

Gauging the boundary in field-space

Henrique Gomes^{*a,b}

^a Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo,
ON N2L2Y5, Canada

^b Trinity College, Cambridge, CB2 1TQ, United Kingdom

January 5, 2019

Abstract

Local gauge theories are in a complicated relationship with boundaries. Whereas fixing the gauge can often shave off unwanted redundancies, the coupling of different bounded regions requires the use of gauge-variant elements. Therefore, coupling is inimical to gauge-fixing, as usually understood. This resistance to gauge-fixing has led some to declare the coupling of subsystems to be the *raison d'être* of gauge [Rov14].

Indeed, while gauge-fixing is entirely unproblematic for a single region without boundary, it introduces arbitrary boundary conditions on the gauge degrees of freedom themselves—these conditions lack a physical interpretation when they are not functionals of the original fields. Such arbitrary boundary choices creep into the calculation of charges through Noether's second theorem, muddling the assignment of physical charges to local gauge symmetries. The confusion brewed by gauge at boundaries is well-known, and must be contended with both conceptually and technically.

It may seem natural to replace the arbitrary boundary choice with new degrees of freedom, for using such a device we resolve some of these confusions while leaving no gauge-dependence on the computation of Noether charges [DF16]. But, concretely, such boundary degrees of freedom are rather arbitrary—they have no relation to the original field-content of the field theory. How should we conceive of them?

Here I will explicate the problems mentioned above and illustrate a possible resolution: in a recent series of papers [GR17, GR18, GHR18] the notion of a connection-form was put forward and implemented in the field-space of gauge theories. Using this tool, a modified version of symplectic geometry—here called ‘horizontal’—is possible. Independently of boundary conditions, this formalism bestows to each region a physically salient, relational notion of charge: the horizontal Noether charge. Meanwhile, as required, the connection-form mediates an irenic composition of regions, one compatible with the attribution of horizontal Noether charges to each region. The aims of this paper are to highlight the boundary issues in the treatment of gauge, and to discuss how gauge theory may be conceptually clarified in light of a resolution to these issues.

*gomes.ha@gmail.com, hgomes@perimeterinstitute.ca

Contents

1 Gauge and boundary: a complicated relationship.	3
1.1 What is gauge?	3
<i>On perspectives</i> –	3
<i>Rovelli’s relationalism: “Why gauge”</i> –	5
1.2 A multitude of perspectives, not a position in the world	7
1.3 Gauge-fixing and boundaries in field theory	9
<i>Gauge-fixing in field theory without boundaries</i> –	9
<i>Gauge-fixing in field theory with boundaries</i> –	10
<i>New degrees of freedom at the boundary</i> –	12
1.4 Boundaries and charges: a brief intro.	13
<i>The teeth of gauge: Noether theorems</i> –	13
<i>Gauge-variant ‘edge-modes’</i> –	15
2 Horizontal geometry and the connection-form	17
2.1 A summary of the construction	17
2.2 The geometry of field-space	18
<i>Mathematical preliminaries</i> –	18
2.3 Connection-forms: the formal construction	21
3 Deploying connection-forms	22
3.1 Examples of connection-forms	23
<i>Selecting frames: the gravity example</i> –	23
<i>The Singer-DeWitt connection</i> –	24
<i>The Higgs connection</i> –	26
<i>Barbour-Bertotti’s best-matching</i> –	27
3.2 Gluing	28
<i>Establishing communication: gluing different regions</i> –	28
<i>Quantum considerations</i> –	30
3.3 Horizontal (or relational) Noether charges	31
<i>Symplectic geometry</i> –	31
<i>Horizontal symplectic geometry</i> –	32
<i>Conserved global charges</i> –	34
<i>The point of confusion</i> –	34
<i>‘Slippage’: an illustrative example</i> –	35
<i>Conserved global charges: an explicit example</i> –	35
4 Conclusions	37
<i>Perspectives</i> –	37
<i>Gauge-fixing and boundaries</i> –	38
<i>Noether charges</i> –	38
<i>The relational connection-form</i> –	38
<i>Open questions</i> –	39

