Found Sci (2016) 21:225–236
DOI 10.1007/s10699-015-9416-0

Belief Systems and Partial Spaces

Otávio Bueno

Published online: 22 January 2015

Abstract One important role of belief systems is to allow us to represent information about a certain domain of inquiry. This paper presents a formal framework to accommodate such information representation. Three cognitive models to represent information are discussed: conceptual spaces (Gärdenfors in *Conceptual spaces: the geometry of thought*. MIT Press, Cambridge, 2000), state-spaces (van Fraassen in *Quantum mechanics: an empiricist view*. Clarendon Press, Oxford, 1991), and the problem spaces familiar from artificial intelligence. After indicating their weakness to deal with partial information, it is argued that an alternative, formulated in terms of partial structures (da Costa and French in *Science and partial truth*. Oxford University Press, New York, 2003), can be provided which not only captures the positive features of these models, but also accommodates the partiality of information ubiquitous in science and mathematics.

Keywords Problem space · Conceptual space · State-space · Partial space · Belief system

1 Introduction

Belief systems are significant in the information they allow us to represent. Several techniques have been devised, in cognitive science and artificial intelligence, to accommodate this issue, ranging from the use of predicate logic and production rules to semantic nets and frames (see, for instance, Rich and Knight 1991, pp. 103–304). This is also, of course, a crucial problem in the philosophy of science. In this context, one addresses the problem by examining the structure of scientific theories, in which such theories encode certain beliefs and accepted information about the world. At this level we find the various formulations of the semantic conception of science, according to which empirical information is conveyed by specifying a family of structures, the models of the theory, and relating these models to appropriate structures which represent empirical data, the models of phenomena (see van Fraassen 1989, 1991).

O. Bueno (✉)
Department of Philosophy, University of Miami, Coral Gables, FL 33124, USA
e-mail: otaviobueno@mac.com

The problem of information representation is of course important for several reasons. Information representation allows us to articulate heuristic strategies of problem solving. Scientists and other cognitive agents are always searching for better ways of solving problems; but in order for heuristic principles to be developed and applied, we need first to have a proper representation of the information about which we are trying to obtain new results. In other words, given the importance of heuristics, it is not surprising that information representation becomes so relevant. It is the first step to be taken before any serious work in heuristics can be developed.

A significant issue in this context concerns the *partiality of information* we usually have about a given domain of inquiry. If we take seriously the actual epistemic situation of cognitive agents, we soon realize how extraordinarily idealized the extant models of information representation are: they usually assume that we have full, complete information about the domain of knowledge we are investigating. But this requirement can hardly be met in practice: partial, incomplete information is the norm, rather than the exception of scientific research (da Costa and French 2003).

In order to accommodate this partiality, what is required is a more open-ended framework in which central features underlying the construction of the extant models of information representation can be generalized. I will argue that the partial structures approach provides a suitable framework for this task (see da Costa and French 1989, 1990, 1993, 2003). As we shall see, this proposal introduces two new components in the debate: a broader notion of structure (partial structure), which is adequate to accommodate the partiality of information, and a weaker notion of truth (quasi-truth), which generalizes the standard Tarskian formulation of truth to contexts involving partial information. In my view, by exploring the resources of this proposal, a new approach to information representation and, in particular, to heuristics can be devised. Moreover, this approach is better equipped to accommodate our actual epistemic situation than the extant views. The aim of the present paper is to argue for these claims.

2 Conceptual Spaces, State-Spaces and Problem Spaces: Information Representation and Heuristics

I start by considering three formal frameworks to represent information about the world: Peter Gärdenfors's *conceptual spaces* (see his 1990), Bas van Fraassen's *state-spaces* (see his 1989 and 1991), and the *problem spaces* familiar from artificial intelligence (see Rich and Knight 1991). There are striking similarities between them, although the differences are also telling. All of them provide means of representing information in *structural* terms, that is, in terms of configurations in the relevant space: whether it is a space of concepts (Gärdenfors), a space of states of a physical system (van Fraassen), or of solutions of a given problem (Rich and Knight). In each case, *geometrical* properties of the accompanying spaces are used as a source of information about the particular domain of inquiry. As we shall see, Gärdenfors stresses the importance of convex sets; van Fraassen, following a long-standing tradition in the foundations of physics, emphasizes the role of certain subspaces; and depending on the particular problem under consideration, the problem space provides information about the acceptable trajectories (those which lead to the solution of the problem). Although the *topics* dealt by each of these proposals are different (concepts, states of physical systems, and solutions to certain problems), they provide remarkably similar strategies to accommodate information. In this respect, they can be grouped together as cognitive models, as belief systems, broadly understood. Moreover, as we shall see, clear relationships can be established between them.

