

Conditionals: between language and reasoning

Class 1 - Introduction

www.ivanociardelli.altervista.org/conditionals2019
ivano.ciardelli@lrz.uni-muenchen.de

Class:

- ▶ Tuesday at 12.15
- ▶ Please ask questions at any time!

Class:

- ▶ Tuesday at 12.15
- ▶ Please ask questions at any time!

Material

- ▶ See course webpage for program, slides, and other information
- ▶ Relevant papers will be uploaded to LSF

Class:

- ▶ Tuesday at 12.15
- ▶ Please ask questions at any time!

Material

- ▶ See course webpage for program, slides, and other information
- ▶ Relevant papers will be uploaded to LSF

Assessment

- ▶ Term paper: see the guidelines on the webpage
- ▶ A list of topic suggestions will be provided later
- ▶ Deadline: 23/9.

Class:

- ▶ Tuesday at 12.15
- ▶ Please ask questions at any time!

Material

- ▶ See course webpage for program, slides, and other information
- ▶ Relevant papers will be uploaded to LSF

Assessment

- ▶ Term paper: see the guidelines on the webpage
- ▶ A list of topic suggestions will be provided later
- ▶ Deadline: 23/9. Strict:

“Ein späterer Abgabetermin sollte den Studierenden nicht eingeräumt werden (dies kann im Fall von BA/Master-Studierenden regulär nur das Prüfungsamt auf Grundlage eines eingereichten ärztlichen Attestes o. Ä. tun).”

What is a conditional?

- ▶ A conditional is a sentence of the form **if A, C**, where A (the antecedent) and C (the consequent) are sentential clauses:
 - (1)
 - a. If kangaroos had no tails, they would topple over.
 - b. If kangaroos have no tails, they've been fooling us all this time.

What is a conditional?

- ▶ A conditional is a sentence of the form **if A, C**, where A (the antecedent) and C (the consequent) are sentential clauses:
 - (1)
 - a. If kangaroos had no tails, they would topple over.
 - b. If kangaroos have no tails, they've been fooling us all this time.

- ▶ This is too narrow. There are many other ways to express conditionals:
 - (2)
 - a. Had I known about the strike, I would have stayed home.
 - b. No Martini, no party.
 - c. Pay him and he'll tell you everything.
 - d. In the event of rain, the party will be canceled.
 - e. Need a job? We can help you.

What is a conditional?

- ▶ A conditional is a sentence of the form **if A, C**, where A (the antecedent) and C (the consequent) are sentential clauses:
 - (1)
 - a. If kangaroos had no tails, they would topple over.
 - b. If kangaroos have no tails, they've been fooling us all this time.
 - (2)
 - a. Had I known about the strike, I would have stayed home.
 - b. No Martini, no party.
 - c. Pay him and he'll tell you everything.
 - d. In the event of rain, the party will be canceled.
 - e. Need a job? We can help you.
- ▶ The common core of these is the fact that they involve:
 - ▶ the creation of a **hypothetical context/situation** where A holds;
 - ▶ the claim that C holds in this hypothetical context.

- ▶ Most literature focuses on conditional **statements**.

- ▶ However, conditionalization is a more general phenomenon, which applies to **questions** and **imperatives** as well:
 - a. If you invite Alice, how will Bob react?
 - b. If you invite Alice, don't tell Bob that you did.

Conditionals are probably the most widely studied of all linguistic constructions.
Why is that?

Conditionals are probably the most widely studied of all linguistic constructions.
Why is that?

- ▶ Because they are difficult.

Today, no fully satisfactory general theory of conditional exists.

Conditionals are probably the most widely studied of all linguistic constructions.
Why is that?

- ▶ Because they are difficult.

Today, no fully satisfactory general theory of conditional exists.

- ▶ Because they are important.