Roadmap: In section 1, I criticize the usual construal of gauge-invariance as descriptive redundancy, by building on Rovelli’s view that gauge variables are needed to describe the relations between subsystems, and thus, in a field-theory, to describe boundaries. Which gauge variables are to be kept, and how they couple is

not made explicit by Rovelli. I identify the very flexible notion of a ‘perspective’ as the right tool to meet Rovelli’s requirements. The challenge then becomes one of keeping the perspectives while still being able to distinguish physical processes from pure gauge ones.

In section 2, this challenge is met. I elaborate on Rovelli’s intuition, by providing specific, regional, gauge-covariant ‘handles’, used to couple different regions. These handles are given by the *relational connection-form*. If the standard finite-dimensional connection-form arises by augmenting spacetime derivatives to be covariant under spacetime-dependent gauge transformations, the relational connection-form arises from augmenting field-variations to be covariant under field-dependent gauge transformations [GR17]. The point is that, at boundaries, the distinction between physical and gauge degrees of freedom is muddled, and therefore gauge-transformations there must adopt field-dependence.

Alternatively, the guiding requirement for the construction of the relational connection-form is a harmonious melding of regional and global perspectives. Most importantly, retaining gauge-variant elements *in this manner* is compatible with the attribution of a (relational) physical status to processes, and this is true both globally and within bounded regions, and with or without imposing boundary conditions. Lastly, I show that the ensuing notions of regional relationalism are different from other attempts at resolving the problem posed by gauge symmetries for bounded regions. The distinguishing criterion is what I consider to be the ‘acid test’ of local gauge theories in bounded regions: does the theory license only those regional charges which depend solely on the original field content? In a satisfactory theory, the regional charges should only depend on the original field content. In section 3, I introduce explicit examples of relational connection-forms, and show that the ensuing horizontal symplectic geometry passes this ‘acid test’.

1 Gauge and boundary: a complicated relationship.

In section 1.1 I argue that gauge is a matter of alternative descriptions, and that keeping these different alternatives in play is indispensable when we recognize that the system described is not the entire cosmos. These different descriptions are the ‘perspectives’.

In section 1.2 I argue that we need a more precise definition than Rovelli provides of what types of “gauge-variant handles” are needed to couple regions. Harmonizing regional and global perspectives could do the trick.

In section 1.3, using gauge-fixings, I show how in the field-theory context boundaries retain the difficulties about coupling identified by Rovelli. I show that the underlying reason is that boundaries drive a wedge between the notion of gauge as: i) a matter of alternative descriptions of a system and as ii) a matter of applications of an abstract group of gauge transformations. It is this wedge that makes room for the—erroneous, in my view—introduction of extra physical degrees of freedom at boundaries, an issue tackled in 1.4. In that subsection, we take stock of our argument so far, then move on to the consequences of this argument for Noether charges, and Donnelly and Freidel’s proposal of physical edge-modes [DF16].

1.1 What is gauge?

On perspectives –

“Like moths attracted to a bright light, philosophers are drawn to glitz. So in discussing the notions of ‘gauge’, ‘gauge freedom’, and ‘gauge theories’, they have tended to focus on examples such as Yang–Mills theories and on the mathematical apparatus of fibre bundles. But while Yang–Mills theories are crucial to modern elementary particle physics, they are only a special case of a much broader class of gauge theories. And while the fibre bundle apparatus turned out, in retrospect, to be the right formalism to illuminate the structure of Yang–Mills theories, the strength of this apparatus is also its weakness: the fibre bundle formalism is very flexible and general, and, as such, fibre bundles can be seen lurking under, over, and around every bush.” – J. Earman [Ear03]

I am one of those who is attracted to the “glitz” of fiber bundles. However, I believe there is more to the glitz than glamorous display; the generality of fiber bundles captures a defining feature of gauge theories. In simple terms, fiber bundles tell us a ‘gauge system’ is one that allows different descriptions of itself—different perspectives. Relating these perspectives is then the job of gauge transformations. This broad characterization doesn’t make reference to action functionals or equations of motion; they come in later in the definition of the dynamics of the system.

