)}$$

   $$\Rightarrow P(E|H) \times \mathbf{P}(\mathbf{H}) = P(H|E) \times P(E)$$

   $$\Rightarrow \frac{P(E|H)}{\mathbf{P}(\mathbf{E})} \times P(H) = P(H|E)$$

[1] The value for $P(R)$ (the absolute prior probability of there being rain) is not simply $\frac{1}{3}$, but would have to be calculated including a probability for $P(raning|\overline{cloudy})$. Setting this value plausibly low to $\frac{1}{1000}$:

$$P(R) = P(C) \times P(R|C) + P(\overline{C}) \times P(R|\overline{C}) = \frac{1}{2} \times \frac{2}{3} + \frac{1}{2} \times \frac{1}{1000} = \frac{1}{3} + \frac{1}{2000} \approx \frac{1}{3}$$

4. To check this once again with our *rain* example, we insert the respective values in the new definition of $P(H|E)$. The result is what we started with as prior for $P(R|C)$:

$$P(R|C) = \frac{P(C|R)}{P(C)} \times P(R) = \frac{1}{\frac{1}{2}} \times \frac{1}{3} = \frac{2}{3} = 0.\overline{6}$$

5. To sum up: The last formulation of $P(H|E)$ takes the probability that H and that E and the probability that E when H to calculate the probability of H when E.

6. This gives us a nice measure to check the 'general import' E has on H. Technically, we take the ratio of how much E influences H and how much it influences $\bar{H}$ (i.e., the posteriors for $P(E|H)$ and $P(E|\bar{H})$): **If the fraction calculates as greater than** 1**, then** E **favors** H**, and below** 1**, it favors** $\bar{H}$**.**

7. (1.4) on p. 4 is easily derivable by transformation.

8. The *log* formation is just a means to compress the values between -1 and 0 (favors $\bar{H}$) and 0 and higher (favors H).

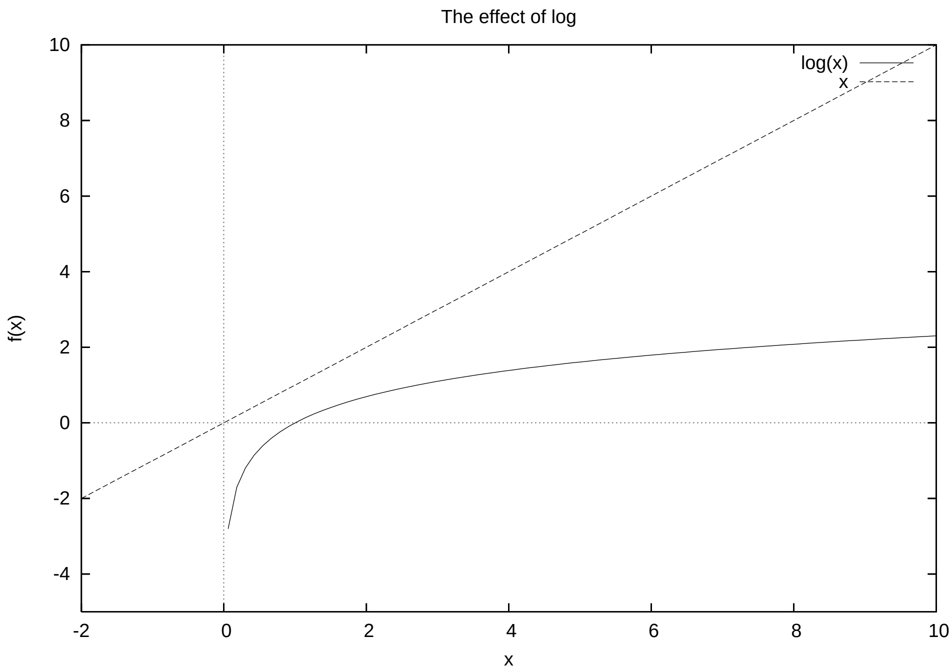

5 Relevance in decision situations

The measure of relevance in terms of informativity as described above is not adequate in all kinds of situations. If preferences over outcomes play a role, then receiving a piece of information which favors some events or 'states of affairs' over others might not be helpful in making a decision. The situation described on p. 5 is such an example.

1. Kermit has to choose between drinking a glass of wine (action a_1) or a beer (action a_2).

2. He knows that Miss Piggy has either gone shopping at Edeka, tegut, Plus, or Real – so, he is in one of four possible worlds which are characterized by which store Miss Piggy has chosen: v_1, v_2, v_3, v_4.

