Philosophical Studies
https://doi.org/10.1007/s11098-024-02171-1

Locative grounding harmony

Samuel Baron[1] · **Kristie Miller**[2] · **Jonathan Tallant**[3]

Accepted: 5 May 2024

Abstract
In this paper, we explore locative grounding harmony, according to which the location of the grounds mirrors the location of the grounded. We proceed in three stages. First, we clarify the notion of locative harmony and describe different locative harmony principles. Second, we offer two arguments for the claim that grounding between physically located entities obeys principles of locative harmony. Third, we consider and respond to a range of cases that seem to show that grounding relations between physically located entities do not obey such principles. We conclude that grounding between such entities obeys locative harmony.

Keyword Grounding · Harmony · Location

1 Introduction

Grounding harmony is the idea that there is a 'mirroring' between grounding and location. Mereological harmony is oft discussed, grounding harmony not so.[1] This is surprising. There is a substantial literature that concerns itself with determining what kind of relation grounding might be. Whether or not grounding is harmonious fits naturally into this investigation.

[1] For discussion of mereological harmony see Casati and Varzi (1999), Donnelly (2004), Saucedo (2011), Uzquiano (2011) and Varzi (2007).

✉ Jonathan Tallant
Jonathan.tallant@nottingham.ac.uk

Samuel Baron
s.baron@unimelb.edu.au

Kristie Miller
kristie.miller@sydney.edu.au

[1] School of Historical and Philosophical Studies, University of Melbourne, Parkville, Australia

[2] School of Philosophical and Historical Inquiry and The Centre for Time, The University of Sydney, Camperdown, Australia

[3] Department of Philosophy, University of Nottingham, Nottingham, England

Published online: 01 July 2024

There are two ways to spell out the notion of grounding harmony. The first option closely mimics mereological harmony—the idea that the mereological structure of objects is reflected in the mereological structure of locations. According to one prominent principle of mereological harmony, x is a part of y only if x's location is a part of y's.

Structural grounding harmony is the analogous idea; the grounding structure of objects is reflected in the grounding structure of their locations. So, we might say: x grounds y only if x's location grounds y's.

The second notion of grounding harmony departs substantially from the mereological notion. *Locative grounding harmony* is the idea that the locations of the grounds mirror the locations of the grounded and says nothing about grounding relations between locations. In this paper, we focus on locative harmony. Henceforth, when we talk of grounding harmony or of grounding being a harmonious relation, it is locative harmony that we have in mind.

Our aim is to first clarify, and then defend, locative grounding harmony as a feature of grounding relations between physically located entities. Note that our claim is not that locative grounding harmony is a feature of all grounding relations; our arguments are restricted to what might be called 'physical grounding': grounding between physical entities. In Section Two we outline a number of different principles of locative grounding harmony. In Section Three we provide two arguments in favour of the claim that grounding between physically located entities obeys principles of locative grounding harmony. In Section Four we consider a range of purported cases of grounding, which appear to be counterexamples to locative grounding harmony for physical entities. We argue that these cases fail to undermine the harmony of grounding. We thus conclude that grounding between physically located entities obeys locative harmony.

2 Harmony principles

This section of the paper contains a number of steps. We start by drawing a distinction between different kinds of grounding, and by clarifying how we are thinking about grounding in general. We move then to the articulation of different *kinds* of locative grounding principles, and draw out some implications. These enable us to narrow in on one type of locative principle in particular. We then use this to state a number of more specific principles.

Here, then, is the promised distinction: we differentiate whole from proper partial grounding. On our view, x wholly grounds y when x, on its own, is fully sufficient for grounding y. By contrast, we assume in this paper that x is a proper partial ground of y when x on its own is not sufficient for y, but x plus some plurality P is sufficient to ground y.[2] We will use '/' to represent whole or full grounding, and '//'

[2] It is worth noting that some take partial grounding to not be augmentable, and so deny that a partial ground is always a part of a plurality that constitutes a full ground (see Trogdon and Witmer (2021)). Nothing we say here hangs on how we specify partial grounding. It is possible to formulate the locative harmony principles for grounding that we outline below in terms of the Trogdon and Witmer account.

to represent proper partial grounding. Thus, 'x/y' is to be read as 'x wholly grounds y' and 'x//y' is to be read as 'x is a proper partial ground of y'. We assume that full grounding forms a strict partial order over the totality of existing things, and it is therefore asymmetric, transitive and irreflexive.

Next: the promised clarification vis a vis grounding. First, there is substantial disagreement over whether grounding is a relation. On some views, grounding is a sentential connective, and not a relation. Since what we say here focuses on the relata of grounding relations, it does not apply straightforwardly to a sentential conception of grounding. That being said, there is scope to extend what we say here by introducing a second sentential connective for location, one that is analogous to grounding.[3] Then our discussion of the relationship between grounding and location in terms of relata can be transposed into a sentential key. For present purposes, though, we will work with the idea that grounding is a relation, since it is admittedly more obvious how grounding so construed might connect to location (which typically is also construed as a relation). Second, even within those who take grounding to be a relation, there is disagreement about what the relata of the grounding relation are. We do not take a firm stand on this issue. However, as we note in Sect. 2.1, we are interested only in cases of grounding where both ground and grounded are physically located. That focus means that some putative relata fall outside of our purview (such as grounding relations between abstract objects).

2.1 Kinds of grounding harmony

Broadly, principles of grounding harmony can take one of three forms. These general forms can be specified in terms of the 'R'-relation, where being 'R-related' is just a placeholder for some more specific locative relationship between the location of the grounds and the location of the grounded, such as overlap or identity. This relationship could be non-symmetrical (as in, x's location is a sub-region of y's) or asymmetrical (as in x's location is a proper sub-region of y's) or symmetrical (as in x's location and y's location wholly overlap) and so we won't provide any further

Footnote 2 (Continued)

However, we focus on the augmentable account as it makes the ensuing discussion a bit neater. It is also worth noting that one way of defining a notion of partial grounding uses proper partial grounding: x is a partial ground of y when either x is a full ground of y or x plus some plurality P is a full ground of y. This is not the only definition available, however. For instance, some use 'partial ground' to refer to what we call proper partial ground. This terminological issue makes no difference in what follows.

[3] The norm in the literature on location is to treat the predicates 'is exactly located at' and 'is weakly located at' as key, treat one as primitive, and define one in terms of the other. In proposing a sentential connective for location, we suspect that we would adopt a scheme that adopted both of '…exactly where …' and '….weakly where…' as connectives and treat one as defined in terms of the other. To illustrate, 'x φs exactly where y φs' would be the canonical form. To be suitably permissive, the connective would have to be treated as allowing instances of the form 'x exists exactly where y exists'. This would enable us to capture the thought that an entity is exactly located at a spacetime point. The logic for the connective would need to be developed, but if we are right about the rest of what we say about grounding then we would expect to be able to develop work done in the case of grounding by the likes of Schneider (2011) and Correia (2010, 2016) to operate in the case of location.

constraints on the R-relation yet. We will develop more specific principles below, but here are the general forms:

(i) If x/y then x and y are located and their locations are R-related.
(ii) If x/y and either x is located or y is located then their locations are R-related
(iii) If x/y and x and y are located then their locations are R-related.

Principles like (i) and (ii) have implications for grounding that we may wish to avoid. In particular, locative principles of these types may be incompatible with putative instances of grounding.

Given a principle like (i), a grounding relation can only obtain between x and y if both x and y are located. This potentially rules out abstract objects standing in grounding relations to one another; for instance, it may rule out the existence of the singleton set {2} being grounded in the existence of the number 2, assuming both are abstract. It also rules out cases of grounding between located and non-located objects. Among other things, it may rule out impure set membership (as when we say that the existence of {Socrates} is grounded in the existence of Socrates) and truthmaking (since this involves a grounding relation between the truth of a proposition and some concrete aspect of reality). Whether these cases are in fact ruled out depends on whether abstract objects have locations, and (if so) what the relationship might be between their locations and the locations of physical objects. For now, we will just assume that abstract objects are not located (though see below the discussion of colour spaces).

Principles in the form of (ii) are similar though a little more forgiving. When x and y are not located the antecedent condition (which requires grounds or grounded to be located) cannot be satisfied, and so the principle is trivially true. As such, these principles allow for x to ground y when both x and y are abstract objects assuming, again, that abstract objects lack locations. However, principles in the form of (ii) still rule out grounding relations between located and non-located objects. For if, say, the grounds are located and the grounded is not, then the disjunction in the antecedent condition will be satisfied, while the consequent won't (given that the grounds and grounded cannot be locatively related if one of them lacks a location). Thus, a principle like (ii) would potentially rule out located objects from grounding abstract objects and vice versa, as we might expect to find in cases of impure set membership and some cases of truthmaking.

Principle (iii) is more forgiving still. It is trivially true for grounding relations between non-located objects. Moreover, it is trivially true in the case of grounding relations between located and non-located objects. Locative harmony principles like (iii) are thus constraints only on objects that are located. For any located objects that stand in a grounding relation, the locations of those objects must also stand in some further relation like overlap or identity.