The electromagnetic and isospin fields—which originally motivated Yang and Mills to introduce the local non-Abelian gauge groups [YM54]—are good examples; admitting descriptions under different Lorentz frames, and under different labeling of particles with isotopic spin, respectively.¹

Underlying this simple definition lies a subtle distinction between the gauge (or symmetry) group, and the perspectives. The perspective is the description of the state from a given viewpoint. The perspectives emphasize how the same system is viewed within different frames, whereas the symmetry groups emphasize the abstract ‘rotations’ between frames, or viewpoints. A ‘frame’ would be the characterization of the coordinates themselves, whereas the perspective characterizes the system as described in that frame. As we will see, a change of frames is identified with a group element and generically corresponds to a change in perspectives

Although relations between perspectives could be said to be more concrete than the abstract definition of gauge symmetries, there are more important, objective, albeit subtle, differences between ‘frame’ and ‘perspective’. Indeed, there are at least two instances in which the notions detach from each other.² Both of these instances require the consideration of bounded regions.

This touches on one of the main points I want to bring to the fore of philosophy of physics with this work: the importance of treating bounded regions in our aim to understand the true nature of gauge systems. Hitherto, discussions of gauge have steered away from questions of regionality. But it is only once we consider the infinite-dimensional fiber bundles appropriate to the treatment of field theories, and the matter of gluing together the corresponding regional mathematical structures, that we will see just how indispensable gauge degrees of freedom truly are.

¹ But so is the lamp on my desk, which can be visually described by several different perspectives. See chapter 3, pages 66-75, of Van Fraassen [Fra08] for a philosophical discussion connecting visual perspectives, Cartesian frames of reference and Lorentz-invariant descriptions. Such analogies can lend at least a smidgeon of plausibility to this broad painting of gauge, even if now gauge systems may “lurk under every bush”. Admittedly, a strict adherence to this characterization of gauge systems might overextend its usefulness by including too many types of systems under the umbrella of ‘gauge’. Rest assured, we will stay dry; I won’t stray from standard examples.

²For field theories: in gauge-fixing in the presence of boundaries, seen in section 1.3; and in the computation of Noether charges, in section 3.3.

Gauge degrees of freedom are often introduced for practical purposes. For example: the totality of gauge-invariant degrees of freedom cannot be localized, as evidenced by Wilson loops.³ The lack of a local description of all the gauge-invariant degrees of freedom compels us to use a particular (but arbitrary) perspective. In other words, for some types of system, there is no *tractable* ‘god’s eye view’—unveiling all the physics without requiring a perspective.

Still, retaining all these perspectives to describe the same system seems uneconomical. Gauge-fixings are usually taken to rid our descriptions of the superfluosity in a tractable, albeit undemocratic, manner. Fixing a gauge amounts to choosing a selective class of perspectives whose members are asked to satisfy certain conditions. Starting from an arbitrary perspective in the class, the conditions should be just strong enough to allow for a unique second perspective, related to the first by a change of frame, to satisfy these conditions. If two gauge-fixed perspectives differ, they *really are* descriptions of physically distinct states; they really do concern different states of affairs. Gauge-fixing thus restricts us to perspectives that use the same language to describe their world. Gauge-fixing avoids the comparison of “apples to oranges”, as the saying goes; using it we compare oranges to oranges and protons to protons.⁴

Therefore, one might think, if we can always fully gauge-fix, the plurality of perspectives is surely superfluous; a surplus structure that we can shave off by quotienting or by agreeing on a common language [Dew17]. But a fly lands in this ointment once we countenance a merely provincial access to the World.