Gärdenfors's conceptual spaces have been devised as a way of representing *conceptual information* about a certain domain (Gärdenfors 1990, 2000). This is done in terms of certain *quality dimensions*, which are meant to represent the qualities of the objects in this domain. For instance, color, weight, temperature, mass, time, and the three dimensions of ordinary space are all examples of quality dimensions in Gärdenfors's sense. Each dimension is endowed with a certain geometrical structure, which provides the resources to represent the relevant objects: ordinary space is often conceived of as a three-dimensional Euclidean space, and weight dimension is measured on a scale isomorphic to the positive real numbers. More formally, according to Gärdenfors (1990, p. 85, and 2000, Chapter 1), a conceptual space S is a class $D_1, \ldots, D_n$ of quality dimensions. Moreover, a point in a conceptual space S is represented by a vector $s = \langle d_1, \ldots, d_n \rangle$ with one index for each dimension. In this way, an object receives a complete description of its qualities by being assigned to a point in the space. In this way, its color, spatial position, temperature, weight etc. are represented.

Crucial about this representation are the geometrical configurations formed in the conceptual space, for in terms of them the notion of a property can be characterized. According to Gärdenfors, a property can be defined as a *region* of the conceptual space (1990, pp. 87–88, and 2000, Chapter 3). For instance, in the time dimension, the point that represents 'now' divides the space into two regions, corresponding to the properties 'past' and 'future'. Moreover, if we consider the geometrical features of such regions, we can distinguish natural from unnatural properties. A property is *natural*, in Gärdenfors's view, if the region H it describes in the conceptual space is *convex*, that is, for every pair s_1 and s_2 of points in H, all points between s_1 and s_2 are also in H (*ibid*., p. 88). In other words, a convex region is closed in a uniform way, which suggests that the property it represents is natural (in some sense).

This is a fairly weak characterization of natural property, since how a convex region will look like depends on the particular conceptual space one considers (in particular, it depends on how the notion of 'between' has been spelled out). Despite this, Gärdenfors claims that this characterization is strong enough to provide an account of induction. After all, it allows us to identify pathological, unnatural properties that cannot be taken as the basis for any inductive generalization. For example, a property like 'grue' (green if observed before the year 3000, and blue if observed afterwards) does not determine a convex region in the two-dimensional conceptual space formed by the color spectrum and time, since it has a discontinuity in the year 3000. Therefore, because it is not a natural property, no inductive inferences should be made in terms of it. As a result, in Gärdenfors' view, since conceptual spaces provide a clear way of representing information, we can examine the problem of induction in a new and more fruitful way.

Moving now to van Fraassen's proposal, the issue of how to represent empirical information in science is similarly crucial. On his view, this is one of the main tasks of a scientific theory. Following the semantic approach, to present a theory is to specify a class of structures, its models (van Fraassen 1989, 1991, 2008). These models represent, in particular, the *states* that a certain physical system is assumed to have. The states are characterized by physical magnitudes (observables) that pertain to the system and can take certain *values* (van Fraassen 1991, p. 26). In classical mechanics, for instance, the history of a system (its evolution in time) can be represented as a trajectory in the space of possible states of the system (its state-space). Such a trajectory is a map from time to the state-space. As an example, consider a classical mechanical system (see Varadarajan 1968). Its states can be completely specified by a $2n$-tuple $(x_1, \ldots, x_n, p_1, \ldots, p_n)$, where $(x_1, \ldots, x_n)$ represents the configuration and $(p_1, \ldots, p_n)$ the momentum vector of the system in a given time instant. Thus, the possible states of the system are represented by points of the open set $S^n \times S^n$ (where S^n is the n-dimensional space of n-tuples of real numbers). The system's evolution is then determined by its Hamiltonian. In other words, if the state of the system at the time t is represented by $(x_1(t), \ldots, x_n(t), p_1(t), \ldots, p_n(t))$, the functions $x_i(t)$, $p_i(t)$ (with $1 \leq i \leq n$) satisfy the equations:

$$dx_i/dt = \partial H/\partial p_i, \text{ and } dp_i/dt = -\partial H/\partial x_i (\text{with } 1 \leq i \leq n). \quad (1)$$

The set of all possible states of the system satisfying (1) gives its state-space.

