Absolute representations and modern physics

Caspar Jacobs[*] and James Read[†]

Abstract

Famously, Adrian Moore has argued that absolute representations of reality are possible: that it is possible to represent reality from no particular point of view. Moreover, Moore believes that absolute representations are a *desideratum* of physics. Recently, however, debates in the philosophy of physics have arisen regarding the apparent *impossibility* of absolute representations of certain aspects of nature in light of our current best theories of physics. Throughout this article, we take gravitational energy as a particular case study of an aspect of nature that seemingly does not admit of an absolute representation. There is, therefore, a *prima facie* tension between Moore's *a priori* case on the one hand, and the state-of-play in modern physics on the other. This article overcomes this tension by demonstrating how, when formulated in the correct way, modern physics admits of an absolute representation of gravitational energy after all. In so doing, the article offers a detailed case study of Moore's argument for absolute representations, clarifying its structure and bringing it into contact with the distinction drawn by philosophers of physics between coordinate-freedom and coordinate-independence, as well as the philosophy of spacetime physics.

Contents

[*]University of Leiden, NL. c.a.jacobs@phil.leidenuniv.nl

[†]University of Oxford, UK. james.read@philosophy.ox.ac.uk

1 Introduction

> [I]t is the business of physicists, as it is the business of no other enquirers, to find some minimal set of concepts that can be used for the indirect endorsement of any true representation: evidence that the concepts physicists currently employ are inadequate for these purposes is evidence that they have further work to do. — Adrian Moore [1997, p. 75]

The search for 'absolute' representations of reality has been a staple in the history of philosophy, and it continues to be a central issue in contemporary philosophy of science.[1] Briefly, absolute representations are representations that are not from any point of view. In physics, absolute representations are often identified as coordinate-independent or coordinate-free ones—that is, representations that do not depend on a specific choice of reference frame or coordinate system.

In his monumental book *Points of View* [1997], Adrian Moore has given a very general argument in favour of the possibility of absolute representation, which states that because representations have to 'answer to reality' there must always exist an account of what makes a representation true that is not itself from any point of view. Furthermore, Moore, following Quine [1978], believes that it is up to physics to produce such absolute representations, as the quotation at the start of this article illustrates. But the possibility of absolute representations of certain aspects of nature continues to be a matter of dispute.

One particularly important case that has generated recent discussion is the status of gravitational stress-energy in general relativity (GR): momentum and energy carried by

[1] Much of the contemporary discussion goes back to Bernard Williams (*Descartes* [1978] and *Ethics and the Limits of Philosophy* [1985]); see also Thomas Nagel's *The View from Nowhere* [1986]. The idea of a metaphysically perspicuous representation of the world is related; on this, see Sider's *Writing the Book of the World* [2011].

the gravitational field itself, rather than by the matter located within spacetime. On the one hand, there are good *prima facie* reasons to believe that gravitational energy exists—for example, to secure a notion of local energy conservation. On the other hand, there seems to be no 'geometric object' on the spacetime manifold that represents this quantity.[2] The best option available is to represent gravitational energy by a so-called 'pseudotensor', but such an object is coordinate-dependent in a vicious sense to be made precise. If Moore is correct that it is the business of physicists to find absolute representations of reality, then this is surely an embarrassment for physics. Put differently, the absence of an absolute representation of gravitational energy is evidence that physicists (or indeed philosophers of physics!) "have further work to do".

In this article we propose a way out of the dilemma by way of a different, coordinate-independent representation of gravitational energy known as the 'Sparling form'. Although this representation was already proposed by László Szabados in 1991, it has received virtually no attention in the philosophical literature on gravitational energy. We argue that this proposal both offers a new solution to an important puzzle in the philosophy of physics, and (in turn) affords the resources with which to overcome any perceived tension here with Moore's *a priori* argument for absolute representations. Since modern physics is rife with other non-geometric objects of substantial importance (most notably spinors, which we discuss in §8), this article should be taken to constitute just one case study in a broader investigation into absolute representations and modern physics—and, therefore, as an invitation to a great deal of "further work".

The structure of the article is as follows. In §2, we present a detailed analysis of Moore's argument in *Points of View*. This analysis also highlights a lacuna in the argument, which leads to a slight weakening of its conclusion. In §3, we connect the notion of absolute representations to that of coordinate-independence in physics. First, we show that a representation need not be coordinate-*free* to represent nature absolutely, but rather only need be coordinate-*independent*. Then, we explain that such coordinate-independence is achieved if the mathematical representations of physical quantities are so-called 'geometric objects'. Crucially, the gravitational stress-energy pseudotensor is *not* a geometric object, as we discuss in §4. This leads to the central dilemma of this article. In §5, we present our solution to this dilemma, based on the work of Szabados. This proposal also helps to solve a distinct but related problem with the gravitational stress-energy pseudotensor, namely that it is not unique, as we show in §6. The proposal of this article provides an absolute representation of gravitational stress-energy, but it has certain counter-intuitive metaphysical implications. These implications are discussed in §7. §8 concludes.

[2] 'Geometric object' is a technical term, which we will define precisely later in the article.

2 Moore's Argument

Moore [1997] argues that absolute representations of reality—that is, representations of reality that are from no particular point of view—are invariably possible. Following Moore, by a 'representation' we mean anything which has content and which is true or false in virtue of the content that it has. Here is a reconstruction of Moore's argument for the claim that absolute representations are always possible:[3]

1. There exists a set C of possible representations that are integrable by simple addition such that for any pair of true possible representations r_1 and r_2, there exists an $R \in C$ part of which reveals how r_1 and r_2 are made true by reality.

2. For any representation r_1 from a point of view p_1, there exists another representation r_2 that is from an incompatible point of view p_2.

3. If r_1 and r_2 are from incompatible points of view p_1 and p_2, and part of R reveals how r_1 and r_2 are made true by reality, then R is neither from p_1 nor from p_2.

4. Therefore, there exists a set C of representations that are integrable by simple addition such that no element of C is from any point of view: C is *an absolute representation of reality*.

Let us comment on the premises one-by one. The first premise states that there exists a set C of true representations, such that these representations are "integrable by simple addition". This means that one can 'add' them together to form another true representation. The typical form of simple addition is *conjunction*: if r_1 and r_2 are true sentences of propositional logic, for instance, then their simple addition yields $r_1 \wedge r_2$. It is clear that not all true representations are integrable by simple addition. For example, if r_1 is an utterance of "It is raining" on Monday, and r_2 is an utterance of "It is dry" on Tuesday, then their conjunction "It is raining and it is dry" is *not* a true representation whether uttered on Monday or on Tuesday.

The members of C are supposed to "reveal" how pairs of true representations are "made true by reality". Consider first the weaker claim that for any pair of true representations r_1 and r_2, there is some true representation R that provides an account of the way in which r_1 and r_2 are made true by reality—no claim is yet made as to whether, for different pairs of representations, such accounts are integrable by simple addition. For Moore, this weaker claim just follows from the fact that r_1 and r_2 *are* made true

[3] In personal correspondence, Moore has endorsed this reconstruction.

by reality: "it means nothing to say that each of them is made true by reality unless it is possible, in principle, to produce a representation that reveals how" (p. 69).

The stronger claim that, for different pairs of true representations, the accounts that reveal how they are made true by reality are from the same point of view—and hence integrable by simple addition—follows from the fact that true representations are made true by "the same reality in every case". If that is the case, Moore says, then "not only must it be possible to provide an account of the kind just described for any possible true representation, but the part of this account that is used for the indirect endorsement of the representation must be combinable with every other such part into a single conception of reality—call it C" (p. 69).

So much for the first premise. The second premise is straightforward: if there were a point of view p such that no representation from an incompatible point of view were possible, then a representation from p would just amount to a representation from no point of view.

The third premise, however, is more controversial. Moore defends it as follows. Firstly, to reveal how r_1 and r_2 are made true by reality is to (indirectly) integrate them. Secondly, since r_1 and r_2 are from incompatible points of view, they are not integrable by simple addition—that is just what it means for their points of view to be incompatible. Therefore, in order for a representation R to integrate them somehow, it has to endorse one of them—say r_1—without adopting the associated point of view p_1. But this does not yet establish that R is from *neither* p_1 nor p_2. Moore goes on to claim that "[R's] treatment of r_1 and r_2 will be entirely symmetrical", but offers no further justification for this claim.[4] As it stands, then, the third premise is unjustified.

Here is an example to illustrate the lacuna. Consider again an utterance of "It is raining" on Monday and an utterance of "It is dry" on Tuesday. These are from incompatible points of view, so cannot be integrated by simple addition. Clearly, one could indirectly integrate them in a tenseless way: "'It is raining' is uttered on Monday and it is raining on Monday, and 'It is dry' is uttered on Tuesday and it is raining on Tuesday." But is is *also* possible to indirectly integrate these utterances from, say, the first one's point of view. For example, one can say: "'It is raining' was uttered yesterday, and yesterday it was raining; but 'It is dry' was uttered today, and today it is dry." This latter representation reveals how both utterances are made true by reality, but it does so from a particular temporal perspective.

