Formal Ontology in Information Systems
C. Trojahn et al. (Eds.)

doi:10.3233/FAIA241313

Fiat Surfaces in the Basic Formal Ontology

Michael Rabenberg[a,1] and Werner Ceusters[a]
[a] *University at Buffalo, 77 Goodell Street, Buffalo, NY 14203, USA*
ORCiD ID: Michael Rabenberg, https://orcid.org/0000-0001-9004-5322
Werner Ceusters, https://orcid.org/0000-0002-2676-8689

Abstract. The Basic Formal Ontology (BFO) class *continuant fiat boundary* and its subclasses, including *fiat surface*, are not heavily axiomatized; they have elucidations, not definitions; and the meanings of these elucidations are poorly captured by the relevant BFO axioms. This paper is an effort to make progress in these respects for *fiat surface*. We identify a range of desiderata for a BFO-conformant view of *fiat surface*, argue that the GitHub does not satisfy them, argue that the view of fiat surfaces in Arp et al. (2015) does a better job, and supplement that view in ways that do a still better job. Our discussion allows us to, *inter alia*, present a number of axioms relevant to our topic worthy of consideration for inclusion in BFO or dependent ontologies in specific domains and for specific purposes.

Keywords. Basic Formal Ontology, FOL axiomatization, surface.

1. Introduction

One class of *immaterial entity* in the Basic Formal Ontology (BFO) is *continuant fiat boundary*; its subclasses in the BFO-2020 GitHub repository, henceforth 'the GitHub', are *fiat point*, *fiat line*, and *fiat surface*.[1] Unlike some BFO categories, such as *temporal interval*, these categories are not heavily axiomatized; on the GitHub they have elucidations, not definitions; and (as we shall make clear) the meanings of these elucidations are poorly captured by the relevant axioms. This paper is an attempt to improve on the status quo of these matters, specifically concerning *fiat surface*.

Our interest in this topic derives partly from our interest in the Ontology for General Medical Science (OGMS), some components of which we have begun axiomatizing in first-order logic. Some definitions in the OGMS raise questions about the fiat surfaces of the entities with which these definitions are concerned. Consider, for example, the OGMS's definition of *extended organism*: 'An object aggregate consisting of an organism and all material entities located within the organism, overlapping the organism, or occupying sites formed in part by the organism' (github.com/OGMS/ogms/blob/master/ogms.owl). ('Extended organism' was not a defined term in [2], the article from which the OGMS in its present form was derived.) One question provoked by this definition is whether any entities located entirely *outside of* an organism's outermost material parts, such as micro-organisms on someone's skin, might nevertheless be member-parts of this organism's extended organism. It seems that producing a satisfactory answer to this question requires a thorough examination of *fiat*

[1] Corresponding Author: Michael Rabenberg, Department of Biomedical Informatics, University at Buffalo, 77 Goodell Street, Buffalo NY, 14203, USA rabenbergm@gmail.com.

surface and related universals (e.g., *site*, *continuant fiat boundary*). Furthermore, 'treatment of boundary_of relations (incl. fiat_boundary_of)' is a topic for future work listed in the BFO 2.0 specifications [3, p6]), but has thus far not been addressed; no fiat_boundary_of relation, or any other such relation, is defined in the BFO at present. In addition, the results of an investigation of the sort we have in mind would also help domain ontology developers who might wish to capitalize on BFO to link their classes accurately. One such ontology is SNOMED CT [4, 5], in which the term 'surface' appears in the fully specified names of over 1700 concepts across 19 concept hierarchies, some of which denote entities related only indirectly, if at all, to the sort of immaterial entity intended by the BFO. Some examples include 'Surface (attribute)', 'Surface (qualifier value)', 'Tooth surface loss (disorder)', 'Fall on hard surface (event)', and 'Surface compensator (physical object)'.

We proceed as follows. In Section 2, we present four desiderata for a BFO-conformant view of fiat surfaces. We bring out these desiderata partly via reflection on some remarks by H.H. Price concerning *surfaces* (not *fiat surfaces*) written long before anyone had uttered the phrase 'Basic Formal Ontology'. These remarks, we suggest, help to motivate some of the desiderata we propose. In Section 3, we turn to what the GitHub says about *fiat surface* and argue that the GitHub goes a small way toward satisfying just two of these desiderata at present. In Section 4, we argue that a view of fiat surfaces that emerges from a charitable interpretation of Arp et al. (2015) [6], a text widely treated as a highly authoritative source for information about the BFO, goes much further. Central to this view, as we understand it, is the thesis that *fiat surfaces are continuant parts of material entities*. In Section 5, we supplement this view in ways that allow it to satisfy the desiderata to a greater extent. We also articulate several axioms concerned with *fiat surface* and related BFO categories that we believe ought to be included, or considered for inclusion, in future versions of the BFO. These axioms are summarized in Section 6 and are accessible in two separate files (https://github.com/michaelrabenberg/FOIS-2024-Rabenberg-Ceusters-Paper/tree/main).

2. Four Desiderata for a BFO-Conformant View of Fiat Surfaces

The following is an excerpt from H.H. Price's *Perception* (quoted in Varzi (2023: Section 4) [7], a catalogue of philosophical passages from antiquity to recent times concerned with boundaries, surfaces, and so on). Sub-passages are named and underlined.

[P] [P1] A surface must be the surface of something; and not merely that, [P2] it must be the surface of some material thing. [Price (1932: 106) [8]]

[P] expresses what Price evidently takes to be a commonsensical view of *surfaces*, not, of course, of the BFO class *fiat surface*. Still, reflection on [P] will help us to introduce several desiderata for a BFO-conformant view of fiat surfaces. [P1] seems to have among its intended implications the following two propositions (variables not explicitly existentially quantified are presumed universally quantified):

(P1a) If x is a surface, then x is a surface *of* an entity other than x.
(P1b) If x is a surface *of* y, then no entity other than x is a surface of y.