This representation provides an immediate way of studying logical relations between statements about the system's states. These statements, whose truth conditions are represented in the models of the system, are called *elementary statements* (1991, pp. 29–30). As an example, let A be the statement 'the system X has state s at time t'. Now, A is true exactly if the state of X at t is (represented by a point) in the region S_A of its state-space. In other words, the state-space provides information about the physical system X. The crucial point for van Fraassen is that the family of statements about X receives some structure from the geometric character of the state-space. For instance, in quantum mechanics, if A attributes a pure state to X, the region S_A it determines in the state-space is a *subspace*. In this way, by investigating the structure of the state-space, one investigates the *logic* of elementary statements (van Fraassen 1991, p. 30). Of particular interest are of course the following logical relationships: (i) the elementary statement A *implies* the elementary statement B if S_A is a subset of S_B; (ii) C is the conjunction of A and B if S_C is the g.l.b. (greatest lower bound), with respect to the implication relation, of S_A and S_B.

So, just as conceptual spaces represent information about the properties of certain objects, state-spaces are adequate to study the underlying logic of statements about physical systems. In particular, we can move from conceptual spaces to state-spaces by focusing on a particular kind of quality dimensions: those that are required to describe the states of a physical system (Gärdenfors 1990, pp. 85–86).

But what is the *point* of these representations? In a nutshell, it is to solve problems. This is quite clear in the state-space case, given, in particular, the heuristic role of *symmetry* in theory construction (see van Fraassen 1989, 1991, pp. 21–26). The idea is that in order to solve a problem we have first to devise an adequate *representation* of it: a model of the situation described by the problem. This model introduces certain variables, and what we are looking for is a rule, a function, from certain *inputs* (the data of the problem) to one, and usually only one, *output* (the problem's solution). The representation of the problem brings certain symmetries, which are crucial for its solution. Now, it often happens that we may not know how to solve a problem P_1 (stated in terms of the model we initially devised), but we may well know how to solve a *related* problem P_2. One way of solving P_1 is then to show that it is essentially P_2. This requires that we provide a transformation (let's call it T) from P_1 into P_2 that leaves P_1 essentially the same. What this means is that the structure that characterizes P_1 (the relationships between its variables) is preserved by T. The underlying heuristic move is then clear: the same problems have the same solutions (van Fraassen 1991, p. 25). The symmetries of the problem, being preserved by T, allow this heuristic move to get off the ground.

These considerations can then be straightforwardly represented in terms of problem spaces (see Rich and Knight 1991). Similarly to the two kinds of spaces we have just discussed, problem spaces are also devised to represent information—the one relevant to the solution of a given problem. The model of the problem provides a configuration of *acceptable states* (in a chess game, for instance, positions of the pieces on the board), and it includes several *trajectories* (legal chess moves), which are devised in such a way that the *goal state* is reached (the board position in which the opponent does not have any legal move and his or her king is being attacked). The solution to the problem is a trajectory leading from the initial state to the goal state. In those cases in which we do not know this trajectory, we may try to find a similar problem (represented by a similar problem space), whose winning trajectory we know, and then provide a transformation from the former into the latter preserving the relevant symmetries.

It should now be clear that these three representations have several features in common. In particular, state-spaces provide one of the ways of representing problem spaces (by specifying the states involved in the solution to a problem). Conceptual spaces, in turn, when used to represent a particular problem, can also be taken as problem spaces. In this sense, what distinguishes these representations is a *pragmatic* issue: the uses each of them are put to.

But there is a salient feature that is also shared by all these representations. They all assume that the information to be represented (in a conceptual space, a state-space or a problem space) is *complete*. That is, the relations among the objects of a given domain (whether it is the state of a physical system or a problem-situation) are always defined and specified. No gaps are allowed. A state-space represents *all possible* states, but it has no obvious room to accommodate the *lack* of information about the state of a given system at a particular time. The same holds for a problem space, since it is usually articulated in terms of a convenient state-space. Conceptual spaces are not different. Gärdenfors does note that, given a quality dimension, we may not know its value for a particular object, and in this case, only a partial vector is assigned to the object (1990, p. 85). However, it is not clear how the notion of a partial vector is to be represented in terms of the full, complete conceptual space, since the latter is formulated with the assumption that *all* quality dimensions are spelled out. What is the status of a vector for which only some dimensions are available? It does not seem to be a properly formulated object, given the conceptual spaces framework. What we need, therefore, is a more open-ended framework in which the *partiality of information* can be properly accommodated. To this issue we should now turn.

