University of Groningen

What conceptual spaces can do for Carnap's late inductive logic

Sznajder, Marta

Published in:
Studies in History and Philosophy of Science. Part A

DOI:
[10.1016/j.shpsa.2015.12.001](https://doi.org/10.1016/j.shpsa.2015.12.001)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2016

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Sznajder, M. (2016). What conceptual spaces can do for Carnap's late inductive logic. *Studies in History and Philosophy of Science. Part A*, 56, 62-71. <https://doi.org/10.1016/j.shpsa.2015.12.001>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

What conceptual spaces can do for Carnap's late inductive logic

Marta Sznajder^{a,b}

^a Munich Center for Mathematical Philosophy, Geschwister-Scholl-Platz 1, 80539 Munich, Germany

^b University of Groningen, Netherlands

ARTICLE INFO

Article history:

Received 4 May 2015

Received in revised form

14 December 2015

Available online 25 January 2016

Keywords:

Inductive logic;

Carnap;

Conceptual spaces;

Attribute spaces;

Linguistic frameworks

ABSTRACT

In the last published account of his late inductive logic, the Basic System of Inductive Logic, Rudolf Carnap introduced a new element to the systems of inductive logic, namely the so-called attribute spaces. These geometrical structures model the meanings of the predicates of the object language and have a similar structure as the conceptual spaces employed by cognitive scientists like Peter Gärdenfors. I show how the development of the theory of conceptual spaces helps us to see the addition of attribute spaces as a step forward in explicating the concept of confirmation. I discuss the differences and similarities of the two theories and investigate the possibilities for developing further connections.

© 2016 Elsevier Ltd. All rights reserved.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Science*

1. Introduction

In recent years we can observe two seemingly unrelated phenomena. The first one, related to the growing interest in formal philosophy, is that the interest in Rudolf Carnap's work is picking up even more momentum (Awodey & Klein, 2004; Carus, 2007; Leitgeb, 2011). The second is the emergence and strengthening of the theory of conceptual spaces, advanced mainly in cognitive science and related fields by authors such as Peter Gärdenfors. Conceptual spaces have been by now successfully applied also to classical philosophical problems (Zenker & Gärdenfors, 2015) and thus formal philosophers are beginning to take interest in them.

What is relatively unknown is that Carnap in his late work on inductive logic introduced and made heavy use of a notion very similar to today's conceptual spaces—namely, attribute spaces. The addition of the attribute spaces to inductive logic has not been discussed a lot in the literature since Hilpinen (1973), and the possible connections with the current work done in cognitive science have not received due attention (Gärdenfors (2004) mentions Carnap's work only in passing). This paper lies the foundations for

developing connections between the two theories. In particular, I show how the theory of conceptual spaces can be seen as a factor in the explicative success of inductive logic.

When I mention Carnap's late work on inductive logic, I am referring to "A Basic System of Inductive Logic" (henceforth *the Basic System*), a two-part paper published as (Carnap, 1971a, 1980). Carnap concluded "Logical Foundations of Probability" (1950) with an exposition of the contents of the future second volume of that book—that was never published. In the course of the 1960s it became clear that his plan was not going to be realized and that the scope of another volume would have to be different. The Basic System is the result of those years' work and the system presented there was supposed to serve as a foundation for future inductive logic (see "Introduction" in (Carnap & Jeffrey, 1971)).

One big difference between the Basic System and Carnap's previous inductive logic is the introduction of the aforementioned attribute spaces—geometric representations of concepts—as another element of linguistic frameworks. The idea that such geometric structures could be introduced to model meaning relations between predicates was not entirely novel at the time. Van Fraassen (1967) already introduced a logical space—essentially a set of points representing some (not necessarily all) possible objects (which Carnap would think of as objects of a particular

E-mail address: Marta.Sznajder@lrz.uni-muenchen.de.

investigation)—used to model intensions of predicates. In this setting predicates correspond to subsets of the set of points of the space and hence intensional relations such as synonymy or inclusion of meaning can be read off directly from the space, avoiding the need for meaning postulates. Stalnaker (1981) used similar structures to define properties without the use of possible worlds.

Section 2 introduces the main elements of the Basic System, focusing on the notions of a linguistic framework and a confirmation function. Attribute spaces are then presented separately in Section 3, with the discussion of their role in the system in Section 4. All of the above sets the scene for Section 5 in which the addition of the attribute spaces is evaluated using the criteria for evaluating explications. This leads to the discussion of the theory of conceptual spaces in Section 6, followed by Section 7 which clarifies the possible relations between inductive logic and the theory of conceptual spaces, pointing out to some open questions.

2. The Basic System of Inductive Logic

Linguistic frameworks are a recurring theme in Carnap's writing and much of his work can be seen as a gradual elaboration of this notion. In the most general way the frameworks can be described as different language systems adopted in order to be used as a basis for scientific investigations. The closest Carnap comes in the Basic System to defining a linguistic framework is the following: "a universe of objects and a system of descriptive concepts that characterize the objects" (Carnap, 1971a, 47). The objects that a framework was designed to talk about are thus included as a part of the framework itself.

In the Basic System a linguistic framework is no longer just a language in the sense of a vocabulary and a set of rules governing the formation of expressions; it consists of many additional elements, namely a geometrical system of concepts (possibly, though not necessarily, supplied with a language) to be used for studying and describing a particular set of individuals, together with that set of individuals, and a number of meta-constraints and rules that determine the admissible confirmation functions for propositions about those objects. Later on it will be made clear what is meant by a 'geometrical system of concepts'; for now it suffices to say that in a linguistic framework the concepts are represented geometrically, and that in virtue of this representation the relations between concepts are derived from topological and metric properties of such a geometrical structure. Linguistic frameworks are set up for the purpose of conducting scientific investigations. Carnap describes a paradigmatic case of such activity as follows: "The person X wishes to assign rational credence values to unknown propositions on the basis of the observations he has made" (1980, 106).

2.1. Language and semantics

Although Carnap puts some stress onto the idea that concepts in a framework do not need to belong to a language, it still aids the clarity of exposition to present the formal apparatus of the system using the notion of a formalized language whose predicates correspond to the relevant concepts. "In this article the signs, formulas, and sentences do not play an essential role (...) at some occasions a reference to linguistic entities is made in order to facilitate understanding" (1971a, 47). Hence the notion of a formal language is still present. In the Basic System Carnap works with the same kind of formal language we recognize from his earlier works on probability. $\mathcal{L}$ is a monadic predicate language with (usually infinitely) countably many individual constants ($a_1, a_2, \dots$), and a standard logical vocabulary (connectives, identity, quantifiers, individual variables). The logic of $\mathcal{L}$ is classical, as is the metalogic of the system. The object language is chosen for an investigation and

A family of predicates

e.g. color names

Predicates

R B G ...

A modality

e.g. color

Attributes

Red Blue Green ...

Fig. 1. Predicates and modalities.

therefore the choice of its particular characteristics will depend on the nature of the investigation and its goals.

The basic predicates of the language represent primitive attributes (properties) of the objects: "observable properties of observable objects" (1971a, 43). They are ordered into finitely many families. Two properties belong to the same family if they belong to the same general kind—here called a modality—for instance a modality of color or age¹ (see Fig. 1). The families are disjoint and exhaustive, i.e. it is logically necessary for any object to have at least one of the attributes of the family and it is inconsistent for it to have more than one of these properties. The division of predicates into families was already mentioned in the Logical Foundations (1950, §18c), where Carnap speaks of *families of related attributes*. However, it wasn't until the Basic System that the modalities were introduced as the factor determining the families.