We are interested in cases where grounds and grounded are located. We thus focus on principles like (iii) which are constraints on located entities. We take it that

there will be those who find (ii) to be plausible, or at least, well worth exploring,[4] but for now we set such principles aside. We also focus on cases in which grounds and grounded are physically located, where physical location includes (at least) spatiotemporal location and, perhaps, non-spatiotemporal physical locations (which may potentially arise in non-spatiotemporal theories of quantum gravity).[5] To make this clear, we thus reformulate (iii) under this restriction:

(iii)* If x/y and x and y are physically located then their physical locations are R-related. Does focusing on physically located entities reduce the import of what we have to say? We don't believe so. Grounding relations between physically located entities are an important class of grounding relations. It is plausible that such grounding relations include a large proportion of the grounding relations between physical entities, construed broadly to include material objects, instantiated properties, fields (including, potentially, both matter and metric fields) and events. A great deal of what we might call 'physical grounding', which captures grounding relations between physical entities, is thus subsumed under the broad category of grounding relations that we consider. These grounding relations are important, as they are likely to play a significant role in our best scientific theories, and are likely to contribute to our understanding of the physical universe that we inhabit. Moreover, it is plausible that grounding between such entities will bear some interesting connection to location. So, if one is interested in harmony principles for grounding, as we are, then grounding between physically located entities is a natural place to start.

We recognise that grounding is often used in a broader sense to include relations between facts. If facts are just constituted by physical entities, then our approach extends to grounding relations between such facts as well. If, however, facts are understood to include true propositions, or abstract objects in general (objects that are not physically located), then our approach does not include all facts. Our approach to grounding does not, therefore, extend to grounding between logical facts (see, e.g., Correia (2013) for discussion) and may not extend to normative facts either (see e.g., Väyrynen and Kirchin (2013)), if these are constituted by physical facts. Some cases of truthmaking are also not covered by what we say here, insofar as truthmaking is sometimes a relation between physical entities and abstract objects (namely, propositions, see e.g., Schaffer (2016)), which arguably lack any physical location. Both pure and impure set membership are also not covered, assuming that sets (even impure ones) are abstract objects and assuming, as seems plausible, that abstract objects lack physical locations.

[4] Principles like (ii) have their merits. If grounding obeys principles like (ii), then it would have more bite in metaphysical debates. Consider, for instance, the debate over dualism in the philosophy of mind. Even for a dualist, it is plausible that mental states are grounded in physical states. If mental states are non-physical however, then arguably they cannot have a location (assuming that x is located iff x is physical). Dualism of a certain kind would thus be ruled out by the locative harmony of grounding. Similar considerations hold for views on which abstract objects are grounded in concrete objects. We take it that some people may find (ii) plausible precisely because they are already committed to various metaphysical claims which rule out the sorts of cases that are ruled out by (ii). By contrast, those who find those cases plausible will be inclined to press the moral in the other direction and take, say, the viability of dualism as a counterexample to principles like (ii). Ultimately, the choice between (ii) and (iii) represents an important choice point.

[5] See, for discussion, Huggett and Wöthrich (2013, 2018).

Given that our discussion excludes a great many cases of grounding, should we be interested in harmony principles of the relevant kind? Again, we think so. Our study of grounding is analogous and to some extent inspired by similar studies of composition relations between entities with physical locations. This important class of composition relations includes what Healey (2013) calls 'physical composition', which plays a particularly important role in our best scientific theories, and plausibly includes a great many spatiotemporally located entities. Of course, it may be that properties, propositions, mathematical objects and other abstracta compose one another despite lacking physical locations (see, for instance, Paul (2002)). And so, in this sense, mereology is a broader study than just the study of composition relations between physically located entities. Nonetheless, there is a great deal of work that focuses on the composition relation between physically located entities. Thus, what's good for composition in this case is, we submit, good for grounding. In both cases, the class of physically located entities is an important class to consider.

We should note, however, that there is scope to extend what we say here to grounding in general, and thus beyond the case of grounding between physically located entities. Thus far we have assumed that abstract objects lack locations. While they certainly lack *physical* locations, it is possible that they possess locations in some other sense. For instance, following Wetzel (2009), we might look to the notion of an abstract space. Wetzel (2009: 132) considers the idea of an abstract space, within which universals may reside (2009: 144)—with colour properties residing within a colour space being one possible example. If one were tempted by that approach, then it could be argued that colour properties—in the sense of abstract universals—reside in colour space, and so have location in some sense. Moreover, it seems relatively clear that there can be grounding relations between colour properties. For instance, the determinate crimson seems to be grounded in the determinable redness. Thus, the question of whether locative grounding harmony holds for abstract objects arises.

Full discussion of this kind of case is beyond the scope of this paper. However, let us briefly note two points. First, it may be that locative grounding harmony is false for colour properties in a colour space. If the location of the determinate is disjoint from the location of the determinable then even very weak harmony principles will fail (e.g., the principle we call Partial Locative Overlap below). Second, and on the other hand, harmony principles may be preserved even in this case. It could be, for instance, that colour properties correspond to regions of a colour space, with determinables being located at the sum of the regions of their determinates. In this case, redness and crimsonness properties would not be disjoint. What emerges, then, is an interesting interaction between grounding harmony and the way we think about colour spaces. We flag this interesting issue as a topic for future work.[6] For now we return to the case of physically located entities, and press on.

[6] We are grateful to an anonymous referee for raising this issue.

2.2 Specific principles

In order to state specific principles of locative harmony we need to say more about location. We will start with a notion of exact location. Following Gilmore (2007), we assume that exact location is a primitive relation L(x, y) between an object x and a spatiotemporal region y. We also assume that x is exactly located at y when y shares the same size and shape as x, and stands in the same spatiotemporal relations to other objects as x does. We are thus effectively building physical location into location from the very beginning. We have chosen spatiotemporal location since, so far as anyone has discovered, that is the notion of location relevant for physical entities. That said, we are happy to leave it open that there are non-spatiotemporal forms of physical location as well, in which case our notion of location can be reformulated accordingly.[7] What matters is physical location, not spatiotemporal location per se. Going forward, we won't mention this restriction again unless necessary, leaving it implicit in the way we have understood location.

Using exact location we can then define a notion of weak location—WKL(x,y).[8] This notion can be defined in terms of exact location L(x,y) and a notion of mereological overlap O between locations (i.e., two locations overlap when they share a part in common). Thus, we can say that[9]:

$$WKL(x, y) =_{df} \exists z\left[L(x, z) \& O(z, y)\right]$$

The strongest system of locative harmony involves exact location, and treats the locative relationship between grounds and grounded in terms of identity:

> **Strong Co-location:** If x/y, and x and y are exactly located, then x's exact location = y's exact location.

Strong Co-location allows us to rule it out that x is grounded in y, when x and y fail to share an exact location. But if x and y do share an exact location, the locative relationship itself does not yield further information about priority.

[7] Following on from the discussion of location in a colour space above, one could also formulate location in a very general sense, perhaps along the lines of Correia (2022). Doing this would allow one to look at harmony principles beyond physical location. The principles we state below can be used for that purpose, though it is not yet clear whether any of them are constraints on ground tout court, rather than just what we've called 'physical grounding'.

[8] Weak location was introduced in Parsons (2007), though Parsons takes it as a primitive where, roughly, x is weakly located at a region if the region is not completely free of x.

[9] We take this definition of weak location from Parsons (2007, p. 204). As pointed out to us by a referee, however, it is controversial how weak location should be formulated. For instance, the above definition forces the principle that Correia (2022, p. 557) calls Exactness:

Exactness: Every object that is weakly located is exactly located.

As Correia notes, exactness does seem prima facie plausible for spatiotemporally located entities, which is a large part of what we are interested in here. However, we recognise that exactness may fail for some located entities. If so, then we are happy to shift toward the neutral framework of location outlined by Correia, in which both weak and exact location are defined in terms of entire location. So far as we can tell, this won't make a difference to the arguments we develop in this paper. Still, we recognise that it would be useful to reformulate our harmony principles inside Correia's system, and then study the implications for grounding. That is an interesting project, and one we intend to pursue in future work.

Note that Strong Co-location conditionalizes on exact location. Thus, it is only if x and y both have exact locations that the principle applies non-trivially. This is important since it is controversial as to whether *all* physical entities have exact locations. For instance, it has been argued that some physical entities have at best indeterminate locations, which may mean that they have no exact locations. Perhaps such entities are only located weakly, or perhaps they are not determinately located at all. By formulating Strong Co-location in a conditional sense, we leave it open that there are grounding relations between physical entities that lack determinate locations.[10] This is one of the advantages of formulating the harmony principles in the style of (iii), which conditionalizes on locations. In what follows, then, we will continue to formulate principles in this way so that we don't prejudge any important questions about grounding and indeterminate location.

There are two kinds of weakening that may be applied to the strong principle of co-location identified above. The first retains exact location as the operative notion of location and weakens the demand for identity between the exact locations of the grounds and the grounded, requiring only that the exact locations of grounds and grounded overlap. The second weakening involves trading in exact location for weak location.

2.2.1 Weakening the demand for identity

There are three principles based on the first kind of weakening available. To begin with:

> **Containing Grounded**: If x/y, and x and y are exactly located, then y's exact location is a sub-region of x's exact location.

We call this principle *containing grounded,* since it entails that the grounded is always contained within the region at which the grounds are exactly located. Containing Grounded entails that a grounded entity cannot 'spill outside' the exact location of its grounds, but it allows that there might be some region at which some of the grounds is located, and at which none of the grounded is located. Of course, ground and grounded sharing the same exact location is just the limiting case of the exact location of one being a sub-region of the exact location of the other (an improper sub-region) and hence is a limiting case of containment. Thus, Strong Co-location entails Containing Grounded, but not vice versa.