It is possible to replace the third premise with a weaker version, namely that if r_1

[4]In personal correspondence, Moore has acknowledged that this is a lacuna in the argument as presented in *Points of View* [1997]. Moore offers a brief response to the issue in [Moore, forthcoming, fn. 8], which we discuss further below.

and r_2 are from incompatible points of view p_1 and p_2, and part of R reveals how r_1 and r_2 are made true by reality, then R is either not from p_1 or not from p_2—or from neither. From this premise a weaker conclusion follows:

4'. Therefore, there exists a representation of reality, C, such that all elements of C are from the same point of view.

Notice that this conclusion remains far from trivial! For it establishes that one can make sense of reality—all of reality—from one unified point of view. This perspective need not be privileged; there may exist unified representations of reality from many different, incompatible points of view. But (4') suffices to overcome any radical form of perspectivalism on which true representations from different points of view are fundamentally irreconcilable.

Moore has briefly responded to this objection in recent work [Moore, forthcoming, fn. 8]. In order to reveal how true representations from incompatible perspectives are made true by reality, Moore claims, one also has to reveal how their respective perspectives contribute towards their truth. It may seem as though one could show this from a particular point of view: on Tuesday, I can make sense of the difference in perspectives between "It is raining" uttered on Monday and "It is dry" uttered on Tuesday by appeal to the fact that these utterances were made *one day apart*. But, Moore argues, this account of the difference in perspectives itself presumes a certain perspective. If one were to account for the same difference in perspective on the preceding Sunday, that account would appeal to the fact that the utterances in question will be made one day apart. The availability of distinct explanations of the difference in perspective—that the utterance *were* made one day apart or that they *will be* made one day apart—belies the fact that it is the *same* difference in perspective in each case. Therefore, Moore concludes, it is impossible to truly reveal how these representations are made true by reality from any particular point of view.

We concur with Moore (in conversation) that more needs to be said at this point, but we will not belabour it further here. In any case, we affirm the claim that in order to reveal how representations from incompatible perspectives are made true by reality, one also has to account for the way in which their respective perspectives contribute towards their truth. We will see below that in the case of gravitational stress-energy, it seems as if this is impossible *even from a particular point of view*. In this case, then, one cannot even present a complete representation of physical reality from a unified point of view—let alone from no point of view! This threatens the weakened conclusion (4') and *a fortiori* the possibility of an absolute conception of reality.

3 Absolute Representations, Coordinate-Independence, and Geometric Objects

We now connect these abstract issues to contemporary physics. Firstly, we contrast the notion of absolute representations with the distinction between *coordinate-independence* and *coordinate-freedom* (§3.1). Secondly, we argue that the possibility of absolute/coordinate-independent representations requires that physics employs so-called *geometric objects* (§3.2). In the next section, we present a challenge to the possibility of an absolute representation of reality: the gravitational stress-energy pseudotensor is *not* a geometric object, and hence does not admit of an absolute representation.

3.1 Coordinate-Independence

A key way in which perspectives enter physics is via *coordinate systems*. We understand a coordinate system to be an assignment of tuples of real numbers to spacetime points in a way that respects the structure of spacetime (e.g. smoothness). Each coordinate system codifies a certain perspective. For instance, the 'lab frame' is associated to a coordinate system in which the laboratory is at rest, i.e. in which the spatial coordinates of the lab remain constant over time. The lab frame thus embodies the *perspective* of the lab. Claims that are made with respect to these coordinates—for instance: the ball moves at 50 km/h—are from the lab's perspective. They are incompatible with those from another frame, say that of a car that drives past the lab. From the perspective of the driver of the car, the ball moves at 20 km/h. Both claims are true *from their respective points of view*. But one cannot form another true representation by simple addition: "the ball moves at 50 km/h and also at 20 km/h" is necessarily false in any frame. The use of coordinates in physics thus seems to inhibit the *desideratum* of absolute representation.

The most straightforward way to reveal how a pair of coordinate-dependent representations is made true by reality is to proffer a coordinate-*free* representation that indirectly endorses the coordinate-dependent ones. Call the lab frame x, and the car frame x'. Then v_l, v_c and v_b denote the velocity of the lab, car and ball respectively from the lab's point of view, and v_l', v_c' and v_b' denote these same quantities from the car's point of view. Moreover, let $\mathbf{v}_{lb}$ denote the relative velocity between the lab and the ball, and likewise for other combinations. The boldface indicates that this quantity is *coordinate-invariant*, that is, that it does not depend on a choice of coordinates. In other words, velocity differences $\mathbf{v}$ are from no point of view. The quantity v_b stands to $\mathbf{v}_{lb}$ as the expression "tomorrow" uttered on Monday stands to the expression "the day after Monday" uttered whenever. It is possible to appeal to invariant quantities to reveal how claims from different perspectives are made true by reality:

> "The velocity of the ball is 50 km/h", uttered from the lab's point of view, refers to the velocity difference between the ball and the lab ($v_b \equiv \mathbf{v}_{lb}$); and the latter quantity is 50 km/h ($\mathbf{v}_{lb} = 50$). "The velocity of the ball is 20 km/h", uttered from the car's point of view, refers to the velocity difference between the car and the lab ($v_b' \equiv \mathbf{v}_{cb}$); and the latter quantity is 20 km/h ($\mathbf{v}_{cb} = 20$). The difference in perspective is accounted for by the fact that the car moves at 30 km/h with respect to the lab ($\mathbf{v}_{lc} = 30$).

(Here, of course, we are using the standard formula for non-relativistic addition of velocities.) This account does not depend on any choice of coordinates. Indeed, it does not even *presuppose* a coordinate system; it is truly coordinate-free. In order to vindicate Moore's argument for the possibility of absolute representations, it would seem that physics has to produce a coordinate-free representation of reality.

However, Wallace [2019] has pointed out that coordinate-independence does not require coordinate-freedom: there are coordinate-independent representations that are not coordinate-free. For an example, consider the claim that the velocity of the car is 30 km/h more than the velocity of the ball. This claim presumes a coordinate system, since "the velocity of the car" and "the velocity of the ball" are only well-defined with respect to a specific frame of reference. But it is true in *any* inertial system of coordinates, where inertial coordinates are those adapted to a reference frame that moves inertially. Insofar as inertial frames are concerned, then, this claim is coordinate-independent. Again, by appeal to such representations it is possible to proffer an account that reveals how claims from different perspectives are made true by reality:

> Let x'' denote an arbitrary reference frame. "The velocity of the ball is 50 km/h", uttered from the lab's point of view, refers to the difference between the velocity of the lab and the velocity of the ball ($v_b \equiv v_b'' - v_l''$); and that quantity is 50 km/h. "The velocity of the ball is 20 km/h", uttered from the car's point of view, refers to the difference between the velocity of the car and the velocity of the ball ($v_b' \equiv v_b'' - v_l''$); and that quantity is 20 km/h. The difference in perspective is accounted for by the fact that the difference between the velocity of the car and the lab is 30 km/h ($v_c'' - v_l'' = 30$).

This account, too, does not presume any frame's *particular* point of view—even though it does presume *some* frame's point of view.

There are thus two distinct ways in which one can rid physics of coordinate-dependence in order to obtain absolute representations. The first is to produce coordinate-*free* representations, whilst the second is to produce coordinate-independent representations whose truth does not depend on the adoption of any *particular* frame of reference.

We saw in the previous section that there is a weaker version of Moore's argument that does not establish the possibility of an absolute representation of reality, but only of a representation of reality—all of reality—that is from a one unified point of view. In the language of coordinates, this amounts to a representation from within a particular reference frame that can yet reveal how claims within different reference frames are made true by reality. In our example, one can choose to adopt the lab frame and express all velocities within that frame. It is then possible to account for the claim that the ball's velocity is 30 km/h *from the car's perspective* as follows: the car's velocity is 30 km/h ($v_c = 30$), and the ball's velocity from the perspective of the car is equal to the difference between the ball's velocity and the car's velocity ($v'_b \equiv v_b - v_c$). This is a coordinate-dependent representation, yet it resolves the incompatibility between the car's and the lab's point of view.

Of course, the lab's reference frame is in no way special. If it were then the above account would not *really* reveal how a pair of coordinate-dependent representations are made true by reality, since it does not mention the supposedly privileged role of the lab frame. This is just the upshot of what Bell [1976] calls the 'Lorentzian pedagogy': "the laws of physics in any one reference frame account for all physical phenomena, *including the observations of moving observers*." This piece of pedagogy entails that one can offer a complete representation of reality from *any* point of view within the range of perspectives represented by different choices of coordinates. For example, one could instead adopt the car's frame x' and account for the velocity measurements within the lab from that perspective. Crucially, this only works if none of these frames is in some way privileged.