Conditionals are connected to key issues in a wide range of disciplines:

- ▶ Linguistics/philosophy of language
- ▶ Logic
- ▶ Psychology/cognitive science
- ▶ Philosophy of science
- ▶ Artificial intelligence

Linguistics/philosophy of language

- ▶ The working assumption of most work in natural language semantics is that the semantic content of a sentence lies in its truth-conditions.
- ▶ These determine the proposition expressed, and play a key role in theories of semantic composition, discourse pragmatics, and propositional attitudes.
- ▶ That is, most existing answers to questions such as:
 - ▶ What is the meaning of “probably A”?
 - ▶ What happens when someone asserts A in a discourse?
 - ▶ What is it to believe that A?

refer to the truth-conditions of A, or the proposition expressed.

Linguistics/philosophy of language

- ▶ The working assumption of most work in natural language semantics is that the semantic content of a sentence lies in its truth-conditions.
- ▶ These determine the proposition expressed, and play a key role in theories of semantic composition, discourse pragmatics, and propositional attitudes.
- ▶ That is, most existing answers to questions such as:
 - ▶ What is the meaning of “probably A”?
 - ▶ What happens when someone asserts A in a discourse?
 - ▶ What is it to believe that A?

refer to the truth-conditions of A, or the proposition expressed.

- ▶ But what are the truth-conditions of a conditional?

- ▶ Consider first a counterfactual such as:

(3) If Smith had been elected, he would have cut public spending.

- ▶ What must the world be like for (3) to be true?
- ▶ (3) seems to talk about an unrealized possibility.

- ▶ What about non-counterfactual conditionals like (4)?
(4) If Smith is elected, he will cut public spending.
- ▶ Here, there is a traditional answer: the material analysis.
If A then B is true if A is false or B is true.
- ▶ But, as is well-known, this analysis leads to much trouble.
For instance, (a) is predicted to be equivalent to (b).
 - a. It is not true that if Smith is elected, he will cut public spending.
 - b. Smith will be elected and he won't cut public spending.
- ▶ But clearly someone can believe/assert (a) without necessarily believing that Smith will be elected.

- ▶ In fact, as we'll see, there are serious reasons to doubt that (some) conditionals have truth-conditions at all.
- ▶ If they don't, this calls for a serious revision of our linguistic theories, including:
 - ▶ **compositional semantics**: to be able to interpret sentences like (5), where a quantifier scopes over a conditional.

(5) Most wild boars will attack a predator if threatened.
 - ▶ **pragmatics**: to understand when conditionals can be asserted, and what effects their assertion on a conversational context.
 - ▶ **propositional attitudes**: to understand what it is to believe, want, suppose, ... a conditional.

Logic

- ▶ Standard logical consequence is monotonic: $A \models C$ implies $A, B \models C$.
- ▶ Conditionals don't seem to be monotonic:
 - (6)
 - a. If I had called, I would have lost;
 - b. If I had called and had a poker in my hand, I would have won.
 - (7)
 - a. If Alice invites Bob for dinner, he will go.
 - b. If Alice invites Bob for dinner and then cancels, he won't go.

Understanding the logic of conditionals is tightly linked with understanding non-monotonic reasoning, in particular:

► **Belief revision**

Models how an agent's beliefs S change in response to new information.

- a. $S + \text{Alice invited Bob} \models \text{Bob will go}$
- b. $S + \text{Alice invited Bob} + \text{Alice canceled} \models \text{Bob won't go}$

► **Default logic**

Models defeasible inferences drawn based on what is normal.

- a. $\text{Zack is a Kangaroo} \rightsquigarrow \text{Zack lives in Australia}$
- b. $\text{Zack is a Kangaroo, Zack is in a zoo} \not\rightsquigarrow \text{Zack lives in Australia}$

Cognition

- ▶ Making hypotheses is one of the most common and most important mental processes. It allows us to run “mental simulations”.
- ▶ This ability is crucial to decision making. E.g., suppose you are offered an apartment to rent. You reason:
 - (8) What if I took it? It is far from the center, so it would take me ages to go to work. On the other hand there are no neighbors, so I could practice playing trumpet. The rent is cheap, so I'd save money. . .