That [P1] is meant to imply (P1a) is obvious. (Strictly speaking the letter of [P1] is consistent with the claim that every surface is a surface *of itself*, but [P1] is clearly intended to rule that possibility out.) That [P1] implies (P1b) is suggested by the definite article in '*the surface of something*'. We consider these implications in turn.

Adapted to fiat surfaces, (P1a) suggests what we might call the *relation thesis*, the thesis that every fiat surface bears some *special relation*, which we might call the '*hugs*' relation, to some other entity. Now, this is of course extremely unspecific language, but it brings out an initial desideratum for a BFO-conformant view of fiat surfaces:

(D1) A BFO-conformant view of fiat surfaces ought to contain a plausible version of the relation thesis that explicates the *hugs* relation via extant BFO categories and relations.

Why accept (D1)? For one thing, we think it clear that *some* version of the relation thesis is correct. A world containing for all its history one fiat surface and nothing else to which the fiat surface is related in any special way seems no less absurd than a world containing for all its history one color quality and nothing else to which it is related in any special way. And here is a very plausible at-least-partial explanation of this absurdity: some version of the relation thesis is correct. If that is right, then any plausible view of fiat surfaces will involve some such thesis. Furthermore, we take it that a *BFO-conformant* view of fiat surfaces ought to involve a version of the relation thesis that explicates the *hugs* relation via extant BFO categories and relations. For nothing would be gained that ontologies are meant to provide from the mere assertion of the following thesis:

(Dud) If x is an instance of fiat-surface at t, then there exists a y such that $x \neq y$ and x *hugs* y at t.

(Dud) is pretty much unimpeachable. But absent a definition of the *hugs* relation formulated in terms of extant BFO categories and relations, or at least the specification of some necessary or sufficient conditions for something's bearing *hugs* to something, (Dud) is useless. Hence the need for a version of the relation thesis of the sort demanded by (D1).

Some candidates for versions of the relation thesis responsive to (D1) are promising so far as they go, but they do not go very far. For example, consider the following claim:

(Locus) If x is an instance of fiat-surface at t, then there exists a y such that $x \neq y$ and x is *located-in* y at t.

(Locus) is formulated in BFO-ese, and it expresses the plausible (though resistible) thought that *hugs* entails (in addition to nonidentity) *located-in*. But we do not think that (Locus) goes very far toward satisfying (D1), because we doubt that *hugs* is mere location-in (though it might imply this). By contrast, consider the following claim:

(Part) If x is an instance of fiat-surface at t, then there exists a y such that $x \neq y$ and x is continuant-part-of y at t.

(Part) arguably goes farther than (Locus) toward satisfying (D1), if only because for BFO *continuant-part-of* entails *located-in* but the converse does not hold. Again, though, *hugs*

is not identical to *continuant-part-of* (though, again, it might entail this). We shall return to claims very similar to (Locus) and (Part) later on.
So much, for now, for (D1). Now we shall introduce a second desideratum for a BFO-conformant view of fiat surfaces, which we shall clarify and motivate via reflection on (P1b):

(D2) A BFO-conformant view of fiat surfaces ought to distinguish between *maximal* and *non-maximal* fiat surfaces, and it ought to clarify some of their characteristics, relations to one another, and relations to other entities to which they bear important relations.

Adapted to fiat surfaces and interpreted in light of the relation thesis, (P1b) suggests the *uniqueness thesis*, the thesis that everything hugged by a fiat surface is hugged by just one fiat surface. Now, if every fiat surface is what might be called a *maximal* fiat surface, i.e., a fiat surface that is a proper continuant part of no other fiat surface, then it is at least somewhat plausible that the uniqueness thesis is true. This is because it is at least somewhat plausible—though hardly incontestable—that no entity is hugged by two numerically distinct maximal fiat surfaces at a given time. But presumably not every fiat surface is maximal. Presumably there are *non-maximal* fiat surfaces, i.e., fiat surfaces that are proper continuant parts of other fiat surfaces. But if there are non-maximal fiat surfaces, then presumably no version of the uniqueness thesis is true as applied to them. To illustrate: It is plausible that, right now, one fiat surface, *lefty*, hugs Barry and is located at the exterior of Barry's left arm and a distinct fiat surface, *righty*, hugs Barry and is located at the exterior of Barry's right arm, but each of these is a proper part of a third fiat surface (e.g., one that hugs Barry *in his entirety*, so to speak).

These considerations bring out that a BFO-conformant view of fiat surfaces ought at least to accommodate the maximal/non-maximal fiat surface distinction; it ought not, for example, be an implication of BFO that there are no non-maximal fiat surfaces. Furthermore, there are questions that can be asked about these different sorts of fiat-surface, such as how many non-maximal fiat surfaces can be proper continuant parts of a given maximal one at a time, which entity or entities a given non-maximal fiat surface ought strictly speaking to be said to hug at a given time (for example, do *lefty* and *righty* hug Barry alone or do they also hug (proper continuant parts of) Barry's left and right arms, respectively?), and so on. Hence (D2).

Now turn to [P2], the second underlined sub-passage in [P]. [P2] seems clearly to have the following proposition among its intended implications:

(P2a) If x is a surface, then x is a surface of a *material entity* that is not x.