Because one can supply these coordinate-dependent accounts from the perspective of any arbitrary frame, it is also possible to construct a coordinate-independent (yet not coordinate-free) account from them. The idea is to take the full collection of accounts from any frame of reference (their 'equivalence class') and present *that* as a unified representation of reality. Since this account does not privilege any particular frame, it is coordinate-independent; but since each of the accounts is from some frame, it is not coordinate-free. This is the essence of what is known as the 'Kleinian' approach to geometry, which identifies geometrical structures in terms of whatever remains invariant under a group of well-defined transformation rules (such as coordinate transformations).[5] The upshot of this approach is that one can 'ascend' from a unified conception of reality from any *arbitrary* perspective to a unified conception of reality from *no* particular perspective.

Perhaps Moore would criticise this construction on the basis that any such account does not fully explain the role of the difference in perspective between the lab and the car. Take the story from the lab's perspective. This account's explanation of the contribution of perspective appealed to the fact that the velocity difference between the lab and the car is 30 km/h ($v_c - v_l = 30$). But this fact itself is represented from the perspective of the lab. Yet the difference made by their respective perspectives transcends the lab's point of view: it is the same difference *whatever* perspective one adopts. This, of course, is just the lacuna between the weaker conclusion (4'), and Moore's intended conclusion (4).

Whether or not one believes that (4) or only (4') is justified, however, the difference here is moot: for the case in which we are interested, it seems impossible to offer a consistent account of the difference in perspective even from any arbitrary perspective. This means that it is *a fortiori* impossible to construct a truly absolute representation by the Kleinian approach. As we show below, this is the case when a theory is partially formulated in terms of non-geometric objects. When a theory posits non-geometric objects, and those objects are interpreted as physically real, it seems that the possibility of an absolute representation of physical reality is foreclosed. This is the challenge to Moore's argument—whether the argument's conclusion is taken as (4) or only (4')—that this article aims to answer.

3.2 Geometric Objects

We turn now to the definition of these geometric objects: a mathematical notion making its first appearance in the literature in [Nijenhuis, 1952, Schouten, 1954, Trautman, 1962, 1965]. It turns out that there are two definitions of a geometric object available in the literature: a 'traditional' and a 'modern' one. Since the former is somewhat more intuitive, we will rely on it in what follows.[6]

---

[5] For further details on this, see Wallace [2019].

[6] On the modern definition, geometric objects are defined as sections over natural bundles on a manifold. Let M be a base manifold, P a bundle over this manifold with projection map π, and let d be a diffeomorphism of M onto itself. Then P is a natural bundle iff d induces a unique diffeomorphism ϕ of P such that $\phi \circ \pi = \pi \circ d$. Let σ be a section of P, i.e. a function $\sigma : M \to P$ such that $\pi(\sigma(p)) = p$. Then σ is a geometric object iff P is a natural bundle. This means that to any diffeomorphism d of the base manifold, there is associated a unique transformation $\sigma \to \phi^*\sigma$, where ϕ^* is the pullback map of ϕ. Generally, on the modern conception a geometric object is defined over a natural bundle with a well-defined pull-back under diffeomorphisms of the base space. [Kolář et al., 1996]

The kinds of object in which we are generally interested in contemporary theoretical physics are *fields*. Broadly, a field assigns a value to each point on a manifold. The latter usually represents spacetime. Note that neither the manifold nor the field involve coordinates: both are typically characterised 'intrinsically'.[7] But suppose that one has defined a local coordinate system, x, on the manifold. One can then express the values of the field within these coordinates: call the result the field's *components* in a coordinate system. The field may generally have different components in different coordinate systems. For a slightly contrived example, consider a 'velocity field' which assigns a velocity value to each point in spacetime. The components of this field would depend on the velocity of the coordinate system itself (with respect to some arbitrarily chosen standard).

For some arbitrary point $p \in M$, consider a pair of arbitrary coordinate systems around p. According to the traditional definition, a geometric object consists of

1. a set of components (a set of N real numbers) for each coordinate system, and
2. a well-defined rule relating the components in the one coordinate system to the components in the other.

The transformation rule in question is 'well-defined' in the region of overlap only if it forms a group: (1) there is an 'identity' transformation that leaves every set of components the same; (2) for each coordinate transformation, there is an 'inverse' transformation that undoes the first; (3) coordinate transformations are transitive, so the successive application of well-defined coordinate transformations is itself a well-defined coordinate transformation; and (4) coordinate transformations are associative, so it doesn't matter whether one evaluates successive coordinate transformations 'from the left' or 'from the right'. For example, let x, x' and x'' denote three coordinate systems. O is a geometric object only if the transformation $O \to O''$ defines the same object as the transformation $O \to O' \to O''$. If this is not the case, then the object is *non-geometric*. To put this succinctly, let $(O')'$ denote the result of first applying a transformation to O from x to x', and then a transformation from x' to x''; and let O'' denote the result of applying a transformation to O from x directly to x''. Then for a geometric object, $(O')' = O''$.[8] We will show below that it is the failure of this property for non-geometric objects that precludes their representation from no particular point of view. This is important because physics does sometimes involve non-geometric objects

[7] Wallace [2019] has rightly pointed out that coordinates still lurk in the traditional, 'intrinsic' characterisation of a manifold. We'll set this aside here.

[8] For a clear discussion of this property of geometric objects, see [Duerr, 2019b].

that *prima facie* seem to represent physically real quantities—such as the gravitational stress-energy pseudotensor discussed in the next section.

Most fields used in contemporary physics are geometric objects: any vector or tensor field is geometric (this includes the metric tensor that determines the geometry of spacetime in relativity theory); in addition, however, certain non-tensorial objects, such as the Christoffel symbols, are also geometric. We would further like to point out that certain *bona fide* physical quantities are represented by *non*-geometric objects. The Yang-Mills field in fibre bundle formulations of electromagnetism, for example, is a non-geometric object, because (briefly) the $U(1)$ bundle is not soldered to the base space [Dewar, 2020]. Another example, relevant to particle physics, is that of spinor fields (on which see [Pitts, 2012]). We therefore believe that many recent commentators are too quick to presume without much justification that geometric objects are the *sine qua non* of modern physics.

Nevertheless, in many cases the presence of non-geometric objects is problematic. In particular, we claim that when physics involves non-geometric objects, and those non-geometric objects are taken to represent a physically real quantity, it becomes impossible to offer a unified account of physical reality—inclusive of the difference made by the perspective of different observers—from a single frame of reference *à la* the Lorentzian pedagogy. But physics does seem to involve non-geometric objects, for instance the gravitational stress-energy pseudotensor we discuss in detail below. This poses a two-fold challenge to Moore's claim that physics is in the business of finding absolute representations. Firstly, it seems to present a counterexample to the weaker conclusion (4') from the previous section. If it is not even possible to offer a unified account of reality from any one point of view, then radical perspectivalism looms. Secondly, we saw in the prevous subsection that one can construct a coordinate-independent account of reality from an equivalence class of coordinate-dependent ones that are from an arbitrary point of view. But if such an account of reality is impossible, then this construction likewise fails. Therefore, the presence of non-geometric objects also stands in the way of a certain recipe for the construction of a coordinate-independent representation of reality.

Let us explain the latter in more detail. Suppose that one has representations of O in x and x'. In order to offer an account that reveals how both representations are made true by reality, one could adopt a third point of view, x''. Part of that account must 'translate' the representation of O from x and x' respectively to x''. This results in a pair of representations which we will denote $(O')'$ and O''. The first challenge occurs here: because O is a non-geometric object, generally $(O')' \neq O''$. Therefore, even from the perspective of an arbitrary coordinate system x'', the representation of O in x and the representation of the same object O in x' seem to present reality differently—they make different claims about the value of O. To make this concrete: it is possible that, considered from x'', O as represented from x vanishes in a certain region whereas O as represented from x' does not vanish in that region. It seems that there is no way to reconcile such representations, and so one cannot reveal how both are made true by reality from an arbitrary point of view. It is therefore also impossible to take an equivalence class of coordinate-dependent accounts of reality in order to construct a coordinate-independent representation. This recipe fails because the transformation rules between coordinate representations of a non-geometric object are not well-behaved: they are not mutually coherent; they are not made true "by the same reality in every case". This is the second challenge.