Then you assess the outcome and compare it with the alternatives.

▶ Conditionals are used to describe the outcomes of such process:

(9) If I took the apartment, it would take forever to go to work.

- ▶ Conditionals are used to describe the outcomes of such process:
 - (9) If I took the apartment, it would take forever to go to work.

- ▶ And to discuss them (such simulations can be wrong!)
 - (10) Actually no, there's a bus which runs directly from there. It would take you only about 20 minutes.

- ▶ Conditionals are used to describe the outcomes of such process:
 - (9) If I took the apartment, it would take forever to go to work.
- ▶ And to discuss them (such simulations can be wrong!)
 - (10) Actually no, there's a bus which runs directly from there. It would take you only about 20 minutes.
- ▶ Conditionals gives us a window into human hypothetical thinking. They give us an empirical handle on how this key cognitive process works.

The ability to think conditional thoughts is a basic part of our mental equipment. A view of the world would be an idle, ineffectual affair without them. There's not much point in recognising that there's a predator in your path unless you also realise that if you don't change direction pretty quickly you will be eaten.

(Edgington 1995)

Philosophy of science

- ▶ What is a law? How is it different from an accidental generalization?
 - a. Every human under stress produces adrenaline.
 - b. Every human in this room is under 40.

Philosophy of science

- ▶ What is a law? How is it different from an accidental generalization?
 - Every human under stress produces adrenaline.
 - Every human in this room is under 40.
- ▶ Only laws, and not accidental generalizations, support counterfactual conditional claims:
 - If my uncle was now under stress, he would produce adrenaline.
 - #If my uncle was now in this room, he would be under 40.

This connection was exploited by various theorists in both directions:

⇒ Give a theory of counterfactuals and use it to explain causation.

A caused B := “if A had not occurred, B would not have occurred”

⇐ Analyze causality and use it to explain counterfactuals.

“if A then B” true ⇔ B causally follows from A + relevant facts

Either way, (counterfactual) conditionals and (causal) laws are tightly linked.

Another link: dispositional properties

Many properties that are part of our understanding of the world seem to be really counterfactual properties:

- ▶ x is fragile \iff if x were struck with some force, x would shatter.

Another link: dispositional properties

Many properties that are part of our understanding of the world seem to be really counterfactual properties:

- ▶ x is fragile \iff if x were struck with some force, x would shatter.
- ▶ x is red \iff if x were exposed to white light, it would reflect only frequencies in the range. . .
- ▶ Things can be red even in the dark!

Another link: dispositional properties

Many properties that are part of our understanding of the world seem to be really counterfactual properties:

- ▶ x is fragile \iff if x were struck with some force, x would shatter.
- ▶ x is red \iff if x were exposed to white light, it would reflect only frequencies in the range. . .
- ▶ Things can be red even in the dark!
- ▶ The problem of how to test these properties is the problem of assessing the truth of certain counterfactual conditionals.

A satisfactory definition of scientific law, a satisfactory theory of confirmation or of disposition terms [...] would solve a large part of the problem of counterfactuals. Conversely, a solution to the problem of counterfactuals would give us the answer to critical questions about law, confirmation, and the meaning of potentiality.

I am not at all contending that the problem of counterfactuals is logically or psychologically the first of these related problems. It makes little difference where we start if we can go ahead.