Ordinary language connotations notwithstanding, (P2a) is neutral about whether some surfaces are material entities. Applied to fiat surfaces, (P2a) suggests a relation thesis logically stronger than the one suggested by (P1a), which we might call the *material-entity relation thesis*, according to which every fiat-surface hugs a *material entity* other than itself. The material-entity relation thesis is nearly as plausible as the already extremely plausible bare relation thesis suggested by (P1a). Presumably, if the material-entity relation thesis were false, then there could be a world containing nothing but fiat surfaces and immaterial Cartesian egos hugged by them for all of its history. This seems barely less absurd than a world containing nothing but fiat surfaces, period, for all of its history. Furthermore, if there are any objections to the material-entity relation thesis from

the point of view of BFO specifically, then they will have to appeal in some way to immaterial entities that occupy spatial regions but that are not themselves continuant fiat boundaries, such as sites. Sites are the only such entities explicitly recognized by BFO, though BFO does not entail that there are no other such things. But even sites (e.g., the cargo hold of a ship) do not seem to be hugged by *fiat surfaces*. Rather, sites seems to be related in some special manner to *non–fiat-surface* continuant fiat boundaries that occupy two-dimensional spatial regions. So, we see no good reason to resist the material-entity relation thesis. Thus the following desideratum supersedes (D1):

(D1+) A BFO-conformant view of fiat surfaces ought to contain a plausible version of the material-entity relation thesis that explicates the *hugs* relation via extant BFO categories and relations.

As it happens, the GitHub elucidations that we shall discuss in the next section imply that every fiat surface is a continuant fiat boundary and that, consequently, every fiat surface has its 'location…determined in relation to some material entity.' [1]. Thus the GitHub elucidations themselves seem to reinforce (D1+).

This concludes our reflection on [P] and our derivation of desiderata fairly directly from it. In the rest of this section, we shall introduce two additional desiderata for a BFO-conformant view of fiat surfaces.

The first is provoked by (Locus), entertained above. (Locus) seems plausible in its own right, and some plausible views of certain *non-spatial* relations between fiat surfaces and other things entail it. For example, if one holds that every fiat surface is a continuant part of a material entity, then one gets (Locus) *gratis*. But, as we briefly indicated earlier, (Locus) is resistible. For example, one could defend what might be called the *halo view* of fiat surfaces, according to which every fiat surface is located *just outside* of some material entity. The halo view is not easily reconcilable with (Locus), and is surely inconsistent with its spirit. Thus these considerations yield another desideratum:

(D3) A BFO-conformant view of fiat surfaces ought to say something about where fiat surfaces are located relative to salient other entities.

In particular, in light of our already-defended view that *fiat surfaces hug material entities*, the 'salient other entities' at issue in (D3) most obviously include the material entities that fiat surfaces hug.

Finally, we shall introduce a desideratum that concerns a topic of great importance to BFO, namely change over time. Consider the following case:

> ***Lone Star Bamboo.*** At t, a perfectly cylindrical bamboo stalk, *bamby*, has height 10cm and diameter 2cm. Also at t, *ricky*, an ontologist, draws a small closed figure in the shape of Texas on an uninteresting location on Bambi's side and says, 'I dub this fiat surface, the one enclosed at present by this figure, '*texy*'.' Between t and some later time t* *bamby* grows to 15cm in height and 3cm in diameter, seriously distorting the figure that *ricky* had drawn.

At t, there is (let us assume) a maximal fiat surface that hugs *bamby*. Call this surface '*maxy*.' (In saying that *maxy* hugs *bambi* at t, we do not mean to imply that *maxy* does not hug *bambi* at times other than t.) *Lone Star Bamboo* provokes several questions about fiat surfaces and their ability or inability to persist through changes. Some examples:

Does *maxy* exist at t*? Does *texy* exist at t (or does *ricky* falsely believe that there is a fiat surface enclosed by the Texas-shaped figure at t that he can dub)? Assuming *texy* exists at t, does *texy* exist at t*? If either of these fiat surfaces exist at t*, does it hug *bamby* at t*? Whatever the answers are to these questions, why are these answers the correct ones? And so on. Thus *Lone Star Bamboo* motivates the following desideratum:

(D4) A BFO-conformant view of fiat surfaces ought to say something about the circumstances under which fiat surfaces do and do not persist through changes (in particular dimensional changes) undergone by salient other entities.

We have thus far proposed four desiderata for a BFO-conformant view of fiat surfaces: (D1+), (D2), (D3), and (D4). In the next section, we argue that the GitHub's view of fiat surfaces goes a small way toward satisfying just two of these four.

3. The GitHub

We first consider the two GitHub *elucidations* most directly relevant to *fiat-surface*. We consider relevant *axioms* later in this section. Here are the elucidations:

(E-cfb) b is a continuant fiat boundary means: b is an immaterial entity that is of zero, one or two dimensions, which is such that there is no time t when b has a spatial region as continuant part at t, and whose location is determined in relation to some material entity.

(E-fs) A fiat surface is a two-dimensional continuant fiat boundary that is self-connected. [1]

These elucidations have some important implications about fiat surfaces (e.g., that every fiat surface is an immaterial entity, that no fiat surface has a spatial region for a continuant part), but they do not imply anything very responsive to our desiderata.

Consider (D1+) and (D3). Because (E-cfb) and (E-fs) imply that every fiat surface has its 'location…determined in relation to some material entity,' these elucidations are somewhat responsive to (D1+) and (D3). But they leave entirely unanswered what, exactly, this location-determination amounts to. For example, for all (E-cfb) and (E-fs) imply, this location-determination could be, say, location-in, or lack of even partial location-in, or some very different relation. Furthermore, if the 'location…determined in relation to' language is interpreted literally, then (E-cfb) and (E-fs) seem overinclusive. Suppose Alan imagines a roughly spherical, two-dimensional, self-connected immaterial entity floating 12 feet above his own head, but surrounding nothing but air. There is an obvious sense in which this is an entity 'determined in relation to' a material entity, namely Alan. We assume, though, that this entity is not meant to count as a fiat surface.

