Jaynes and Cox: Practical Inference Under Uncertainty

R. Henderson

January 19, 2012

Outline of Topics

Inferences we know and love

- **Modus ponens**
  $\{A \rightarrow B, A\} \Rightarrow B$
  "If Digby is a dog, then Digby likes french fries. Digby is a dog. Therefore, Digby likes french fries."

Inferences we know and love

- **Modus tollens**
  $\{A \rightarrow B, \neg B\} \Rightarrow \neg A$
  "If Digby makes his own french fries, then Digby is smarter than the average dog. Digby is not smarter than the average dog. Therefore, Digby does not make his own french fries."

Inferences we know and love (and classical logics hate)

- **Affirming the Consequent**
  $\{A \rightarrow B, B\} \rightsquigarrow A$
  "If John mows Bill's grass, Bill gives him \$30. Bill gave John \$30. Therefore, John mowed Bill's grass."

Inferences we know and love (and classical logics hate)

- **Denying the Antecedent**
  $\{A \rightarrow B, \neg A\} \rightsquigarrow \neg B$
  "If Digby rolls over on command, Ryan give him a treat. Digby didn't roll over on command. Therefore, Ryan didn't give him a treat."

Inferences we know and love (and classical logics hate)

While none of these are valid in propositional logic, they are deeply appealing. They get even better when stick a possibility adverbial in the consequent.

- "If John mows Bill's grass, Bill gives him $30. Bill gave John $30. Therefore, John **probably** mowed Bill's grass."
- "If Digby rolls over on command, Ryan gives him a treat. Digby didn't roll over on command. Therefore, Ryan **probably** didn't give him a treat."

Inferences we know and love (and classical logics hate)

We like even weaker syllogisms!

- **Gratuitously Affirming the Consequent**
  $\{\text{Plausibly}(A \to B), B\} \rightsquigarrow \text{Plausibly}(A)$
  "It's plausible that if a man has recently escaped from prison, he'll be wearing handcuffs and an orange jumpsuit. That man is wearing handcuffs and an orange jumpsuit. Therefore, he plausibly recently escaped from prison."

Inferences we know and love (and classical logics hate)

While none of these are classically valid, they seem fine if our goal is inference to the best explanation and not logically validity.

Common Sense = Weak Syllogisms + Prior Information

Or: If you transport a man wearing handcuffs and an orange jumpsuit across state lines, you do not avoid a federal prison sentence because "$\{$Plausibly$(A \to B)$, B $\} \rightsquigarrow$ Plausibly(A)" is not truth-preserving.

Inferences we know and love (and classical logics hate)

Goal: A theory that tells us how plausible the conclusion of a weak syllogism is given some specified background information.

Syntax

- A denumerable set of propositional letters: $A, B, C, ...$, standing for NL propositions. Any propositional letter is a formula.
- Connectives:
  - If ϕ and ψ are formulas, $\overline{\phi}$, $\phi\psi$, $\phi \vee \psi$, $\phi \rightarrow \psi$ are formulas.
  - Note that due to an accident of ghastly notation:

$\overline{\phi\psi}$ is wide scope negation, i.e., $\neg(\phi \wedge \psi)$, while:

$\bar{\phi}\bar{\psi}$ is narrow scope negation, i.e., $\neg\phi \wedge \neg\psi$

Syntax

Note that so far this is vanilla propositional logic!

- Modeling weak syllogisms doesn't mean making our logic weaker.
- We will retain all propositional logic validities.
- The logics we know and love are not inconsistent with how we reason, just insufficient for modeling the entire space inferences we make.

Syntax

Conditional Probability: If ϕ and ψ are formulas, then $\phi|\psi$ is the conditional probability of ϕ given ψ.

Warning, $\phi|\psi$ is not a formula. Otherwise we could write things like: $\phi|\psi \rightarrow \gamma|\theta|\eta$

We can write things like $\phi \rightarrow \psi|\overline{\theta\eta}$, though.

But do they mean?

Semantics for the PL connectives

It is just what you would expect.

- $V(A) \mapsto \{1, 0\}$
- $V(\phi\psi) = V(\phi) \cdot V(\psi)$
- $V(\phi \vee \psi) = \textbf{max}\{V(\phi), V(\psi)\}$
- $V(\neg\phi) = 1 - V(\phi)$
- $\phi \to \psi = \textbf{max}\{1 - V(\phi), V(\psi)\}$

Semantics for conditional probability

We need to decide what $\phi|\psi$ means.

Desiderata:

- Degrees of plausibility are real numbers, so $\phi|\psi$ should have a real number value.
- Infinitesimal increases in plausibility should correspond to infinitesimal increases in the real number value of $\phi|\psi$.
- ψ in $\phi|\psi$ should not be a contradiction—its value should not be 0.
- The plausibility of $\phi|\psi$ should accord with common sense (remember that ϕ and ψ are arbitrary PL formulas)

Interlude

There is a significantly more beautiful way to define the interpretation of a language in a probability model. See, for instance, Kaufmann's 2005 paper "Conditional Predictions" in Linguistics and Philosophy.

Semantics for conditional probability

We need to decide what $\phi|\psi$ means.

Some Common-Sense Requirements:

- if $(A|C') > (A|C)$ and $(B|AC') = (B|AC)$, then $(AB|C') \geq (AB|C)$
  "If we get information that A is more plausible, but the plausibility of $B|A\{C, C'\}$ is unchanged, then AB can be no less plausible than it was before we learned what we learned."