(Goodman '47)

Artificial intelligence

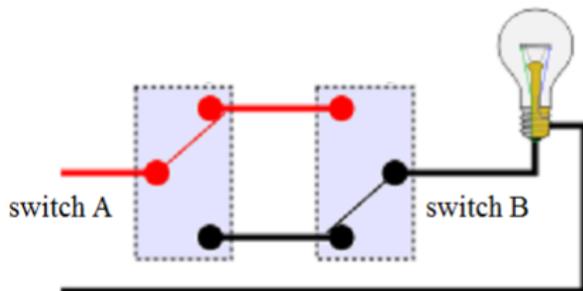
- ▶ In order to act intelligently, a robot needs to be able to reason about the consequences of its actions:
 - ▶ to come up with strategies to achieve goals;
 - ▶ to choose which strategy to pick among several options.
- ▶ It needs to be able to reason not only about how the world currently *is*, but also about how it *would be* if certain actions were performed.
- ▶ That is, it needs to be able to carry out conditional reasoning.
- ▶ Classical logic alone is not sufficient to this purpose.

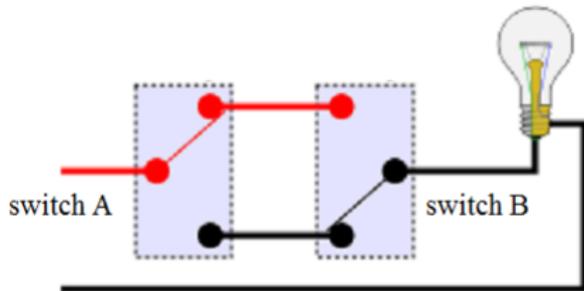
An example:

- ▶ A light bulb is controlled by two switches, A and B.
The light is on iff the switches are in the same position.

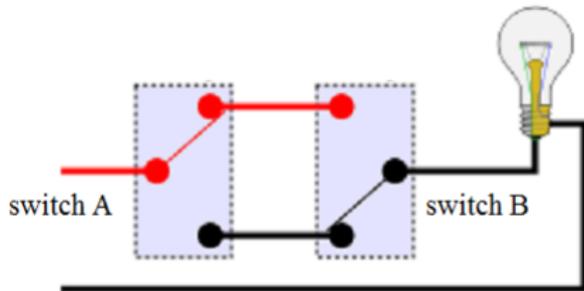
On \leftrightarrow (A-up \leftrightarrow B-up)

- ▶ Currently switch A is up and switch B is down.





- ▶ Now consider a robot whose task is to turn on the light: it should be able to see that this can be achieved, e.g., by toggling switch B.
- ▶ I.e. it should infer that **if B were switched up, the light would turn on.**



- ▶ Now consider a robot whose task is to turn on the light: it should be able to see that this can be achieved, e.g., by toggling switch B.
- ▶ I.e. it should infer that **if B were switched up, the light would turn on.**
- ▶ But this doesn't follow by classical logic alone:

$$\text{On} \leftrightarrow (\text{A-up} \leftrightarrow \text{B-up}), \text{ B-up} \not\leftrightarrow \text{On}$$

- ▶ After all, the law is compatible with the following alternative conclusion:
if B were switched up, the light would stay off and A would switch down.
- ▶ Making the right predictions about conditionals is both crucial to AI,
and a non-trivial logical task!

So, the study of conditionals is tied to fundamental questions about the workings of:

- ▶ natural language
- ▶ hypothetical thinking
- ▶ default reasoning
- ▶ causality and causal reasoning

which are crucial for a number of scientific disciplines:

- ▶ linguistics
- ▶ logics
- ▶ psychology
- ▶ AI
- ▶ philosophy of science

How many conditionals?

Conditionals come in many varieties, differing by mood (indicative vs. subjunctive) and tense.

▶ Indicatives:

- a. If John gets invited, he will go. (future-oriented)
- b. If John was invited, he went. (past-oriented)

▶ Subjunctive:

- a. If John were invited, he would go. (future-oriented)
- b. If John had been invited, he would have gone. (past-oriented)

Tense indicates the temporal positions of the events.

But what is the role of the indicative/subjunctive distinction?