About (D2), (E-cfb) and (E-fs) seem silent. Nothing in them implies either that there are maximal and non-maximal fiat surfaces or (say) that there are only maximal ones. Certainly they imply nothing about special relations between maximal and non-maximal fiat surfaces, or between fiat surfaces of these sorts and salient other entities (beyond what has already been said about the relevance of (E-cfb) and (E-fs) for (D1+) and (D3)).

Finally, (E-cfb) and (E-fs) seem to imply nothing about the conditions (if any) under which a given fiat surface persists through the sorts of change with which (D4) is concerned, so the elucidations seem to go no way toward satisfying (D4).

So much for the elucidations. Before we proceed, we note that the propositions in the BFO 2.0 specifications document most closely corresponding to (E-cfb) and (E-fs) do not satisfy our desiderata any more than (E-cfb) and (E-fs) themselves do. (A specifications document for BFO 2020 is currently in progress; we do not discuss it because it is unpublished at the time of this writing.) Here are the propositions:

> DEFINITION: *b* is a *continuant fiat boundary* = Def. *b* is an *immaterial entity* that is of zero, one or two dimensions, which does not include a *spatial region* as part, and whose location is defined in relation to some material entity. [3, p42])
>
> ELUCIDATION: a *two-dimensional continuant fiat boundary* (surface) is a self-connected fiat surface. [3, p43])

Because the first proposition is identical to (E-cfb), except that it is presented as a definition rather than as an elucidation, it goes no further than (E-cfb) in satisfying the desiderata. Furthermore, the second proposition is conspicuously an elucidation not of *fiat surface* but of *two-dimensional continuant fiat boundary* in terms of *fiat surface*. *Fiat surface* itself receives no elucidation or definition in the specifications document. Nor, incidentally, is *fiat surface* explicitly defined in Arp et al. (2015). Consequently, the second proposition rather trivially does not go further than (e-fs) does toward satisfying the desiderata.

Now we shall discuss what bearing, if any, the GitHub axioms, hereafter called '*BFO-2020*' (italicized), might have on the desiderata. To explore this matter, we used an in-house developed reasoner, similar in function to familiar theorem provers such as Prover9, to test what implications could be deduced from *BFO-2020* when combined with various relevant input assumptions. For example, one test combined *BFO-2020* with the following assumptions (which we reproduce here in the CLIF dialect):

(t1a) (instance-of surface-1 fiat-surface time-1)
(t1b) (forall (m t) (not (instance-of m material-entity t)))

This test yielded a contradiction, because (t1b) is inconsistent with *BFO-2020*. By contrast, as we had anticipated on the basis of a prior manual inspection of what seemed to us to be the relevant axioms of *BFO-2020*, the following test yielded no contradiction:

(t1a) (instance-of surface-1 fiat-surface time-1)
(t1c) (instance-of time-1 temporal-instant time-1)
(t1d) (forall (m t) (if (instance-of m material-entity t) (not (= time-1 t))))

We take these tests to support the conclusion that *BFO-2020* goes only a small way toward satisfying (D1+). For although *BFO-2020* rules out the possibility of a world in which there is a fiat surface and *never* a material entity, it does not rule out the possibility of a world in which there is a fiat surface at some temporal instant but no material entity at that temporal instant. Adequately satisfying (D1+) would, we take it, rule out this second possibility. We also take these tests to support the conclusion that *BFO-2020* goes no way toward satisfying (D3). Given the second possibility just described, it is *a fortiori*

possible for there to be a fiat surface at some temporal instant that at that temporal instant bears no interesting locational relation to any material entity. But satisfying (D3) would, we assume, rule out this possibility, as well.

Our testing concerning (D2) (see https://github.com/michaelrabenberg/FOIS-2024-Rabenberg-Ceusters-Paper/tree/main) supported the conclusion that *BFO-2020* is consistent with the conjunction of the following propositions:

(i) At some time, some fiat surface is not a proper continuant part of any fiat surface.
(ii) At some time, some fiat surface is a proper continuant part of some fiat surface.
(iii) At some time, there is a fiat surface with exactly one fiat surface among its proper continuant parts.
(iv) At some time, there is a fiat surface, s1, with a fiat surface for a proper continuant part, and for any fiat surface, s2, that is a proper continuant part of s1, there is a fiat surface, s3, such that s3 is a proper continuant part of s1 and s2 is a proper continuant part of s3.

That *BFO-2020* is consistent with (i) and (ii) implies that *BFO-2020* is neutral on the existence of maximal fiat surfaces and on the existence of non-maximal ones. That *BFO-2020* is consistent with (iii) and (iv) implies that *BFO-2020* is consistent with the existence of fiat surfaces that are 'mereologically discrete' with fiat surfaces and with the existence of fiat surfaces that are 'mereologically dense' with fiat surfaces; that *BFO-2020* is consistent with the existence of fiat surfaces of both of these sorts supports the further conclusion that *BFO-2020* implies nothing about *how many* fiat surfaces are located in a given fiat surface. These findings led to conclude that *BFO-2020* is neutral about the issues with which (D2) is concerned.

Finally, concerning (D4), we did two things. First, we ran *BFO-2020* along with all of the following propositions through our reasoner:

(t1a) (instance-of surface-1 fiat-surface time-1)
(t1e) (instance-of material-entity-1 material-entity time-1)
(t1d) (instance-of surface-1 fiat-surface time-2)
(t1g) (precedes time-1 time-2)
(t1h) (not (exists (m) (instance-of m material-entity time-2)))
(t1i) (instance-of time-2 temporal-instant time-2)

This test yielded no contradiction, supporting the conclusion that, for all *BFO-2020* implies, a fiat surface that exists alongside a material entity at one time can exist at a later instant at which no material entity exists any longer. In addition, a manual inspection of *BFO-2020* provided no evidence that it entails that if a material entity is contemporary with a fiat surface and then the material entity undergoes some change (dimensional or otherwise), then a fiat surface undergoes some corresponding change. These investigations make us confident that *BFO-2020* is silent about (D4)-relevant matters.

