Meta-, Anti-, Induction

Christian J. Feldbacher-Escamilla

Joint Work with Paul D. Thorn

Spring 2019

Project Information

Publication(s):

- Feldbacher-Escamilla, Christian J. and Thorn, Paul D. (submitted). "Meta-Induction and its Assumptions: The case of convexity". In: *manuscript*.

Talk(s):

- Feldbacher-Escamilla, Christian J. (2019a-06-05/2019-06-07). *Meta-, Anti-, Induction*. Conference. Presentation (contributed). IACAP 2019. Universidad Nacional Autónoma de México: The International Association for Computing and Philosophy.
- Feldbacher-Escamilla, Christian J. (2019b-05-16). *Meta-, Anti-, Induction*. Research Seminar. Presentation (invited). Summer term 2019. LMU Munich: MCMP.

Project(s):

- DFG funded research unit *Inductive Metaphysics* (FOR 2495); subproject: *Creative Abductive Inference and its Role for Inductive Metaphysics in Comparison to Other Metaphysical Methods*.

Introduction

C.D. Broad:
"Induction is the glory of science and the scandal of philosophy."

The "scandal" is about failing to account for:
Hume's Problem of Induction and Goodman's New Riddle

In this talk we want to discuss two accounts to these problems:
Meta-Induction and Conceptual Spaces

We will see that both approaches make an important convexity-assumption.

Problem: How to justify this assumption?

Contents

Hume's Problem and the New Riddle of Induction

1. A Brief History of Induction

(384–322 BC)

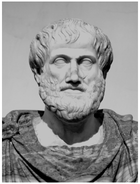

Aristotle

(1564–1642)

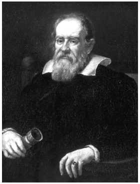

Galilei

(1561–1626)

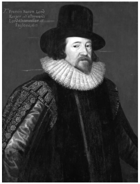

Bacon

(1642–1726/27)

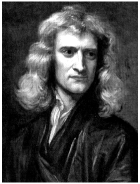

Newton

(1711–1776)

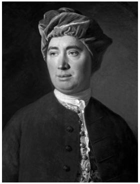

Hume

(1806–1873)

Mill

(1891–1953)

Reichenbach

(1891–1970)

Carnap

(1902–1994)

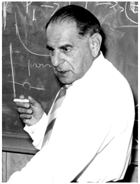

Popper

(1906–1998)

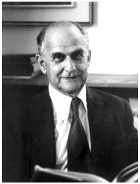

Goodman

2. Hume's Problem of Induction

"There is no object, which implies the existence of any other if we consider these objects in themselves, and never look beyond the ideas which we form of them.

[. . .]

We have no other notion of cause and effect, but that of certain objects, which have been always conjoined together. . . . We cannot penetrate into the reason of the conjunction.

[. . .]

All our reasonings concerning causes and effects are derived from nothing but custom; and that belief is more properly an act of the sensitive, than of the cogitative part of our natures."

(Hume, Enquiry, 1748)

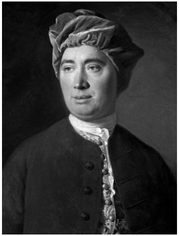

3. Approaches to Hume's Problem

How to justify induction?

- Foundationalism: Induction *as* Foundation (Bacon, 1620)
- Naturalised Epistemology: Induction = Custom (Hume 1748)
- Infinitism: Induction via Uniformity via Induction via ... (Mill 1843)
- "Logicism": Inductive Logic (Keynes 1921 and Carnap 1950)
- Eliminativism: Falsificationism (Popper 1934)
- "Pragmatism": Induction as a means for success (Reichenbach 1938)

4. The New Riddle of Induction

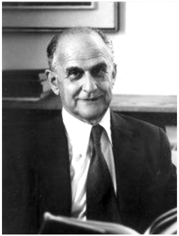

Justification of Induction $\Rightarrow$ Justification of Anti-Induction

Example: (cf. Goodman 1946, 1955/1983, chpt.3)
We start with enumerative induction of the form:

$$Pa_1, \ldots, Pa_{n-1} \mid\!\sim Pa_n$$

$a_1, \ldots, a_{n-1}$ ... observed, a_n ... unobserved, P ... *green*
We can define a property Q ... *gr*ue (and $\neg Q$ ... *bl*een):

$$Qx \quad \Leftrightarrow_{df} \quad (Px \ \leftrightarrow (x = a_1 \vee \cdots \vee x = a_{n-1}))$$