Modalities are vital for the structure of the framework, since the division of the language into families is governed by them. As to what they are, Carnap only says that modalities are "general kinds" to which predicates of the same family belong, and that they can be qualitative (e.g. color, shape, substance) or quantitative (e.g. age, height). The best way to think about modalities is as those respects in which objects can be judged similar or different. While the attributes subsumed under a modality are comparable (e.g. we can say that red and orange are more similar to each other than red and blue), the modalities themselves are not (e.g. it is not immediately meaningful to say that color is more similar to shape than color is to sound).

Models for the object language are interpretation functions assigning attributes (properties) to individuals from the chosen domain. The domain is assumed to be always countable, and usually infinite, and for every individual there is exactly one constant denoting it. The initial set of all models—that is, of all the possible ways of assigning predicates to the individuals—can be further restricted by meaning postulates assumed for a given investigation (different kinds of postulates are discussed in Section 5 and an example of the restriction procedure can be found on pp. 82–83 of (Carnap, 1971a)). As will be seen shortly, the set of models for a given investigation becomes the sample space over which conditional probability functions representing confirmation will be assigned.

Propositions are sets of models, that is events on the sample space for the confirmation measures. Atomic propositions on $\mathcal{L}$ correspond to atomic sentences of $\mathcal{L}$, which are of the form $P_j a_i$ (where a_i is an individual constant). Hence atomic propositions are sets of those models in which $P_j a_i$ holds. The class of all propositions, denoted $\mathcal{E}_{\mathcal{L}}$, is a σ -field on the set of all models, generated by the collection of atomic propositions (1971a, 37). The latter means that $\mathcal{E}_{\mathcal{L}}$ is closed under complementation and finite and countably infinite unions. $\mathcal{E}_{\mathcal{L}}$ is very rich: for each class of sentences of the language, there is a corresponding proposition in $\mathcal{E}_{\mathcal{L}}$ (1971a, 60).

¹ In modern linguistics such families of predicates are referred to as domains, see (Langacker, 1987, chap. 4).

(Second argument of $\mathcal{C}(H|E)$ always assumed nonempty)

A1	Axiom of the lower bound	$\mathcal{C}(H E) \geq 0$
A2	Axiom of self-confirmation	$\mathcal{C}(E E) = 1$
A3	Axiom of the complement	$\mathcal{C}(H E) + \mathcal{C}(-H E) = 1$
A4	General multiplication principle	If $E \cap H$ is possible, $\mathcal{C}(H \cap H' E) =$ $= \mathcal{C}(H E)\mathcal{C}(H' E \cap H)$
A5	Axiom of regularity	$\mathcal{C}$ is regular
A6	Axiom of symmetry	$\mathcal{C}$ is symmetric (wrt individuals)
A7	Axiom of convergence	Let $\mathcal{C}$ be a symmetric $\mathcal{C}$ -function. Then $\lim_s [\mathcal{C}(P_j a_{s+1} E_s(Z)) - r_j^s(Z)]$ exists and is 0 for all Z in $\mathbf{Z}$.

Fig. 2. Axioms of the Basic System.

However, many propositions in $\mathcal{L}$ do not correspond to any single sentence of $\mathcal{L}$. For instance, in case there are infinitely many objects in the domain, the proposition that specifies for each object which property from a given family that object has (also known as a state description) does not correspond to any sentence in $\mathcal{L}$, since it corresponds to an infinite conjunction of formulas of the form $P_j a_i$. Such richness of the algebra of propositions is needed in order for it to serve its main purpose, that is serving as a base for making predictions. Notice that because the class of propositions is defined purely set-theoretically on the basis of the set of atomic propositions, the system of concepts can be viewed as being independent of the syntax of the chosen language.

2.2. Confirmation functions

Setting up a linguistic framework, as described so far, occurs with one goal in mind: to perform scientific investigations. The kind of activity that this formal setup aims to model consists of gathering observational data and forming beliefs regarding future observations. Confirmation functions are supposed to be formal representations of “rational credibility functions”—that is (conditional) credence functions that are rational to have before any evidence is accumulated. According to Carnap, while credences may reflect temporary opinions of an agent, a credibility, or initial credence, function is “a trait of his underlying permanent intellectual character” (1971b, 19). Inductive logic, then, is meant to restrict the range of possible confirmation functions to include only those that correspond to rational credibility functions. Therefore each of the axioms of inductive logic reflects some rationality condition on credibility.

In the Basic System confirmation functions ($\mathcal{C}$ -functions) are conditional probability assignments of real values to pairs of propositions, with the second proposition nonempty. These values are meant to represent degrees to which the second proposition (the evidence) confirms the first (the hypothesis)²: a formula $\mathcal{C}(H|E) = 1/3$ says that E confirms H to degree $1/3$. Unconditional probability ($\mathcal{P}$ -functions) is defined in a standard way as probability conditional on tautological evidence, $\mathcal{P}(A) = \mathcal{C}(A|\top)$. The

axioms of the Basic System, summarized in Fig. 2, are the basic axioms for Carnap functions plus symmetry with respect to individuals and regularity. Reichenbach's axiom of convergence is also included in the system.

The first four axioms define conditional probability functions (see also (Zabell, 2011, 267)). Given the subjective interpretation of probability functions mentioned above, it comes as no surprise that the justification given for the probability axioms relies on Dutch book arguments (Carnap, 1971a, chap. 8). The rationality condition that gives rise to the first four axioms is coherence: having a probabilistically incoherent credence is irrational, because it commits the agent to accepting bets that guarantee a loss.

The Carnapian statement of the axiom of regularity differs slightly from what is more commonly known under that label, namely the requirement that only logical falsehoods are assigned probability 0 (and, as a consequence, only logical truths can have probability 1). Kemeny's formulation (Kemeny, 1955, 266) of regularity in terms of conditional probability assignments, rephrased in the language of the Basic System, would read: $\mathcal{C}(H|E) = 1$ iff $E \subseteq H$, that is: the only hypotheses confirmed to degree 1 by the evidence are the logical consequences of the evidence (and the only hypotheses with degree of confirmation 0 are the ones that are logically incompatible with the evidence). In the Basic System, though, Carnap allows for additional cases in which $\mathcal{C}(H|E) = 0$ (D7-1 of (Carnap, 1980)). First, this can happen when the propositions in question are mutually exclusive, which amounts to Kemeny's regularity. Second, when H and E both fail to be molecular propositions (which means they cannot be constructed from atomic propositions by complementation and finite unions). Third, when some family of predicates involved in those propositions is infinite. The latter two cases cover the case of propositions that correspond to quantified sentences and sentences that in a sense are like quantified sentences.

Carnap's rationale for assuming regularity is based on a Dutch book style argument for strict probabilistic coherence as a rationality requirement for credences. An agent whose credence function is not strictly coherent—that is, not regular—can be presented with a finite series of bets each of which she will accept, the bets being such that no possible outcome will result in a gain for the agent; only losses and zero gains are possible. According to Carnap, an agent accepting such a system of bets is irrational because “It is clearly unreasonable to enter a contract involving the risk of a loss which is not balanced by the possibility of a gain” (1980, 114).