A natural idea (e.g., Stalnaker '68):

- ▶ indicative and subjunctive conditionals have the same truth-conditions;
- ▶ mood is used to express an attitude towards the antecedent, to indicate whether it is being treated as possible (indicative) or not (subjunctive).
 - a. If John went to Paris, he won't be back until Monday.
↪ he might have gone
 - b. If John had gone to Paris, he wouldn't be back until Monday.
↪ he didn't go

This is not the case, as illustrated by this example from Adams 1970:

Suppose that on a given occasion three persons, A, B and V, are isolated in a room which is sealed off from the outside. During this time, the third person, V ('the victim'), is murdered by being shot.

The circumstances are such that only A and B could have done the shooting, though both deny it and accuse the other, and no one else witnessed the murder. An investigation is therefore instituted.

It establishes that A had in fact a very strong motive for wanting V dead, and furthermore had a gun which might well have been the murder weapon. It also shows that B had no known motive to murder V, and it fails to show that B might have had a gun at the time. On this basis it is concluded that A was the murderer.

Given that B appear to have had neither motive nor means to shoot V, it seems false to say:

(11) If A hadn't shot V, B would have.

On the other hand, since only A and B were in the room with V, the following seems true:

(12) If A didn't shoot V, then B did.

Given that B appear to have had neither motive nor means to shoot V, it seems false to say:

(11) If A hadn't shot V, B would have.

On the other hand, since only A and B were in the room with V, the following seems true:

(12) If A didn't shoot V, then B did.

Similarly, we judge the following differently:

- (13) a. If A hadn't shot V, then V would still be alive.
b. If A didn't shoot V, then V is alive.

True
False

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion.

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion. Consider:

(14) If Oswald hadn't shot Kennedy, someone else would have.

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion. Consider:

(14) If Oswald hadn't shot Kennedy, someone else would have.

F

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion. Consider:

- (14) If Oswald hadn't shot Kennedy, someone else would have. F
- (15) If Oswald didn't shoot Kennedy, someone else did.

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion. Consider:

- (14) If Oswald hadn't shot Kennedy, someone else would have. F
- (15) If Oswald didn't shoot Kennedy, someone else did. T

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion. Consider:

- | | | |
|------|---|---|
| (14) | If Oswald hadn't shot Kennedy, someone else would have. | F |
| (15) | If Oswald didn't shoot Kennedy, someone else did. | T |

Conclusion:

indicative and subjunctive conditionals do not just signal different attitudes, but can come apart in truth-value, and therefore express different things.

A concrete version of the example. Context:

- ▶ On November 22nd, 1963, John Kennedy was shot in Dallas.
- ▶ The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- ▶ Suppose we trust this conclusion. Consider:

- | | | |
|------|---|---|
| (14) | If Oswald hadn't shot Kennedy, someone else would have. | F |
| (15) | If Oswald didn't shoot Kennedy, someone else did. | T |

Conclusion:

indicative and subjunctive conditionals do not just signal different attitudes, but can come apart in truth-value, and therefore express different things.

Let us examine more carefully why this difference arises.

The ontic interpretation

- ▶ What if Oswald had not shot Kennedy?
- ▶ In supposing this, we are keeping our actual beliefs fixed (Oswald did kill Kennedy, and acted alone) and we are manipulating the world.
- ▶ We are imagining an alternative history where things went differently, and drawing conclusions about what would have happened.
- ▶ Since we think Oswald was the only shooter, in the alternative history where Oswald doesn't shoot, no killing takes place.
- ▶ Since we manipulate the world, we call this an **ontic** reading.

The epistemic interpretation

- ▶ What if Oswald did not shoot Kennedy?
- ▶ In supposing this, we are holding the world fixed (we are still talking about the actual world) and manipulating our beliefs about it.
- ▶ We are supposing that Oswald *in fact* didn't shoot (the official conclusion was wrong!) and trying to square this with the rest of our beliefs.
- ▶ Since we have decisive evidence that the shooting took place, we conclude that someone else must have shot.
- ▶ Since we manipulate our beliefs, we call this an **epistemic** reading.