By help of enumerative induction we can infer:

$$\underbrace{Qa_1, \ldots, Qa_{n-1}}_{Pa_1, \ldots, Pa_{n-1}} \mid\!\sim \underbrace{Qa_n}_{\neg Pa_n}$$

5. The More General Problem: Language Dependency

Synonymity: Φ and Ψ are synonymous iff they have a *common definitional extension*, i.e. there are $\mathcal{D}_\Phi, \mathcal{D}_\Psi$ such that:

1. $\mathcal{D}_\Psi$ contains exactly one definition for each descriptive symbol in Ψ in terms of descriptive symbols of Φ only, and
2. $\mathcal{D}_\Phi$ contains exactly one definition for each descriptive symbol in Φ in terms of descriptive symbols of Ψ only, and
3. $\Phi \cup \mathcal{D}_\Psi \vdash\dashv \Psi \cup \mathcal{D}_\Phi$

Language Dependency: $\mid\!\sim$ is language dependent iff there are $\Phi_1, \ldots, \Phi_n$ and $\Psi_1, \ldots, \Psi_n$ such that

1. Φ_1, Ψ_1 and $\ldots$ and Φ_n, Ψ_n are synonymous (given a common definitional extension), and
2. $\Phi_1, \ldots, \Phi_{n-1} \mid\!\sim \Phi_n$ and $\Psi_1, \ldots, \Psi_{n-1} \not\mid\!\sim \Psi_n$.

If $\mid\!\sim$ excludes anti-induction, then $\mid\!\sim$ is language dependent.

6. Approaches to the Problem of Language Dependency

Language dependency is a general problem: It is not only about *induction*, but also *truthlikeness*, *simplicity*, *causality*, etc.

One can approach this problem by ...

- ... living with this relativism of epistemic notions. (Carnap 1928)*
- ... excluding "alien" languages. (Barnes 1991)
- ... excluding "alien" translations. (Tichý 1976)
- ... arguing that lang-invariance constraints are fishy. (cf. my 2019)

7. Epistemic Engineering

Epistemic Means-End Principle:

$$\overbrace{(\mathcal{O}(A)}^{\text{ulterior epistemic ends}} \quad \& \quad \underbrace{\Box(A \rightarrow B) \vee (B \circ\!\!\rightarrow A))}_{\text{epistemic engineering}} \quad \Rightarrow \quad \overbrace{\mathcal{O}(B)}^{\text{derived epistemic ends}}$$

E.g. Ought Implies Can Heuristics:

$$\neg\exists B(\Box(B \rightarrow A) \vee (B \circ\!\!\rightarrow A)) \;\Rightarrow\; \neg\mathcal{O}(A)$$

The Meta-Inductive Approach to Hume's Problem

1. The Framework: Prediction Games

Let us consider a series of events $e_1, e_2, \ldots$ with outcomes in $[0, 1]$.

Now, consider prediction methods for the event outcomes: $pred_1, \ldots, pred_n$ of the form $pred_i(e_t) \in [0, 1]$

A simple prediction method for binary events would be, e.g., a binarized likelihood method: $pred(e_t) = 1$ if $\frac{E_1 + \cdots + E_{t-1}}{t-1} \geq 0.5$ otherwise $pred(e_t) = 0$

	e_1	e_2	e_3	e_4	e_5	e_6	e_7	$\ldots$
E_i	0	0	1	1	1	1	0	
$pred_1$	1	0	0	0	1	1	1	

Now, assume that past predictions and event outcomes (E's) are available.

Then we can evaluate prediction methods according to their success.

2. Hume's Problem of Induction in this Setting

How can we justify inductive methods like the likelihood method?

In this framework:
problem of induction $\approx$ we against nature $\approx$ problem of Cartesian daemon

Here is why:	nature	sets	$\boxed{e_t : pred_{i,t}}$	is set by	us
Best case:	true outcome		$=$		$pred_i$
Worst case:			**max. distance**		$pred_i$

is most sceptic . . . what is its logic?

3. The Impossibility of Solving Hume's Problem

If there are no constraints for , then any $pred_i$ fails:

can always maximise $pred_i$ loss.