The axiom of symmetry requires confirmation functions to be invariant under finite permutations of individuals in the domain. Assuming individual symmetry amounts to saying that it does not matter for the confirmation function whether a_1 or a_3 have the property P —it only matters how many individuals have that property. A symmetric confirmation function then will assign the same values to the propositions with the same corresponding structure description (given the same evidence). The reason for accepting the symmetry axiom stems from the fact that confirmation functions are supposed to model initial credences. Since no evidence is available that could inform us about the differences between particular individuals, we should treat them equally when it comes to beliefs about them; in particular, we shouldn't expect any of them to have a given property any more than other objects to have that same property (Carnap, 1971b, 23). Thus, the axiom of symmetry is one way of applying the principle of indifference—which says that whenever we do not have information distinguishing between a set of given alternatives, we should assign the same probability to each alternative—in formulating inductive logic.

Finally, there is the axiom of convergence (Reichenbach's axiom). Informally, the axiom can be spelled out in the following

² This is Carnap's measure of absolute, as opposed to incremental, confirmation (see (Carnap, 1950, §87)).

way: as the evidence grows, the expectation for observing a P_j -object, $P_j a_{s+1}$, given evidence consisting of s observations, should converge to the relative frequency of P_j type observations among the total evidence. In the notation of Fig. 2, $E_Z(Z)$ is the proposition corresponding to the sentence that says which property P_i each of the first s individuals ($a_1, \dots, a_s$) has in the model Z . Furthermore, $r_j^s(Z)$ is the relative frequency in Z of P_j objects among the first s individuals. The axiom says that for each model Z , in the limit those two values are equal. The axiom of convergence ensures that, loosely speaking, regardless of what the world is like, e.g. which objects are in fact P , the expectation for observing a P thing will converge to the actual proportion of the P s in the observed sample, as the sample grows larger and larger. In other words, experience will always eventually trump any beliefs held prior to obtaining the evidence; the agent cannot indefinitely resist learning from experience.

3. Attribute spaces

Once the set of individuals and the concepts or predicates that describe them are given, the investigator faces one more choice to make: the choice of attribute spaces that accompany the concepts. Before the Basic System, in the inductive logic familiar from “Logical Foundations of Probability”, the language and its semantics as described above would provide the only structure that accompanied the confirmation functions. The Basic System introduces a new element: attribute spaces. They are structures that model the meaning of predicates, providing a (partial) interpretation for the language (Carnap, 1971a, 56). However, Carnap does not explain why he chose this particular way of representing meanings of the predicates; the examples he provides in the text suggest that he was familiar with the developments of similar techniques in psychology.

There are two uses of these structures within the Basic System. First, they encode some assumptions about the meaning relations between predicates in a family, like the fact that predicates within a family are supposed to be mutually exclusive. In this capacity attribute spaces can encode some of the meaning postulates. The second use of the attribute spaces is in informing the choice of some of the values of the confirmation functions.

An attribute space is “an abstract, logical space whose points represent the elementary (i.e., most specific) properties of the modality in question” (Carnap, 1971a, 43). Each family of predicates of the language is assigned its own space. For example, in the case of the family of color names, its space will consist of points that are thought of as the most specific colors. To each predicate a whole region of the space is assigned (there are many points that can be examples of *red*). Since the families consist of mutually exclusive and logically exhaustive predicates, every family determines a partition of its space, called an attribute partition.

Consider one of the simpler examples of an attribute space, a version of which was provided by Carnap (1971a, sec. 2): an interval of real numbers ranging from 0 to 120. For an actual investigation about, say, the population of Munich in 2014, it can be used as a one-dimensional space for a family of predicates containing the predicates *younger than 18*, *18–24 years old*, *24–55 years old* and *at least 55 years old*. It is clear that the “resolution” of these predicates is much less fine-grained than that of a real interval; therefore to each of the predicates there will correspond an interval of our one-dimensional space. Care must also be taken to ensure that the predicates are mutually exclusive and exhaustive—which in this case would amount to, e.g. deciding where the cutoff points between pairs of neighboring predicates lie.

Another, more elaborate example of an attribute space is the well-known three-dimensional color space, already discussed by Carnap in the *Aufbau* (1967, §90 and §118). Every color can be assigned a value on each of the three dimensions: hue, brightness, and saturation. The color space is, roughly, a three-dimensional double cone. Its two vertices are pure black and pure white colors, the vertical axis of the double cone therefore corresponding to the brightness dimension. The circular cross-section in the middle is the familiar color circle. The hue value of each color is given by its angle on the circle, while the distance from the vertical axis is the saturation value. Here, again, individual points in the space correspond to specific shades of specific colors. No scientifically useful set of color predicates will have enough names to cover the space point by point (not only for reasons of cardinality, but also because of the properties of human perception and the available methods and devices for measurement)—so again the names of colors assumed for a given investigation will be assigned regions in the space. For example, the predicate *red* will be assigned a region roughly neighboring to the region corresponding to *orange*.

While Carnap does not give a specific recipe for choosing the shape of attribute spaces, the number of dimensions of an attribute space seems to depend on what kind of modality that space corresponds to, and since “the classification of attributes into different modalities is not unique but can be made in various ways” (1971a, 71), there can be multiple ways of dividing an attribute space into dimensions. For instance, the color space can be either seen as a three-dimensional space based on the modality of color (which is the most common interpretation), or as three one-dimensional spaces based on three modalities of hue, brightness and saturation (1971a, 71). This picture in fact shows a very Carnapian approach to such issues: choose the shape of your spaces as it suits your scientific goals.

The regions corresponding to predicates are characterized by two magnitudes: their size and the distances between pairs of them. The available notion of distance between points of the space can be qualitative, although quantitative, normalized distance measures—metrics—are to be preferred (1980, 10). A quantitative distance measure, i.e. one specifying for each pair of points an exact number representing the distance between them, provides a more straightforward translation of its relevant values into probability assignments with the use of the γ and η rules discussed in Section 4.

The concept of distance in an attribute space is tightly connected to similarity judgments. Similarity is a relation that is, as a matter of degree, assumed to hold between attributes in the attribute space, that is between the properties studied in a given investigation. Therefore statements about how similar two attributes are cannot be a part of the object language $\mathcal{L}$ (since the latter can say only which individuals in the domain have which properties). On the basis of the distance function on points of the attribute space, Carnap defines distances between regions, i.e. predicates (Carnap, 1980, 25). Note, however, that the assumption that similarities between predicates behave like distances is a strong one, for it follows from it that the numerical representations of similarities satisfy the conditions satisfied by distance functions, like symmetry or the triangle condition (an assumption challenged by e.g. Tversky (1977)).

Where do the similarity judgments come from then? Unlike the *Aufbau*, where the basic relation R_s of (recollection of) similarity holds between elements of the “autopsychological” domain of private sense experience (Carnap, 1967, §78), the Basic System does not mention similar restrictions on the origin of the statements of similarity used in the construction of an attribute space. In the most basic case similarity seems to be similarity as perceived by a single agent, but Carnap does not seem to exclude cases of aggregation of judgments of groups of people (1980, 11).