Rovelli’s relationalism: “Why gauge” – The diluted physical meaning usually attributed to pure gauge degrees of freedom is a consequence of the monopoly of physical meaning that is usually awarded to gauge-invariant objects.

In [Rov14], a dissonant view is put forward. It is a fundamentally relational view, and attributes an important role to perspectives. It implies that when we are in possession of only partial subsystems—and therefore lacking the complete set of relations physically characterizing the entire system—it makes sense to keep gauge-variant information in-hand.

For Rovelli, gauge-variant⁵ objects are necessary to couple certain types of subsystems. As he points out, “gauge non-invariant quantities [...] represent, in a sense, handles through which systems can couple.” To illustrate the idea, he describes two squadrons, each made up of N spaceships, which are just coming into contact with each other. Given q_j^1 as the position of the j -th ship in the first squadron, the Lagrangian for the first squadron is given by

$$L_1 = \frac{1}{2} \sum_{i=1}^{N-1} (\dot{q}_{i+1}^1 - \dot{q}_i^1)^2 \quad (1.1)$$

which has a time-dependent displacement symmetry acting as $\delta_1 q_i^1 := q_i^1 + f_1(t)$ (*mutatis mutandis* for the second squadron, with ‘2’ in place of ‘1’). For each squadron we can find gauge-invariant variables by taking the difference in position

³Using Wilson loops as basic variables is also in practice problematic for a variety of reasons [Hea07].

⁴This is in reference to the (approximate) isospin symmetry. This feature of gauge-fixings is well-known in the standard local gauge theory case, but it also applies to the colloquial use of ‘perspectives’ to discuss vision: fixing one’s representation of, for example, a desk lamp in terms of characteristics of the lamp—e.g. ‘the chord is in direct view’—enables one to unambiguously tell lamps apart [Fra08].

⁵‘Gauge-variant’ is preferable to the double-negative ‘gauge non-invariant’, as Rovelli and others refer to the property.

of two ships $\bar{q}_i^1 = q_{i+1}^1 - q_i^1$, $i = 1, \dots, N-1$ where the barred variables, \bar{q}_i^1 are gauge-invariant. Although there are N spaceships in each squadron, we can describe the subsystems individually in a gauge-invariant manner by using the gauge-fixed $N-1$ variables; of course, we could have chosen many different parametrizations of the gauge-invariant variables.

Now the second squadron appears and fighting begins. Interaction terms between the two squadrons might no longer be independently gauge-invariant under δ_1, δ_2 . They still could be, if expressed in terms of the barred variables, but for the example Rovelli gives,

$$L_{\text{int}} = \frac{1}{2}(\dot{q}_1^1 - \dot{q}_N^2)^2$$

they are not, and invariance requires $f_1(t) = f_2(t)$; only global time-dependent displacements act as a symmetry. There are now $2N-1$ relative distances, and therefore the number of degrees of freedom for the coupled system exceeds the sum of degrees of freedom from the individual systems by one: $\#(1 \cup 2) - (\#1 + \#2) = 1$. Gauge-invariant degrees of freedom which cross the subsystem boundary, e.g. $q_1^1 - q_N^2$, and we are unable to describe the coupling of the two squadrons merely by using the relational gauge invariant variables of each subsystem; we require their gauge variant degrees of freedom for coupling.

