Abstracts: International Workshop "Learning conditionals"

Peter Collins (Birkbeck)

Goal framing and utility conditionals

Imagine a health campaign which informs people that 'If you have regular HIV tests, you will detect an HIV infection early'. Will that campaign be more or less effective than one which informs people that 'If you don't have regular HIV tests, you won't detect an HIV infection early'? That is, does the effectiveness depend on the way goals are framed. A large literature has explored this question using predictions from the psychology of decision making, and has yielded inconsistent data (see, e.g., Gallagher & Updegraff, 2012; Levin, Schneider & Gaeth, 1998; O'Keefe & Jensen, 2007, 2009). Although such conditionals generally feature utilities – some 89% of studies in an informal survey – the role of utilities has not yet been systematically investigated. In contrast, utilities feature prominently in the psychology of argumentation (e.g. Corner, Hahn & Oaksford, 2011; Evans, Neilens, Handley & Over, 2008; Ohm & Thompson, 2004). Bonnefon (e.g. Bonnefon, 2009, 2012, 2016) has drawn such work together under the umbrella of his theory of utility conditionals. We report the results of experiments which use Bonnefon's theory of utility conditionals to predict the effectiveness of goal framing. We discuss the relevance of the data for the psychology of decision making, for the psychology of argumentation, and for the theory of utility conditionals.

Ben Eva (LMU Munich)

Imaging Uncertainty

The technique of imaging was first introduced by Lewis (1976), in order to provide a novel account of the probability of conditional propositions. In the intervening years, imaging has been the object of significant interest in both AI and philosophy, and has come to be seen (in some quarters) as a plausible alternative to Bayesian conditionalisation. In this paper, we draw attention to a major challenge that faces any proponent of imaging. In particular, we introduce a new philosophical criterion that any update rule should satisfy, and use it to evaluate a range of different approaches to generalising imaging to situations involving uncertain evidence. We show that none of the currently prevalent approaches to imaging allow for such a generalisation, although a lesser known version of imaging, introduced by Joyce (2010), is able to deal satisfactorily with uncertain evidence. In light of these results, we conclude that Joyce's approach to imaging is the only one that provides a genuine alternative to Bayesian conditionalisation.

(Joint work with Stephan Hartmann)

Lupita Estefania Gazzo Castañeda (Giessen)

How phrasing influences conditional reasoning

Conditionals are everywhere. We use them to describe causalities, hypotheses, inducement and advice, or even to describe temporal relationships. It is therefore not surprising how much research has been conducted to understand how people interpret, understand, and reason with conditionals. However, there is still more research needed on how differences in the phrasing of conditionals can affect their interpretation and the inferences drawn from them. Does the specificity of the agent in a conditional matter? Do we reason differently if the agent is a specific (e.g., Jack) instead of an unspecific person (e.g., a person)? What happens if conditionals are preceded by quantifiers such as "all" or "some"? And does the phrasing of the consequent affect inferences? The aim of this talk is to show that differences in the phrasing of conditionals do affect how people interpret and reason with them. These differences are also very important when trying to understand what people learn from conditionals: if differences in the phrasing of conditionals affect inferences, then they probably will also affect what people extract and learn from them.

Mario Günther (LMU Munich)

Learning Conditional Information by Jeffrey Imaging on Stalnaker Conditionals

We propose a method of learning indicative conditional information. An agent learns conditional information by Jeffrey imaging on the minimally informative proposition expressed by a Stalnaker conditional. We show that the predictions of the proposed method align with the intuitions in Douven (2012)'s benchmark examples. As a generalisation of Lewis's imaging, Jeffrey imaging on Stalnaker conditionals can also capture the learning of uncertain conditional information, which we illustrate by generating predictions for the Judy Benjamin Problem. The proposed learning method manages to clearly discern between conditional, causal and conjunctive information.

Ulrike Hahn (Birkbeck / LMU Munich)

A failure to consider learning has hampered understanding of the conditional

The talk details evidence from very simple experiments that illustrate how and why present theories of the conditional are incompatible with how our beliefs change upon hearing a conditional. It is argued that a failure to consider what is learned on communication of natural language (indicative) conditionals has hampered development of an adequate theory of conditionals.

Stephan Hartmann (LMU Munich)

Bayesian Argumentation and the Value of Logical Validity

We often make arguments based on uncertain premises. In such cases, the conclusion of the argument does not follow with certainty, even if the underlying argument pattern is deductively valid. This raises the questions `what is so special about deductively valid arguments?', and `what advantage do we gain by using them?'. We will provide a novel answer to this question. In doing so, we will introduce the distance-based approach to probabilistic updating. Unlike the many updating rules that have been considered in the philosophical literature (such as conditionalization, imaging, and Leitgeb-Pettigrew updating), this approach takes seriously the idea of updating on non-propositional evidence. We conclude that the distance-based approach is the only probabilistic updating method that is able to provide a philosophically satisfactory account of arguments with uncertain premises.

(Joint work with Ben Eva)

Gabriele Kern-Isberner (TU Dortmund)

Reasoning, Revising, and Learning based on Conditionals

Conditionals are fascinating objects of reasoning in nearly all domains where rational logic-based thinking is important. As plausible rules or rules with exceptions, they help us not only to manage our everyday lives successfully, but also to support a physician's diagnostic capabilities, or a scientist's findings of good hypotheses. Interpreting conditionals via two-valued material implications is often done for the sake of logical clarity but cuts off their power considerably. Conditionals are understood and dealt with much better as at least three-valued entities of knowledge or belief that encode often semantical relationships between antecedent and consequent. This expressiveness of conditionals has been well appreciated in the probabilistic environment, where conditional probabilities and Bayes' rule have been offering a rich framework for reasoning for centuries. However, counterparts of these methods in more qualitative environments are much less well known. With strong involvement of philosophers, the fields of nonmonotonic reasoning and belief revision emerged in the 80s of the last century, and here conditionals proved to be particularly useful both for encoding both nonmonotonic inferences and revision strategies, continuing their probabilistic success story in more qualitative domains. In particular researchers like Ernest Adams and Judea Pearl emphasized the role of conditionals to link qualitative and quantitative reasoning.