Compare this to the case of velocities discussed before. The key fact here is that velocity transformations are additive: if the car frame x' moves at 30 km/h with respect to the lab frame x, and if an arbitrary third frame x'' moves with (say) 10 km/h with respect to the car frame, then the latter frame moves with $30 + 10 = 40$ km/h with respect to the lab frame. So, if v_b is 50 km/h, then v_b'' is 10 km/h; but equally, if v_b' is 20 km/h, then $(v_b')'$—the velocity of the ball with respect to the car *with respect to an arbitrary third frame*—is also 10 km/h. The result is the same in each case. Therefore, the transformation rules for velocities *are* coherent: one can freely switch from one perspective to another without contradiction. The representations from these different perspectives are made true by the same reality in each case.

To sum up the story so far, before we present our main case study: we started with Moore's argument for the possibility of an absolute representation of reality—that is, a representation of reality that is from no point of view. We noted that the argument as presented in Moore's book *Points of View* [1997] contains a lacuna, but that it is at least possible to establish the weaker conclusion that there is a representation of reality, which includes an account of the difference made by a representation's perspective, from one unified point of view. In physical terms, an absolute representation is coordinate-independent, and one (but not the only) way to obtain coordinate-independence is coordinate-freedom. But even a coordinate-dependent representation may still offer a unified representation of reality in the weaker sense that it can account for any perspective—this is the Lorentzian pedagogy. Moreover, because such an account is possible from any arbitrary perspective, it is in fact possible to construct a coordinate-*in*dependent account from their equivalence class. But this is only true if physics employs geometric objects. It is impossible to offer a unified representation of a non-geometric object from any particular perspective, because no such representation can account coherently for the way in which that same non-geometric object is represented from other perspectives. Consequently, it is also impossible to construct a coordinate-independent (hence absolute) representation from an equivalence class of coordinate-dependent ones. This poses a dilemma: either non-geometric objects are not physically real, or absolute representations from physics are not possible. In the case of the gravitational stress-energy, to be discussed in the next section, the former horn is objectionable on physical grounds. But since Moore's argument is *a priori*, the second horn must be false too. The aim of the remainder of this article is to offer a way out of this dilemma.

4 The Gravitational Stress-Energy Pseudotensor

This section presents the *gravitational stress-energy pseudotensor*. This is a non-geometric object that nevertheless seems to represent a physically real quantity in the framework of general relativity (GR). If this is indeed the case, then modern physics would seem to deal with irreducibly non-absolute representations. This would clearly conflict with Moore's view of physics as being in the business of the discovery of absolute representations of reality. Ultimately, we will argue that one *can* represent gravitational stress-energy absolutely, although this requires a revision of both its mathematical and metaphysical nature. But in this section we first discuss why the standard approach invoking the pseudotensor is not up to the job.[9]

The idea that both energy and momentum are conserved quantities is familiar already from classical mechanics, and carries over straightforwardly to the special theory of relativity (SR). In SR, the density and flux of the energy and momentum of *matter* (which here includes non-gravitational fields such as the electromagnetic field) is represented by the so-called 'stress-energy tensor', $T^{\mu\nu}$, which is a *bona fide* geometric object. The conservation of energy and momentum is expressed in terms of this tensor as follows

$$\partial_\mu T^{\mu\nu} = 0. \tag{1}$$

This equation states that the total flux of stress-energy through a point vanishes—i.e., that stress-energy is conserved at any point in spacetime.[10]

[9]The (non-)existence of gravitational stress-energy in general relativity has received much attention in the recent philosophical literature (see e.g. [Lam, 2011, Hoefer, 2000, Curiel, 2019, Read, 2018, Duerr, 2019a, Pitts, 2010]). Here, we will present only the details of the debates regarding gravitational stress-energy which are important for our purposes.

[10]There are three points to make here. (1) As is standard, in the above Greek indices are used to denote the components of an object in some coordinate system. (2) Strictly, (1) holds only in the frames of references adapted to the structure of Minkowski spacetime of SR—see e.g. [Read, 2018] for discussion. (3) Equation (1) can be converted into a conservation law through a region by integrating and applying Stokes' theorem.

In GR, on the other hand, it seems at first as if it is not the case that energy is conserved. In particular, the lack of global symmetries of a (curved, dynamical) spacetime in GR means that

$$\partial_\mu T^{\mu\nu} \neq 0, \tag{2}$$

which would seem to imply that material stress-energy is not a conserved quantity in GR. But perhaps this should not surprise us! For it is often—though not universally—claimed that, in GR, the gravitational degrees of freedom *also* carry energy, for example in gravitational waves.[11] It might still be the case, then, that the *total* energy of a system—material *and* gravitational—is conserved.

In order to formalise this claim, one can introduce a gravitational stress-energy pseudotensor, $t^{\mu\nu}$. (The reason this is called a *pseudo*tensor becomes clear below). The claim then is that the total stress-energy $\mathfrak{T}^{\mu\nu} := T^{\mu\nu} + t^{\mu\nu}$ is conserved, i.e. that

$$\partial_\mu \mathfrak{T}^{\mu\nu} := \partial_\mu (T^{\mu\nu} + t^{\mu\nu}) = 0. \tag{3}$$

This is indeed the case: in GR, there always exists a $t^{\mu\nu}$ such that (3) is satisfied. Therefore, it seems that one finds continuity between classical mechanics, special relativity and general relativity: in each theory, the total energy of the universe is locally conserved. The only difference is that in GR, unlike in its predecessors, the gravitational field itself carries energy.

There are weighty physical reasons to believe that the gravitational field really does carry stress-energy, which is represented by $t^{\mu\nu}$. Firstly, there is the above-mentioned continuity with classical and special-relativistic mechanics: in both theories, total energy is conserved. Moreover, such conservation principles are certainly not 'idle posits' of those theories. The scientific realist would therefore do well to preserve the principle of conservation of energy by positing the conservation of $\mathfrak{T}^{\mu\nu}$. Secondly, quite apart from past physics the conservation of energy and momentum are generally regarded as fundamental principles in contemporary physics. For example, Lange [2007] thinks of such principles as 'meta-laws' that modally constrain the laws of nature. In this way, conservation principles have explanatory import. The availability of such explanations is another reason to posit $t^{\mu\nu}$. Finally, GR simply does seem to describe a universe in which entities associated with gravitation carry energy. This is most clearly exemplified by the phenomenon of gravitational waves: fluctuations in the gravitational field that can have very real effects, as recently detected in the Nobel Prize-winning LIGO experiment.[12]

[11] For an article questioning this orthodoxy, see [Duerr, 2019c].

[12] For a recent defence of this view, see [Gomes and Rovelli, 2023]). Note, though, that these authors embrace the mainstream position that gravitational energy is non-localisable in GR.

However, the fact that $t^{\mu\nu}$ is a *pseudo*tensor—which indicates *inter alia* that it is a non-geometric object!—has led several philosophers to object to the claim that it represents a physically real quantity. In fact, there are at least three distinct problems related to $t^{\mu\nu}$ that are discussed in the literature:

1. In GR, there are in fact *infinitely many* distinct pseudotensorial quantities that are candidates to represent gravitational stress-energy: from Noether's first theorem, one for each rigid symmetry of the Lagrangian density [Pitts, 2010];

2. Conservation laws such as (3) are closely associated to (trivial?) mathematical identities ([Brading, 2005]);

3. Pseudotensors are not geometric objects: they don't have associated well-defined transformation laws (Duerr [2019b]).

The first problem refers to the fact that (3) does not define $t^{\mu\nu}$ uniquely. Rather, (3) is satisfied for any object such that

$$t^{\mu\nu} = \partial_\lambda U^{\mu\lambda\nu} - \frac{1}{8\pi} G^{\mu\nu}, \tag{4}$$

where $G^{\mu\nu}$ is the Einstein tensor (appearing on the left-hand side of the Einstein equations) and $U^{\mu\lambda\nu} = U^{\mu[\lambda\nu]}$ is a so-called 'superpotential': different superpotentials lead to different pseudotensors.[13] This raises the question: which of the infinitely many pseudotensors (associated to one of infinitely many superpotentials) is the 'real' one that represents the gravitational field's 'true' stress-energy? Relatedly, the second problem is based upon the thought that each of these infinitely many conservation laws is in fact a mathematical identity, which it would seem cannot have any non-trivial physical content. We will return to these first two problems—and especially the non-uniqueness problem—in §6.

For our present purposes, however, it is clearly the third problem that is most relevant. The fact that the pseudotensor is not a geometric object means, as explained in §3, that it is viciously coordinate-dependent.[14] This in turn means that one cannot construct a coordinate-independent representation of it, either in a entirely coordinate-free fashion or by way of the Kleinian approach discussed above. There is therefore a direct conflict between the fact that gravitational energy is represented by a non-geometric

[13] See [Trautman, 1962] for a classic discussion of this ambiguity, and [de Haro, 2022] for a more recent source.