3. If she went to Edeka or tegut, he would fancy a glass of wine, because he knows that Miss Piggy buys good wine at these stores. However, if she went to Plus or Real he would avoid the heartburn from the cheap wine they sell there and just have a beer. Additionally, his absolute favorite drink is the Chateau from Edeka.

4. His preferences look like this then:

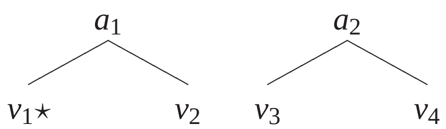

5. If all four worlds have the same prior probability, then learning that Miss Piggy went to tegut or Plus (v_2 or v_3) will equally favor both possible actions, and the Good/Merin relevance will be 0 with respect to the question of what Kermit should do.

6. However, his actions might be heavily influenced by learning that he won't get his favorite Chateau tonight.

- A more useful measure might take into account **how much an agent values the outcomes of action in certain possible worlds**.

- We can formulate such preferences by a *von Neumann-Morgenstern* utility measure U. It takes a possible action and a world and gives a rela number corrsponding to the preference of some agent w.r.t. the pair of action and world.

- For Kermit an his drinks maybe: $U(v_1,a_1) > U(v_2,a_1) = U(v_3,a_2) = U(v_4,a_3) > U(v_2,a_1) = U(v_1,a_2) = ...$

- Formally, we now have a a **decision problem** as $\langle\langle\Omega,P\rangle,\mathcal{A},U\rangle$: a tuple of a probability space (a set of possible worlds Ω and their propabilities) and a set of actions $\mathcal{A}$ and their utilities U). P, Ω, $\mathcal{A}$, U are all **sets**.

- **Maximization** of expected utilities means that an agent calculates for some possible action the utility in each world mutiplied with the probability of that world...and then he adds up all those values. The action that is the one with the best outcomes in the most probable worlds will be favored by him.

- In math:

  $$EU(a) = \sum_{v\in\Omega} P(v) \times U(v,a) = P(v_1) \times U(v_1,a) + P(v_2) \times U(v_2,a) + \ldots$$

  choose the action with the highest EU value: $\max_{b\in\mathcal{A}} EU(b)$

- Turning to the job interview example, review calculation on p. 6.

- When new information A comes in, we can just calculate a new max value with updated probabilities:

  $$EU(a) = \sum_{v\in\Omega} P(v|A) \times U(v,a)$$

- $P(v|A)$ is calculated as described above (Bayesian Update). If A does not make any world more or less probable, nothing will change (because learning A does not affect the utilities, but only the probabilities of worlds. If A, for example, tells the agent that one of his preferred worlds for some action is (as opposed to hi prior beliefs) highly improbable, then $EU(a|A) < EU(a)$, and a might not be chosen by the agent.

- **Relevance of A, heuristics I**: If the utility (EU value) of new best action a after learning A is greater than the utility of the action the agent would have chosen before learning A, called a^*, given that A is the case, then A was relevant because it makes the agent choose a over a^*. If, however, given that A is the case, the new best action a and the previous best action a^* would have the same utility, then A was not relevant because $EU(a) = EU(a^*)$ which means that the agent chooses the same action as before.

- In math, we define the utility value UV of a proposition A:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a|A) - EU(a^*|A)$$

- **Relevance of A, heuristics II**: We just compare the value (EU) of the best action before learning A with the one after learning A. If the new one is better, relevance (as maximization of expectations) was positive, otherwise negative, or nothing has changed.

- In math:

$$UV'(A) = \max_{a \in \mathcal{A}} EU(a|A) - \max_{a \in \mathcal{A}} EU(a)$$

- Thus, in the employment example, Adam can say things that change Eve's beliefs both in a negative sense and in a positive sense. Both cases are relevant, though for Adam only positive relevance is of interest, and he should try to choose his utterances such as to maximize positive relevance for Eve.

- Putting the right side in absolutes gives us a measure of pure relevance, no matter whether it was positive (boosted expectations) or negative (lowered expectations):

$$UV''(A) = |\max_{a \in \mathcal{A}} EU(a|A) - \max_{a \in \mathcal{A}} EU(a)|$$

- (Travel to France example to show that reduction to at least more promising worlds is also informative, even if we don't single out the one maximally informative world.)

6 Games, domination, and two forms

- Games are decision problems where decision/actions of others play a crucial role.