The converse principle to Containing Grounded says that the grounds must be where the grounded are but allows that the grounded can be where the grounds is not. This principle can be stated as follows:

> **Containing Grounds**: If x/y, and x and y are exactly located, then x's exact location is a sub-region of y's exact location.

[10] For discussion, see Calosi & Wilson (2019, 2021), Calosi & Mariani (2021), Calosi (2021), Darby (2010), Glick (2017) and Williams (2008).

Containing Grounds entails that if x grounds y, then everywhere we find x, we find y, but allows that y might 'spill outside' the region at which we find x. So, the grounded can overflow the region at which the grounds are exactly located. For example, if one thinks that holes are grounded in hole-linings, but that holes are located at places that hole-linings are not (i.e. that holes are what is contained inside the hole-linings), then one will think there can be such cases.

If one thinks both that grounds can spill outside the exact location of the grounded, and that the grounded can spill outside the exact location of the grounds, then one will have reason to adopt a still weaker principle which countenances both possibilities. For instance, one might think this if one thinks that holes are grounded in hole-linings, but that not only is the hole located at places that the hole-lining is not (as per the above suggestion), but that also the hole is not located at some places that the hole-lining is (because, one might think, holes are not located where matter is, and a hole-lining is located at such a place).

Locative Overlap: If x/y, and x and y are exactly located, then x's exact location and y's exact location overlap.

Each of the three weaker principles just outlined have straightforward analogues for the case of proper partial grounding. The proper partial grounding principles are, respectively, as follows:

Strong Partial Containing Grounds: If x//y, and x and y are exactly located, then x's exact location is a proper sub-region of y's exact location.

Strong Partial Containing Grounded: If x//y, and x and y are exactly located, then y's exact location is a proper sub-region of x's exact location.

Partial Locative Overlap: If x//y, and x and y are exactly located, then x's exact location and y's exact location overlap.

2.2.2 Weakening by trading exact for weak location

As noted, there are two ways to weaken the strong principles of locative harmony. The second involves trading exact for weak location. This yields the following principles:

Weak Co-location: If x/y, and x and y are weakly located, then for at least one of x's weak locations, z, and for at least one of y's weak locations, w, z = w.

Weak Partial Co-location: If x//y, and x and y are weakly located, then for at least one of x's weak locations, z, and for at least one of y's weak locations, w, z is a proper sub-region of w.

Note that these principles quantify existentially over weak locations. The reason for this approach is that entities usually have multiple weak locations, in contrast to exact locations, which are typically thought to be unique. Entities generally have multiple weak locations, because weak location (at least, on the definition used above) only requires partial overlap with some region, and entities partially overlap many regions. Exact location is a more demanding requirement than partial overlap, likely requiring at least whole overlap with some region, which is why entities don't have multiple exact locations.

2.2.3 Combining both approaches

These two kinds of weakening are compatible, so we can produce a very weak system of locative harmony by simultaneously weakening both the notion of location and the relationship between locations. These twice-weakened principles can be stated as follows:

Weak Containing Grounded: If x/y, and x and y are weakly located, then for at least one of y's weak locations, z, and for at least one of x's weak locations, w, z is a proper sub-region of w.

Weak Containing Grounds: If x/y, and x and y are weakly located, then for at least one of x's weak locations, z, and for at least one of y's weak locations, w, z is a proper sub-region of w.

Weak Locative Overlap: If x/y, and x and y are weakly located, then for at least one of x's weak locations, z, and for at least one of y's weak locations, w, z and w overlap.

Weak Partial Containing Grounds: If x//y, and x and y are weakly located, then for at least one of x's weak locations, z, and for at least one of y's weak locations, w, z is a proper sub-region of w.

Weak Partial Containing Grounded: If x//y, and x and y are weakly located, then for at least one of y's weak locations, z, and for at least one of x's weak locations, w, z is a proper sub-region of w.

Weak Partial Locative Overlap: If x//y, and x and y are weakly located, then for at least one of x's weak locations, z, and for at least one of y's weak locations, w, z and w overlap.

Each of these principles requires that the grounds and the grounded either share some weak location, or that they have weak locations that overlap. So long as the grounds and the grounded are physically located at all, they should satisfy these principles: weak locations of the relevant kind are very easy to come by. For

example, suppose that x/y and that x is exactly located at a, and that y is exactly located at b and that a and b are regions of spacetime. Then there is always some region of spacetime c, such that c ≠ a, c ≠ b and c overlaps both a and b. Since c just is a weak location for x and y, it follows that x and y share a weak location. There will also, in general, be a spacetime region d such that d ≠ a, d ≠ b and d overlaps y and has a sub-region that overlaps x (or vice versa).

Basically, for any pair of spacetime regions, we can always find a weak location of the kind that will satisfy the weak principles of locative harmony introduced above. This is true even if the two spacetime regions at issue are at opposite ends of the universe. For locative grounding harmony to be substantive, then, we need to be thinking in terms of exact locations, since exact locations are much harder to come by (and, depending on one's views about multi-location, each physical entity only has one). Thus, grounding obeys locative harmony in an interesting sense when at least one of: Strong Co-location, Strong Partial Co-location, Containing Grounds, Containing Grounded, Locative Overlap, Strong Partial Containing Grounds, Strong Partial Containing Grounded or Partial Locative Overlap is true.[11]

In what remains, we make the case that grounding between physically located entities *is* constrained by interesting principles of locative harmony, though to some extent we leave it open as to what these might be. We will thus set aside weak location and focus entirely on exact location. Our view is that grounding between physically located entities obeys *at least* Partial Locative Overlap. This is an important principle to test: if Partial Locative Overlap is false then all of the other principles are false as well (since they rely on there being at least some overlap between the location of the grounds and the location of the grounded). Thus, establishing that grounding between physically located entities obeys at least this principle shows that grounding in such cases is tied to location at some level. We leave it open as to whether grounding between physically located entities obeys a stronger principle of harmony. We suspect it does, but we recognise that some of the stronger principles conflict with existing views in metaphysics. For instance, as has been pointed out to us by a referee, Strong Partial Containing Grounds may conflict with the idea that wholes ground their parts.

Despite leaving much open, our argument still delivers a substantive outcome. It places non-trivial constraints on how grounding between physically located entities works. Indeed, one of our arguments in favour of Partial Locative Overlap shows that giving up even this weak constraint has untoward implications. What we say here thus provides a starting point for further discussion about the relationship between grounding and location for physical entities. Given that the principles we have identified have the potential to conflict with existing metaphysical claims, the further exploration of grounding harmony is important, and so the foundation we lay

[11] As a referee has pointed out to us, the weak principles may not be entirely trivial. For instance, these principles seem to rule out grounding relations between objects located in distinct worlds (within a concrete modal realist framework). That's because while all spatioemporally located objects in a single world will plausibly share a single weak location (and so grounds and grounded will always be weakly co-located in some sense), spatiotemporally located objects in distinct worlds won't share a single weak location, since the boundaries of weak location are specified by the boundaries of a single spacetime. We agree that there is potentially more to say here.

here will provide a useful basis for future discussion. The rest of the paper therefore proceeds in two stages. First, we argue directly in favour of Partial Locative Overlap. Second, we consider and respond to a range of apparent counterexamples to the idea that grounding between physical entities is constrained by locative harmony in this sense.

3 Arguments for locative harmony

As we see it, there are two arguments that speak in favour of the broad idea that grounding between physically located entities obeys some system of locative harmony principles.

3.1 Grounding and building relations

Grounding is intimately connected to a range of other 'building' relations. Candidate relations to which grounding is related include: part/whole relations, determinate/determinable relations, material constitution relations, functional realization, micro-based determination and emergence. There are, broadly, two approaches to understanding the intimate connection between grounding and other building relations. The first approach takes the connection to be that of unification. On this view, building relations are unified insofar as they are all instances of grounding. A second approach takes the connection between grounding and the building relations to be weaker. On this picture, some of the building relations are tracked by grounding, but not all of them are. Grounding is one broad type of building among a group of building relations. On either view, grounding is intimately connected to all or some building relations. We would thus expect any general patterns across building relations to transfer to grounding. Importantly, when we consider each of the relations we find that they all appear to obey a principle of locative harmony. This provides preliminary evidence that grounding does as well.

A point of clarification before we press on. Below we discuss a range of different building relations and note how each appears to bear an important connection to location. We assume, in each case, that both relata of the relations we discuss are physically located. It is thus compatible with our argument that building relations do not obey harmony for the case of non-physically located objects. This is fine for our purposes. What matters is that the building relations are harmonious in the case of physically located entities, since this can form an inductive base for inferring that grounding for physically located entities obeys a similar constraint. It is worth noting, also, that it is controversial as to whether all of the relations we consider are in fact instances of grounding. Indeed, for some of these relations, this has been denied in some quarters. However, this is only a problem if grounding is taken to be a unifying relation. But, as noted, our argument does not hang on this way of thinking about grounding. So long as there is a family resemblance between the relations we discuss, such that they can all reasonably be thought of as building relations, we

have some reason to suppose that general features possessed by a number of such relations will carry over to grounding as well.