3 A Formal Framework: Partial Structures and Quasi-Truth

The partial structures approach relies on three main notions: partial relation, partial structure and quasi-truth.[1] One of the main motivations for introducing this proposal comes from the need for supplying a formal framework in which the openness and incompleteness of information dealt with in scientific practice can be accommodated in a unified way (da Costa and French 2003). This is accomplished by extending, on the one hand, the usual notion of structure—in order to model the partialness of information we have about a certain domain (introducing then the notion of a *partial* structure)—and on the other hand, by generalizing the Tarskian characterization of the concept of truth for such partial contexts (advancing the corresponding concept of *quasi-truth*).

In order to introduce a partial structure, the first step is to formulate an appropriate notion of *partial relation*. When investigating a certain domain of knowledge Δ, we formulate a conceptual framework that helps us in systematizing and organizing the information we obtain about Δ. This domain is tentatively represented by a set D of objects, and is studied by the examination of the relations holding among D's elements. However, given a relation R defined over D, often we do not know whether all the objects of D (or n-tuples thereof)

[1] This approach was first presented in Mikenberg et al. (1986). For further developments, see da Costa and French (1989, 1990, 1993, 2003) and Bueno (1997).

are in the extension of R. This is part and parcel of the incompleteness of our information about Δ, and is formally accommodated by the concept of partial relation. More formally, let D be a non-empty set. An n-place *partial relation* R over D is a triple $\langle R_1, R_2, R_3\rangle$, where R_1, R_2, and R_3 are mutually disjoint sets, with $R_1 \cup R_2 \cup R_3 = D^n$, and such that: R_1 is the set of n-tuples that (we know that) belong to R, R_2 is the set of n-tuples that (we know that) do not belong to R, and R_3 is the set of n-tuples for which we do not know whether they belong or not to R. (If R_3 is empty, R is a usual n-place relation which can be identified with R_1.)

But in order to represent the information about the domain under consideration, we need a notion of *structure*. The following characterization is meant to supply a notion that is broad enough to accommodate the partiality usually found in scientific practice. Partial relations, of course, do the main work. A *partial structure* S is an ordered pair $\langle D, R_i\rangle_{i\in I}$, where D is a non-empty set, and $(R_i)_{i\in I}$ is a family of partial relations defined over D.[2]

Two of the three basic notions of the partial structures approach are now defined. In order to spell out the last one, quasi-truth, an auxiliary notion is required. The idea is to use, in the characterization of quasi-truth, the resources supplied by Tarski's characterization of truth. However, since this characterization is only formulated for full structures, we need to introduce an intermediary notion that links full to partial structures. This is the first role of those structures that extend a partial structure A into a full, total structure (which are called A-normal structures). Their second role is purely model-theoretic, namely, to put forward an interpretation of a given language and, in terms of it, to characterize basic semantic notions. But how are A-normal structures to be defined? Let $A = \langle D, R_i\rangle_{i\in I}$ be a partial structure. We say that the structure $B = \langle D', R'_i\rangle_{i\in I}$ is an *A-normal structure* if (i) $D = D'$, (ii) every constant of the language in question is interpreted by the same object both in A and in B, and (iii) R'_i extends the corresponding relation R_i (in the sense that each R'_i, supposed of arity n, is defined for all n-tuples of elements of D'). Note that, although each R'_i is *defined* for all n-tuples over D′, it holds for some of them (the R'^1_i-component of R'_i), and it doesn't hold for others (the R'^2_i-component).

As a result, given a partial structure A, there are *several* A-normal structures. Suppose that, for a given n-place partial relation R_i, we don't know whether $R_i a_1 \ldots a_n$ holds or not. A way of extending R_i into a full R'_i relation is to look for information to establish that it *does* hold, another way is to look for the contrary information. Both are *prima facie* possible ways of extending the partiality of R_i. But there are many different ways in which such extensions can go. In fact, there are *too many* possible extensions of the partial relations that constitute A. Thus, we need to provide constraints to restrict the acceptable extensions of the relevant partial structure.

In order to do that, we need first to formulate a further auxiliary notion (see Mikenberg et al. 1986). A *pragmatic structure* is a partial structure to which a third component has been added: a set of accepted sentences P, which represents the accepted information about the structure's domain. (Depending on the interpretation of science that is adopted, different kinds of sentences are introduced in P: realists will typically include laws and theories, whereas empiricists will add empirical regularities and observational statements about the domain in question.) A *pragmatic structure* is then a triple $A = \langle D, R_i, P\rangle_{i\in I}$, where D is a non-empty set, $(R_i)_{i\in I}$ is a family of partial relations defined over D, and P is a set of

[2] The partiality modeled here is due to the incompleteness of the information about the domain under investigation. With additional information, a partial relation may eventually become total. Thus, the partiality is not understood as a property of the world, but it is the result of our lack of information about it. The framework is concerned with an epistemic, not an ontological, partiality.

accepted sentences. The idea, as will become clear, is that P introduces constraints on the ways that a partial structure can be extended.