Rovelli’s succinct explanation for the existence and importance of gauge-variant variables follows:

“The gauge invariance of [one squadron] is invariance under an arbitrary time-dependent displacement of all spaceships. The position of the individual ships is redundant in the theory insofar as we measure only relative distances among these. But in the physical world, each ship has a position nevertheless: this becomes meaningful with respect to one additional ship, as this appears. [this additional ship plays the role of a coupling between the two squadrons]. In other words, *the existence of a gauge [symmetry] expresses the fact that a reference is needed to measure the position of a ship. It expresses the fact that we measure position relationally.* It suggests that this is the general case for all gauge systems. The fact that the world is well described by gauge theories expresses the fact that the quantities we deal with in the world are generally quantities that pertain to relations between different parts of the world, that is, which are defined across subsystems. The example shows how a gauge quantity typically describes an individual component of relative observables.” [my italics]

To cement the ingredients we will need from Rovelli’s discussion, let me offer another, more ‘modern’ example. Suppose Alice describes regional spacetime events in her Lorentz frame. She has fixed Lorentz invariance—that ambiguity no longer acts on *her own frame’s* description of events. The gauge has been fixed, and a gauge-fixed system is a gauge-invariant system. If Alice’s region comprises the entire spacetime, we lose nothing by shedding information characterizing Alice’s frame itself. Her coordinate description of the Universe encodes all the gauge-invariant physics and we can forget about Lorentz invariance.⁶ For a single region, nothing is lost by gauge-fixing.

⁶This does not endow this frame with a privileged status, a point emphasized by John Bell [Bel76]: “we need not accept Lorentz’s philosophy [of the ether] to accept a Lorentzian pedagogy. Its special merit is to *drive home the lesson that the laws of physics in any one reference frame account for all physical phenomena.*” [my italics]

But, if the region does *not* comprise the entire Universe, we must relate Alice’s description of her region to Bob’s description of his. To accomplish this, we need to keep some information characterizing the perspectives themselves. Rovelli’s example at least recognizes the importance of this information.

Because we have no local description of the physical degrees of freedom of gauge theories, we use particular (but arbitrary) perspectives. But the multitude of perspectives is usually seen as superfluous, since the assumption is that we can always gauge-fix to a single selective class (and so secure a gauge-invariant description). However, without global access to the system, i.e. when we know ourselves to be provincial in the World, gauge-fixing is premature. Instead, we need to stay flexible about perspectives, which entails keeping gauge-variant elements in the theory.

As an analogy, imagine a given region in the World, where an edict obliges the population to speak Brazilian Portuguese. Since dictionaries there are no longer useful, they are all destroyed. Later on, this region comes into contact with another region where only German is spoken. In the analogy, a gauge-transformation is effected by inter-language dictionaries, and a gauge-fixing is the settling on one language, once and for all. Thus a regional gauge-fixing abjures the possibility of a larger World with other languages, rendering dictionaries entirely redundant. But if we don’t have the full picture—if there is or could be other regions—we need to keep some flexibility of language, some language-variance if you will. In other words: just as the lack of a tractable ‘god’s-eye-view’ of the system obliges us to start using perspectives, the lack of a global view obliges us to keep them.

In sum: Rovelli implicitly attacks a naive approach to all theories with local gauge symmetry. According to this naive approach, gauge degrees of freedom are at bottom just redundancies in our descriptions. We only need these descriptive redundancies so as to write the theory locally, but then we cull them by fixing the gauge, and in the end a gauge-fixed theory carries just as much ontological content as the non-fixed one. All agree that no objections to this approach arise, *if* we have access to the entire Universe. But they do arise, if we don’t have such access.

1.2 A multitude of perspectives, not a position in the world

Although Rovelli has identified one aspect of the importance of gauge-variance which we will later exploit, Rovelli’s construction is not fully perspectival in the sense that I want to implement here. For note that he requires the ‘bare position’ of one spaceship in each of the two squadrons—a gauge-variant quantity in each subsystem—and this only insofar ‘as the other squadron appears’. These two gauge-variant elements are then combined together, forming one more gauge-invariant quantity. But what if we had chosen a different basis as the gauge-invariant variables of each squadron? Or what if instead of the ‘bare position’ of a ship, we would like to use a different gauge-variant parameter for the first squadron, e.g.: the center of mass of the squadron? Or if the interaction term had been different? Without knowledge of which kind of second subsystem will appear over the horizon, it could be a non-trivial task to couple the two new parameters in a gauge-invariant fashion; the two choices must be not only gauge-variant, but gauge-covariant in the right manner.