And for any $pred_i$ the average (w.r.t. all possible event series) loss/success rate is the same: .5 (there is *no free lunch* Wolpert 1996).

> *"A learner that makes no a priori assumptions regarding the identity of the target concept has no rational basis for classifying any unseen instances." (cf. Mitchell 1997, p.42)*

Hence, there is no absolute justification for any prediction method.

Hence, by *ought implies can*: $\neg\mathcal{O}(\textit{absolute success})$

4. Epistemic Re-Engineering: Reichenbach's Vindication

Instead of $\mathcal{O}$(*absolute success*), aim at $\mathcal{O}$(*relative success*)

Skyrms (2000, p.46): "If no method is guaranteed to be successful, then it would seem rational to bet on that method which will be successful, if any method will."

Lightbulb-Example

- You have to bet on some colour.
- Possible states:
  1. No light turns on.
  2. The green light turns on.
  3. All lights turn on.

Predicting is not sufficient for success, but necessary: Whenever you are successful with your prediction, you would have been also with predicting .

Induction: might fail, but if we are successful, then also by induction

5. Meta-Induction: The Mechanism I

How to cook up $pred_{MI}$:

- We measure the past success of a method by inverting the loss ℓ.

E_i	0	0	0
$pred_1$	1	0	1
$pred_2$	0	0	1

$\Rightarrow$

success
0.33
0.66

- We measure the attractivity of a method for the *MI*-method ($pred_{MI}$) by cutting off worse than *MI*-performing methods.

$pred_{MI}$	0.66
$pred_1$	0.33
$pred_2$	0.66

$\Rightarrow$

attractivity
0.0
0.66

- We calculate weights out of the attractivities.

	attractivity
$pred_1$	0.0
$pred_2$	0.66

$\Rightarrow$

weight
0.0
1.0

- We define $pred_{MI}$ by attractivity-based weighting of predictions $pred_i$.

6. Meta-Induction: The Mechanism II

$$success(pred_i, t) = \frac{\sum_{k=1}^{t} 1 - \ell(pred_i(e_k), E_k)}{t}$$

$$attractivity(pred_i, t+1) = \begin{cases} success(pred_i, t), & \text{if } success(pred_i, t) \geq \\ & success(pred_{MI}, t) \\ 0, & \text{otherwise} \end{cases}$$

$$weight(pred_i, t+1) = \frac{attractivity(pred_i, t+1)}{\sum_{k=1}^{n} attractivity(pred_k, t+1)}$$

$$pred_{MI}(e_{t+1}) = \sum_{k=1}^{n} weight(pred_k, t+1) \cdot pred_k(e_{t+1})$$

7. Meta-Induction: Hume's Problem* Solved

Main result of the meta-inductive research programme: long-run optimality; In the long run $pred_{MI}$'s performs at least as good as any other method, if loss ℓ is convex:

$$\lim_{t \longrightarrow \infty} success(pred_{MI}, t) - success(pred_i, t) \geq 0, \quad \text{for all } 1 \leq i \leq n$$

By this, success-based induction is justified (*per comparationem*).

Hence, given the past success of inductive methods as, e.g., the so-called *straight rule*, a success-based choice of these methods is also justified.

Proviso: garbage in $\Rightarrow$ garbage out, $pred_{MI}$ is "parasitical".

8. The Assumptions

Optimality of $pred_{MI}$ holds only ...

- ... a comparison with accessible prediction methods,
- ... consideration of the long run,
- ... a set of continuous (not discrete) predictions,
- ... a finite number of object-methods,
- ... a loss ℓ that is convex.

9. The Approach in a Nutshell

1. Hume's problem is about an absolute justification of induction. (Hume)
   Schema: $\mathcal{O}(\textit{absolute successful induction})$
2. There is no means for such an absolute justification. (no free lunch)
   Schema: $\neg\exists B(\Box(B \to \textit{absolute successful induction}))$
3. Hence: We need new epistemic ends. (ought implies can)
   Schema: $\neg\mathcal{O}(\textit{absolute successful induction})$
4. We aim at a relative justification. (from Reichenbach 1938)
   Schema: $\mathcal{O}(\textit{relative successful induction})$
5. There is a means for such a relative justification. (meta-induction)
   Schema: $(\textit{meta-induction} \circ\!\!\to \textit{relative successful induction})$
6. This justification depends on several assumptions. (meta-induction)
   Schema: $\textit{convexity} \;\to\; \textit{meta-induction}$

Slogan

Hume's Problem* Solved*!