Consider, first, mereology. For some it is a necessary truth about the part/whole relation that it obeys a group of harmony principles (see, e.g., Schaffer (2009)). As noted, however, these harmony principles are strictly stronger than the notion of locative harmony at issue, since they impute a relationship between the structure of objects and the structure of locations. Nonetheless, such principles do tend to suggest locative harmony as well. The whole is generally located where the parts are located, at least in cases of composition for physical entities. And, indeed, this seems intuitively correct. It is odd indeed to suppose that a physical entity has a part, without in at least some sense being located where the part is located.

Similar considerations apply to material constitution. A statue that is entirely constituted from a lump of clay appears to be located wherever the lump of clay is located; nowhere else. It is not the case, for instance, that the statue is located in a region that is completely free of any clay. Similarly, the lump of clay appears to be located wherever the statue that it entirely constitutes is located. Even if one does not believe that the location of the statue is exactly reflected in the location of the clay, it is difficult to imagine there being no locative relationship between the two, and so some notion of harmony seems to constrain material constitution.

Consider, next: determinate/determinable relations. Consider a square of crimson cloth that a pigeon is dutifully pecking. The instance of redness seems to be located where the instance of crimsonness is located: namely where the cloth is.[12] Can the location of the determinate and the determinable come apart? Well, clearly there can be some red where there is no crimson. But it is hard to see how there could be some crimson located somewhere without there being some redness located there as well.[13]

[12] Here, we suppose that the crimsonness of the cloth is at least partially located in the cloth itself. Something like this claim follows on most theories of properties, and something even stronger follows for others, like trope theory. On each of the main theories of properties, property instances of physical entities like pieces of cloth are located wherever the cloth is located. What theories of properties disagree about is whether properties construed as types are located where their instances are. An Aristotelian says 'yes' a Platonist says 'no'. We will set aside this debate, however, and focus just on property instances throughout this paper, and the grounding relations between them.

[13] Above we noted that determinates may not be located where their determinables are located. However, that discussion related to properties qua abstract objects. The case here relates to properties qua property instances. Our claim is that property instances of determinates are co-located, to some extent, with property instances of their determinables. It is compatible with this that determinates and determinables have disjoint locations in an abstract property space or that they locatively overlap. The locations of property instances in physical space is, so far as we can tell, orthogonal to the location of properties in a property space.

Functional realization also seems to be harmonious. Suppose that some version of physicalism about the mind is true and suppose that functionalism about mental states is also true.[14] On such a view, pain is given a functional characterisation in terms of the pain role: a network of causal inputs, outputs and interactions with other mental states. Now, suppose that what realizes the functional role for pain is a neural state (or a sequence of neural states over time). Where is the pain located? Again, it seems natural to suppose that the pain is located wherever its realizer is located and nowhere else.[15] For example, suppose that Martian physiology is not capable of realizing pain. It would be odd, then, to say that there are pains on Mars despite the absence of any realizers for pain.

Next, consider the connected notions of micro-based determination and emergence. Very roughly, micro-based determination involves the determination of some macro state by an underlying micro state, where a 'state' in this sense just refers to one or more material objects plus the properties they possess. For instance, the temperature of a body of water is determined by the mean kinetic energy of the H_2O molecules that make it up. Cases of micro-based determination seem to involve a strong locative relationship between the relevant macro and micro states. In the water case, it is natural to assume that the macro state—a body of water with a certain temperature—is located wherever the H_2O molecules that determine that state are jointly located. It is hard to imagine that the H_2O molecules are located in some region that is completely free of the relevant body of water. Similarly, it would be odd for the body of water to be located in some region that is completely free of H_2O molecules.

Some uses of the term 'emergence' pick out something like micro-based determination. Thus, what we have said for micro-based determination goes for those uses of emergence as well. But there is another notion of emergence that is sometimes used in the philosophical literature, and that goes by the name of 'strong emergence'. The difference between strong emergence and micro-based determination is, roughly, that in micro-based determination a description of the higher-level property or phenomenon is, in principle, deducible from a description of the lower-level property or phenomenon. In strong emergence, there is a failure of deducibility, and so the emergent property is generally thought to be something 'over and above' whatever the lower-level phenomenon might be.

[14] We use the term 'mental state' here loosely to include any number of physicalist approaches. Thus, a mental state could be a material object, a material object plus instantiated properties, a set of instantiated properties, an event and so on. Different accounts of what mental states are will have different implications for *where* mental states are located. However, by virtue of being physical objects mental states will always have some location (assuming, as seems plausible, that everything physical is located). What we say about the location of mental states holds regardless of what mental states are. No matter what mental states are, their locations seem to accord with the locations of the physical states to which they correspond in some sense (though what that sense is may be up for debate).

[15] Of course, one might resist the view that pains are mental states altogether, in favour of the view that pains are things that inhere in body parts (the so-called common-sense conception of pain). But then it will still be that the grounds of the pain (i.e. the body damage) and the pain itself, are located in the same place (namely the relevant body part).

It is controversial whether there are any genuine cases of strong emergence. One candidate case, however, is consciousness. Some believe that consciousness cannot be deduced from the underlying neural properties of brains and a set of functional characterisations of mental states and so there is an explanatory gap between neural realizers and conscious mental states. One response to the explanatory gap, of course, is to deny that conscious mental states are physical states, and thus that this is a case in which physical states come to possess non-physical properties. Some philosophers, however, are happy to accept that conscious mental states are physical states that cannot be deduced from underlying physical states and are thus strongly emergent. If one believes that the strongly emergent states are physical, it is reasonable to suppose that they are located in space and time. The question, then, is whether the location of the emergent states can come apart from the location of the underlying states from which they arise. Intuitively, again, the answer would seem to be 'no'.

In sum: part/whole relations, determinate/determinable relations, material constitution relations, functional realization, micro-based determination and emergence all appear to obey locative harmony in the case of physical entities. Because grounding is intimately connected to the building relations, we can expect this general pattern to extend to grounding relations as well. Exactly which one of the principles considered in Sect. 2 should we conclude constrains grounding? Because the argument is general in nature, it does not provide much guidance here. Thus, we think that *at best* a conservative conclusion is warranted; grounding for physically located entities obeys the weakest of the harmony principles that we've outlined which, as discussed, is Partial Locative Overlap.

We can also push the moral in the other direction. According to Koslicki (2015, p. 315), there aren't any features that all building relations have in common and so it is implausible to suppose that they are all instances of a single core relation (see also Wilson (2014)). This is taken as evidence against the idea that grounding unifies the building relations. On the contrary, there does seem to be a feature common to grounding relations between physically located entities: locative harmony. Thus, not only is it plausible that grounding is harmonious because the building relations are, harmony can potentially give us a better handle on what unifies building relations in the first place, and thus provides a preliminary answer to Koslicki's concern. Perhaps they are all unified for the case of physical entities at least by virtue of satisfying (at least) Partial Locative Overlap. This point is strongest if we take grounding relations to unify the class of building relations between physical entities. However, it applies even if grounding is just a member of a broad family of building relations between physical entities. For then there is still the question of in virtue of what the building relations constitute a family. Aside from being building relations, being harmonious in the case of physically located entities, may be one further defining feature.

3.2 Action at a distance

The second argument in favour of locative harmony for grounding focuses on action at a distance. The argument, in brief, is that if the weakest principle of grounding harmony is false, namely Partial Locative Overlap, then we get action at a distance on the cheap. So, we should accept at least this principle. Consider the following sufficient condition for action at a distance outlined by Frisch (2010, p. 666):

> (AD) If an event c lies in the causal past of an event e and there is a spacetime shell S surrounding e but not c such that the occurrence of e counterfactually depends on the occurrence of c even when we keep the physical state of S fixed (except for events on S that lie in the causal future of e), then c is a cause acting at a distance on e.

A spacetime shell around an event is a cross-section of the backwards light-cone of that event. Crucially, the shell around e will include one or more spacetime points, but won't include the spacetime point at which c is located. The physical state of the shell is just a specification of any objects within the spacetime region enclosed by S along with the properties of those objects. Thus, we can think of the shell as a slice of e's past light-cone that doesn't include c. In this way, a shell is supposed to act as a screen for c. If, when we hold the shell fixed, and intervene on c, something happens to e, then the causal action of c on e is not being mediated by events that are held fixed within the shell and in this way c is acting at a distance on e.

To see why the failure of Partial Locative Overlap would give rise to action at a distance, it is useful to first consider the stronger principle, Locative Overlap, which involves full grounds. This helps to warm us up to the example. It will then be straightforward to consider the case of Partial Locative Overlap. Suppose, then, that we have two events: e and c. Suppose, further, that c is a whole ground of e, but that Locative Overlap is false. If Locative Overlap is false, then it is possible for e and c to be exactly located at disjoint regions of spacetime a and b (i.e., Lea and Lcb and it is not the case that Oab). Let us suppose this is so and let us also stipulate that c is in e's causal past. Finally, let us assume that there is a shell S around e that does not include c and, moreover, that c does not ground e via anything in S. Rather, c is a *direct* ground of e, despite being located outside of a shell around e.

Now, consider an intervention on c that annihilates it. Given that c is a full and direct ground of e, e is thereby prevented, on the assumption that nothing can obtain in the absence of whatever grounds it. But then it seems that the following counterfactual is true:

[CF1] If c had not occurred, then e would not have occurred.

This counterfactual will be true even when we hold fixed the shell around e. For when we hold fixed the shell around e, we hold fixed none of the grounds of e, since all of the grounds of e are located at a region that is outside the relevant shell (namely, at b, where c is located) and c does not ground e via anything in S. So, it is not the case that by holding fixed the shell around e we are holding fixed something that could screen off the impact of removing c.