[14] This terminology is also used by Duerr [2019b].

object on the one hand, and Moore's claim that physics must provide absolute representations of physical reality on the other. It is this conflict that leads to the dilemma raised above: either gravitational stress-energy is not a physically real quantity after all, or it is but one cannot represent it absolutely. Although many philosophers have taken the first horn in response to this dilemma,[15] we believe that it is undesirable for physical reasons as set out above. But this leaves us with the latter horn, which seems to conflict with the (*a priori*) case for the possibility of an absolute representation of reality by physics.

5 An Absolute Representation of Gravitational Stress-Energy

Fortunately, there is a way out of this dilemma. It is in fact possible to construct a *bona fide* geometric—indeed, tensorial!—object that can represent gravitational stress-energy. By taking this object to represent gravitational stress-energy in GR, the vicious coordinate-dependence associated with non-geometric objects is thereby avoided. The aim of this section is to put forward this alternative representation of gravitational energy—first proposed in the physics literature by Szabados [1991, 1992]—by means of which one can overcome the dilemma between the real existence of gravitational energy and the possibility of an absolute representations of reality.

We will see that our solution does require a significant revision to the concept of gravitational stress-energy. In particular, the geometric object that represents such stress-energy is not defined over the spacetime manifold itself, but rather over the so-called *bundle of linear frames*: a mathematical structure that consists of all possible 'choices of basis' for vector spaces at each point of spacetime. If taken seriously, this picture implies that gravitational energy is not a field that 'lives' on spacetime, but rather within this bundle of frames. We will comment on the metaphysical implications of this picture in §7, but overall we believe that taking an object on the bundle of linear frames to represent gravitational stress-energy is a price worth paying in order to reconcile its existence with the possibility of its absolute representations.

Szabados [1992] frames the problem of the vicious coordinate-dependence of pseu-

[15] See e.g. [Duerr, 2019b]. Sometimes, the non-geometric nature of pseudotensors is conflated with the fact that their components can be made to vanish in some coordinate system—see e.g. [Lam, 2011]. These are distinct features of an object: for example, Christoffel symbols are *geometric*, but have components which can be made to vanish at a point in some coordinate system. And conversely: some pseudotensors, e.g. that of Møller, are not geometric objects, but cannot be made to vanish at a point—our thanks to Brian Pitts for pointing this out this latter example to us.

dotensors as a contradiction with the principle of 'general covariance', which (in his words) states that nature is most appropriately described in terms of geometric objects.[16] We have seen that this does not seem to be the case for the gravitational stress-energy pseudotensor. Szabados proposes a solution as follows:

> However, if the geometric objects [...] were not required to be geometric objects on the spacetime manifold, but they were allowed to be geometric objects on the manifold of frames of the spacetime; i.e. on the bundle of linear frames $L(M)$ over M, and if the previous coordinate and/or gauge dependent quantities and formulae could be reformulated in terms e.g. of differential forms on $L(M)$, then the contradiction with the principle of general covariance would be resolved. [Szabados, 1992, p. 2522]

In brief, then, Szabados' idea is that one could reformulate the offensive pseudotensor as a geometric object defined over the bundle of linear frames. Since this is object is geometric, it allows for an absolute representation of gravitational stress-energy. The remainder of this section is devoted to a more detailed explication of Szabados' proposed solution. The reader already familiar with the concept of a bundle of linear frames may skip ahead to §5.2.

5.1 Mathematical Preliminaries

Let us start with the bundle of linear frames, denoted $L(M)$ by Szabados. In order to understand this concept, we first require the notion of a vector space. Generally, a vector space consists of a set V—the set of vectors—and a pair of binary operations: vector addition and scalar multiplication. The first operation allows one to add vectors to form another vector; the second to multiply a vector by a real number to obtain another vector. These operations must satisfy certain conditions, such as associativity and commutativity. If the vectors are two-dimensional, one can represent them as arrows on a fixed plane. Then vector addition is carried out by the familiar parallelogram rule, whereas scalar multiplication affects the length of a vector.

Although it is customary to think of vectors as tuples of real numbers, the elements of a vector space are distinct from those tuples; the latter are said to *represent* the former. The ordered pair (2, 3), for example, can be taken to represent a vector v that points in the direction 'two steps east and three steps north'. But that is not the only way to represent this vector. The coordinates (2, 3) only represent a vector in the 'two-east, three-north' direction conditional on the choice of the 'east'-vector and the 'north'-vector as

[16] In this article, we avoid couching the issue in terms of 'general covariance', for that term itself is notoriously fraught in GR: see [Norton, 1993] for background.

basis vectors. Given this basis, it is possible to express any other vector in terms of them: this many steps east, and that many steps north. But one could equally well have chosen a basis that consists, say, of a vector $\sqrt{2}$ units in the 'north-east' direction and a unit vector in the east direction. In terms of this basis, the same vector v is now represented as (2, 1): two steps of $\sqrt{2}$ in the north-east direction equals two steps to the north and two steps to the east; add another step to the east and one arrives at the same point as before. The crucial point here is that (2, 3) and (2, 1) are different representations of the *same* vector, v, in different bases. Thus, the numerical representation of a vector depends on an antecedent choice of basis vectors.

Generally, for an n-dimensional vector space, a basis consists of a choice of n linearly independent vectors (that is, vectors that are not sums of multiples of one another). Indeed, one can as well define the dimensionality of a vector space as the number of bases vectors required to 'span' it. The choice of basis in effect consists of a conventional choice for which direction to call the x, y, z, etc. axes, and for the unit in each direction. Once one has chosen a basis, each further vector can be expressed as a sum of multiples of these basis vectors. Moreover, it is also possible to consider the set of *all* possible bases for a certain vector space. This set itself has a non-trivial structure, since different bases are related by elements of the so-called *general linear group* (which is represented by the set of $n \times n$ matrices). The elements of this group act on any element within the set of bases to obtain another basis. We will return to the set of bases of a vector space below.

First, however, we will discuss the idea of a *tangent* vector space. Given a point p on a manifold M, one can always define the vector space tangent to that point. Think of a point on a curve: one can construct a tangent line parallel to the curve at that point; likewise, one can construct a tangent plane parallel to a point on a curved surface. If one further endows such a line or plane (or higher-dimensional generalisation thereof) with the structure required for a vector space, then one obtains the tangent vector space of that point. For example, the tangent vector space at a point on a surface consists of the arrows within the plane parallel to the surface at that point, together with the operations of vector addition and scalar multiplication. We will let T_p denote the tangent vector space of a spacetime point p.

If spacetime is flat, then there is a unique (or 'canonical') map between the tangent spaces at different points. Intuitively, the planes that are parallel to different points of a flat surface fully overlap, so one can identify them as the same plane. The more technical sense in which this is the case is that for any pair of points p, q on a flat space, one can transport a vector from T_p to T_q along any arbitrary path between p and q, and obtain the same result every time. There is therefore a sense in which one can speak of a sole vector space that is the same for every point. But GR tells us that spacetime is not flat, but curved. And in a curved spacetime, the transportation of a vector from T_p to T_q along different paths between p and q need not always yield the same result. If spacetime is curved, then, each point has a unique tangent vector space. Consequently, there is no path-independent way to transport a vector from one point to another.

To illustrate this latter case, consider the surface of the earth. Suppose John points in some direction on the North Pole, and Jane points in some direction on the equator. Do they point in the *same* direction along the surface of the Earth? The question is ill-posed insofar as the plane parallel to the earth's surface at the North Pole is different from the plane parallel to the earth's surface at the equator. Of course, Jane could move to the North Pole, *all the while keeping to point in the same direction*: if she points in the same direction as John once she has arrived at the North Pole, then they were pointing in the same direction all along. But the outcome of this 'experiment' crucially depends on the path Jane takes. It is possible that if Jane were to move from the equator to the North Pole along the shortest path, she would end up pointing in a different direction than John; yet if she were to move some distance along the equator first and *then* travel northwards, she would end up pointing in the same direction as John. Thus, there is no determinate answer to the question whether John and Jane pointed in the same direction or not: it depends on an arbitrarily chosen path between them. This illustrates the sense in which, on a curved spacetime, there is no 'global' vector space within which one can compare directions. Each point of spacetime has its own tangent vector space, and there is no canonical map between them.