However, intersubjective similarity is also deemed possible. For some investigations the similarity judgments may be informed by a background theory—especially when it comes to highly theoretical enterprises—although such a situation presupposes a move into inductive logic for a theoretical language (which in the Basic System is only mentioned as future work): “The basis of IL_O is different for a person of normal color vision and one with a specific color blindness (...). In contrast, theoretical physics supplies an intersubjective basis” (1971a, 52). When the underlying modality is itself quantitative and measurable, like age or length, the metric on such a space is immediately given by the standard scale used to measure the magnitude in question. In this case the statements about similarity of attributes are not needed as a pretheoretical basis for the quantitative concept of distance, on the contrary: they can be straightforwardly derived from statements about distances.

4. Additional rationality constraints

One of the uses that the Basic System puts attribute spaces to is constraining some values of rational confirmation functions. Once the whole linguistic framework is in place—that is, a language, its semantics, and the attribute spaces related to its families—there are some values of the confirmation functions that will be determined by the framework itself. That means: given your linguistic framework, there are certain beliefs that are rational for you to hold.

In order to say which values of the confirmation functions can be deduced from the attribute spaces and what the form of the rationality constraints is, some further formal concepts need to be introduced. In “The Continuum of Inductive Methods” (Carnap, 1952) the prior expectation for observing one object of any given kind depended only on the number of predicates in the language. The values of the unconditional probability assignments $\mathcal{M}$ for all propositions $P_j a$ (the belief of observing an object of the kind P_j prior to acquiring any evidence) were taken to be the same and to be equal to $1/k$ for predicates belonging to families of k predicates: $\mathcal{M}(P_j a) = 1/k$. In the Basic System that assumption was lifted, and the values of the unconditional probability function for single observations are now tied to the geometrical structure of the attribute space underlying a family of predicates and do not have to be equal for every predicate.

Carnap introduces two new sets of parameters— γ and η . The two kinds of parameters are the values of the confirmation function for some specific, simple propositions. The γ_j for a predicate P_j is the unconditional probability of observing a P_j object—the prior expectation of an outcome of the type P_j : $\gamma_j = \mathcal{M}(P_j a)$. Parameters of the type η_{jl} concern pairs of predicates P_j, P_l ; η_{jl} measures how much the knowledge that an object of type P_l was observed influences the belief that a P_j will be observed. Hence η_{jl} is the ratio of the probability of observing a P_j thing given the evidence consisting of a single outcome of the type P_l , and $\mathcal{M}(P_j a)$:

$$\eta_{jl} = \frac{\mathcal{C}(P_j a | P_l b)}{\mathcal{M}(P_j a)} = \frac{\mathcal{M}(P_j a \wedge P_l b)}{\mathcal{M}(P_j a) \mathcal{M}(P_l b)}$$

The η_{jl} is supposed to measure the strength of the analogical effect between P_j and P_l : whether and how much our expectation for a P_j observation changes after observing a P_l object when no other evidence is available. Another way to think about the η 's is as a coherence measure (Shogenji (1999) proposes the same formula to measure pairwise coherence of beliefs). Under this interpretation η_{jl} measures the extent to which $P_j b$ and $P_l a$ cohere with each other,

in other words: how surprised we are to see $P_j b$ after observing $P_l a$; the higher the η_{jl} , the less surprising the second observation. Note that once the values of all η_{jl} 's are given, other values of the confirmation function need not be uniquely specified.

The new parameters are used to describe a wider class of confirmation functions: the λ -system known from the Continuum turns out to be a special case within the Basic System. Now the λ -parameter can be specified only for languages in which the η_{jl} values are the same for all pairs of predicates, leading to a $\lambda - \gamma$ continuum, a further generalization of the λ continuum.

In the Basic System Carnap proposes two rules connecting certain features of attribute spaces with the values of the γ and η parameters. Both of these rules can be argued to presuppose thinking of collecting evidence as observing points in the space, i.e. particular qualities of an object, rather than as simply recording what is seen using just an atomic proposition. The first, called the γ rule, relates the size of the region corresponding to a predicate P_j with the prior probability $\mathcal{M}(P_j a)$ of observing an object of this type. The larger the region, the higher the probability. In the case of there being available a normalized width function (i.e. a function, based on the distance function on the space, that to each region in the space assigns its size expressed as a fraction of the size of the whole space), the probabilities are taken to be equal to the sizes of the relevant regions (Carnap, 1980, 34). For instance, if the region corresponding to the predicate *Red* spans over a quarter of the whole space, the recommended probability of observing a red object conditional on empty evidence is 0.25.

The second rule, called the η rule, is supposed to enable the confirmation functions to capture inductive reasoning that uses analogy by similarity between predicates belonging to the same family. The general idea behind such reasoning is that if two predicates are similar to each other, observing an object of the first kind will to some extent raise not only our expectation of observing another one of the same kind (instantial relevance), but also the expectation of observing an object of the second kind. For instance, consider a family of predicates *Hippie*, *Hipster*, and *Businessman*, partitioning the set of possible humans. If I take hippies to be more similar to hipsters than to businessmen, upon seeing a lot of hippies people in a bar I will expect to see more hipsters than businessmen in there as well. Notice, however, that if similar predicates were to receive such an “analogy boost” every time a relevant observation was made, we would run a risk of violating the axiom of convergence. Carnap's solution to this problem was to limit the analogy effect to cover only the cases in which one observation is made—hence the introduction of the η_{jl} parameters³—and to let the effect fade out as more evidence is considered (the rate at which it would diminish was left as an open question (Carnap, 1980, 45)).

Carnap's way of spelling out this notion of similarity of predicates proceeded through the use of attribute spaces. Recall that what the distances in a space are capturing is exactly the similarity between different qualities. Hence analogical reasoning and distances between predicates could be somehow tied to each other. However, while the γ rule gives a straightforward recipe for establishing the value of each γ_j , in the case of connecting conditional probabilities to distances between regions, introducing a straight numerical constraint on the values of $\mathcal{C}(P_j b | P_m a)$ given the distances between regions for P_j and P_m seemed too arbitrary to Carnap. Therefore he formulated only a qualitative constraint (the η

³ It is worth noting that some authors choose to give up convergence rather than restrict analogy (Costantini, 1983; Huttegger, 2015). See also Niiniluoto (1981), who however retracted that position (Niiniluoto, 1988).

rule): if the attribute P_m is closer to P_j than P_n is, evidence of the form $P_m a$ will have a greater influence on the posterior for $P_j b$ afterwards than observing $P_n a$. Thus the rule does not say how much our expectations should be influenced given the distance between two predicates. Rather, it gives a comparative restriction: different distances should lead to different analogical effects.⁴

In the λ -continuum the γ_j values for all predicates P_j were taken to be equal. This was motivated by indifference considerations: Carnap believed that invariance of the confirmation function under permutation of predicates within a family belonged to the “valid part of the principle of indifference” (Carnap, 1963, 975). One way of spelling out the principle of indifference (Keynes, 1921, 42) is that we should assign equal probabilities to alternatives of which there is no known reason to prefer some over the other. In the system of the Continuum there was nothing in the language that would differentiate between different predicates of the same family, and hence there was no reason for unequal γ s. In the Basic System, however, there is extra information on the predicates, namely the position and size of their corresponding regions in the attribute space. So applying the principle of indifference on the level of predicates is no longer an option.