- ▶ Does the (syntactic) indicative/subjunctive distinction reliably mark the (semantic) ontic/epistemic distinction?
- ▶ In other words, is it generally true that that:
 - ▶ indicative conditionals express epistemic readings;
 - ▶ subjunctive conditionals express ontic readings.

- ▶ Does the (syntactic) indicative/subjunctive distinction reliably mark the (semantic) ontic/epistemic distinction?
- ▶ In other words, is it generally true that that:
 - ▶ indicative conditionals express epistemic readings;
 - ▶ subjunctive conditionals express ontic readings.
- ▶ This is a debated topic in the literature.
- ▶ I myself think that this is plausible for past-tense conditionals, but that future-tense indicative conditionals can have both readings.

(16) If Alice goes to the party, Bob will not go.

- ▶ The difference between the two interpretations becomes hard to detect when assumptions about the world are not counterfactuals.

Future-oriented indicatives can express ontic conditionals

Consider this example by Edgington:

“Don’t go in there”, I say, “If you go in you will get hurt”. You look sceptical but stay outside, when there is large crash as the roof collapses.

“You see”, I say, “if you had gone in you would have got hurt. I told you so.”

Future-oriented indicatives can express ontic conditionals

Consider this example by Edgington:

“Don’t go in there”, I say, “If you go in you will get hurt”. You look sceptical but stay outside, when there is large crash as the roof collapses. “You see”, I say, “if you had gone in you would have got hurt. I told you so.”

- ▶ Here, the meaning of the initial indicative conditional is later re-cast, accurately, by means of a subjunctive past conditional.
- ▶ The reading of “if you had gone in you would have got hurt” is ontic: we are imagining a world in which things went differently than they did.
- ▶ The original conditional “if you go in you will get hurt” must thus share this ontic reading. This is plausible, as there is a causal, and not merely epistemic, connection between going in and getting hurt.

Future-oriented indicatives can express epistemic conditionals

Another example, also inspired by Edgington:

There are two prisoners, Smith and Jones, and we have reliable information that one of them will try to escape tonight. Jones is a docile, unadventurous chap; Smith is just the opposite.

I say:

- (17)
- a. It is probably Smith who will try to escape.
 - b. But if Smith doesn't try to escape, Jones will.

Future-oriented indicatives can express epistemic conditionals

Another example, also inspired by Edgington:

There are two prisoners, Smith and Jones, and we have reliable information that one of them will try to escape tonight. Jones is a docile, unadventurous chap; Smith is just the opposite.

I say:

- (17) a. It is probably Smith who will try to escape.
 b. But if Smith doesn't try to escape, Jones will.

As a matter of fact, Smith does try to escape. Can I now assert:

- (18) If Smith had not tried to escape, Jones would have.

Future-oriented indicatives can express epistemic conditionals

Another example, also inspired by Edgington:

There are two prisoners, Smith and Jones, and we have reliable information that one of them will try to escape tonight. Jones is a docile, unadventurous chap; Smith is just the opposite.

I say:

- (17) a. It is probably Smith who will try to escape.
 b. But if Smith doesn't try to escape, Jones will.

As a matter of fact, Smith does try to escape. Can I now assert:

- (18) If Smith had not tried to escape, Jones would have.

No: if Smith had not tried to escape, it would have been a quiet night.

The reading of (17-b) is clearly epistemic: the claim is not based on a causal connection between Smith's non-escape and Jones's escape. Accordingly, it is not legitimate to cast (17-b) as a subjunctive.

- ▶ In the literature, it is common to draw the line between subjunctive and indicative conditionals, with the former referred to as “counterfactuals”.
- ▶ I prefer the semantic distinction between ontic and epistemic conditionals, since this seems to track two different modes of hypothetical thinking.
- ▶ Ontic and epistemic conditionals pose different theoretical challenges; the relevant literature is also different, though with an overlap.
- ▶ We will devote the first part of the course to ontic conditionals (counterfactuals) and the second to epistemic conditionals (indicatives).