We now have all the ingredients in place to define the bundle of linear frames $L(M)$ over a manifold M. Firstly, each point $p \in M$ has its own associated vector space T_p. Secondly, *each* such vector space has a set of bases. For any point p, let F_p denote the set of bases for T_p (or, more accurately, a structured 'fibre' of bases). Just like there is no canonical map between the vector spaces themselves, there is no canonical map between their respective bases. Put differently, one simply cannot say whether a basis in F_p is the same as a basis in F_q for $p \neq q$. Finally, the bundle of frames is the disjoint union of all F_p, for any $p \in M$: $L(M) := \bigcup_{p \in M} F_p$. The elements of $L(M)$ are therefore ordered pairs (p, f) such that $f \in F_p$. For a given p, the set of elements $(p, f) \in F$ are just the different vector bases for T_p. In other words, the 'points' of the bundle of linear frames are local vector bases. The bundle of linear frames, $L(M)$ thus contains the possible bases for *all* tangent vector spaces of a manifold.[17]

The final concept that is helpful in what follows is that of a *section* of $L(M)$: this is a continuous function from M to $L(M)$ such that each point p in M is mapped onto

[17] In technical terms, $L(M)$ is a principle fibre bundle over M with structure group $GL(n, \mathbb{R})$, where n is the dimension of M.

a point (p, f) of $L(M)$, i.e. a unique basis for T_p. A section thus specifies a choice of basis (not necessarily orthonormal) at each point of M in such a way that this choice of basis varies continuously as one moves across M. This is also known as a *vielbein*. This notion of a section does not play a direct role in this section, but it will help us to solve the non-uniqueness problem in §6.

5.2 The Sparling Form

With this set-up in place, Szabados [1991, 1992] defines two important objects: the 'Nester-Witten form', u_i and the 'Sparling form', t_i.[18] Both of these are defined on $L(M)$. Just as a field over a manifold M assigns a field-value to each point of M, a field over the bundle of frames $L(M)$ assigns a field-value to each point of $L(M)$, that is, to each choice of basis for each tangent vector space. Although such a field is mathematically well-defined, it may seem odd physically: we can make easily sense of a field on spacetime, but what does a field defined over a manifold of vector bases represent? What does it mean to assign a value to a choice of basis? We will comment on the metaphysical interpretation of such fields in the next section; here we restrict ourselves to their mathematical definition.

Although the precise definitions of u_i and t_i are not crucially important for what follows, we provide it here for completeness. Those not interested in mathematical details may skip ahead to the theorem below.

Let $L(M)$ denote the linear frame bundle over M, $\{\delta_i\}$ $(i = 1, \ldots, m)$ the standard basis for $\mathbb{R}^m$, i.e. $\delta_i = (0, \ldots, 1_i, \ldots 0)$, and $\theta = \theta^i \delta_i$ the canonical $\mathbb{R}^m$-valued 1-form on $L(M)$. For any $r = 0, 1, \ldots, m$, let

$$\Sigma_{a_1 \ldots a_r} := \frac{1}{(m-r)!} \epsilon_{a_1 \ldots a_r e_{r+1} \ldots e_m} \theta^{e_{r+1}} \wedge \ldots \wedge \theta^{e_m}, \tag{5}$$

where ϵ denotes the Levi-Civita symbol and $\wedge$ denotes the wedge product on differential forms.[19] Letting $\omega^a{}_b$ be a spin connection on $L(M)$, the Cartan equations for torsion Ξ^a and curvature $\Omega^a{}_b$ are as usual given by

$$\Xi^a = d\theta^a + \omega^a{}_b \wedge \theta^b, \tag{6}$$

$$\Omega^a{}_b = d\omega^a{}_b + \omega^a{}_c \wedge \omega^c{}_b. \tag{7}$$

[18] Following Szabados [1992], Latin indices denote objects on $L(M)$.

[19] For background on differential forms, see e.g. [Burke, 1985]. To reassure those unfamiliar with differential forms: the details of these objects and their constructions will not matter for our purposes!

The Nester-Witten form is then defined as:

$$u_i := -\frac{1}{2}\omega^{ab} \wedge \Sigma_{iab}. \tag{8}$$

The exterior derivative of the Nester-Witten form is

$$du_i = -\frac{1}{2}\Omega^{ab} \wedge \Sigma_{iab} + \frac{1}{2}\Xi^c \wedge \omega^{ab} \wedge \Sigma_{iabc} + t_i, \tag{9}$$

where

$$t_i := -\frac{1}{2}\left(\omega^c{}_i \wedge \omega^{ab} \wedge \Sigma_{cab} + \omega^a{}_c \wedge \omega^{cb} \wedge \Sigma_{iab}\right) \tag{10}$$

is the Sparling $(m-1)$-form.

So much for the mathematical construction of these objects. The significance of the Nester-Witten form and the Sparling form lies in the following theorem:[20]

Theorem 1 (Sparling–Dubois-Violette–Madore)**.** *For any $\mathbb{R}^{m*}$-valued horizontal $(m-1)$ form T_i satisfying $DT_i := dT_i - \omega^c{}_i \wedge T_c = 0$ and $\kappa \in \mathbb{R}$, the following statements are equivalent:*

1. *$\omega^a{}_b$ is torsion free, $\Xi^a = 0$, and $\frac{1}{2}\Omega^{ab} \wedge \Sigma_{iab} + \kappa T_i = 0$;*
2. *$\kappa T_i + t_i = du_i$;*
3. *$d\left(\kappa T_i + t_i\right) = 0$.*

The first condition expresses the fact that a metric connection is torsion-free and satisfies Einstein's equations, so (1) is satisfied whenever the theory's equations of motion are satisfied.[21] The theorem states that this is the case if and only if the Nester-Witten form and the Sparling form jointly satisfy condition (2), if and only if the latter form satisfies condition (3). Szabados [1992] notes that "in Einstein's theory (3) looks like as [*sic*] a conservation equation, while (2) gives us the 'superpotential' for the conserved quantity $\kappa T_i + t_i$: it is just the Nester-Witten form." In particular, (3) states that the sum $\kappa T_i + t_i$ is a conserved quantity, where T_i is the matter stress-energy tensor. If t_i

[20]The following is the theorem as stated by Szabados [1992]; for original sources, see [Sparling, 1982, Dubois-Violette and Madore, 1987].

[21]The pullback of T_i along a local section is independent of the choice of section, and yields a geometric object apt to represent material stress-energy—see [Szabados, 1991, p. 24]. Note, though, that as-yet no energy conditions are imposed on said object: for the *locus classicus* on energy conditions, see [Curiel, 2017].

is interpreted as a representation of gravitational stress-energy, then (3) states that the sum of matter and gravitational stress-energy is conserved—as desired. Therefore, on this interpretation Theorem 1 tells us that the total sum of stress-energy is conserved whenever the torsion-free metric connection satisfies the Einstein equations.

The second condition further deepens the parallel between the Sparling form and the gravitational stress-energy pseudotensor, since it shows that the Nester-Witten form acts as a 'superpotential' to t_i. This provides further legitimacy to the suggestion that the latter quantity represents gravitational stress-energy.

Of course, however, the crucial difference is that u_i and t_i are geometric objects! Insofar as geometric objects are required for the possibility of absolute representations, then, Theorem 1 proves that it *is* possible to provide an absolute representation of gravitational stress-energy. This means that it is possible to escape the dilemma raised before. Recall that the dilemma seemed to force a choice between the reality of gravitational energy and the possibility of an absolute representation of physical reality. The conservation of energy provides a physical reason to choose the first horn, but Moore's *a priori* case for absolute representations necessitates the second. We have shown, however, that it is possible to represent gravitational energy in a way that does not lead to a vicious form of coordinate-dependence. Therefore, it does not spoil the possibility of an absolute representation of it.

6 The Non-Uniqueness Problem

It turns out that recourse to differential forms on $L(M)$ also allows us to solve another problem associated with pseudotensors: their non-uniqueness (recall that this was the first problem with pseudotensors presented in §4). To see how such a solution proceeds, first recall that we can write a pseudotensor $t^{\mu\nu}$ on M in terms of a superpotential $U^{\mu\lambda\nu}$ as per (4). But the pseudotensor satisfies this equation for any superpotential that is antisymmetric in its second and third indices. Thus, different choices for the superpotential lead to different—but, it would seem, equally valid—pseudotensors.

The situation here is somewhat analogous to the freedom to add a constant to the total energy in classical mechanics. But the freedom here is worse, in the sense that each pseudotensor is associated with a *distinct* charge via Noether's theorem [de Haro, 2022]. Put in terms of the distinction between determinables and determinates, adding a constant to the potential in classical mechanics simply changes the determinate value of the same determinable quantity, but adding a superpotential to the pseudotensor seems to result in a different determinable quantity altogether.

However, it seems that we can use features of the coordinate-independent framework developed in the previous sections in order to at least alleviate—if not totally dissolve—the non-uniqueness problem. We will make use of three facts:

1. We can write any superpotential in terms of a particular superpotential called the von Freud superpotential [Trautman, 1962];

2. The von Freud superpotential is just the pullback of the Nester-Witten form along a particular coordinate section [Szabados, 1992];

3. Different superpotentials can be obtained from different pull-backs of the Nester-Witten form.

Combining these facts, we get that any superpotential is related to the von Freud superpotential, which has a coordinate-independent formulation in terms of the Nester-Witten form on $L(M)$. Moreover, different superpotentials also correspond to the pullback of the Nester-Witten form along different sections. This provides a sense in which the distinct superpotentials—and hence their associated pseudotensors—are simply different reflections on spacetime of the same geometric object on the bundle of linear frames. Put differently, the Nester-Witten form *unifies* the myriad choices of superpotential: "This reformulation may yield a unification of the different pseudotensorial and rigid-basis-dependent approaches into a single manifest gauge invariant formalism" [Szabados, 1992, p. 2522].