4. Arp et al. (2015) [6]

In this section, we argue that Arp et al. (2015) can be interpreted as endorsing a view of *fiat surface* that goes further than the GitHub does toward satisfying our four desiderata.

Start with (D1+). Although no passage in Arp et al. (2015) unambiguously satisfies (D1+), the following passage and Arp et al.'s subsequent discussion of it can be interpreted as being responsive to (D1+), in a way that we shall explain below:

[A1] a boundary *a* of an object *b* is boundary-dependent on this object if and only if it is necessarily such that it can exist only if either *b* exists or there exists some part of *b* that includes *a* as part. [6, p109, underlining ours]

To illustrate what they mean by [A1], Arp et al. present a case involving a cuboid block of marble that contains a hole in its interior, itself containing a 'supremely powerful corrosive acid' that is causing the hole to grow (*Figure 1)*.

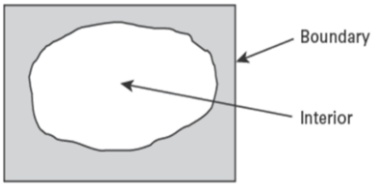

Figure 1. Block of marble with growing interior hole. [6, p109]

Arp et al. say that, at least until the acid has penetrated the exterior of the block, the block's 'boundaries', by which they appear to mean the outer continuant fiat boundaries hugging the block, remain in existence: 'They [i.e., the block's boundaries] will continue to exist for just as long as there is at least *some* remaining part of the block that includes them as part' [6, p109, underlining ours]. A seeming implication of this remark is that at least *before* the acid breaks through the block's exterior (i.e., as long as things are as Figure 1 depicts them), *the block's boundaries are continuant parts of the block*. But Arp et al. of course accept BFO's classification of all fiat surfaces as continuant fiat boundaries, and they are clear that, in their view, an entity that BFO would call a 'fiat surface' is among the block's boundaries. It is natural, then, to attribute to Arp et al. an endorsement of the view that *hugs entails continuant-parthood*, and thus also of this claim, a strengthened version of one introduced in Section 2:

(Part+) If x is an instance of fiat-surface at t, then there exists a y such that $x \neq y$, y is an instance of material-entity at t, and x is continuant-part-of y at t.

And this position is of course responsive to (D1+), even if it does not satisfy it completely.

Here a clarification is necessary. Arp et al. use the term 'fiat surface' on only a small number of occasions (just two, to our finding; see Arp et al. (2015: 110, 186)), and they never propose a definition of *fiat surface*. (On the contrary, like the BFO 2.0 specifications document, Arp et al. (2015: 110, 186) define *two-dimensional continuant fiat boundary*, a class that does not appear in the GitHub axioms, definitions, or elucidations, partly in terms of *fiat surface*.) So, in attributing (Part+) to Arp et al., we are engaging in some speculative reconstruction of their view in light of the present GitHub materials (which were produced a few years after Arp et al. (2015) was published). Our aim in this section is to suggest what Arp et al. *might* say if they were to try, to the extent possible, to rearticulate the view they endorse in their book, at least in broad outline, using current BFO language. This endeavor on our part requires some deviation from their own exact language, but we take such deviations as are required not to do violence to the substance of their position.

Now turn to (D2). Arp et al. often speak of 'the' surface of this or that material entity; or refer to an entity's 'surface' in a way that strongly suggests *uniqueness*; or refer to 'the' fiat object boundary of a material entity, which they equate with 'its maximally

connected two-dimensional surface,' where the "its" also strongly suggests uniqueness; see for example Arp et al. (2015: 92, 108, 109, 110, 116, 117, 135, 186). This language leads us to believe that Arp et al. would endorse some version of what we have called the uniqueness thesis. And a natural version of the uniqueness thesis would be the following claim, which is at least partly responsive to (D2): For each entity that a *maximal* fiat surface hugs at a given time, *exactly one* maximal fiat surface hugs this entity at this time. Thus Arp et al.'s view can be understood as at least somewhat responsive to (D2).

However, to our finding Arp et al. say nothing explicit about what we have called *non-maximal* fiat surfaces. In particular, here are two natural questions about them to which Arp et al.'s discussion seems unresponsive: (i) Is a given non-maximal fiat surface located in a given maximal fiat surface a *continuant part* of this maximal fiat surface? (ii) *How many* non-maximal fiat surfaces can hug a given material entity at a time? We suggest answers to these questions consistent with Arp et al.'s view in Section 5.

Now turn to (D3). First, because every continuant part of a material entity is located in this material entity, and we have attributed (Part+) to Arp et al., we also attribute to them this thesis, which is a strengthened version of a thesis introduced in Section 2 and which is responsive to (D3):

(Locus+) If x is an instance of fiat-surface at t, then there exists a y such that $x \neq y$, y is an instance of material-entity at t, and x is located-in y at t.

Of course, (Locus+) is not *very* responsive to (D3). For all (Locus+) says, it could be that every fiat surface is shaped like Texas and located in an undistinguished location entirely within some material entity. However, Arp et al. also say that 'a *continuant fiat boundary* is a boundary of some material entity that exists exactly where that object meets its surroundings' [6, p108] . Here one could get the impression that Arp et al. think that the only material entities *of which* continuant fiat boundaries can be continuant fiat boundaries are objects. But this is presumably not the case; presumably fiat points, say, can be continuant fiat boundaries of, for example, non-object fiat object parts. The 'meets its surroundings' locution, combined with (Locus+), suggests that, according to Arp et al., every fiat surface is located in some material entity at the 'exterior,' so to speak, of the material entity. This is of course a plausible suggestion, and it is surely at least close to true; but we shall present a wrinkle for it in Section 5.