This gives us everything we need for action at a distance in Frisch's sense: c lies in e's causal past, and is such that, holding fixed a shell around e, if c had not occurred, e would not have occurred because c fully and directly grounds e. It follows that c is acting at a distance on e.

We can extend this to proper partial grounding as well. Suppose, as before, that c is in e's causal past, and that c and e are located at disjoint locations a and b. Suppose, further, that c is a partial ground of e, in this sense: if c does not exist, then neither does e.[16] Next, suppose that Partial Locative Overlap is false and, moreover, that the locations of c and e are in fact disjoint. As before, we can define a shell around e, S, which screens off e's causal past and does not include c. We can also suppose that there is no, d, such that c grounds e via d and d is located in the shell. Then, once we remove c, we must thereby remove e because the loss of c implies the loss of e. Moreover, this is true even when we hold fixed the shell, because c's grounding of e does not go via anything in S.

So far, we've shown that, if Partial Locative Overlap is false *and* the partial grounds and the grounded are located at disjoint locations, then we get a case of action at a distance. But, of course, it does not follow from Partial Locative Overlap being false that the grounds *always* occupy a location that is disjoint from the grounded. What it does do, however, is open up this possibility, and thus opens up the possibility of action at a distance. This possibility conflicts with relativity theory. The causal structure of spacetime apparently forces causation to be local, in this sense: causation between x and y is always mediated by a spatiotemporally local chain of events, where that chain is a continuous, timelike curve (see, e.g., Weinstein (2006) for discussion). If a cause is not spatiotemporally local to its effect—where we can understand this in the terms outlined by Frisch, as involving action on some e despite holding fixed a cross-section of e's backwards lightcone—then what we have is a case in which there is causation that is not mediated by a spatiotemporally local chain of events, and so does not proceed along a continuous, timelike curve. For if there were such a causal chain, then it would go through every cross-section of e's backwards lightcone between e and c. When holding fixed a shell on e between e and c, there should thus not be any sense in which changing c affects e, since all influence from c is mediated by events in the shell being held fixed. In the case we have considered, changing c does affect e despite holding a shell on e fixed, so there must be some influence from c that is not mediated by any event in e's backwards lightcone. This kind of action at a distance is not physically possible as it requires causal paths through spacetime not generally allowed by relativity.

One might remain unconvinced: doesn't quantum mechanics require action at a distance? If so, isn't action at a distance something we may be forced to accept?

[16] In line with our discussion of partial grounding in fn. 3, we do not require that the partial grounds are augmentable. What matters for our argument here is that e's existence implies the existence of one of its partial grounds. Though, as we shall see below, the relationship doesn't even have to be as strong as implication. If the destruction of c might lead to the destruction of e, that is sufficient for action at a distance to be possible.

But it is extremely controversial whether quantum mechanics requires causal action at a distance. Granted, the theory does seem to require non-locality in some sense. But it is not widely agreed that quantum non-locality is causal action at a distance, precisely because that would conflict in a deep way with relativity (and quantum mechanics is compatible with, at least, special relativity, in which we have fixed relativistic causal structure).

One might still be unconvinced: doesn't even relativity require causal action at a distance in the form of gravitational influence? The short answer is 'no'. Arguably, gravitation in a Newtonian context can be thought of as a form of action at a distance. But in relativity, this is implausible. Gravitation is always mediated by spatiotemporally local influences. For instance, gravity waves—which seem to travel at the speed of light—provide one example of how gravitational influence propagates locally, and does not involve action at a distance.

We anticipate three further objections: that this is grounding at a distance rather than causation at a distance; that we're over-reliant on a particular model of action at a distance; that we're depending upon a controversial notion of grounding as necessitation. We take each in turn.

First, objection: the kind of action at a distance we get is *grounding* at a distance, rather than *causation* at a distance, and so is not something we should be worried about. This, however, is not so: the action at issue is causal. To bring out the causal nature of the case, let us suppose that there is a possible cause x that brings about a change in the grounds, c. Let us also suppose that there is an effect, y, which is brought about by a change in the grounded entity, e. Specifically, let x be a cause that destroys the grounds c, and y be an effect that is brought about by the destruction of the grounded entity e. It follows that x causes y via the grounding chain between c and e. Moreover, the causation between x and y in this case is not mediated via any events that are local to y. To see this, assume that y is in the timelike future of e. Then simply extend the shell on e (the one that we held fixed in the example above) to produce a cross-section of y's backwards lightcone. Holding this shell on y fixed, x still makes a difference to y. That's because x switches c off, which switches e off, which switches y off, and this is true even when we hold fixed a shell on y between x and y (so long as we don't hold fixed a shell containing e, which we can do in this case).

The causal action is thus mediated by a direct grounding chain. One might respond that causation is not mediated by grounding chains in this way, and so the example is not probative. But that is too quick. Consider a diamond and its molecular structure. The diamond is hard, and this hardness is grounded in the tight lattice structure of the diamond's molecules. It seems clear that we can bring about an effect on the diamond's hardness via causation on the diamond's molecular structure. For instance, suppose we shoot a high-energy stream of particles at the diamond's molecular lattice, destabilising it. This results in the diamond becoming brittle, and thus no longer hard. Here the cause, which operates on the diamond's molecular structure, is mediated by the grounding chain between the molecular structure and the diamond. Nonetheless, we still consider this to be a case of causal action. Thus, in the example of action at a distance via grounding considered above, the mere use of the grounding chain does not reclassify the case as an instance of

grounding at a distance. It remains causal action at a distance mediated by direct grounding.

Second objection: one might argue that the way in which we've managed to identify causal action at a distance hangs on Frisch's particular definition. Surely, there are *other* definitions that don't have the same consequence. In fact, our argument works for a range of definitions. Here are four drawn from the literature on the topic:

1. … the state on a spacelike patch Σ determines the state on the future domain of dependence $D+(\Sigma)$ consisting of the spacetime points p such that every past inextendible causal (i.e., timelike or null) curve through p intersects Σ. This determination indicates that the state of any field at $p \in D+(\Sigma)$ cannot be influenced by events so far away that an influence from them to p would have to be superluminal. (Butterfield, 2007, pp. 301–302)
2. What is action-at-distance? That, it turns out, is a difficult question, but here is a sufficient condition (SC): we have a case of it if we have a cause, at least one of whose effects is not connected to it via any spatiotemporally causal chain. (Hall, 2002, p. 276)
3. If A and B are spatially distant (space-like separated) things, then an external influence on A has no immediate effect on B. (Berkovitz, 1998, pp. 204–205).
4. The idea is that if an external influence on A is to have any effect on B, that effect must propagate from A to B via some continuous physical process. Any such mediation between A and B must occur via some definite temporally ordered sequence of stages of this process. Such mediation cannot be instantaneous, and so an effect on B cannot occur at the same time as the external influence on A. (Healey, 1994, p. 352)

The case that we have already considered violates definitions 2 and 4. In definition 2, a 'spatiotemporally causal chain' means a sequence of events leading from the cause to the effect, such that any adjacent events in the sequence a and b are timelike related, i.e. a continuous timelike chain. In our example, we have a cause x that destroys the grounds c for an event e which brings about an effect y, *without* influencing any of the events in a shell around y. We know this because we can hold the shell fixed and still get the effect on y. The causal influence of x on y is thus not proceeding via a continuous timelike chain of events leading to y since if there were such a chain, then it would pass through every shell on y between x and y and so holding fixed a shell on y should wipe out x's influence on y.

Similar considerations lead to a violation of definition 4. In our example, when x destroys c leading to the destruction of e and thus the causation of y, there is no continuous physical process that propagates between x and y, because there is no temporally ordered sequence of stages that can constitute such a process. Again, such a process would require a continuous timelike chain of causally related events. For there to be such a chain, x would need to make some difference to the events in every slice across y's backwards lightcone between x and y. But x doesn't make such a difference, because there is always one shell on y that is indifferent to the influence

of x. Rather, x brings about y simply by wiping out the grounding chain between e and c.

The case we have outlined does not straightforwardly lead to a violation of definitions 1 and 3. The reason for this is that our case involves a kind of action at a temporal distance (Adlam, 2018). The idea being that c is the grounds for e that lies in the past of e, and is such that a cause on x wipes out c, leading to an effect y that follows from the destruction of e, without operating via a continuous timelike chain of events between x and y. It is, however, straightforward to build a case that involves action at a spatial distance, which would violate definitions 1 and 3 as well.

Here's the case: suppose that c directly grounds e and c is spacelike separated from e. Now, suppose a cause x wipes out c. This immediately destroys e, which we may suppose has a follow-on effect of causing y, which we can set up to be also spacelike separated from c and x. In this case, x causes y. But because x and y are at a spacelike distance, any influence between them would have to be superluminal. This case of causation thus violates definition 1. This might not be obvious: when an event x destroys the grounds c (leading to the destruction of e which triggers event y), it is not clear that we 'send' anything from x to y in a way that we might intuitively associate with something travelling faster than light. But this thought interprets the constraint on superluminal influence too literally. The constraint is not really about sending things from x to y. Rather, the constraint is just another way of saying that causation between x and y should be mediated by a continuous, timelike chain of causes. In our example, if x causes y, then any mediating chain between x and y would need to be along the kind of path that would be traced by something travelling faster than light, since x and y are not timelike related. The point, then, is that because nothing travels faster than light, there are no such paths, and so there is no way for there to be a continuous chain of causes that mediates the action of x on y.