Given a *pragmatic* structure A, what are the necessary and sufficient conditions for the existence of an A-normal structure? These conditions can be formulated as follows (see Mikenberg et al. 1986): Let $A = \langle D, R_i, P \rangle_{i \in I}$ be a pragmatic structure. For each partial relation R_i, we construct a set M_i of atomic sentences and negations of atomic sentences, such that the former correspond to the n-tuples that satisfy R_i, and the latter to those n-tuples that do not satisfy R_i. Let M be $\cup_{i \in I} M_i$. Therefore, a pragmatic structure A admits an A-normal structure if, and only if, the set $M \cup P$ is *consistent*.[3]

We can now formulate the concept of quasi-truth. A sentence α is *quasi-true* in a pragmatic structure $A = \langle D, R_i, P \rangle_{i \in I}$ according to an A-normal structure $B = \langle D', R'_i \rangle_{i \in I}$ if α is true in B (in the Tarskian sense). If α is not quasi-true in A according to B, we say that α is *quasi-false* (in A according to B). Moreover, we say that a sentence α is *quasi-true* if there is a pragmatic structure A and a corresponding A-normal structure B such that α is true in B (according to Tarski's account). Otherwise, α is *quasi-false*.

Intuitively speaking, a quasi-true sentence α does not necessarily describe, in complete detail, the whole domain to which it refers, but only an aspect of it, the one modeled by the relevant partial structure A. For there are several different ways in which A can be extended to a full structure, and in some of these extensions, α may not be true. Thus, the notion of quasi-truth is strictly weaker than truth: although every true sentence is (trivially) quasi-true, a quasi-true sentence is need not be true (since it can be false in certain extensions of A).

In order to clarify the concept of quasi-truth, let me consider three objections that could be raised against it. (1) It may be argued that because quasi-truth has been defined in terms of full structures and the standard notion of truth, there is no gain with its introduction. In my view, there are several reasons why this is *not* the case. First, as we have just seen, despite the use of full structures, quasi-truth is weaker than truth: a sentence that is quasi-true in a particular domain—that is, with respect to a given partial structure A—may not be true if considered in an extended domain. Thus, we have here a sort of underdetermination, involving distinct ways of extending the same partial structure, which makes the notion of quasi-truth especially appropriate for empiricists. Second, one of the points of introducing the notion of quasi-truth, as da Costa and French (2003) have argued in detail, is that in terms of this notion, a formal framework can be advanced to accommodate the openness and partialness typically found in science. Bluntly put, the actual information at our disposal about a certain domain is modeled by a *partial* (but not full) structure A. Full, A-normal structures represent ways of extending the actual information that are possible according to A. In this respect, the use of full structures is a semantic expedient of the framework (in order to provide a definition of quasi-truth), but no epistemic import is assigned to them. Third, and crucially, full structures can be ultimately dispensed with in the formulation of quasi-truth, since this concept can be characterized independently of the standard Tarskian account of truth (see Bueno and Souza 1996). This provides, of course, the strongest argument for the dispensability of full structures (as well as of the Tarskian account) *vis-à-vis* quasi-truth. Therefore, full, A-normal structures are entirely inessential; their use here is only a convenient device.

(2) The second objection raises the problem of whether quasi-truth is a *truth* notion at all. If 'quasi' means something like 'approximate', the objection goes, what we have is, at most, an *epistemic* notion, but one that has nothing to do with truth. After all, truth is radically *non*-epistemic: it is *not* constrained by evidence. If it is a requirement for an adequate conception of truth that truth is taken in a non-epistemic way, then it may seem that quasi-truth is

[3] For further discussion, see Bueno (1997), section 3.1, and Bueno and Souza (1996).

not a truth concept; at least in the sense that quasi-truth is always relative to the available information encoded in a partial structure. But, in the end, it will all depend on how that information is interpreted. If it is interpreted as information about the world, rather than about our knowledge of it, then quasi-truth is a perfectly non-epistemic notion. If, however, the information is interpreted as about our knowledge of the world, then it is epistemic in nature. The formalism *per se* need not settle that issue.

Having said that, it should be pointed out that there are those who defend an *epistemic* notion of truth. For example, with his model-theoretic argument against metaphysical realism, Putnam indicated some difficulties that are faced by a *non*-epistemic account of truth (see Putnam 1976, 1980). But whatever the fortunes of this argument, the notion of quasi-truth need not be interpreted epistemically. The fact that quasi-truth is constrained by evidence—in the sense of being relative to the information encoded in a given partial structure—rather than being a difficulty is in fact the strength of this notion. After all, it allows us to accommodate the dependence of our truth claims on the evidence at our disposal, given a partial structure. In this respect, our judgments of quasi-truth reflect the current state of information about the domain we are concerned with.