Finally, consider (D4). Arp et al. clearly hold that some continuant fiat boundaries, including some fiat surfaces, persist through certain changes in the entities that they hug. Recall the example of the block of marble: on their telling, at least as long as the acid has not broken through the exterior, at least some of the block's continuant fiat boundaries persist through the growth of the block's interior hole (which growth is of course a change undergone by the marble block). However, to our finding Arp et al. are silent about cases, such as *Lone Star Bamboo*, in which a material entity undergoes changes in its *external* size and shape. But changes of this sort, like the one in *Lone Star Bamboo*, raise the most challenging questions about the persistence or non-persistence of a fiat surface. We shall make some suggestions in the next section relevant to such cases.

To review: we believe that Arp et al. (2015) endorse a view that goes a significant way toward satisfying (D1+), (D2), and (D3), with at least two questions pertinent to (D2) remaining. Their view does not seem to satisfy (D4) to a significant extent.

5. Filling out the View

In this section, we 'run with' the view we attributed in Section 4 to Arp et al. and propose some additions and modifications to it to make it more responsive to the four desiderata, and to make it capable of answering a challenge that we shall present in this section.

First, consider the two questions raised in Section 4 pertinent to (D2): (i) Is a given non-maximal fiat surface located in a given maximal fiat surface a *continuant part* of this maximal fiat surface? (ii) *How many* non-maximal fiat surfaces can hug a given material entity at a time?

We suggest that the most natural answer to (i) is 'Yes, with an important qualification.' Suppose that at some time two material entities, m1 and m2, 'touch' one another, so that there are two maximal fiat surfaces, s1 and s2, hugging m1 and m2, respectively, where s1 has a non-maximal fiat surface continuant part, s1-, and s2 has a non-maximal fiat surface continuant part, s2-, such that *s1- and s2- occupy the same spatial region*. For example, m1 and m2 could be two perfect cubes, with one of m1's faces pressed perfectly against one of m2's faces. (If you deny the possibility of such a case, ignore the present issue.) Now, s1- and s2- are each located in s1 and in s2. Yet s1- is a continuant part of s1 and not of s2, and s2- is a continuant part of s2 and not of s1. This shows that we cannot say without qualification that if a non-maximal fiat surface, s-, is located in a maximal fiat surface, s, at t, then s- is a continuant part of s at t. Rather, such cases recommend the following thought: every non-maximal fiat surface is a continuant part of the maximal fiat surface *that hugs the same material entity as the non-maximal fiat surface does*. Because s1 and s2 hug different entities, this claim allows the case just described to raise no special difficulties.

As for (ii), we think the most natural answer is 'as many as the character of space allows.' If space is continuous, then there are uncountably many. If space is discrete, then there are finitely many - in particular, one for every path connecting two points along the relevant maximal fiat surface.

Now we shall present a challenge to Arp et al.'s view that is relevant to (D3). Let a material entity's 'outer' continuant parts be those continuant parts of it that (to reuse Arp et al.'s phrase) *meet the surroundings* of the material entity. Your skin is plausibly among your outer parts; your liver presumably is not. Now consider a material entity with a *site* among its outer parts. For example, imagine some particular canoe—call it '*canny*'—and let '*hully*' refer to *canny*'s hull, i.e., the hollowed-out, empty area in which a canoeist sits. Suppose *hully* is a site that is among the proper continuant parts, and thus among the *outer* proper continuant parts, of *canny*. (Some will reject this supposition. Bear with us for now.) Finally, suppose for *reductio ad absurdum* that there is a fiat surface, *surfy*, that hugs *canny* and is located wherever and only wherever *canny* meets its surroundings. Then part of *surfy*, in particular the part of surfy that is located precisely on the exterior of *hully*, "floats" above the solid walls and floor of *canny*. But we think this is an implausible result. We think it much more natural to hold that any fiat surface that hugs *canny* is located where the solid walls and floor of *canny* meet *their* surroundings. This suggests that a fiat surface that hugs a given material entity is not in all cases located in a place where *the material entity* meets its surroundings, but rather is located where (so to speak) *the matter* of the material entity meets *its* (i.e., the matter's) surroundings.

Now, one could avoid this challenge by insisting that *hully* is not a part of *canny*, or by saying that there is no site at all where we have claimed *hully* is located. But unless one holds that no material entities have sites among their outer parts, one will have to allow that there are some cases relevantly like that of *canny* as we have described it.

What is it for something to be *the matter of* a material entity? 'Matter,' though not a BFO category, appears in the GitHub elucidation of *material entity*: 'A material entity is an independent continuant that at all times at which it exists has some portion of matter as continuant part.' [1] It might be tempting to define the *is-matter-of* relation roughly as follows: x is-matter-of material entity y at t iff x is the maximal continuant part of y with no immaterial entities for continuant parts at t. But this will probably not work, because every continuant part of a material entity plausibly has some immaterial entities, such as fiat points, for continuant parts.

Better perhaps would be the following definition:

(M) x is-matter-of material entity y at t iff x is y's maximal continuant part with no *sites* for continuant parts at t.