Apparently oblivious to this complication, Rovelli justifies the retention of the gauge-variant variable by resorting to our intuitions about the ‘real world’: “But in the physical world each ship has a position nevertheless: this becomes meaningful with respect to one additional ship, as this appears.” But why single out a ‘bare

position’ of a single ship as the gauge-variant handle? How to choose which ship? And can these choices only make sense ‘as the other squadron appears over the horizon’? After the second squadron has appeared we can in any case, at least in Rovelli’s example, describe everything in terms of gauge-invariant observables for the entire system.

No: the handle must be, as it were, flexible—able to accommodate multiple choices of parametrization—and it must be ready beforehand, ready to be “grabbed by” any *potential* second system.⁷ Thus I will here propose a more general framework, where such a handle automatically embodies covariance, without recourse to the bare position of a ship or to the nature of a second subsystem.

The tenor of the construction is simple. First, we will keep the description of the system through an arbitrary frame; we will keep all perspectives. Then, for any given process, we will extricate the physical component of this process from the effects of a change of perspective. This can be done with respect to each perspective. Although both the physical and the gauge components of the process are gauge-covariant, one component has a bijective correspondence to the gauge-invariant degrees of freedom of the region, and the gauge component tells us how to couple to other regions. We will go through the corresponding mathematics shortly, in the field-theoretic context.

Accordingly, in this paper we will jettison knowledge about the ‘bare location’ of one ship with respect to another squadron, in favor of maintaining arbitrary perspectives. Accommodating this plurality of views—different descriptions of the subsystems by arbitrary frames—moves our discussion to a big, extremely redundant configuration space, usually described by principal fiber bundles. For our purposes here, a principal fiber bundle is essentially a space in which all the different descriptions of the same physical configuration organize themselves tidily in subspaces, called ‘the fibers’. Taking the equivalence class of each fiber would give you a (purely abstract) representation of the physical content of that configuration—the ‘view from nowhere’. The space of all equivalence classes is usually called ‘base space’, and in many field-theoretic cases it is an unwieldy “cubist” landscape, in which no standard notion of localization is possible.

⁷This is reminiscent of a discussion of charts and atlases in the characterization of differentiable manifolds [Lan99]; and this line of thought can be pursued to reach conclusions encompassed by ours [DF16, GR17], by analogously keeping, in addition to the chart description, the information of the embedding maps of the chart into the manifold (see also [BK95, IK85, BK60]). Explicitly: this parametrization, generalizing Rovelli’s ‘position in the real world’, would correspond to an embedding map of a chart. In other words, in the differential geometry language, one could argue that if a coordinate description of the real world M is complete, i.e. there is a global chart, whose inverse we write as $\phi : U \subset \mathbb{R}^n \rightarrow M$ such that $\phi(U) = M$, then this description in $U \subset \mathbb{R}^n$ exhausts the ontological content of the real world, and we lose nothing by confining our talk of reality to elements of U . But if $\phi(U) \neq M$, then we must confer reality on the abstract manifold M , in which the domain of the chart is embedded. In this case, the embedding map itself is significant irrespective of whether there is another intersecting chart. Nonetheless, Rovelli’s position is tenable to the extent that specific properties of ϕ are only relevant once we are required to use transition functions, i.e.: once there is an intersection with another chart. Then, the “position of the ship in the real world” (i.e. the embedding of U) becomes indispensable, for it needs to be compared to another perspective (whereby we obtain the transition $(\phi')^{-1} \circ \phi : \phi^{-1}(\phi'(U') \cap \phi(U)) \subset U \rightarrow U'$). This distinction—between the reality of the manifold M vs. the reality of transition maps—mirrors that between objective realism and Rovelli’s relationalism. Rovelli’s view is part of a larger theoretical framework with which he approaches quantum mechanics, i.e. “relational quantum mechanics”, in which a realistic interpretation of quantum subsystems only makes sense in relation to other subsystems [Rov96].