(3) Finally, it may be argued, given that quasi-truth is formulated in terms of Tarski's account of truth-in-a-structure, and since this account does *not* provide an epistemic notion, it is unclear how quasi-truth *can be* epistemic. In reply, note again that, as opposed to Tarski's account, quasi-truth is always *relative* to a given *partial* structure. And given that such a structure represents our information about a given domain, not only our judgments of quasi-truth become relative to the extant information, but the same goes for quasi-truth itself. After all, if a sentence α is quasi-true it will remain forever as such. If later we discover that α is indeed true, it will still be quasi-true. Moreover, α's quasi-truth depends upon the partial structure we consider. In a *different* partial structure, α might not be quasi-true. This indicates the sense in which quasi-truth is an epistemic notion: it is relative to the information we have, that is, relative to a partial structure.

Far more could be said about quasi-truth, but I hope these remarks are enough to provide an idea of its main features and of the main components of the partial structures approach. We can now return to the main topics of our discussion, examining how this framework can be used to explore them in a new way.

4 The Role of Partial Spaces: Geometry and Heuristics

By using partial structures and quasi-truth, a new way of representing partial information in a belief system can be devised by means of what I call a *partial space*. The main idea is to reformulate the space representations discussed in Sect. 2 in terms of partial structures.

A *partial space* is a family of partial structures $(A_i)_{i \in I}$, where each A_i is of the form $\langle D_i, R_i \rangle_{i \in I}$, and R_i is a partial relation. Note that a partial space incorporates several domains D_i. This is meant to accommodate different quality dimensions of a conceptual space. In fact, each partial structure A_i determines a quality dimension, except for the fact that we now have a clear way of representing incomplete information about the objects in D_i: the n-tuples for which we do not know whether they belong or not to a relation R_i are found in the R_3-component of R_i. In this way, the lack of information about the relations in a conceptual space is straightforwardly accommodated in a partial space.

How about Gärdenfors' account of properties? It turns out that it also has a counterpart in a partial space. For partial relations also determine corresponding regions in this space: the regions determined by the objects that belong to R_i (R_1-component) and the complement determined by the objects that do *not* belong to R_i (R_2-component). The new feature is, of course, the R_3-component. It introduces open regions in the space, which represent our lack of information about the domain D_i. Now, depending on how such regions are filled in, either becoming full members of R_1 or of R_2, the shape of the resulting region will change. This process of filling in the regions corresponds, of course, to the construction of a full A-normal structure, and *that* structure has relations that determine full regions. For example, although we may not have enough information to say that a region H in a partial space is convex, since there may well be open subregions in it (those determined by the R_3-components), we can still claim that H is *quasi-convex*. This is so if (i) for every pair s_1 and s_2 of points in the region H, all points between s_1 and s_2 that are not in the open subregion R_3 are also in H, and (ii) there is an extension of the open subregion R_3 for which H is convex. In other words, as far as our current information is concerned, the region H may well be convex (nothing thus far precludes this possibility), and we know what kind of further information should be obtained in order to establish that this is actually the case.

This illustrates at once the heuristic role of partial spaces. The process of filling in open regions provides a clear heuristic strategy: we should look for the information that is required to extend the open regions in such a way that the resulting regions are convex. If we manage to achieve this, we can then employ Gärdenfors' analysis of properties; otherwise, we refrain from drawing inferences on the basis of those regions—unless we are trying to determine only the quasi-truth of the relevant situation, for which a quasi-convex region will be, of course, relevant.

These considerations also illuminate the role of quasi-truth in a partial space. Van Fraassen emphasized the importance of elementary statements as providing information about a physical system. We have something similar here. But instead of claiming that elementary statements are *true*—we often do not have enough information to establish such claims—we study the conditions in which they are *quasi-true*. An elementary statement S about a physical system X is *quasi-true in a partial structure* A if the state of X is represented by a point in the region of the partial space that is determined by A. Now, since we are dealing with a *partial* space, there may be open regions in the space (those determined by the R_3-components). So, again we have to study the possible extensions of these regions. If among these extensions there is one for which S is true (the state of X is a point in that region), then S is quasi-true. The point here is that our knowledge of the states of the system is often incomplete: we do not know all the system's states. In order to overcome these gaps, we have to explore ways of extending the partial relations, which determine at best open regions in the partial space, to full regions. This also illustrates the heuristic role of a partial space: it represents not only our current information about a physical system, but also our lack of information. The open regions are those that have to be explored further.