Clearly, these results require some further shoring up. Firstly, although Szabados shows that all *extant* pseudotensors are derivable from the Sparling form, he doesn't quite give a proof that this is possible for *all* pseudotensors. Secondly, we have not proven that the Sparling form itself is unique, so one might worry that the same issue reappears at the level of $L(M)$. Addressing these issues will have to remain mathematical tasks for another day: our philosophical point here is simply that appeal to geometric objects on $L(M)$ has the potential to address not only the third problem for pseudotensors as presented in §4 (non-geometric status), but also the first (non-uniqueness).

Moreover, even these limited results already offer us a way to understand in more detail the relation between the coordinate-dependent representation of gravitational stress-energy in terms of the pseudotensor and our coordinate-independent representation of the same quantity by the Sparling form. For they suggest that these quantities are indeed representations of the *same* quantity; or, more precisely, it suggests that a pseudotensor is no more than a particular way to pull back the Sparling form onto the spacetime manifold. The different pseudotensors are, as it were, the spatiotemporal 'shadows' of the same object on the bundle of linear frames, seen from different 'perspectives' (i.e. sections/vielbeins). This result is particularly important for the realist, since it is typically seen as a desideratum to be able to explain the way in which successor theories relate to their predecessors. Think, for example, of the way in which classical mechanics reduces to special relativity in the limit $c \to \infty$.[22] In our case, the Sparling form is not a successor theory, but rather a successor quantity. But the same point applies: it is a desideratum for the realist to show that the Sparling form in some sense reduces to the stress-energy pseudotensors. The fact that the latter is the pullback of the former shows that this is indeed the case. The close connection between the Sparling form and the equivalence class of pseudotensors therefore shows that the former is indeed a candidate to represent gravitational stress-energy, and at the same time explains the success (albeit limited) of the latter in accounting for the conservation of total stress-energy in GR.

7 The Metaphysics of Frame Bundles

Mathematically, it all works out: the Sparling form is a *bona fide* geometric object, and it has the credentials to represent gravitational energy in GR. Metaphysically, however, our proposal lands us in strange waters. The gravitational stress-energy is no longer a field on spacetime, but rather on the bundle of linear frames. The pseudotensors defined over the spacetime manifold are at best one-sided reflections of the Sparling form, which is the real deal.

In response to our claim that the Sparling form represents gravitational stress-energy, one might raise the following worry: by considering objects defined on $L(M)$ when addressing issues regarding gravitational energy in GR, we have dodged the question by moving to a new space, thereby implicating us in an enriched set of ontological commitments over and above 'standard' GR. In fact, there are two closely related objections here: (1) that the introduction of a novel space, the bundle of frames, makes GR less parsimonious; and (2) that it is difficult to even make sense of a physically real field defined over this bundle.

In response to the first objection, we would say this: there is a straightforward sense in which we are *already* committed to $L(M)$ when we posit a manifold M, since the bundle of linear frames is *definable* from the standard manifold structure. Weatherall [2015] forcefully makes this point: "[When formulated in terms of a frame bundle] no additional structure has been added to the theory. Any manifold gives rise, in a canonical way, to an associated frame bundle. Thus there is a straightforward sense in which a relativistic spacetime (M, g_{ab}) always comes equipped with a principal [frame] bundle over it; we just have little occasion to mention it in ordinary applications of relativ-

[22] Or $v/c \to 0$, or whichever limit one prefers.

ity theory." A similar point about definability and ontological commitment (which, of course, we endorse) is made by Barrett [2017]; it is also part of the lesson of the functionalism of (for instance) Lewis [1970], first expounded in the context of the philosophy of mind, but more recently taken up by Butterfield and Gomes [2020] and others in the philosophy of space and time. The basic idea here is that if one has explicitly posited a certain fundamental structure, then any further structure that one can define from the former 'comes for free'. Therefore, it does not add to a theory's ontology to take this further structure as seriously as the initially posited fundamental structure.

To apply this line of thought to the case at hand: when one takes different sections of the Sparling and Nester-Witten forms defined on $L(M)$, one can derive different pseudotensors and superpotentials (respectively) on M via the 'pull-back' construction [Szabados, 1992]. Now, we claim (but do not here rigorously prove!) that t_i and u_i can be defined from just the structure of standard GR (i.e. the usual geometric objects of the metric and material fields on M, together with the Einstein equation). For the union of all pseudotensorial conservation laws is equivalent to the Einstein equation [Pitts, 2010], but if one has all such conservation laws, then one has the associated conservation law—((3) in Theorem 1)—for the associated objects on $L(M)$. So, the structure of standard GR seem to allow one to define a suitable conservation law for suitable objects on $L(M)$, without the need to posit additional structure. If this is correct, then it makes sense to say that it is only if one chooses a *particular* section of $L(M)$ that one is introducing additional structure to GR—tantamount to taking the perspective of a preferred frame. But the point of our proposal is exactly that one does not have to take any such perspective; one can take the view from nowhere on the bundle of frames.

So much for the first objection. It is harder to see how to respond to the second objection—that it is difficult to conceive of a physical field defined on $L(M)$, and for that object to play the role of gravitational stress-energy. True, this is an unintuitive metaphysical picture. But so, of course, is much else in GR: that spacetime is curved and that it expands; the possibility of singularities and black holes; and indeed the very notion that the gravitational field can carry energy. We have, to some extent, become accustomed to these ideas. We believe that the notion of gravitational stress-energy on the bundle of frames is simply another novel implication of GR.[23]

The objection might be sharpened: it is not simply that a field on $L(M)$ is beyond the pale, but rather that it seems odd to identify this field with gravitational stress-energy. After all, gravity concerns the curvature of spacetime itself, and does not seem

[23] Conversely, in Yang-Mills theories it is absolutely standard to represent physical fields, such as the electromagnetic field, on a principal bundle defined over the spacetime manifold. This provides further reason to believe that a physical field on the bundle of frames is not unacceptable either.

conceptually close to a local choice of basis. We admit that our proposal does not conform to these intuitions—but it is not unjustified! In particular, Theorem 1 tells us that the Sparling form certainly plays the *functional role* of gravitational energy, insofar as it satisfies the conservation equation $d\left(\kappa T_i + t_i\right) = 0$ [Read, 2018]. Moreover, the Theorem also relates this conservation equation to the Einstein equation in a natural way, so the Sparling form is not just an idle posit but an integral part of the theory. Finally, as explained in the previous section, there is a natural way to understand the gravitational stress-energy pseudotensor as the 'pull-back' of the Sparling form onto the spacetime manifold. Insofar as the pseudotensor was a natural candidate for gravitational stress-energy, then, and bar any worries about the fact that it is a pseudotensor, the Sparling form is just as natural—if not more so, due to the fact that it can unify different pseudotensors as explained in the next section.

Therefore, it is our view that the Sparling form *does* represent gravitational stress-energy. It just turns out that gravitational stress-energy lives on $L(M)$, not M. This is, perhaps, a metaphysically radical conclusion. But since it is justified on the basis of Theorem 1, we believe that this is simply where GR leads. In the physicist's pursuit of absolute representations of nature, pre-theoretical metaphysical intuitions just don't carry much weight.

8 Conclusion

Let's sum up. Moore [1997] has given an abstract and *a priori* argument to the effect that an absolute representations of reality is possible: that is, that it is possible to represent the world 'from no point of view'. He has also stated that it is the business of physics to find absolute representations. But this seems to stand in conflict with modern physics, in which various non-geometric objects—such as the gravitational stress-energy pseudotensor—at least *prima facie* seem to have physical content. One way to resolve the tension is in fact to denude all such objects of representational significance. In this article, however, we have sought to explore how one may instead secure an absolute representation of whatever it is that said objects purport to represent: again, our case study has been gravitational energy, in which case we have appealed to geometric objects defined on the bundle of linear frames. Since such objects are geometric, they *do* admit of absolute representations—thereby, any tension between Moore's argument and modern physics (at least in this particular case!) is resolved.

Clearly, what we have presented here is but one example of a potentially much broader methodology, which it would be worth exploring in the context of other non-geometric objects in modern physics which one would like to invest with representational import. The most significant such case is that of spinors, which are taken to represent all fermionic matter in the universe. Since spinors are non-geometric objects, it again appears that they stand in tension with Moore's argument: insofar as one seeks a representation of physics 'from no point of view', it becomes a pressing task to investigate how this might be secured in this case also. Although Pitts [2012] has undertaken some admirable and beautiful work in this direction, as of yet he has shown only that such objects are 'almost geometric' (in a technical sense); given this, in our view, there remains more work to be done when exploring how well (or otherwise) spinors and other non-geometric objects sit with Moore's argument.