- **static game**: Only one decision is made simultaneously by every player (*rock-paper-scissors*, normal form).

- **dynamic game**: There is a sequence of moves/decisions made by the players (*chess*, extensive form).

- **cooperative game**: Players can commincate before the game and form coalitions/agreements (vs. **non-cooperative**).

- The already known Prisoner's Dilemma describes a static, non-cooperative game.

- **zero sum game**: A higher payoff for one user means a lower payoff for the other users, payoff is distributed exclusively.

- **game of pure coordination**: If player A receives some payoff, player B receives the same amount of payoff.

- Formally, a game is described as consisting of:

  1. a set of n **players**: $N = \{1, 2, \dots, n\}$ (only 2 in PD)
  2. one **set of possible actions** per player: $\mathcal{A}_i$ (= {defect,cooperate} for PD)
  3. **action profiles**, possible sets of one action (in static games) per player: $\langle a_i, a_{i+1}, \dots, a_n \rangle$, $a_n \in \mathcal{A}_i$
  4. **strategy profiles** (in dynamic games): sets of tuples of situations (in the game) plus actions (for each player)
     (coincides with action profiles in static games)
  5. player preference over strategy profiles: $(s_1, \dots, s_n) \prec (s'_1, \dots, s'_n)$
  6. payoff function for strategy profiles per player *i* (equivalent to previous): $u_i(s_1, \dots, s_n) < u_i(s'_1, \dots, s'_n)$
  7. one mapping of strategies to payoffs per player: U (type $\mathcal{A} \rightarrow \mathbb{R}^n$), i.e. it will give us for each action profile each user's utility for the profile. It is the function we use to draw payoff-matrices.

- Notational twist: We sometimes want to compare one strategy $(s_1, \dots s_n)$ with another one, where the second one is distinguished from the first only in that one player chooses a different action/strategy. We first take the player *i*'s old strategy s_i away: $s_{-i} = (s_1, \dots s_{i-1}, s_{i+1}, \dots s_n)$, then add the new strategy for player i s'_i: $\langle s'_i, s_{-i} \rangle$

Payoff matrices for PD with negative numbers for prison years:

	player 2 cooperates	player 2 defects
player 1 cooperates	$(u_1(1c, 2c); u_2(1c, 2c)) =$ $(-2; -2)$	$(u_1(1c, 2d); u_2(1c, 2d)) =$ $(-10; 0)$
player 1 defects	$(u_1(1d, 2c); u_2(1d, 2c)) =$ $(0; -10)$	$(u_1(1d, 2d); u_2(1d, 2d)) =$ $(-8; -8)$

- **row player**: **player 1**

- **column player**: **player 2**

- Strict Domination comes from the following reasoning for player 1:

  1. **Player 1** checks column-wise what payoff he gets when player 2 cooperates or defects:
  2. …should player 2 cooperate, I get -2 if I cooperate and 0 if I defect. $\rightarrow$ **I should defect if 2c.**
  3. …should player 2 defect, I get -10 if I cooperate and -8 if I defect. $\rightarrow$ **I should defect if 2d.**
  4. $\longrightarrow$ **Rationally, I can only choose to defect!** (= "defect" strictly dominates "cooperate" for 1.)

- Strict Domination comes from the following reasoning for player 2:

  1. **Player 2** checks row-wise what payoff he gets when player 1 cooperates or defects:

2. …should player 1 cooperate, I get -2 if I cooperate and 0 if I defect. → **I should defect if 1c.**
3. …should player 1 defect, I get -10 if I cooperate and -8 if I defect. → **I should defect if 1d.**
4. ⟶ **Rationally, I can only choose to defect!** (= "defect" strictly dominates "cooperate" for 2.)

- Rationally, both players can only defect!

Why being compassionate is also rational…your payoffs simply change:

	player 2 cooperates	player 2 defects
player 1 cooperates	$(-4;-2)$	$(-10;0)$
player 1 defects	$(-10;-10)$	$(-16;-8)$

The extensive form

- It can model dynmic games.
- It shows how a game develops as a sequence of states and actions (states also as outcomes of actions).
- For games where decisions are made simultaneously, the extensive form shows a sequential notation, but that does not mean that the second player knows about the move of the first player.
- This is captured by putting states into large ovals. They represent **information sets**. The player who makes a move in such an information set is in one of the states covered by the set, but he doesn't know in which he is.
- (Discuss fig. 1.1, p. 15 and fig. 1.2, p. 17. Do the PD in extensive form.)