Let us consider an example. A mathematical investigation usually starts when the researcher is unsure whether a certain conjecture holds or not. At this stage, the mathematician tries both to prove and to refute the conjecture in question (for an insightful analysis, see Lakatos 1976). This move can be modeled very naturally in terms of partial spaces. We can think of the various possible outcomes of such an inquiry as represented in a *partial problem space*. The inquiry itself can be represented as a *trajectory* in such a problem space. The trajectory, in turn, is represented in terms of the relations that hold between the objects of the problem space. If such relations are *partial*, they are not defined for every n-tuple of objects of the domain. This means that the mathematician is still uncertain about the actual truth-value of the conjecture that depends upon such relations, which may either belong to the relevant domain (in which case, let us say, the conjecture would be true), or not belong to the domain (in which case the conjecture would be false). In order to determine which of them is the case, the mathematician will try both to prove and to refute the conjecture. Given the openness of the actual epistemic situation, this strategy is quite appropriate.

Moreover, we can also accommodate concept stretching, in Lakatos' (1976) sense, in terms of partial spaces. When a concept is 'stretched' in mathematics, that is, extended to accommodate some novel information beyond its original extension, the range of structures under consideration changes, since new conditions are given or lifted from the relevant domain. As a result, the structures that satisfy these conditions change as well. In the present framework, we can accommodate this shift by changing the *partial structure* with which we have been investigating the domain. The *new* concepts that are eventually introduced by 'stretching' the previous ones correspond to new kinds of relations introduced in a new partial structure. In this way, this important component of the dynamics of mathematical research, beautifully described by Lakatos, can be accounted for.

But this framework can also accommodate a further heuristic strategy. It was the strategy used by Abraham Robinson in the formulation of non-standard analysis (see his 1974/1996; see also Luxemburg 1996). As is well known, the early formulation of the calculus, due to Leibniz and Newton, was heavily dependent on infinitesimals, which were crucially employed in the derivation of the rules of Newton's method of fluxions. However, given the lack of a precise mathematical definition, infinitesimals were received with harsh criticism (in particular, by Berkeley). What Leibniz was trying to devise was a program of construction of numbers that would include infinitesimals. The idea was to introduce the latter, by appropriate arithmetic rules, as ideal numbers into the system of real numbers, in such a way that the resulting system would have the same properties as the real number system. However, given that neither Leibniz nor his followers managed to produce such a system, infinitesimals gradually fell into disrepute, and were eventually eliminated in the classical theory of limits elaborated in the nineteenth century.

But with sufficient heuristic resources every program can be revived. And at the end of the 1950s, with the work of Robinson, Leibniz's program was brought back. In fact, what Robinson realized (1974/1996, p. xiii) is that the model-theoretic techniques developed in the twentieth century provided the adequate framework in which Leibniz's intuitions could be properly articulated and vindicated. He showed (Luxemburg 1996, p. viii) that the ordered fields that are non-standard models of the theory of real numbers could be thought of, in the metamathematical sense, as non-archimedean ordered field extensions of real numbers **R**, and they included numbers behaving like infinitesimals with regard to **R**. Moreover, since these ordered fields are *models* of **R**, they have the same properties as **R** does. As a result, on Robinson's view, Leibniz's problem was solved.

It took so long for Leibniz's intuitions about infinitesimals to be developed because what was lacking was a proper *representation* of the problem. Given Leibniz's actual epistemic situation, with the lack of technical resources that were available to Robinson (especially model-theoretic techniques), he had at best a *partial* representation. But the crucial step in Robinson's remarkable achievement was to employ the methods of model theory. These methods provided the framework in which the partial representation of Leibniz's problem could be extended to a *full* one: by studying properties of non-standard models of real numbers, infinitesimals could be introduced in a precise and straightforward way.[4]

This case illustrates the importance of having an adequate *representation* of a problem in order to devise a solution for it. It also highlights the heuristic role of partial information,

[4] I have concerns that the reconstruction of Leibniz's program in terms of nonstandard analysis is historically accurate, since the Leibnizian continuum did not include the model-theoretic techniques that are needed in order to obtain the relevant nonstandard models. But I won't press this point here.

and *ipso facto* of partial spaces in which this information is represented: to indicate where further research should concentrate if we are to solve the problems we want to.