. . . by Meta-Induction based on Convex ℓ.

The Conceptual Spaces Approach to the New Riddle

1. The Conceptual Spaces Approach

For the following, cf. (Gärdenfors 1990, 2000, sect.6.3.1).

Main idea: (AI) Problems related to knowledge are about knowledge representation; should be not in form of propositions, but by geometrical terms.

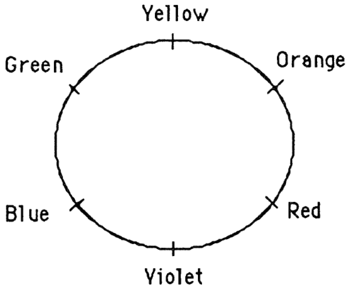

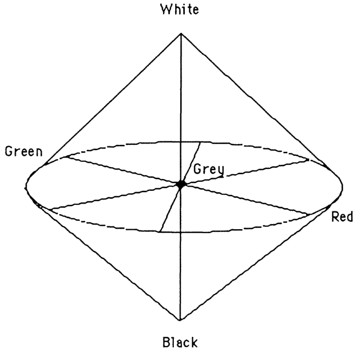

⇒ knowledge problems: metric properties instead of logical structure

2. Projectibility and Naturalness

- **Projectibility:** "[I] outline a theory of conceptual spaces. I shall argue that by representing knowledge in terms of conceptual spaces one can rather easily identify the projectible properties [with the natural ones]." (p.79)
- **Naturalness:** "The topological properties of the dimensions now allow us to introduce the notion of a natural property, which we have seen to be a central task for a theory of induction. The definition is simply that a property, that is, a region of a conceptual space, is natural only if the region is *convex*." (pp.87f)

$$\text{projectible/natural} \;\Rightarrow\; \textit{convex}$$

"Carnap [...] finds it "useful" to consider only *connected* [$<$ convex] regions of attribute spaces when looking for rules of inductive logic." (cf. p.88)

3. Conceptual Spaces: Convexity

Convex Space:

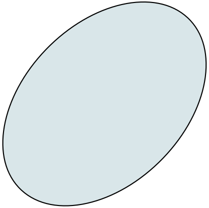

convex

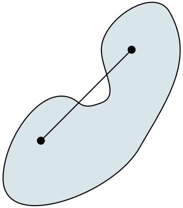

non-convex

Example of a conjecture (cf. p.88): Colours are natural kinds.

If o_1 and o_2 are said to have the colour C, then any object o_3 with a colour which lies between the colour of o_1 and o_2 will also be described by C.

4. Tackling the New Riddle

Given our standard representation of colours, *green* and *blue* are natural properties, while *gr*ue and *bl*een are not. (cf. pp.88f)

*Gr*ue presumes two dimensions, colour and time, for its description.

Even if we consider the cylindrical space that would be generated by taking the Cartesian product of the time and hue dimensions, *gr*ue would not represent a convex region, but rather be discontinuous.

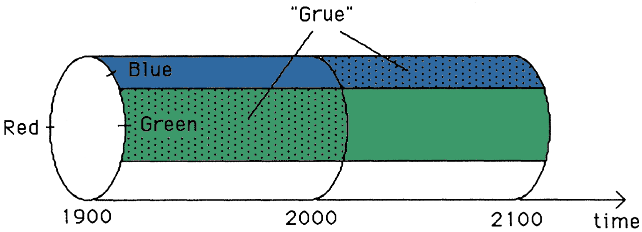

5. A Problem

"Even if predicates like *grue* and its ilk do not correspond to natural properties in our standard conceptual space it is conceivable that such predicates would correspond to natural properties in another conceptual space where, consequently, our predicates *green* and *blue* would denote non-natural properties." (p.90)

What counts as natural is *dependent* on the underlying conceptual space.