However, in the way that Carnap conceived of the γ rule, it presupposes the idea that indifference considerations should be applied on the level of points of the space rather than on that of the predicates of the language. That is, the set of alternatives which should be considered equiprobable because of no information distinguishing between them, is the set of points of the space. And once the initial probability mass is distributed uniformly over the whole space, the probabilities for regions will indeed be equal to their sizes.

This move can be seen as a way of addressing some of the well-known objections to the principle of indifference (see Gillies (2000, chap. 3)) concerning the problem of there being too many alternative ways to carve up the problem space. The attribute space is supposed to offer a privileged set of indivisible alternatives—i.e. particular qualities associated with a given modality—which can be later on divided into a family of predicates in various ways via attribute partitions. For example, the color space provides all the possible ways in which an object can have a color—be colored, so to speak—the ways that are considered possible within the chosen linguistic framework. In this sense attribute spaces are a constitutive a priori (in the sense discussed by Friedman (2001)) element of a linguistic framework: they define what is empirically possible at all. And in that sense they provide *all* the alternatives—relative to the choice of the linguistic framework. The space then can be divided into an attribute partition in any way, the flat prior over the points translating into different assignments of probabilities for different predicates.

The η rule also inspires thinking of observations as points of the space. The attribute space provides us with the respects in which the future will resemble the past; the modality is the inductively salient respect, and the space itself gives us an interpretation of what “similar” will mean in this context.⁵ We have the notion of distance—and hence similarity—not only between predicates in

the space, but primarily between the points. If the outcomes so far were concentrated in one region of the space, then it seems reasonable that also in the future observations that kind of concentration will occur. Therefore we will expect future outcomes to appear close to the region into which the previous observations tended to fall as well, regardless of where the borders between predicates lie (Carnap, 1980, 40). In Section 7 the idea that evidence should be thought of in terms of points in the space rather than just atomic sentences will be revisited.

5. Attribute spaces as a part of an explicatum

After investigating the use that Carnap puts attribute spaces to, we can ask more generally whether the addition of the spaces brings his inductive logic project closer to its goals. Some authors, most notably Hill and Paris (2013, 1296), criticize Carnap, stating that the addition of attribute spaces meant introducing nonlogical elements into inductive logic, therefore abandoning the project of developing pure inductive logic based on an uninterpreted language. However, instead of thinking of inductive logic as an attempt to establish the purely logical theory of confirmation, in this section I focus on looking at the Basic System as proposing an explication of the concept of confirmation and on how the addition of attribute spaces influenced that project.

Carnapian inductive logic was from its start a project in explication (Carnap, 1950, §1). The main target of the explication was the concept of confirmation, and through clarifying that concept also related concepts, like probability and inductive reasoning, were to be analyzed. It seems fitting then to ask the following question: if and in what ways did the addition of attribute spaces to the system of inductive logic accord with the goals that inductive logic was conceived of for itself? The talk about the increase in the success of Carnap's inductive logic project can be a red flag to some, since Carnapian inductive logic had been subject to a lot of criticism aimed at its very core, voiced for instance by Lakatos (1968). Here is not the place to get involved in this discussion; the aim of this section is much more modest. I am not trying to say whether Carnapian inductive logic was successful as a whole; right now I am only evaluating what role attribute spaces can be seen to play in the explication that inductive logic aims at.

To this end, we can evoke Carnap's own four criteria for evaluating explications, i.e. similarity, exactness, fruitfulness, and simplicity (Carnap, 1950, §3). The explicandum in this case is the “informal”, “everyday” concept of confirmation (as used by scientists rather than laymen). The explicatum—the formal concept of confirmation—is built using a whole range of notions, like linguistic framework or models for a formal language. And in that sense attribute spaces are part of the explication of confirmation. The Basic System tells us that confirmation is a relation between propositions spelled out within a linguistic framework; and that the meaning relations between the terms in the framework can be represented geometrically as attribute spaces.

To start with, consider the desideratum of simplicity. At the first glance one might think that it was negatively affected by the introduction of attribute spaces into the system: even more confirmation functions are now deemed rational by the system than it was the case with the Continuum, which required an addition of more parameters. Moreover, the notion of a linguistic framework was made more complicated by the addition of attribute spaces. Nevertheless, there is an obvious trade-off with the positive impact of the changes to the system on the other desiderata, and the inevitability of such trade-offs had been already acknowledged by Carnap himself (1950, 7). Given what the additions to the system are achieving in terms of making it more precise

⁴ It needs to be stressed here that Carnap's proposal was not the final word on the topic of analogy by similarity of predicates. Among others, Kuipers (1984) explores the idea even further, also using the idea of distances between attributes. Moreover, he proposes a more general principle of analogy, diverging from Carnap in some crucial aspects. The study of the relations between Carnap's and Kuiper's systems, while lying outside the scope of this paper, remains a worthwhile project.

⁵ Notice, however, that Carnap doesn't say much in the Basic System about how to make a decision which respects or similarities should be taken as inductively relevant. How pressing this issue is was of course pointed out by Goodman (1955, 1972).

and more general, they are actually a surprisingly simple means to do so.

Although Carnap (1963, 936) mentions precision as possibly the most important of the requirements for explication, Justus (2012, 168) argues that the importance of precision lies in the fact that it tends to increase fruitfulness. Hence in cases when introducing more precise concepts doesn't increase fruitfulness, the requirement of precision can be relaxed. Carnap's mentioning how a precise, numerical statement of the η rule would be undesirable is a good example of such a possibility of tension between precision and fruitfulness.

A fruitful explicatum is one that facilitates the formulation of "universal statements" (Carnap, 1950, 7), i.e. empirical laws or (in case of purely formal explicata) logical laws. Does the addition of attribute spaces increase the Basic System's power to formulate general statements about confirmation? Yes, since the γ and η rules are a generalization of sorts, spelling out a universal connection between meaning relations among predicates and features of the confirmation functions. Moreover, the addition of attribute spaces opens inductive logic up into a number of new themes and applications, making possible the formulation of general principles concerning extra concepts like analogy. Attribute spaces are also a more principled framework for meaning postulates than simple sentences of the object and metalanguage of inductive logic.

The desideratum of exactness demands the explicatum to be formulated in such a way that it fits "into a well-connected system of scientific concepts" (Carnap, 1950, 7). Outside of purely formal sciences, this requirement does not mean fitting into an axiomatic system, but rather as connecting to other established scientific theories, as well as increasing the precision of the target concept (Justus, 2012, 165). Speaking most generally, explication is understood as replacing an imprecise concept with a precise one. How the attribute spaces can be seen to increase the precision of the Basic System is most immediately seen in the case of the γ rule. The γ rule specifies exactly how to choose a priori probabilities for each predicate. The connection between the meaning of terms used in inductive reasoning and the confirmation of (at least some of) the statements of the language has been made more precise, while the scope of the explicatum is now also broader because of the acceptance of uneven γ values for different predicates.

When it comes to connecting the explicatum to a broader network of scientific concepts, it turns out that attribute spaces serve inductive logic well again. If there was a well-established scientific theory that employed structures similar to attribute spaces and, of course, employed them to the same purpose, i.e. modeling the meaning of terms, then the exactness of the Basic System could be said to have increased through the addition of the spaces. And in fact the theory of conceptual spaces, which we will turn to shortly, does provide a model of meaning in terms of geometric representation of concepts.