As we saw, partial spaces do suggest heuristic moves: by exploring the open regions of the space, and filling them in, we would be able to uncover new information—we will find out ways of moving R_3-components of the partial relations into R_1- or R_2-components. Furthermore, when we formulate a partial structure, and try to extend it to a full one, we are *representing* the information we have about a domain D. In this sense, we abstract some features of D, or we idealize them, by suggesting which features of D are taken as relevant, and which are not, or by adjusting them as needed by the relevant descriptions (see also French and Ladyman 1998). This is another heuristic feature of this framework: it helps us understand heuristic moves that have been made on the basis of the partiality of information about D.

5 Conclusion

By developing the concept of partial spaces, it is possible to represent a significant feature of belief systems: the inherent incompleteness of the information they encode. Given the way in which such spaces have been constructed, we can also preserve, when dealing with complete information, the salient features found in the different representations discussed in the beginning of this paper: from conceptual spaces through state-spaces to problem spaces. In this way, we can accommodate the incompleteness of belief systems together with the significant heuristic role that it plays.

References

Bueno, O. (1997). Empirical adequacy: A partial structures approach. *Studies in History and Philosophy of Science*, *28*, 585–610.

Bueno, O., & de Souza, E. (1996). The concept of quasi-truth. *Logique et Analyse*, *153–154*, 183–199.

da Costa, N. C. A., & French, S. (1989). Pragmatic truth and the logic of induction. *British Journal for the Philosophy of Science*, *40*, 333–356.

da Costa, N. C. A., & French, S. (1990). The model-theoretic approach in the philosophy of science. *Philosophy of Science*, *57*, 248–265.

da Costa, N. C. A., & French, S. (1993). Towards an acceptable theory of acceptance: partial structures and the general correspondence principle. In French, & Kamminga (Eds.), pp. 137–158.

da Costa, N. C. A., & French, S. (2003). *Science and partial truth*. New York: Oxford University Press.

French, S., & Kamminga, H. (Eds.). (1993). *Correspondence, invariance and heuristics: Essays in honour of Heinz Post*. Dordrecht: Reidel.

French, S., & Ladyman, J. (1998). A semantic perspective on idealization in quantum mechanics. In Shanks (Ed.), pp. 51–73.

Gärdenfors, P. (1990). Induction conceptual spaces and AI. *Philosophy of Science*, *57*, 78–95.

Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. Cambridge, MA: MIT Press.

Lakatos, I. (1976). *Proofs and refutations*. Cambridge: Cambridge University Press.

Luxemburg, W. A. J. (1996). *Foreword*, in Robinson (1974/1996), pp. vii–xi.

Mikenberg, I., da Costa, N. C. A., & Chuaqui, R. (1986). Pragmatic truth and approximation to truth. *Journal of Symbolic Logic*, *51*, 201–221.

Putnam, H. (1976). *Realism and reason*, in Putnam (1978), pp. 123–140.

Putnam, H. (1978). *Meaning and the moral sciences*. London: Routledge and Kegan Paul.

Putnam, H. (1980). Models and reality. *Journal of Symbolic Logic*, *45*, 464–482.

Rich, E., & Knight, K. (1991). *Artificial intelligence* (2nd ed.). New York: McGraw-Hill.

Robinson, A. (1974/1996). *Non-standard analysis*. (First edition, 1966.) Amsterdam: North-Holland. (Reprinted by Princeton University Press, 1996).

Shanks, N. (Ed.). (1998). *Idealization in contemporary physics. Poznan studies in the philosophy of the sciences and the humanities*. Amsterdam: Rodopi.
van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: Clarendon Press.
van Fraassen, B. C. (1991). *Quantum mechanics: An empiricist view*. Oxford: Clarendon Press.
van Fraassen, B. C. (2008). *Scientific representation: Paradoxes of perspective*. Oxford: Clarendon Press.
Varadarajan, V. S. (1968). *Geometry of quantum theory* (vol. 1). Princeton: van Nostrand.

Otávio Bueno is Professor of Philosophy and Chair of the Philosophy Department at the University of Miami. His research concentrates in philosophy of science, philosophy of mathematics, philosophy of logic, and epistemology. He has published widely in these areas in journals such as: Noûs, Mind, British Journal for the Philosophy of Science, Philosophical Studies, Philosophy of Science, Synthese, Journal of Philosophical Logic, Studies in History and Philosophy of Science, Analysis, Studies in History of Philosophy of Modern Physics, Erkenntnis, Monist, Metaphilosophy, Studies in History and Philosophy of Biological and Biomedical Sciences, Ratio, History and Philosophy of Logic, and Logique et Analyse. He is the author or editor of several books, and editor in chief of Synthese. In his free time, he enjoys running ultramarathons.