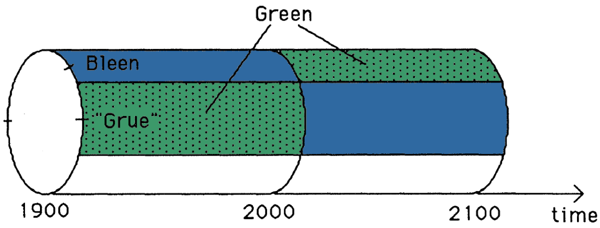

6. A Naturalised Solution

Gärdenfors (1990, pp.91f):

- "Human beings, to a remarkable extent, agree on which properties are projectible."
- "This far-reaching agreement suggests that the psychological conceptual spaces of humans are, at least in their fundamental dimensions, close to identical."
- "Why our way of identifying natural properties accords so well with the external world[?]"
- "The answer, it seems to me, comes from evolutionary theory. Natural selection has made us all develop a conceptual space that results in inductions that are valid most of the time and thus promote survival."

This "evolutionary" approach to *projectibility* is in the tradition of Peirce (1994).

7. The Approach in a Nutshell

1. The problem of induction consists in the problem of projectibility: Induction is valid only, if it is about projectible properties X. (from Goodman 1946)
   Schema: $\mathit{Induction}(X) \Rightarrow \mathit{Projectible}(X)$

2. A property is projectible iff it is natural. (classicism of natural kinds)
   Schema: $\mathit{Projectible}(X) \Leftrightarrow \mathit{Natural}(X)$

3. A property is natural only, if it(s extension) is convex. (conceptual spaces)
   Schema: $\mathit{Natural}(X) \Rightarrow \mathit{Convex}(X)$

4. *Green* is convex, whereas *Gr*ue is not. (fact of natural language)
   Schema: $\mathit{Convex}(\mathit{Green})$ & $\neg\mathit{Convex}(\mathit{Gr}\mathrm{ue})$

5. So, if at all, then induction is valid only with regards to *Green*, but not with regards to *Gr*ue. (1–4)
   Schema: $\Diamond\mathit{Induction}(\mathit{Green})$ & $\neg\mathit{Induction}(\mathit{Gr}\mathrm{ue})$

Slogan

Goodman's Riddle Solved*!

. . . by Convex Conceptual Spaces.

Meta-Induction and Convexity

1. Recall the Assumptions of Meta-Induction

Optimality of $pred_{MI}$ presupposes ...

- ... a comparison with accessible prediction methods,
- ... consideration of the long run,
- ... a set of continuous (not discrete) predictions,
- ... a finite number of object-methods,
- ... a loss ℓ that is convex.

How to justify the **convexity** assumption ?

2. Convexity of the Loss Function

$$\forall a, b, r, w : \ \ell(w \cdot a + (1 - w) \cdot b, r) \ \leq \ w \cdot \ell(a, r) + (1 - w) \cdot \ell(b, r)$$

Example: squared loss

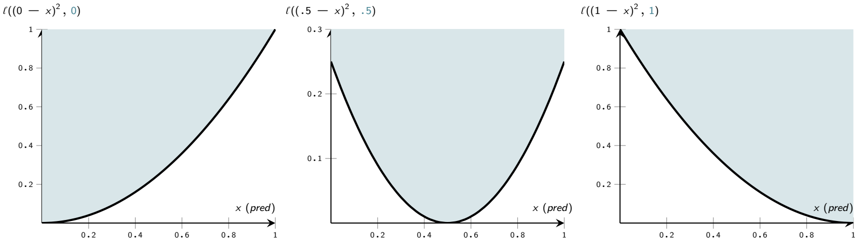

3. Convexity of the Loss Function

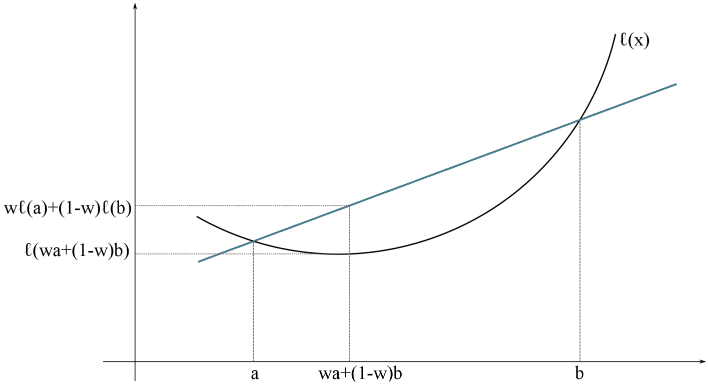

4. Different Forms of Non-Convexity

Non-Convexity:

$$\exists a, b, r, w: \ \ell(w \cdot a + (1-w) \cdot b, r) \ > \ w \cdot \ell(a, r) + (1-w) \cdot \ell(b, r)$$