Finally, the question of how similar the explicatum is to the concept being explicated, that is the concept of confirmation proper, and the auxiliary notions like linguistic framework, or, most importantly, attribute spaces. When it comes to confirmation proper, the introduction of the attribute spaces into the system does work towards making the explicatum more similar to the "everyday" concept of confirmation. Attribute spaces provide an exact way to measure relevance between predicates, therefore enabling the Basic System to approach the problem of analogical reasoning. This means that the new explicatum can cover more cases in which the explicandum might have been used; the old explicatum (i.e. the one in the Foundations) deemed statements concerning confirmation via analogy as improper. This means that the similarity of the new explicatum to the "informal" concept of

confirmation is greater than the one of the explicatum that can be found in Logical Foundations or even the Continuum. A similar thing can be said of the introduction of the γ rule and the idea that the initial probabilities can be different for different predicates. The addition of this rule allows the system to cover cases in which scientists have reasons to start with uneven initial probabilities, for instance as a result of earlier studies.

The similarity of linguistic frameworks in the Basic System to the thing they are supposed to represent could be lower if the resulting concept of confirmation was very satisfactory. This is because similarity is not a very tight desideratum for confirmation; if precision and fruitfulness of the explicatum less similar to the explicandum are higher, then the less similar explicatum can be preferred. That comes from the fact that the explicandum is always a vague or imprecise concept, and sometimes also an incoherent one, and a very tightly fitting explicatum might inherit those undesirable qualities.

However, if there was a proof or a good reason to think that linguistic frameworks of the Basic System were a particularly good model of the conceptual structures used by actual scientists, then it would be an argument in favor of greater similarity of the concepts of the Basic System with what they are supposed to explicate. It would be the case even more so if we found that the way that the attribute spaces predict the concepts used by agents relate to each other is actually a good representation of how they relate in real life.

Not only similarity would be influenced then; we have noted above that other desiderata would also be positively affected if there was a reason to believe that actual linguistic frameworks look somewhat similar to what Carnap proposes, or that human concepts are structured somewhat like attribute spaces, and if there was a well-established theory where notions like attribute spaces were already put to use. And as a matter of fact, we do have such a reason, provided by recent developments in cognitive science, under the theory of conceptual spaces.

6. Conceptual spaces

Conceptual spaces are a tool for modeling cognitive representations advocated in recent years by Peter Gärdenfors (2004, 2014). As we have seen in the previous section, there is a way in which the existence of an empirically successful theory of concepts that employs structures similar to attribute spaces can be seen as relevant to the normative, explicative project of Carnap. Attribute spaces connect inductive logic to a broad and empirically driven enterprise of modeling concepts geometrically, a connection which strengthens Carnap's own explication.

On the face of it the Carnapian attribute spaces and Gärdenfors' conceptual spaces are the same kinds of entities: geometrical structures consisting of a set of points organized along a number of dimensions, meant as a representation of concepts. However, the details of the proposals differ, as well as the use that is made of the structures. Some differences result from different goals that the projects pursue. Others stem from the fact that the theory of conceptual spaces is more developed than the attribute spaces are within the Basic System. This section reviews the similarities and differences, and together with the next section offers ways in which interaction with the theory of conceptual spaces can help strengthen and advance the Basic System.

Conceptual spaces are meant to provide a middle-level model for cognitive representations, lying in between the symbolic and the subconceptual approaches. The approach is motivated by the failures of the symbolic and associationist models for representing information when it comes to explaining phenomena like concept formation and acquisition. Gärdenfors' goal is to provide a platform

for describing cognitive representations: “The epistemological role of the theory of conceptual spaces (...) is to serve as a tool in modeling various *relations* among our experiences, that is, what we perceive, remember, or imagine” (Gärdenfors, 2004, 5). Hence one of the uses of conceptual spaces is to form empirically testable predictions regarding (not only) human cognition. They are also to be used as a tool for constructing cognitive representations of artificial agents (2004, 30). These two goals correspond to the two goals of cognitive science in general, that is description and modeling of existing cognitive agents, and constructing artificial ones.

While for Carnap an attribute space is primarily a simple collection of points divided into an attribute partition, conceptual spaces are primarily collections of domains—that is sets of dimensions that relate to each other in a particular way: “A conceptual space is defined here as a set of quality dimensions with a geometrical structure” (2004, 24). It seems more important that a space spans over a number of dimensions, and what dimensions those are, than the points it is built from. Carnapian attribute spaces can have any number of dimensions and, unlike the cognitive scientists, Carnap is not concerned with the issue of whether the dimensions of attribute spaces have a clear interpretation, although for all of the examples he discusses it is in fact the case that the dimensions can be interpreted in a natural way. For Gärdenfors the dimensions correspond to “the different ways stimuli are judged to be similar or different” (2004, 6), which is what was earlier remarked about the Carnapian modalities as well. However, there is a slight difference here. For Carnap a modality can be assigned an attribute space of many dimensions (recall the three-dimensional color space), and so modalities and dimensions are not exactly the same. So in that sense Carnapian modalities are a little more general than the dimensions of cognitive science.

Recall that the attribute spaces were divided into regions by attribute partitions in such a way that every point belonged to a region corresponding to a predicate of the relevant family. Conceptual spaces are not supposed to have such a tight connection with the language used to refer to the concepts represented in them—and quite naturally, since they have been proposed as an alternative to a symbolic, i.e. linguistic, approach to concept representation. The dimensions and relations between them are what constitutes a space, and concepts do not necessarily have to fully partition the spaces: Carnap being open to artificial language construction in science, his assumption might be fitting as a requirement for scientific concepts, but not necessarily as a model for natural language predicates. For instance, there can be regions in the color space that don’t fall under any, say, English common color term.

The regions in conceptual spaces that correspond to concepts are usually considered to be nonhomogeneous, i.e. some points are assumed to be more prototypical within a region; that insight was brought in by studies in cognitive linguistics (Gärdenfors, 2004, 3.8, 3.9). The division of the space into regions is then determined by the arrangement of such prototypical points—every point in the space is assigned the region corresponding to the prototype that is the closest, and hence most similar, to it (this way of dividing a space is known as a Voronoi tessellation). The idea of prototypical points can be developed even further, into prototypical regions, in order to account for vagueness of predicates (Douven, Decock, Dietz, & Egré, 2013). Concepts in the Carnapian attribute spaces did not have any of such extra structure to them; an attribute partition simply divides the space into sharp regions.

The tight connection between similarity and distances is assumed both in the attribute and the conceptual spaces: “the similarity of two objects can be defined via the distance between their representing points in a conceptual space” (Gärdenfors, 2004,

44). However, Gärdenfors (2004, 110) stresses that similarity measurements should always be understood as relative to a particular cognitive model. There are no similarity judgments that would be absolute in the sense of not being formed assuming a particular (model of) conceptual representation. This resonates well with Carnap’s general stance on linguistic frameworks—even if not explicitly discussed in connection with attribute spaces in the Basic System—and what kinds of questions can be asked only within a framework rather than outside of it.

At the time when Carnap was introducing them, the attribute spaces might have seemed a bit ad hoc, in spite of the earlier use of them by van Fraassen. The recent developments in cognitive science provide us with an argument that as far as the explicative project of Carnapian inductive logic goes, attribute spaces do make a positive contribution within it. In spite of the differences between conceptual and attribute spaces, there is a clear correspondence which allows us to draw the earlier conclusion about the explicative success of the Basic System brought about by the inclusion of attribute spaces. However, when establishing this and further connections, we have to be careful when stating our goals and we have to make sure to not collapse the normative and descriptive claims of the two theories. And it is the elaboration of these further connections that I will now turn to.