The spacelike case also violates definition 3. That definition tells us that there is no immediate causal action between spacelike separated events. What this means, effectively, is that any causal action between spacelike separated events must be mediate, and in particular mediated by a continuous timelike chain of events. However, as we've already seen, there is no such chain. Rather, the causal action is mediated by a direct grounding chain. Indeed, we can fix that there is no chain between x and y other than the grounding chain between c and e and *still* we get causal action. So, definition 3 is violated: there is immediate action of x and y, in the sense of there being no continuous timelike chain of events that mediates the action.

Other definitions in the literature are similarly violated if grounding is unconstrained by location (see, e.g., the definitions discussed in Lange (2002, p. 15)). Is there some as yet undiscovered definition that does better? Perhaps, but we seriously doubt it. Definitions of action at a distance typically aim to capture a certain kind of violation to relativistic causal structure. In particular, such definitions aim to capture cases in which causation occurs between two events x and y but without the presence of a continuous timelike chain of causal intermediaries connecting x and y. If we aim to preserve this relativistic structure—and we submit one should because of the massive empirical success of relativity—then there must be *some* locative constraint on grounding between physical entities. For without such a constraint, there

can be grounding chains that are not continuous timelike chains but that can nonetheless mediate causal action, resulting in action at a distance.

This brings us to the third objection against our argument. Our argument, one might contend, relies on grounding necessitation, which is controversial. Grounding necessitation is the idea, roughly, that, if x grounds y, then necessarily, if x exists then y exists. But as far as we can see our argument doesn't presuppose this claim. The cases we considered are ones in which removing the grounds leads to the loss of the grounded, rather than cases in which removing the grounded leads to the loss of the grounds (which would rely on necessitation). Our approach relies only on counterfactual dependence: when x grounds y and x exists, then had x not existed, y would not have existed. Moreover, we do not require this in general. We only require there to be some possible cases in which this is true, and in which the grounds and the grounded are physically located entities that do not share locations.

Still, even this counterfactual connection might be deemed contentious. But, ultimately, we may not even require anything this strong. Suppose that when x grounds y the destruction of x can lead to the destruction of y. Even if this weaker claim is true, we can still build a case that violates the structure of relativity. For the mere fact that destroying the grounds can lead to the destruction of the grounded can be used to produce the possibility of action at a distance and, as noted, not only does relativity seem to imply that there is no causal action at a distance, it also seems to imply that there can't be any.

4 Counterexamples

We have outlined a system of locative harmony principles for grounding and have provided two arguments in favour of taking grounding between physically located entities to be constrained by harmony in some sense. We now consider some potential counterexamples to the idea that grounding between physically located entities obeys Partial Locative Overlap. A successful counterexample is one where the location of a grounded physical entity and the location of its whole physical grounds are completely disjoint. For then Partial Locative Overlap is falsified for physical entities. We call this a case of *disjoint grounding* and use it as a basis for testing the relationship between grounding and physical location.

4.1 Swiss cheese, flag poles and planets

The first case we will consider involves a piece of Swiss cheese. It is commonly thought that the holes in a piece of Swiss cheese depend on the cheese itself. This is an example of a more general grounding relation between boundaries and whatever they bound. In this case, the cheese is the boundary and the hole is the bounded. Intuitively, the hole is located wherever the Swiss cheese is not located and vice versa. More generally, boundaries appear to occupy locations that are disjoint from the bounded. At first glance, this appears to be a case of disjoint grounding.

However, while it is true that the Swiss cheese and its holes don't occupy the same spatial location, they must nonetheless share a location in another sense: they must share a temporal location. It cannot be, for instance, that the Swiss cheese is located only in 1922 and its holes are located only in 1985. The cheese and its holes are always temporally co-located. At best, then, this is a case of grounding at a *spatial* distance, i.e., grounding between entities that have disjoint spatial locations.[17] It is not a case of disjoint grounding proper.[18]

A similar response is available for two other cases that have been discussed in the literature (see Baron et al. (2019)). Consider, first, the flagpole and its shadow. Plausibly, the flagpole's shadow is grounded in the flagpole, plus the fact that the flagpole is illuminated, where the grounding is governed by laws concerning the rectilinear propagation of light. On the face of it, however, the location of the flagpole and the location of the shadow are completely disjoint. Thus, it seems we have a case of disjoint grounding. But the flagpole is, at best, spatially disjoint from the shadow: there is no time at which the flagpole's shadow exists, and the flagpole doesn't. So, the flagpole is not completely disjoint from the shadow in a locative sense.

Next, cast your mind back to when Pluto was a planet. Pluto's being a planet appears to be a fact about Pluto, which is located out in space. This fact, however, is grounded in a range of conventions laid down by the International Astronomical Union concerning what it takes to be a planet. The relevant conventions exist here on Earth (assuming, as seems plausible, that conventions are grounded in minds and thus people). It therefore appears that what grounds Pluto's planethood in space is located at a region disjoint from the location of what is grounded: namely, Pluto itself.

As with the flagpole and Swiss cheese case, however, the location of Pluto's planethood is, at best, spatially disjoint from the conventions that ground it. The planethood and the conventions have to be temporally coincident in order for Pluto to be a planet. Indeed, as soon as we change the conventions—which is precisely what actually happened—we end up preventing Pluto from being a planet. The conventions at issue need to exist at the same time as Pluto's planethood for grounding to occur.

[17] It may also be worth noting that a hole in a piece of Swiss cheese is surrounded by the matter making up the cheese—at least, this is true of instances of genuine holes, as opposed to indents in the surface of the cheese. Call such phenomena 'internal cavities'. An internal cavity of x would then be defined as a set of connected interior points of x. It may be possible to formulate grounding principles in terms of this and similar notions. We thank a referee for this point.

[18] A referee suggests that our reply violates a plausible principle on boundaries, namely:

Boundaries depend on the plurality of the interior parts (or on the interior, construed as the sum of such parts).

We don't see that this principle is violated. We agree that the boundary of a hole in the piece of Swiss cheese depends on the *spatially located* parts at each time that the Swiss cheese exists. Moreover, we agree that the boundary and the interior are located at disjoint spatial locations. What we deny is that this is all there is to the case: the boundary and the interior also have temporal locations, and these are never disjoint. That's compatible with the above principle and is all we need for our solution to work.

We anticipate the following reply. We have claimed that in the Swiss cheese, flagpole and Pluto cases the temporal locations of the grounds and the grounded overlap, despite there being no spatial overlap. One might argue, however, that this division between space and time is not properly scientific. We should be thinking only in terms of an invariant spatiotemporal interval. With respect to spatiotemporal regions, however, one might argue that we cannot even state the response given above, because we can't talk meaningfully of temporal locations as distinct from spatial ones.

There are three things to say here. First, if we fix a frame of reference it is still possible to make sense of spatial and temporal location. Once we fix a frame of reference, however, there will be temporal overlap between, for example, the Swiss cheese and its holes. If one is inclined to take 'location within a frame of reference' as not a genuine form of location, then this response won't work. We worry, however, that if we take 'location within a frame of reference' as not a genuine form of location, then we make trouble for a range of locative claims that seem perfectly reasonable, like: x and y happened at the same time. Granted, there's no frame-independent sense in which this is true, but there's a frame-dependent sense, and we often care about the frame-dependent notion.

Second, once we take account of the full grounds within each of the three cases discussed, it is far from clear that the location of the grounds is fully disjoint from the location of the grounded, even once locations are properly construed as spatiotemporal regions. Take the cheese and its holes. A hole is grounded in the cheesy boundary, to be sure, but it is also partly grounded in the empty region of spacetime that the boundary encloses (no region, no holes). The grounded hole is located at the empty spacetime region that grounds it. Assuming, as seems plausible, that spacetime regions are trivially located at themselves, it follows that the location of the hole overlaps the location of some of its partial grounds. Similarly, in the case of a shadow, the shadow needs to be grounded in the existence of a piece of non-reflective material that is located where the shadow is (the shadow just is the fact that this piece of material is dark). The shadow, then, partly overlaps its full grounds, even if some of its partial grounds are disjoint. Finally, Pluto's planethood is, presumably, partly grounded in the hunk of matter out there in space in addition to some spatially distant conventions. But this means that at least some of its grounds occupy a location that overlaps the grounded, even if not all do.

While these cases are not counterexamples to the locative harmony of grounding for physical entities, they do tend to suggest that a strong principle of co-location won't do, whereby exact physical locations between grounds and grounded must be identical. Rather, it seems that to take account of such cases, a looser locative principle is required. In particular, what we may need to do is weaken Partial Locative Overlap even further to require that some but not all of a physical entity's partial grounds locatively overlap the entity. However, we don't yet see a need to weaken the principle this far since, as noted, we think there is always temporal overlap in the above cases. But we offer it as a potential fall-back position.

4.2 Churches, plumbers and swamp people

The next group of cases are harder to deal with, since they seem to involve both temporal and spatial dislocation. Consider the case of a consecrated church. A building is a church only if there was a past time at which the relevant rites and rituals were performed on it. So, the ground for the *present* fact that a certain building is a church, has an exact location that is in the past and is thus temporally disjoint from the exact location of the building's current instantiation of church-hood. We can easily imagine that the consecration act happens at a different spatial location to the building (suppose it happens in the Vatican) thereby making the case one in which both spatial and temporal locations of grounds and grounded are disjoint.