This, in turn, invites a range of broader mathematical questions: can all non-geometric objects be understood as pullbacks along local sections of geometric objects defined on fibre bundles? And: are all such bundles implicitly definable as in the case of the frame bundle, or not? And there are also further philosophical questions here: to what extent can Lewisian functionalism be assimilated to implicit definability, as we have done in this article? Is Moore's conception of an absolute representation the same as that of Williams [1978, 1985] (who indeed inspired Moore's work, at least in part), and if not does any tension with modern physics in that latter case manifest itself in the same way? And so forth.

The general point—which was also raised in [Read, 2022]—is this: contemporary physics presents a broad zoology of 'fantastic beasts' [Duerr, 2019a]; it is incumbent upon the metaphysician to investigate the metaphysical significance of such objects, as well as how such objects sit with any *a priori* arguments regarding the world which they might elect to muster. Our investigations in this article into how well Moore's argument sits with non-geometric objects from physics can, in this sense, be taken as a call-to-arms to all naturalistically-inclined metaphysicians and philosophers to engage in a substantially broader research programme.

Acknowledgements

We are grateful to Patrick Dürr, Baptiste Le Bihan, Sam Fletcher, Henrique Gomes, Eleanor March, Adrian Moore, Brian Pitts, and Chris Timpson, as well as to audiences in Bristol, Oxford, Geneva, and Boston, for helpful discussions. J.R. thanks the Leverhulme Trust for their support.

Statements and declarations

There are no conflicts of interest associated with this work.

References

Thomas William Barrett. What do symmetries tell us about structure? *Philosophy of Science*, (4):617–639, 2017. doi: 10.1086/699156.

John S. Bell. How to teach special relativity. *Progress in Scientific Culture*, 1, 1976.

Katherine Brading. A note on general relativity, energy conservation, and Noether's theorems. In A. J. Kox and Jean Eisenstaedt, editors, *The Universe of General Relativity*, pages 125–135, Boston, MA, 2005. Birkhäuser Boston. ISBN 978-0-8176-4454-3.

William L. Burke. *Applied Differential Geometry*. Cambridge University Press, 1985. doi: 10.1017/CBO9781139171786.

J. Butterfield and H. Gomes. Functionalism as a species of reduction, 2020.

Erik Curiel. A primer on energy conditions. In Dennis Lehmkuhl, Gregor Schiemann, and Erhard Scholz, editors, *Towards a Theory of Spacetime Theories*, pages 43–104. Birkhäuser, 2017.

Erik Curiel. On geometric objects, the non-existence of a gravitational stress-energy tensor, and the uniqueness of the Einstein field equation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 66: 90–102, 2019. doi: 10.1016/j.shpsb.2018.08.003.

Sebastian de Haro. Noether's theorems and energy in general relativity. In James Read and Nicholas J. Teh, editors, *The Philosophy and Physics of Noether's Theorems: A Centenary Volume*, page 197–256. Cambridge University Press, 2022. doi: 10.1017/9781108665445.010.

Neil Dewar. General-relativistic covariance. *Foundations of Physics*, 50(4):294–318, 2020. doi: 10.1007/s10701-019-00256-0.

Michel Dubois-Violette and John Madore. Conservation laws and integrability conditions for gravitational and Yang-Mills field equations. *Communications in Mathematical Physics*, 108(2):213 – 223, 1987.

Patrick M. Duerr. Fantastic beasts and where to find them: Local gravitational energy and energy conservation in general relativity. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 65:1–14, 2019a. doi: 10.1016/j.shpsb.2018.07.002.

Patrick M. Duerr. Against 'functional gravitational energy': A critical note on functionalism, selective realism, and geometric objects and gravitational energy. *Synthese*, 199(S2):299–333, 2019b. doi: 10.1007/s11229-019-02503-3.

Patrick M. Duerr. It ain't necessarily so: Gravitational waves and energy transport. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 65:25–40, 2019c. doi: 10.1016/j.shpsb.2018.08.005.

Henrique Gomes and Carlo Rovelli. On the analogies between gravitational and electromagnetic radiative energy, 2023.

Carl Hoefer. Energy conservation in GTR. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 31(2):187–199, 2000. doi: 10.1016/s1355-2198(00)00004-6.

Ivan Kolář, Jan Slovák, and Peter Michor. *Natural Operations in Differential Geometry*. 09 1996. ISBN 978-3-642-08149-1. doi: 10.1007/978-3-662-02950-3.

Vincent Lam. Gravitational and nongravitational energy: The need for background structures. *Philosophy of Science*, 78(5):1012–1024, 2011. doi: 10.1086/662260.

Marc Lange. Laws and meta-laws of nature: Conservation laws and symmetries. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(3):457–481, 2007. ISSN 1355-2198. doi: https://doi.org/10.1016/j.shpsb.2006.08.003. URL `https://www.sciencedirect.com/science/article/pii/S1355219806000943`.

David Lewis. How to define theoretical terms. *Journal of Philosophy*, 67(13):427–446, 1970. doi: 10.2307/2023861.

A. W. Moore. *Points of View*. Oxford, GB: Clarendon Press, 1997.

Adrian Moore. The possibility of absolute representations. In James Conant and Jesse M. Mulder, editors, *Reading Rödl on Self-Consciousness and Objectivity*. forthcoming.

Thomas Nagel. *The View From Nowhere*. Oxford University Press, 1986.

A. Nijenhuis. *Theory of the Geometric Object*. University of Amsterdam, 1952. URL `https://books.google.co.uk/books?id=UpeUQAAACAAJ`.

John D. Norton. General covariance and the foundations of general relativity: Eight decades of dispute. *Reports of Progress in Physics*, 56:791–861, 1993.

J. Brian Pitts. Gauge-invariant localization of infinitely many gravitational energies from all possible auxiliary structures. *General Relativity and Gravitation*, 42(3):601–622, 2010. doi: 10.1007/s10714-009-0874-y. URL `https://doi.org/10.1007/s10714-009-0874-y`.

J. Brian Pitts. The nontriviality of trivial general covariance: How electrons restrict 'time' coordinates, spinors fit into tensor calculus, and $\frac{7}{16}$ of a tetrad is surplus structure. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 43(1):1–24, 2012. doi: 10.1016/j.shpsb.2011.11.001.

W. V. Quine. Facts of the matter. *Southwestern Journal of Philosophy*, 9(2), 1978.

James Read. Functional gravitational energy. *British Journal for the Philosophy of Science*, 71(1):205–232, 2018. doi: 10.1093/bjps/axx048.

James Read. Geometric objects and perspectivalism. In James Read and Nicholas J. Teh, editors, *The Philosophy and Physics of Noether's Theorems*, pages 257–273. Cambridge: Cambridge University Press, 2022.

J. A. Schouten. *Ricci-Calculus: An Introduction to Tensor Analysis and Its Geometrical Applications*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1954. ISBN 9783662129272. URL `https://books.google.co.uk/books?id=rf_uCAAAQBAJ`.

Theodore Sider. *Writing the Book of the World*. Oxford, England: Oxford University Press, 2011.

G. A. J. Sparling. Twistors, spinors and the Einstein vacuum equations. 1982.

L. B. Szabados. Canonical pseudotensors, Sparling's form and Noether currents. *Hungarian Academy of Sciences*, 1991.

L B Szabados. On canonical pseudotensors, sparling's form and noether currents. *Classical and Quantum Gravity*, 9(11):2521, nov 1992. doi: 10.1088/0264-9381/9/11/017. URL `https://dx.doi.org/10.1088/0264-9381/9/11/017`.

Andrzej Trautman. Conservation laws in general relativity. In Louis Witten, editor, *Gravitation: An introduction to current research*, pages 169–198. Wiley, 1962.

Andrzej Trautman. Foundations and current problems of general relativity (notes by graham dixon, petros florides and gerald lemmer). In S. Deser and K. W. Ford, editors, *Lectures on General Relativity*. Englewood Cliffs, N.J., Prentice-Hall, 1965.

David Wallace. Who's afraid of coordinate systems? an essay on representation of spacetime structure. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 67:125–136, 2019. doi: 10.1016/j.shpsb.2017.07.002.

James Owen Weatherall. Fiber bundles, yang–mills theory, and general relativity. *Synthese*, 193(8), 2015.

Bernard Williams. *Descartes: The Project of Pure Enquiry*. Routledge, 1978.

Bernard Williams. *Ethics and the Limits of Philosophy*. London: Fontana, 1985.