7. Conclusion: building the bridges

In spite of the Basic System being well-read and often referred to, the attribute spaces themselves have not so far sparked any substantial interest within the inductive logic community. With the connection to the developments in cognitive science, however, a number of new issues worth investigating arise. They can be roughly categorized under two leading questions. First, how would inductive logic have to be modified if we were to plug in conceptual spaces where attribute spaces were in the Basic System. Which properties of conceptual spaces can be useful from the point of view of inductive logic? The second big question is how to conceive of inductive reasoning on the level of conceptual spaces (see (Gärdenfors, 1994) for a statement of related issues). In spite of their apparent similarity, these two questions differ in the direction from which they are looking at the problem; whether they approach conceptual spaces from the point of view of inductive logic, or vice versa. And because inductive logic and the theory of conceptual spaces have very different aims, the resulting questions differ.

Carnapian inductive logic is a normative project; it is not concerned primarily with how people *do* reason inductively, but with how they *should* do it in order to avoid mistakes and further the sciences. From such a perspective a perfect match of the way linguistic frameworks are represented in the Basic System with what we take humans to actually be doing in their daily reasoning is not an absolute must; especially given that Carnap’s original inspiration was scientific, as opposed to everyday, reasoning. However, some descriptive accuracy is needed in order to claim that the system proposed can serve as a guide for inductive reasoners out there (Romeijn, 2005, 40). And in that sense employing conceptual spaces in inductive logic can eventually be worthwhile. For instance, consider the case of vague concepts. In Section 5, when discussing the requirement of fruitfulness, we mentioned the possibility of situations in which introducing more precise concepts does not facilitate the formulation of laws. Therefore there might be good reasons to include a non-sharp concept in our linguistic framework, and the formal model of vague concepts developed within the theory of conceptual spaces might prove useful in setting up attribute spaces for such concepts. Developing the idea of applying confirmation functions to vague concepts could, moreover, prove useful to the independently motivated project of

representing credences using imprecise probabilities (see e.g. Joyce (2011); Bradley & Steele (2014)).

The ambition of the theory of conceptual spaces, on the other hand, is primarily descriptive; it aims to create a model of cognitive representations that will correctly predict known features of (human) reasoning such as categorization or the emergence of concepts. Interpreted in this latter setting, Carnap's postulates such as the γ and η rules, which each prescribe a particular connection between rational expectation and the structure of the attribute space, would turn into empirical statements. Hence, one may ask: is it really the case that human agents match their initial expectations for observing an object to how large the relevant region is in the space?—a question that Carnapian inductive logic would not raise. Finally, when we focus on the constructive goal of cognitive science and think of designing artificial agents who are able to reason inductively in a correct way while using conceptual spaces to represent concepts, the principles they would operate on could come from inductive logic, since here the question of correctness is also applicable.

One way of making changes in inductive logic through taking attribute spaces more seriously is inspired by the discussion in Section 4 of the way in which the γ and η rules can be thought of as deriving from the assumption that observations can be thought of as points in the conceptual space. That this idea is not very far from the way Carnap thought about the attribute spaces is moreover proven by his own example of the η function and the corresponding η curve (Carnap, 1980, §17). In the example Carnap talks about how the initial credence corresponds to a flat density over the points of the space, and discusses how making one observation changes that density into a normal distribution, which gives rise to the desired analogical effects. The mean of that normal distribution can be thought of as corresponding to the exact quality that was observed. However, Carnap did not take that idea any further, i.e. into a deliberation how the probability density over the attribute space would change with more observations, and it still remains an open question.

Thinking of observations as points in the attribute space opens up even more connections to existing theories, bringing us back to the considerations of fruitfulness of the Basic System. In particular, there is a connection to an observation that Ferguson (1974) makes in a paper on Dirichlet processes that “One would like to have a prior distribution for P with the property that if X is a sample from P and $X = x$, then the posterior guess at P gives more weight to values close to x than the prior guess at P does” (Ferguson, 1974, 622). He suggests that in a setting with real-valued outcomes some sort of analogical reasoning would be in order, which is very similar to the Carnapian intuition that in attribute spaces—which tend to be continuous, and do have the notion of closeness defined on them—the situation is the same. This connection suggests immediately that the problem of analogical reasoning on attribute spaces might be closely connected with doing inductive logic with value continuum. The latter has been approached within Bayesian statistics by the use of Dirichlet processes, and within inductive logic with the use of Blackwell-McQueen prediction rules (see the discussion by Skyrms (1996)). However, including analogical reasoning within the setting with a continuum of possible outcomes has not been discussed yet, although a paper by Antoniuk (1974) might serve as a relevant point of departure here.

Developing the idea that the confirmation functions could be applied directly to the attribute space rather than to the algebra of propositions would serve two purposes, corresponding to the two questions in approaching the connections between the two theories that were mentioned earlier in this section. If we approach the issue from the side of inductive logic, we can see such a development as a step towards possibly including more kinds of inductive reasoning in our formal model. On the other hand, it could serve as

a first step towards a theory of inductive reasoning for geometric, conceptual representations: the questions raised by Gärdenfors (1994) about how to conceive of and represent inductive reasoning on the level of conceptual spaces, as opposed to symbolic or connectionist models of representation, are still open.

It was mentioned in the previous section that while regions in Carnapian attribute spaces are homogeneous, conceptual spaces (at least for certain kinds of modalities) have been proposed as structured around a set of prototypical points or regions. Such an assumption is very well supported by studies on human categorization (Lakoff, 1987), but it is not entirely obvious whether we should take this as a reason to jump right into reformulating the Basic System to cover this insight.

Again, our course of action will depend on the kind of project that we are involved in. Inductive logic is ultimately interested in conceptual spaces that scientific theories are based on, not necessarily the ones that best approximate human cognition. And for the most part scientifically useful concepts will be sharp and defined with the use of necessary and sufficient conditions rather than by the similarity to a prototype (however, when vagueness is essential to a scientific concept, conceptual spaces with prototypes could still be useful). But when we think of representing inductive reasoning on conceptual spaces, with the aim of describing actual human reasoning, including prototypes in our formal model of induction might be desirable.

Finally, there is what seems to be the biggest open question, at least on the side of inductive logic. Even though inductive logic was meant to ultimately be a tool for scientists, the Basic System is concerned only with an observational language. Inclusion of theoretical concepts in inductive logic still remains to be done. There is some recent work on representing scientific theories and the frameworks in which they are spelled out in conceptual spaces (Gärdenfors & Zenker, 2013; Zenker & Gärdenfors, 2014a, 2014b). However, we know nothing so far about whether and how these insights can be relevant to inductive logic.

Acknowledgments

This research was supported by the Alexander von Humboldt Foundation. I would like to thank Hannes Leitgeb, Jan-Willem Romeijn, Frank Zenker, and the anonymous referees for comments on the various versions of this paper.