Non-Convexity$^{\exists abr \forall w}$:

$$\exists a, b, r, \forall w: \ \ell(w \cdot a + (1-w) \cdot b, r) \ > \ w \cdot \ell(a, r) + (1-w) \cdot \ell(b, r)$$

Concavity:

$$\forall a, b, r, w: \ \ell(w \cdot a + (1-w) \cdot b, r) \ > \ w \cdot \ell(a, r) + (1-w) \cdot \ell(b, r)$$

5. Transferring Gärdenfors' Argumentation

1. The problem of induction consists in the problem of projectibility: Induction is valid only, if it is about something projectible. (from Goodman 1946)
   Schema: $\textit{Meta-Induction}(\ell) \Rightarrow \textit{Projectible}(\ell)$
2. A property is projectible iff it is natural. (classicism of natural kinds)
   Schema: $\textit{Projectible}(\ell) \Leftrightarrow \textit{Natural}(\ell)$
3. A property is natural only, if it(s extension) is convex. (conceptual spaces)
   Schema: $\textit{Natural}(\ell) \Rightarrow \textit{Convex}(\ell)$
4. So, if at all, then induction is valid only with regards to a convex loss ℓ. (1–3)
   Schema: $\textit{Meta-Induction}(\ell) \Rightarrow \textit{Convex}(\ell)$

Problems:

- relativity to underlying conceptual space
- conditional justification of convexity (Hume✓ ⇒ Goodman✓ via Convexity; however, we need Convexity already for Hume✓)

6. Strengthening the Argument by Epistemic Engineering

Recall the Epistemic Means-End Principle:

$$\overbrace{(\mathcal{O}(A)}^{\text{ulterior epistemic ends}} \quad \& \quad \underbrace{\Box(A \rightarrow B) \vee (B \circ\!\!\rightarrow A))}_{\text{epistemic engineering}} \quad \rightarrow \quad \overbrace{\mathcal{O}(B)}^{\text{derived epistemic ends}}$$

Recall the Justification of (Meta-)Induction:

- $\mathcal{O}(\textit{relative success})$
- $(\textit{meta-})\textit{induction} \circ\!\!\rightarrow \ \textit{relative success}$
- $\mathcal{O}((\textit{meta-})\textit{induction})$

We want to argue for:

- Not only: $(\textit{meta-})\textit{induction} \rightarrow \textit{convexity}$
- But also: $\Box((\textit{meta-})\textit{induction} \rightarrow \textit{convexity})$
- Hence: $\mathcal{O}(\textit{convexity})$

7. Convexity as a Necessary Epistemic Means

Main Result:

Assume $pred_{MI}$ is an ℓ-based weighting method (based on $w, \ell, \mathcal{L}$, where $\mathcal{L}(\alpha, n)$ is α's cumulative loss up to n), satisfying

Monotonicity:
$\forall \alpha, \beta, n: \quad \mathcal{L}(\alpha, n) \ \leq \ \mathcal{L}(\beta, n) \quad \Rightarrow \quad w(\alpha, n+1) \ \geq \ w(\beta, n+1)$

Then convexity* of ℓ is necessary for the optimality of $pred_{MI}$.

I.e.: $\Box((\textit{meta-})\textit{induction} \rightarrow \textit{convexity})$

Hence: $\mathcal{O}(\textit{convexity})$

* more specifically: not non-convexity$^{\exists abr \forall w}$

8. Possible Objection

Non-convex loss ℓ allows for the intuition of rewarding decisiveness of predictions. E.g. for $\ell(x, 1)$:

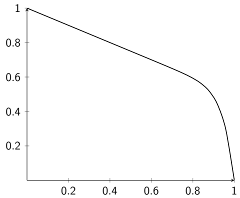

Solution: Why not shifting the decision theoretic part to the utilities?