In this talk, we elaborate on the role of conditionals for qualitative and quantitative logical frameworks, for inductive reasoning from knowledge bases, and for belief revision, and in general, as an excellent belief entity that combines logical quality with semantical expressiveness and commonsense meaning. We propose an algebraic theory of conditionals that is based upon de Finetti's three-valued approach to conditionals and can be used for different semantical frameworks like probability theory, or Spohn's ranking functions. This theory is also apt to establish a link between inductive reasoning with conditionals and learning conditionals from data, addressing the difficult problem of learning structure from data.

KAROLINA KRZYŻANOWSKA (LMU MUNICH)

WHAT DO WE LEARN WHEN WE LEARN A CONDITIONAL?

As many other linguistic expressions, conditionals can be used to make suggestions or promises, to argue, or to persuade others. All these linguistic activities involve some kind of learning: when we try to persuade someone that "if p then q," we want them to incorporate the information that if p then q in their belief system. In the belief revision literature, updating on a conditional is typically modeled as fixing the corresponding conditional probability, Pr(q|p), to a high value, and revising one's degrees of belief accordingly. However, equating the information conveyed by a conditional with high conditional probability is not theory-neutral; different accounts of conditionals can make different predictions on what changes when a conditional is learned. The aim of this talk is to pit the most prominent accounts of indicative conditionals against the recently collected data on how people's beliefs change upon receiving conditional information.

(Joint work with Ulrike Hahn)

MIKE OAKSFORD (BIRKBECK)

LEARNING IN DYNAMIC CONDITIONAL INFERENCE

When modelled as Bayesian conditionalization, many conditional inferences appear to violate the rigidity assumption that conditional probabilities are invariant between old (Pr0()) and new (Pr1(), i.e., post conditionalization) probability distributions. So being told that the car did not start seems to only make sense in a context where the ignition key was turned and hence constitutes a disconfirming observation relative to the claim that if the ignition key is turned (p) the car starts (q). Examples like this suggest that the categorical premise of a conditional inference can lead people to learn a new lower value of Pr(q|p) which is used to assess some inferences. This hypothesis was first tested using causal material. People judged Pr(q|p) with no categorical premise (Plain) and the after being told the categorical premise for MP (p), DA (¬p), AC (q), and MT (¬q). Pr(q|p) was judged lower than Plain for DA and MT but not for MP and AC. Other than for AC, this experiment confirmed the experimental hypothesis. However, the observed pattern of results could be a negation effect. So in Experiment 2, abstract material with explicit and implicit negations (e.g., a letter that is not a vowel (¬A) is a consonant (K)) were combined with a probability manipulation. This experiment replicated Experiment 1 for both explicit and implicit negations but the effects were weaker. Experiment 3 replicated Experiment 2. While showing clear effects in the predicted direction for DA and MT no effect was observed for AC and the value of Pr(q|p) was lower than the best fit values found in previous work. Experiment 4 introduced an inference task in to the design because more coherent probability assignments seem to results from this procedure. Pr(q|p) values were identical to previous best fits and had the same value for DA, AC, and MT but the differences between these inferences and MP and Plain were absent. These results were generally supportive of the revision strategy but some open questions remain.

(Joint work with Simon Hall)

David Over (Durham)

Learning counterfactuals: Towards a psychological study

The psychology of reasoning has totally neglected the topic of how people revise their beliefs when they learn a conditional. However, some experimental results are highly relevant to it. The most important is support for the conditional probability hypothesis: that people judge the probability of an indicative conditional, P(if p then q), to be the conditional probability, P(q|p). We will briefly refer to the results that support this hypothesis and the possible qualifications to it. These results are relevant to the topic of the workshop because they suggest that learning an indicative conditional can be represented as learning a conditional probability. But what can be said about learning a counterfactual conditional? There is a great deal of work in psychology on counterfactuals, but very little, so far, investigates the conditional probability hypothesis for counterfactuals. Moreover, there are relevant psychology studies that still have to be made, at a fundamental level, of people's counterfactual reasoning. For example, an unaddressed question is whether people consider centering, inferring 'if p then q' from 'p & q', valid for counterfactuals. We will ask how these psychological studies can be developed.

(Joint work with Nicole Cruz)

Henrik Singmann (Zurich)

A hierarchical Bayesian implementation of purely Bayesian and Bayesian mixture models of conditional reasoning

I will present a hierarchical Bayesian implementation of different models of conditional reasoning: A variant of the Bayesian model of Oaksford, Chater, and Larkin (2000; Oaksford and Chater; 2007) and the dual-source model (DSM; Klauer, Beller, & Huetter, 2010; Singmann, Klauer, & Beller, 2016). The models assume that the data follow a beta-distribution employing the beta regression approach (Ferrari & Cribari-Neto, 2004). In addition, the models use non-informative priors on the probability distributions underlying antecedent and consequent using a Dirichlet hyperparameter (e.g., Kemp, Perfors, Tenenbaum, 2007). I will then show how these models can be extended to model learning of conditionals using data from the two-stage experimental procedure usually used to fit the DSM.