References

- Antoniuk, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The Annals of Statistics*, 2(6), 1152–1174. Retrieved from <http://www.jstor.org/stable/2958336>.
- Awodey, S., & Klein, C. (Eds.). (2004). *Carnap brought home: The view from Jena*. Chicago and La Salle: Open Court.
- Bradley, S., & Steele, K. (2014). Should subjective probabilities be sharp? *Episteme*, 11, 277–289. <http://dx.doi.org/10.1017/epi.2014.8>.
- Carnap, R. (1950). *Logical foundations of probability*. Univ. of Chicago Press.
- Carnap, R. (1952). *The continuum of inductive methods*. Univ. of Chicago Press.
- Carnap, R. (1963). Replies and systematic expositions. In P. A. Schilpp (Ed.), *The philosophy of Rudolf Carnap* (pp. 859–1012). La Salle IL: Open Court.
- Carnap, R. (1967). *The logical structure of the world*. Berkeley and Los Angeles: Univ. of California Press.
- Carnap, R. (1971a). A basic system of inductive logic, Part I. In R. Carnap, & R. C. Jeffrey (Eds.), *Studies in inductive logic and probability* (Vol. I, pp. 33–165) University of California Press.
- Carnap, R. (1971b). Inductive logic and rational decisions. In R. Carnap, & R. C. Jeffrey (Eds.), *Studies in inductive logic and probability* (Vol. I, pp. 5–31) University of California Press.
- Carnap, R. (1980). A basic system of inductive logic, Part II. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability* (Vol. II, pp. 7–155) Univ. of California Press.
- Carnap, R., & Jeffrey, R. C. (Eds.). (1971). *Studies in inductive logic and probability* (Vol. I) University of California Press.
- Carus, A. W. (2007). *Carnap and twentieth-century thought. Explication as enlightenment*. Cambridge: Cambridge University Press.

- Costantini, D. (1983). Analogy by similarity. *Erkenntnis*, 20, 103–114. <http://dx.doi.org/10.1007/BF00166497>.
- Douven, I., Decock, L., Dietz, R., & Egré, P. (2013). Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic*, 42(1), 137–160. <http://dx.doi.org/10.1007/s10992-011-9216-0>.
- Ferguson, T. S. (1974). Prior distributions on spaces of probability measures. *The Annals of Statistics*, 2(4), 615–629. Retrieved from <http://www.jstor.org/stable/2958401>.
- Friedman, M. (2001). *Dynamics of reason*. CSLI Publications.
- Gärdenfors, P. (1994). Three levels of inductive inference. In D. Prawitz, B. Skyrms, & D. Westerståhl (Eds.), *Logic, methodology and philosophy of science IX* (pp. 427–449). Amsterdam: Elsevier Science.
- Gärdenfors, P. (2004). *Conceptual spaces: The geometry of thought*. MIT Press.
- Gärdenfors, P. (2014). *Geometry of meaning: Semantics based on conceptual spaces*. MIT Press.
- Gärdenfors, P., & Zenker, F. (2013). Theory change as dimensional change: Conceptual spaces applied to the dynamics of empirical theories. *Synthese*, 190, 1039–1058. <http://dx.doi.org/10.1007/s11229-011-0060-0>.
- Gillies, D. (2000). *Philosophical theories of probability*. London and New York: Routledge.
- Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.
- Goodman, N. (1972). Seven strictures on similarity. In *Problems and projects* (pp. 437–446). Indianapolis and New York: The Bobbs-Merrill Company, Inc.
- Hill, A., & Paris, J. B. (2013). An analogy principle in inductive logic. *Annals of Pure and Applied Logic*, 164, 1293–1321.
- Hilpinen, R. (1973). Carnap's new system of inductive logic. *Synthese*, 25(3/4), 307–333. <http://dx.doi.org/10.1007/BF00499684>.
- Huttegger, S. (2015). *Analogical predictive probabilities* (manuscript available from the author).
- Joyce, J. (2011). A defense of imprecise credence in inference and decision. *Philosophical Perspectives*, 24, 281–323. <http://dx.doi.org/10.1111/j.1520-8583.2010.00194.x>.
- Justus, J. (2012). Carnap on concept determination: Methodology for philosophy of science. *European Journal of Philosophy of Science*, 2, 161–179. <http://dx.doi.org/10.1007/s13194-011-0027-5>.
- Kemeny, J. G. (1955). Fair bets and inductive probabilities. *The Journal of Symbolic Logic*, 20(3), 263–273. <http://dx.doi.org/10.2307/2268222>.
- Keynes, J. M. (1921). *A treatise on probability*. Macmillan, 1963.
- Kuipers, T. A. F. (1984). Two types of inductive analogy by similarity. *Erkenntnis*, 21, 63–87. <http://dx.doi.org/10.1007/BF00176183>.
- Lakatos, I. (1968). Changes in the problem of inductive logic. In I. Lakatos (Ed.), *The problem of inductive logic* (pp. 315–417). North-Holland Publishing Company.
- Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago, IL: The University of Chicago Press.
- Langacker, R. W. (1987). *Foundations of cognitive grammar* (Vol. 1) Stanford, CA: Stanford Univ. Press.
- Leitgeb, H. (2011). Logic in general philosophy of science: Old things and new things. *Synthese*, 179, 339–350. <http://dx.doi.org/10.1007/s11229-010-9776-5>.
- Niiniluoto, I. (1981). Analogy and inductive logic. *Erkenntnis*, 16, 1–34. <http://dx.doi.org/10.1007/BF00219640>.
- Niiniluoto, I. (1988). Analogy and similarity in scientific reasoning. In D. H. Helman (Ed.), *Analogical reasoning* (pp. 271–298). Kluwer.
- Romeijn, J. W. (2005). *Bayesian inductive logic*. Ph.D. thesis. Rijksuniversiteit Groningen.
- Shogenji, T. (1999). Is coherence truth conducive? *Analysis*, 59(4), 338–345. <http://dx.doi.org/10.1093/analys/59.4.338>.
- Skyrms, B. (1996). *Carnapian inductive logic and Bayesian statistics*. Lecture Notes–Monograph Series 30: Statistics, Probability and Game Theory (pp. 321–336).
- Stalnaker, R. (1981). Anti-essentialism. *Midwest Studies of Philosophy*, (4), 343–355.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327–352. <http://dx.doi.org/10.1037/0033-295X.84.4.327>.
- Van Fraassen, B. C. (1967). Meaning relations among predicates. *Noûs*, 1(2), 161–179. Retrieved from <http://www.jstor.org/stable/2214583>.
- Zabell, S. (2011). Carnap and the logic of inductive inference. In D. M. Gabbay, S. Hartmann, & J. Woods (Eds.), *Inductive logic: Vol. 10. Handbook of the history of logic* (pp. 265–309) Elsevier.
- Zenker, F., & Gärdenfors, P. (2014a). Communication, rationality, and conceptual changes in scientific theories. In F. Zenker, & P. Gärdenfors (Eds.), *Synthese library: Vol. 359. Applications of conceptual spaces. The case for geometric knowledge representation* (pp. 259–277) Springer.
- Zenker, F., & Gärdenfors, P. (2014b). Modeling diachronic changes in structuralism and in conceptual spaces. *Erkenntnis*, 79(8), 1547–1561. <http://dx.doi.org/10.1007/s10670-013-9582-9>.
- Zenker, F., & Gärdenfors, P. (Eds.). (2015). *Synthese library: Vol. 359. Applications of conceptual spaces. The case for geometric knowledge representation* Springer.