We think it intuitive that for any material entity, x, the entity that is-matter-of x at t is the maximal continuant part of x none of whose own continuant parts are immaterial entities that 'crowd out' other material-entity continuant parts of x at t. Sites seem to 'crowd out' the material-entity continuant parts of the material entities of which they are continuant parts. That is, it seems that if some site, s, is a continuant part of some material entity, x, at t, then at t there is no material entity that is a continuant part of x and that is located *in s*. This is of course not to say that sites, or even sites that are continuant parts of material entities, crowd out material entities *generally*. The cargo hold of a ship, say, can of course have cargo boxes, which are material entities, located inside of it. But these cargo boxes are not continuant parts of the ship. Similarly, the ship can of course have material-entity continuant parts exactly abutting the cargo hold, e.g., the walls, floor, and ceiling circumscribing the cargo hold; but these things are not *located in* the cargo hold. By contrast, continuant fiat boundaries do not have the same 'crowding out' effect. This is why, intuitively, it is no objection to the claim that some particular board is a part of *canny*'s matter that this board has among its continuant parts some continuant fiat boundaries (e.g., some fiat points), whereas it is an objection to the claim that *canny* is a part of *canny*'s own matter that *canny* has *hully* for a continuant part. Hence (M)'s plausibility.

But (M) faces a difficulty of its own—in fact it faces a number of difficulties, but we shall discuss just one here. Consider some material entity, x, that at the first moment of its existence, t, has *no* sites for continuant parts. Given (M), it follows from BFO axiom [scr-1] that either x and that which bears is-matter-of to x, m(x), are identical or x and m(x) are each a continuant part of the other at t. But it seems hard to imagine that x and m(x) might each be a continuant part of the other at t unless x = m(x). So, it appears that x = m(x). But now suppose at some later time, t+, a process occurs in which a site, s, comes into existence, such that we would ordinarily say that x has *gained* s for a continuant part. How can we characterize such a situation? If x has s for a continuant part at t+, then, given that x = m(x), m(x) also has s for a continuant part at t+, but, given (M), this is impossible. So, what can be said instead? All available options have drawbacks. For example, one could just insist that x does not have s for a continuant part at t+. But this seems unmotivated in the absence of a general background assumption that material entities never have sites for continuant parts, and this background assumption would on its own handle the difficulty raised by *canny* and *hully*. Alternatively, one could reject [scr-1], and hold (say) that x is not identical to m(x) *simpliciter* but rather is identical to m(x) *at t* and is not identical to m(x) at t+, but this would be a very radical proposal. Or one could say that at t+, neither x nor m(x) exists,

and instead new entities, x* and m(x*) exist, such that x* has s for a continuant part at t+ and x* has m(x*) for its matter at t+. But this too would be a radical proposal, requiring one to hold that under many circumstances entities do not persist through time that seem to do so.

We are not sure how best to define *is-matter-of*. Nevertheless, we think that (M) is at least close to true, and that fiat surfaces are best described as located where a material entity's matter meets its surroundings, not where the material entity meets its surroundings. We help ourselves to an *is-matter-of* relation (as well as to a *matter* universal, entirely parasitic on *is-matter-of* and introduced purely for the sake of convenience) in the formulation of some of our axioms. We leave the task of rigorously defending a definition of *is-matter-of* to future work.

One final point relevant to (D3). Some might agree with us that fiat surfaces are located where a material entity's matter meets its surroundings, but might say that it would nevertheless be helpful to have a term for a *continuant fiat boundary* located where a *material entity* meets its surroundings. Perhaps 'fiat material-entity boundary' would work as a term for such an entity. (We adapt this term from 'fiat object boundary', a term used by Arp et al. (2015: 108), seemingly to designate something at least very similar to what we have called a 'maximal fiat boundary.')

Finally, we turn to (D4). We make three brief suggestions about *Lone Star Bamboo* relevant to (D4). The first is fairly straightforward: We think that *maxy* exists at t*. In fact, we find the following principle, which would explain *maxy*'s existence at t*, very plausible: If some maximal fiat surface, s, hugs x at t, then s hugs x at *every* time at which x exists. Things don't "swap out" maximal fiat surfaces for other ones, or lose them altogether (without ceasing to exist). Second, we think that *texy* exists at t. To deny this would, it seems, require denying that there is anything very "fiat" about fiat surfaces. But third, we doubt that *texy* exists at t*. This is because we doubt that 'gerrymandered' fiat surfaces such as *texy* have any features that enable them to persist through changes of the sort at issue in *Lone Star Bamboo*. (By contrast, maximal fiat surfaces do have such a feature, namely *maximality*.) This is not to say that *no* non-maximal fiat surfaces have persistence-enabling features. Take a fiat surface exactly 'on top of' a person's birthmark at a given time. Arguably, *continuant parthood relative to exactly the top of the birthmark* enables persistence by this fiat surface through underlying dimensional transformations after t (such as the increase or decrease in size of the body part with the birthmark). We suggest that persistence through time by a fiat surface requires something enabling persistence. For gerrymandered fiat surfaces, such as *texy*, we just can't think of anything suitable to do the trick.

These remarks relevant to (D4) could be challenged, and in future work we will elaborate upon them. But we hope they at least help to bring out some issues connected to (D4) worthy of future exploration.

6. Axioms and BFO Proposals

On the basis of the considerations articulated in sections 2–5, we have produced two CLIF files (see https://github.com/michaelrabenberg/FOIS-2024-Rabenberg-Ceusters-Paper/tree/main) containing axioms concerned with fiat-surface, continuant-fiat-boundary, and some related BFO universals. The file 'boundary-yes.cl' contains axioms that we believe can safely be included in BFO, whether as such or as theorems of some other axioms worthy of inclusion. For example, here is an axiom in boundary-yes.cl:

```
(cl:comment "if something is a continuant fiat boundary at SOME temporal instant then
something is a material entity at THAT temporal instant [bbb-yes-1a]"
  (forall (a t) (if (and (instance-of a continuant-fiat-boundary t)
                         (instance-of t temporal-instant t))
                    (exists (b) (instance-of b material-entity t)))))
```

This axiom is not a theorem of *BFO-2020*, but we see no reasonable interpretation of (e-cfb) on which it does not turn out to be true, and the considerations of Section 2 independently motivate it. If this axiom is false, then revisions to (e-cfb) (and probably to (e-fs)) are required, and we just have no idea what continuant fiat boundaries (or fiat surfaces) are supposed to be. Hence this axiom's presence in boundary-yes.cl.