- if $(A|C') > (A|C)$, then $(\bar{A}|C') < (\bar{A}|C)$
  "Getting more information that A plausible is also getting more information that $\bar{A}$ is less plausible."

Product Rule

How should we relate the plausibility of $AB|C$ given the plausibilities of A and B separately?

- Note that this is similar to the question what is $V(\phi\psi)$ given $V(\phi)$ and $V(\psi)$.

Product Rule

A bad idea:

$F(AB|C) = F(A|C, B|C)$

"It plausible that the next person you meet has a brown right eye. It is plausible that the next person you meet has a blue left eye. But is is not plausible at all that the next person you meet will have a brown right eye and a blue left eye."

"It is plausible that the next person you meet has blue eyes. It is plausibile that the next person you meet has black hair. It is reasonably plausible that the next person you meet will have blue eyes and black hair."

Product Rule

The right idea:

$F(AB|C) = F(B|C, A|BC) = F(A|C, B|AC)$

"Imagine verifying AB. If AB is true, then B must be true. Thus, we could start by determining the plausibility of $B|C$. Furthermore, if B is true, AB is true only if A is true. Thus, we should determine the plausibility of $A|BC$."

The claim is that this is all we need to determine (AB|C).

Suppose it weren't. Well that would mean we need to determine $A|C$. But, if B is false $A|C$ is irrelevant for determining $AB|C$. Thus, $A|C$ tells us nothing new about $AB|C$ if we already have $B|C$ and $A|BC$.

Product Rule

Multiplication meets the requirements for F if p is a positive continuous monotonic function, essentially if it's a probability function, but Jaynes is coy about this in his paper.

$$p(AB|C) = p(A|BC)p(B|C) = p(B|AC)p(A|C)$$

Modus Ponens

- **MP:** $\{A \rightarrow B, A\} \Rightarrow B$
- **P:** $p(AB|C) = p(A|C) \cdot p(B|AC)$

We can rewrite **P** as:

$$p(B|AC) = \frac{p(AB|C)}{p(A|C)}$$

Assuming $C \iff A \rightarrow B$, then:

$$p(B|AC[A \rightarrow B]) = \frac{p(AB|C[A \rightarrow B])}{p(A|C[A \rightarrow B])}$$

Thus, $p(B|AC[A \rightarrow B]) = 1$ since p(numerator)=p(denominator)

Modus Tollens

- $\{A \rightarrow B, \neg B\} \Rightarrow \neg A$
- **P:** $p(A\bar{B}|C) = p(\bar{B}|C) \cdot p(A|\bar{B}C)$

We can rewrite **P** as:

$p(A|\bar{B}C) = \frac{p(A\bar{B}|C)}{p(\bar{B}|C)}$

Assuming $C \iff A \rightarrow B$, then:

$p(A|\bar{B}C[A \rightarrow B]) = \frac{p(A\bar{B}|C[A \rightarrow B])}{p(\bar{B}|C[A \rightarrow B])}$

Thus, $p(A|\bar{B}C[A \rightarrow B]) = 0$, since the p(numerator)=0

Modus {Ponens, Tollens}

Deductive logic is the limiting form of reasoning under uncertainty!

Affirming the Consequent

- **Affirming the Consequent**
  $\{A \rightarrow B,\ B\} \rightsquigarrow A$
- **P**: $p(A|C) \cdot p(B|AC) = p(B|C) \cdot p(A|BC)$

We can rewrite **P** as:

$$p(A|C)\frac{p(B|AC)}{p(B|C)} = p(A|BC)$$

Assuming $C \iff A \rightarrow B$, then:

$$p(A|C[A \rightarrow B])\frac{p(B|AC[A \rightarrow B)]}{p(B|C[A \rightarrow B)]} = p(A|BC[A \rightarrow B])$$

Thus, $p(A|C[A \rightarrow B]) \leq p(A|BC[A \rightarrow B])$, since p(numerator)=1 and p(denominator)≤1

Denying the Antecedent

- **DA**
  $\{A \rightarrow B, \neg A\} \rightsquigarrow B$
- **P**: $p(\bar{A}|C) \cdot p(B|\bar{A}C) = p(B|C) \cdot p(\bar{A}|BC)$

We can rewrite **P** as:

$$p(B|\bar{A}C) = p(B|C)\frac{p(\bar{A}|BC)}{p(\bar{A}|C)}$$

Assuming $C \iff A \rightarrow B$, then:

$$p(B|\bar{A}C[A \rightarrow B]) = p(B|C[A \rightarrow B])\frac{p(\bar{A}|BC[A \rightarrow B])}{p(\bar{A}|C[A \rightarrow B])}$$

Thus, $p(B|\bar{A}C[A \rightarrow B]) \leq p(B|C[A \rightarrow B])$, since p(numerator)$\geq$p(denominator)

Gratuitously Affirming the Consequent

- **GAC**
  $\{\text{Plausibly}(A \rightarrow B),\ B\} \rightsquigarrow \text{Plausibly}(A)$
- **P**: $p(A|C) \cdot p(B|AC) = p(B|C) \cdot p(A|BC)$

We can rewrite the major premise as:

$$p(B|AC) > p(B|C) = C$$

We can rewrite **P** as:

$$p(A|C)\frac{p(B|AC)}{p(B|C)} = p(A|BC)$$

Thus, $p(A|C) < p(A|BC)$, since p(numerator)>p(denominator)