There are two things to say here. First, as a referee notes, even if the consecration act and the church are spatially and temporally disjoint, it doesn't follow that all of the properly partial grounds of the church are disjoint from the church. For, presumably, being a church is at least partially grounded in the present building, which spatially and temporally overlaps with the church. That being so, at best the case at hand would be one in which some, but not all, of the properly partial grounds of the church are disjoint from the church. This is problem enough, however, since it would still undermine Partial Locative Overlap (though not the fall-back position noted above).[19]

This brings us to our second point. Whatever's going on with the church case, it's not a case of disjoint grounding. For if this were a case of disjoint grounding, then it would give rise to cases of action at a distance, of the kind described in the previous section. To see this, it is useful to spell out the church case in a bit more detail. Let us suppose that in 1990, there is a particular building—St. Brigid's—and that at a certain point in 1990, St. Brigid's has the property of being a church. Let us also suppose that, in 1562, there was a particular event that occurred involving a priest, that we'll call the consecration event, c. The grounding structure of the situation is thus: St. Brigid's being a church in 1990 is wholly and directly grounded in the 1562 consecration event c.

Now, all we need is for there to be some event at e_2 that occurs, in part, because St Brigid's is a church. So, for instance, suppose that two people start saying "I do" at t_1 and then finish at t_2. This then causes them to be married at t_2. Suppose the rules around marriage are such that it only occurs if the relevant 'I dos' occur in a church. Given this, we can intervene on c to bring it about that the two people are not married at t_2, despite holding fixed all of the physical facts at t_1. Thus, we can hold fixed the fact that at t_1 the couple says "I", and we can hold fixed all of the physical features of the building at that moment, and still prevent St Brigid's from being a church at t_2, since it's being a church at that moment does not depend on any of the physical facts at t_1 but, instead, depends entirely on 1562 facts involving consecration.

[19] Similar considerations apply to the case of Sara and Alastair below. There too we may suppose that some of the properly partial grounds are co-located with Sara and Alastair respectively. Still, one might try to argue that not all are, and so we think that there is a bit more to say in these cases.

Can the church case be modified to avoid this kind of outcome so that it still yields a case of disjoint grounding? One might think so. Let us stipulate that the consecration act happens at t_1, when the couple say "I" and then at t_2 the building becomes a church at the very moment that the couple say "do" bringing it about that they are married. In this situation, one might argue, we can't make sense of the core counterfactual underlying Frisch's definition of action at a distance. That definition tells us to hold fixed a spacetime shell around the target event, e (the marriage) and then wipe out whatever is supposed to be acting at a distance. But the grounds in this case occur within the shell around the target event that we are holding fixed, since they occur in the moment immediately prior to the marriage. Since we can't both hold fixed the grounds and intervene on the grounds, we can't get a case of action at a distance. But the grounds are nonetheless disjoint from the grounded, since they occur at a different time, and at a different location (in the Vatican).

That response is not quite right, however. What Frisch's definition tells us to do is hold fixed a spacetime shell around the target event e, which is a cross-section of e's causal past. The only events that are within e's causal past, however, are events that are located at points that can reach e at a speed equal to or slower than the speed of light. If, at t_1, the rites are performed in the Vatican and, at the very next moment, St Brigid's becomes a church at a different spatial location, then the Vatican event won't be in the causal past of the church-hood event.

This is easiest to see if we imagine that St Brigid's is in, say, Indonesia. There is no way to get from the Vatican to Indonesia in a moment without travelling faster than light. Even if the church is just outside the walls of the Vatican, however, there would still be a problem. Even if the church is a fraction of a centimetre from the walls, and the difference between t_1 and t_2 is the smallest possible temporal distance, the consecration event still wouldn't be in the causal past of St. Brigid's becoming a church. As far as we can see, the only way to ensure that the consecration event is in the causal past of the church-hood event is for the two events to occur at the same spatial location. Any two spatially distant events will need more than a moment for causation to occur (given the usual causal structure of spacetime).

Thus, it seems that the only way to set up the church case so that it doesn't yield action at a distance via grounding is to put the grounds at the same spatial location as the grounded. But then the case is really just a reversed version of the Swiss cheese case: the locations of the grounds and the grounded are spatially but not temporally overlapping. But there is still some overlap between the locations, and so we don't have a genuine case of disjoint grounding.

Our analysis of the church case helps us to deal with another case, due to Wilson (2020). Suppose that, in 2021, Alastair is human. Alastair's being human is grounded in his causal history, specifically that he is a descendent of a particular lineage that includes *Homo Sapiens*. Aspects of Alastair's causal history, one might argue, are spatially and temporally disjoint from the current fact that Alastair is human. For instance, his having a great-great-grandfather who is human is a partial ground of Alastair's being human now. On the face of it, then, this is a case of disjoint grounding, since at least some of the grounds are located at regions that don't overlap with Alastair's location.

As with St. Brigid's, we shouldn't view this as a case of disjoint grounding. To see this, suppose that at t_1, Alastair is given a test to see if he's human. The test is 'positive' if Alastair is in fact human, otherwise it will come up negative. At t_2, the test comes up positive. Now, suppose we draw a spacetime shell around Alastair, holding fixed all of the physical facts of the world between t_1 and Alastair's birth. Finally, we intervene on Alastair's ancestral lineage by replacing his great-grandfather with swampman: a being who is not descended from *Homo Sapiens* but who, instead, emerges *ex nihilo* from an old swamp in Oxford. By making this change, we should thereby bring it about that Alastair is not human. This, in turn, alters the outcome of the test. The test will show 'negative' at t_2, instead of 'positive'. In this way, it appears we have counterfactual dependence between the outcome of the test and the grounds of Alastair's humanity.

Of course, one oddity about this case is that if the physical facts at t_1 and t_2 remain unchanged throughout this process, we shouldn't expect the outcome of the test to change. We agree. But what this suggests is that there's something wrong with thinking of Alastair's humanity as being grounded in his causal history, rather than in his physical constitution. Alastair is human because of current physical facts about Alastair, not because of who his grandfather was. His humanity is due to having certain DNA, or a certain morphology *right now*. If that's right, however, then the grounds and the grounded are spatially and temporally co-located, and so there is no sense in which Alastair's humanity involves disjoint grounding.

One more case. Suppose that Sara is a retired dancer. Sara's now being a retired dancer appears to be grounded in a past fact about Sara: namely, that she used to dance, along with a present fact about Sara: namely, that she no-longer dances. We can further stipulate that Sara was a dancer in Spain but now lives in the UK and so the present fact to be grounded is spatially and temporally disjoint from the past fact that does the grounding.

This time let us suppose that at t_1, Sara stands before a machine. If Sara is a retired dancer, then at t_2 the machine will give Sara a million dollars. If Sara is not a retired dancer, then she gets nothing. She gets a million. Now, hold fixed the physical facts about t_1, including that Sara is standing before the machine. Next, alter Sara's past so that she was never a dancer. Then Sara won't get her million, because she does not have the property of being a retired dancer when she stands in front of the machine. Sara's getting the million at t_2 is counterfactually dependent on her past dancing, since that grounds her present status as a retired dancer.

As with the previous case, this one is odd because when we are asked to hold fixed all of the physical facts at t_1 it seems natural to suppose that we also hold fixed everything intrinsic about Sara. We should thereby hold fixed her status as a retired dancer if it's got anything to do with Sara's current intrinsic nature. But, in fact, Sara's being a retired dancer now doesn't depend on Sara's physical make-up now, or even her make-up in the moment at which she stands before the machine, except insofar as she doesn't dance. We can thus hold the physical facts fixed in the required manner (thereby holding fixed that she doesn't dance), and still (somewhat bizarrely) determine whether Sara receives the million dollars or not, by manipulating her past grounds.

What's gone wrong with the case is that we've mistakenly treated Sara's being a retired dancer as an intrinsic property of Sara. Sara's being a retired dancer is a relational fact between her current physical make-up and the past. As a non-intrinsic fact about Sara, there is nothing about Sara now that the machine could detect to give her a million dollars. The machine could still detect her status as a retired dancer. But it would need to have infallible access to the past in order to determine whether she stands in the right relation to it. But that access would likely occur via some causal route. If that's right, however, then when we hold fixed the causal past, we should still hold fixed the machine's way of accessing the past, and so the machine should still gain the information that Sara is a retired dancer even if we change the past facts about Sara. If Sara's being a retired dancer is a relational fact of the relevant kind, then there is no possibility of action at a distance in this case.

If, however, Sara's being a retired dancer is a relational fact, then it is no-longer clear that this is a genuine case of disjoint grounding. For it is quite plausible to suppose that relational facts are located wherever their relata are located. In this case, Sara's being a retired dancer is partly located in the present, and partly located in the past. But then it is not the case that the grounds are locatively disjoint from the grounded: the location of the grounded overlaps the location of its grounds.

These insights about Sara help us to draw out a general moral for putative cases of disjoint grounding. For many such cases, we are mistaking *relational* properties for *intrinsic* properties. Relational properties, however, are not counterexamples to locative grounding harmony, because they are plausibly spread across the locations of their relata. Moreover, for any apparent case involving intrinsic properties, we have good reason to analyse it in relational terms. For once we are dealing with intrinsic properties, it seems all too easy to use disjoint grounding as the basis for action at a distance.