The second axiom file, 'boundary-maybe.cl', contains axioms that we believe ought to be considered for inclusion in future versions of BFO. Some of these axioms both authors of this paper accept; some one or both of us reject; about each of them at least one of us is sufficiently unsure that we do not think it appropriate for boundary-yes.cl. For example, the following axiom - which we take it to be a distillation of the core thought that we have attributed to Arp et al. and that we assumed for argument's sake in Section 5 - is in boundary-maybe.cl, because we take it to be worthy of consideration but are not sufficiently confident of its truth to include it in boundary-yes.cl:

```
(cl:comment "if something is a fiat surface at SOME temporal instant then it is
CONTINUANT PART OF a material entity at THAT temporal instant [bbb-maybe-
2a1]"
      (forall (a t) (if (and (instance-of a fiat-surface t)
                             (instance-of t temporal-instant t))
                        (exists (b) (and (instance-of b material-entity t)
                                         (continuant-part-of a b t))))))
```

Other axioms in boundary-maybe.cl include candidates for domain and range restrictions on the *hugs* relation, a formalized version of (M) (discussed in Section 5), and a variety of theses concerning fiat-surfaces' relations to material entities. Boundary-maybe.cl itself contains detailed descriptions of its contents, and the axioms are organized roughly into unified groups that facilitate inspection of them. We encourage interested readers to consider the axioms of boundary-maybe.cl themselves to see which they find plausible and which they do not. We hope these axioms will generate fruitful reflection and discussion.

Our testing has indicated no inconsistency between the axioms in boundary-yes.cl and *BFO-2020* in its current form, but we do not claim that the boundary-maybe.cl axioms are all collectively, or even all distributively, consistent with *BFO-2020*. We do not even claim that the boundary-maybe.cl axioms are *internally* consistent. But this is no problem for our purposes. Some axioms can be distributively worthy of consideration for inclusion in *BFO-2020* even if they are not consistent with *BFO-2020* in its current form and even if they are not consistent with each other, and we mean to claim about the boundary-maybe.cl axioms only that they are distributively worthy of consideration for inclusion in *BFO-2020*.

One final point. Recall that we suggested in Section 2 that although fiat surfaces hug material entities, non–fiat-surface two-dimensional continuant fiat boundaries—and not fiat surfaces—bear a corresponding relation to sites. (Perhaps it would be innocent to say

that this relation just is the *hugs* relation, and indeed boundary-maybe.cl contains an axiom to this effect, but we refrain from taking such a stance here.) This fact presents a challenge to (e-fs) as it is currently formulated. Consider some ship's cargo hold, CH, and a continuant fiat boundary, B, that relates to CH as a fiat surface hugging Barry and located precisely where Barry's matter meets its boundaries relates to Barry. B seems to meet the description of a fiat surface contained in (e-fs). This suggests that (e-fs) ought to be revised in a way that excludes entities such as B. Furthermore, for similar reasons, it might be that a *two-dimensional self-connected continuant-fiat-boundary* universal ought to be inserted 'between' the *continuant-fiat-boundary* and *fiat-surface* universals in the BFO hierarchy. Interestingly, though, no analogous difficulty seems to apply to *fiat-point* and *fiat-line*. It really does seem that everything that meets the description of a fiat point in the *fiat-point* elucidation is a fiat point, and similarly *mutatis mutandis* for *fiat-line*; and it really does seem that every zero-dimensional self-connected continuant fiat boundary is a fiat point, and similarly *mutatis mutandis* for fiat line. So it would seem to be pointless, and indeed misleading, to add the universals *zero-dimensional self-connected continuant-fiat-boundary* and *one-dimensional self-connected continuant-fiat-boundary* to BFO.

Acknowledgments

The work described is funded in part by grant T15-LM012495 of the NLM and grant UL1 TR001412/TR/NCATS of the NIH.

References

[1] Beverley J, Rudnicki R, Ruttenberg A. BFO-2020 GitHub Repository 2024. Available from: https://github.com/BFO-ontology/BFO-2020.
[2] Scheuermann RH, Ceusters W, Smith B. Toward an Ontological Treatment of Disease and Diagnosis. Summit on Translational Bioinformatics. 2009:116-20.
[3] Smith B. Basic Formal Ontology 2.0: Specification and User's Guide2015.
[4] Ceusters W, Pengput A. Axiomatizing SNOMED CT Disorders: Should there be Room for Interpretation? Formal Ontology in Information Systems (FOIS 2023); July 17- 20, 2023; Sherbrooke, Quebec, Canada.
[5] Schulz S, Case JT, Hendler P, Karlsson D, Lawley M, Cornet R, et al. SNOMED CT and Basic Formal Ontology — Convergence or Contradiction between Standards? The Case of "Clinical Finding". Applied Ontology. 2023;18:207-37.
[6] Arp R, Smith B, Spear AD. Building Ontologies with Basic Formal Ontology. Cambridge, Massachusetts: MIT Press; 2015.
[7] Varzi A. Boundary 2023. Available from: https://plato.stanford.edu/entries/boundary/.
[8] Price HH. Perception. London: Methuen; 1932.