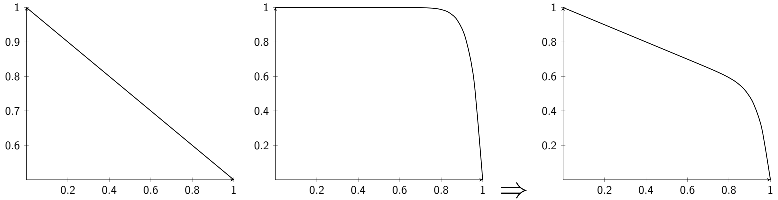

Summary

- Hume's Problem: How to justify Induction?
- Goodman's New Riddle: How to rule out Anti-Induction?
- Gärdenfors' Conceptual Spaces:
  - Highlights Convexity vs. Anti-Induction
- Reichenbach's Vindication $\Rightarrow$ Meta-Induction
  - Highlights relative success of Induction
  - Based on Convexity assumption
- So, both accounts hinge on Convexity
- Problem: How to justify the convexity assumption?
- Gärdenfors' general answer: via Naturalised Epistemology
- Our answer for meta-induction: via Epistemic Engineering

References I

Bacon, Francis (1620/2000). "The New Organon". In: *The New Organon*. Ed. by Jardine, Lisa and Silverthorne, Michael. Cambridge: Cambridge University Press.

Barnes, Eric (1991). "Beyond verisimilitude: A linguistically invariant basis for scientific progress". In: *Synthese* 88.3, pp. 309–339. URL: http://www.springerlink.com/content/n415u420641701m8/fulltext.pdf.

Carnap, Rudolf (1950/1962). *Logical Foundations of Probability*. London: Routledge and Kegan Paul.

— (1928/2003). *The Logical Structure of the World and Pseudoproblems in Philosophy*. Open Court Classics. Illinous: Open Court.

Feldbacher-Escamilla, Christian J. (under revision). *Epistemic Engineering. Uncovering the Logic of Deceivability and Meta-Induction*. book manuscript.

Feldbacher-Escamilla, Christian J. and Thorn, Paul D. (submitted). "Meta-Induction and its Assumptions: The case of convexity". In: *manuscript*.

Gärdenfors, Peter (1990). "Induction, Conceptual Spaces and AI". In: *Philosophy of Science* 57.1, pp. 78–95. DOI: 10.1086/289532.

— (2000). *Conceptual Spaces. The Geometry of Thought*. Cambridge, Massachusetts: The MIT Press.

Goodman, Nelson (1946). "A Query on Confirmation". In: *The Journal of Philosophy* 43.14, pp. 383–385. DOI: 10.2307/2020332.

References II

Goodman, Nelson (1955/1983). *Fact, Fiction, and Forecast*. Ed. by Putnam, Hilary. Fourth Edition. Harvard: Harvard University Press.

Hume, David, ed. (1748/2007). *An Enquiry Concerning Human Understanding*. Oxford World's Classics. (ed. by Millican, Peter). Oxford University Press.

Keynes, John Maynard (1921). *A Treatise on Probability*. London: Macmillan and Co.

Mitchell, Tom (1997). *Machine Learning*. New York: McGraw-Hill.

Hartshorne, Charles, Weiss, Paul, and Burks, Arthur W., eds. (1994). *Collected Papers of Charles Sanders Peirce*. Cambridge, MA: Harvard University Press.

Popper, Karl R. (1935). *Logik der Forschung. Zur Erkenntnistheorie der modernen Naturwissenschaft*. Ed. by Frank, Philipp and Schlick, Moritz. Schriften zur Wissenschaftlichen Weltauffassung. Wien: Springer.

Reichenbach, Hans (1938). *Experience and Prediction. An Analysis of the Foundations and the Structure of Knowledge*. Chicago: University of Chicago Press.

Schurz, Gerhard (2008). "The Meta-Inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem". In: *Philosophy of Science* 75.3, pp. 278–305. DOI: 10.1086/592550.

— (2019). *Hume's Problem Solved. The Optimality of Meta-Induction*. Cambridge, Massachusetts: The MIT Press.

Skyrms, Brian (2000). *Choice and Chance. An Introduction to Inductive Logic*. Fourth Edition. Stamford: Wadsworth.

References III

Tichý, Pavel (1976). "Verisimilitude Redefined". In: *The British Journal for the Philosophy of Science* 27.1, pp. 25–42. DOI: 10.1093/bjps/27.1.25.

Wolpert, David H. (1996). "The Lack of A Priori Distinctions Between Learning Algorithms". In: *Neural Computation* 8.7, pp. 1341–1390. DOI: 10.1162/neco.1996.8.7.1341.