4.3 Shortness of location

There is one final counterexample to consider, suggested to us by a referee. Consider the following four claims:

(i) Property instances are exactly located at the objects that instantiate them.
(ii) (Some) property instances are fully grounded in the objects that instantiate them.
(iii) Objects are exactly located at spacetime regions.
(iv) Objects and spacetime regions are disjoint.

Each of these four claims is defensible. Together, one might argue, they constitute a counterexample to grounding harmony for physical entities. Here's the argument. Let P be a property instance, O be an object and R be a spacetime region. Then:

a. O grounds P (from ii)
b. P is located at O, L(P, O) (from i)

c. O is located at some region R, L(R, O) (from iii)
d. P and O's locations overlap (from i, Locative Overlap)
e. O and R overlap (from b, c, d)
f. O and R don't overlap (from iv)

So, we have a contradiction.[20]

One response to the problem is to note that the combination of (i), (iii) and (iv) violates a plausible constraint on exact location, namely that exact locations are 'short' (Gilmore (2018)):

$$\text{Shortness of Location } \forall x \forall y \forall z((L(x, y) \& L(y, z)) \rightarrow y = z)$$

Consider: if L(P,O) and L(O,R), then by Shortness of Location, O = R. But then (iv) is rather implausible, since it is hard to see how objects can be both identical to regions and disjoint from them. On the face of it, everything should overlap itself completely. Given that the combination of (i), (iii) and (iv) violates shortness of exact location, and given that this principle follows from other plausible constraints on location, we have reason to reject the combination of views that yields a counterexample to Partial Locative Overlap. Put this way, Partial Locative Overlap turns out to cohere with other plausible principles of location, which provides some boost to the plausibility of that principle.

Of course, one may take (i)–(iv) to constitute a counterexample to both Shortness of Location and Partial Locative Overlap. That's fair enough, but now we must consider which package of views we should prefer: a particular approach to property instantiation captured by (i)–(iv) or a view of grounding and location that rules against such a picture. At this point, we can only emphasise the arguments we have already provided in favour of Partial Locative Overlap.

Indeed, we take it to be a positive feature of our account that it rules out *some* metaphysical approaches. For what this shows is that the question of whether grounding between physically located entities is constrained by harmony in even a weak sense has the potential to do substantial work in metaphysics, by ruling out certain positions and cohering with others. This shows that Partial Locative Overlap is a substantive principle, and that the arguments offered here in its favour are worth studying further.

5 Conclusion

It's time to take stock. We have outlined a range of different principles of locative grounding harmony for instances of grounding involving physically located objects, and have identified the most interesting of these (involving exact location). We have also offered two arguments in favour of the view that grounding relations between physical entities obey Partial Locative Overlap: (i) many of the relations unified by

[20] Although we have framed the example in terms of Locative Overlap, let us grant for the sake of argument that we can formulate it in terms of Partial Locative Overlap as well.

grounding are harmonious and so obey at least this principle and (ii) without Partial Locative Overlap, grounding can be used to get action at a distance. We then dealt with a number of potential counterexamples to Partial Locative Overlap. We argued either that such cases are not genuine cases of disjoint grounding or else that they are only if they allow for action at a distance. For the second kind of case we showed how to better understand the grounding structure of the cases in line with Partial Locative Overlap. We conclude that grounding for physically located entities obeys Partial Locative Overlap. This still leaves open the question of whether grounding of this type obeys any stronger harmony principles. It also leaves open the question of whether grounding *in general* (not just between physically located entities) obeys any principles of locative harmony at all.[21]

Acknowledgements The authors would like to thank two anonymous referees for this journal for their helpful and constructive feedback. They would also like to thank Freddie for not eating an earlier version of this paper.

Author contributions This is a joint endeavour, with each author contributing to each section of the paper.

Funding Professor Miller acknowledges the support provided by two grants: ARC DP180100105. ARC FT170100262.

Data availability No data was collected or generated in the production of this paper.

Declarations

Conflict of interest The authors have no competing interests to declare.

Ethical approval There are no ethical issues associated with this paper.

References

Adlam, E. (2018). Spooky action at a temporal distance. *Entropy, 20*(40), 1–20.

Baron, S., Miller, K., & Tallant, J. (2019). Grounding at a distance. *Philosophical Studies*. https://doi.org/10.1007/s11098-019-01374-1

Berkovitz, J. (1998). Aspects of quantum non-locality ii: Superluminal causation and relativity. *Studies in the History and Philosophy of Modern Physics, 29*(4), 509–545.

[21] Many thanks to two referees who helped to reshape the paper.

Butterfield, J. (2007). Reconsidering relativistic causality. *International Studies in the Philosophy of Science, 21*(3), 295–328.
Calosi, C., & Mariani, C. (2021). Quantum Indeterminacy. *Philosophy Compass, 16*, e12731.
Calosi, C., & Wilson, J. (2019). Quantum metaphysical indeterminacy. *Philosophical Studies, 176*, 2599–2627.
Calosi, C., & Wilson, J. (2021). Quantum indeterminacy and the double-slit experiment. *Philosophical Studies, 178*, 3291–3317.
Casati, R., & Varzi, A. (1999). *Parts and Places*. MIT Press.
Correia, F. (2013). Logical grounds. *Review of Symbolic Logic, 1*, 1–29.
Correia, F. (2016). On the logic of factual equivalence. *Review of Symbolic Logic, 9*, 103–122.
Correia, F. (2022). A general theory of location based on the notion of entire loation. *Journal of Philosophical Logic, 51*, 555–582.
Darby, G. (2010). Quantum mechanics and metaphysical indeterminacy. *Australasian Journal of Philosophy, 88*(2), 227–245.
Donnelly, M. (2004). A formal theory for reasoning about parthood, connection, and location. *Artificial Intelligence, 160*, 145–172.
Frisch, M. (2010). Causes, counterfactuals and non-locality. *Australasian Journal of Philosophy, 88*(4), 655–672.
Gilmore, C. (2007). Time Travel, Coinciding Objects, and Persistence. In D. Zimmerman (Ed.), *Oxford Studies in Metaphysics* (Vol. 3, pp. 177–198). Oxford University Press.
Gilmore, C. (2018) Location and Mereology. in E. N. Zalta (ed) *Stanford Encyclopedia of Philosophy*
Glick, D. (2017). Against quantum indeterminacy. *Thought, 6*, 204–213.
Hall, N. (2002). Non-locality on the cheap? A new problem for counterfactual analyses of causation. *Nous, 36*(2), 276–294.
Healey, R. (1994). Nonseperable processes and causal explanation. *Studies in History and Philosophy of Science, 25*(3), 337–374.
Healey, R. (2013). Physical composition. *Studies in History and Philosophy of Modern Physics, 44*, 48–62.
Huggett, N., & Wüthrich, C. (2013). Emergent spacetime and empirical (in)coherence. *Studies in the History and Philosophy of Science Part b: Studies in History and Philosophy of Modern Physics, 44*(3), 276–285.
Huggett, N., & Wüthrich, C. (2018). The (A)temporal emergence of spacetime. *Philosophy of Science, 85*, 1190–1203.
Koslicki, K. (2015). The Coarse-Grainedness of Grounding. In K. Bennett & D. Zimmerman (Eds.), *Oxford Studies in Metaphysics* (Vol. 9, pp. 306–344). Oxford University Press.
Lange, M. (2002). *An Introduction to the Philosophy of Physics*. Oxford University Press.
Parsons, J. (2007). Theories of Location. In D. Zimmerman (Ed.), *Oxford Studies in Metaphysics* (Vol. 3, pp. 201–232). Oxford: Oxford University Press.
Paul, L. A. (2002). Logical parts. *Nous, 36*(4), 578–596.
Saucedo, R. (2011). Parthood and Location. In K. Bennett & D. Zimmerman (Eds.), *Oxford Studies in Metaphysics* (Vol. Six, pp. 223–284). Oxford: Oxford University Press.
Schaffer, J. (2009). Spacetime the one substance. *Philosophical Studies, 145*, 131–148.
Schaffer, J. (2016). Grounding in the image of causation. *Philosophical Studies, 173*, 49–100.
Schnieder, B. (2011). A logic for 'because.' *The Review of Symbolic Logic, 4*, 445–465.
Trogdon, K., & Witmer, D. G. (2021). Full and partial grounding. *Journal of the American Philosophical Association, 7*(2), 252–271.
Uzquiano, G. (2011). Mereological Harmony. In K. Bennett & D. Zimmerman (Eds.), *Oxford Studies in Metaphysics* (Vol. Six, pp. 199–224). Oxford University Press.
Varzi, A. (2007). Spatial Reasoning and Ontology: Parts, Wholes, and Locations. In M. Aiello, I. E. Pratt-Hartmann, & J. van Benthem (Eds.), *Handbook of Spatial Logics* (pp. 945–1038). Springer-Verlag.
Väyrynen, P., & Kirchin, S. (2013). Normativity and grounding. *Proceedings of the Aristotelian Society, 87*, 155–178.
Weinstein, S. (2006). Superluminal signaling and relativity. *Synthese, 148*(2), 381–399.
Wetzel, L. (2009). *Types and Tokens: On Abstract Objects*. MIT Press.
Wilson, J. (2014). No work for a theory of grounding. *Inquiry, 57*(5–6), 535–579.

Wilson, A. (2020). Classifying Dependencies. In D. Glick, G. Darby, & A. Marmodoro (Eds.), *The Foundation of Reality: Fundamentality, Space and Time* (pp. 46–68). Oxford University Press.