1.3 Gauge-fixing and boundaries in field theory

Apart from its undemocratic nature, gauge-fixing is seen by many to resolve issues of redundancy with gauge. In this subsection, our exploration of gauge field theory in a bounded region will dispel the absolute powers of gauge-fixing, or at least fill them with nuance.

Gauge-fixing in field theory without boundaries— We ended the last section with a brief mention of the sort of spaces that accommodate a multitude of perspectives, generally referred to as principal fiber bundles. As evidenced by Earman’s quote [Ear03], the mathematical theory of principal fiber bundles is a staple of gauge theories, be they theories of finite-dimensional systems (like Rovelli’s squadron) or more realistic field-theories (like the standard model of particle physics). The discussion above was focused on Rovelli’s finite-dimensional case. Let us now extend it to the field-theoretic case.⁸

In field-theory, a gauge-fixing is perfectly analogous to the fixing of one frame of reference as discussed above. For simplicity, we start this exposition in the absence of boundaries, i.e. when our patch of the Universe is the whole Universe. So here we are assuming a single chart is able to cover the entire manifold.

To recap, the ‘view from nowhere’ is incapable of *locally* describing the full relational content of the system, and therefore we must start from an arbitrary perspective. For instance, in Yang-Mills theory, in the principal fiber bundle language, we can only describe fields over spacetime using a *finite-dimensional* ‘section’ (see the left hand side of figure 2 for an illustration, which distinguishes the finite-dimensional section from the infinite-dimensional one, representing a gauge-fixing)— we usually denote this by A_μ^a (with μ describing spacetime indices and a the Lie-algebra indices).

Since we now have all these arbitrary perspectives, we resort to a gauge-fixing. A gauge-fixing should allow us to directly compare the physics of two states, originally in arbitrary perspectives, by transforming them into descriptions (perspectives) in the same selective class. A gauge-fixing chooses a selective class of perspectives (or a gauge) by “intransigently demanding” that the field configurations satisfy its conditions.

Such conditions should be *physical*, i.e. they should constrain the frame by the way the system looks in that frame, i.e. by constraining the perspective. For example, in Newtonian mechanics, this could be given by inertial frames, center of mass, and some anisotropy parameters. In the context of electromagnetism, such a condition could be that the electromagnetic potentials of the configuration in the chosen frame are transverse $\partial^\mu A_\mu^f = 0$ (with script f standing for ‘fixed’; it is not a Lie-algebra index, which is absent in electromagnetism). In sum, the condition should be a condition *on the perspective*, not the frame.⁹

Concretely: starting with any potential A_μ , the gauge-fixing procedure requires us to “rotate” our frame (or “translate” your gauge potential) until its condition is met. For example: suppose we are in a spacetime of Euclidean signature, then if the gauge-group acts as $A_\mu \rightarrow A_\mu^\lambda := A_\mu + \partial_\mu \lambda$ for λ an element of the group, one

⁸For the issues that arise in the extension of Rovelli’s argument to the Yang-Mills case, see [Teh15]. Teh discusses the two sorts of couplings for field-theory which we will cover: between different fields and between different regions.

⁹ Indeed, as we will later see, since there is no canonical isomorphism between the gauge group and an orbit of the field under gauge transformations. We cannot implement conditions to fix the frame itself—the field-content must always mediate our choice of frame.

such gauge-fixing condition would require us to solve

$$f(A) := \partial^\mu(A_\mu^\lambda) = \partial^\mu(A_\mu + \partial_\mu\lambda) = 0 \quad (1.2)$$

for λ , obtaining the field-*dependent* $\lambda^f(A)$. In the Yang-Mills case, this is known as the Landau gauge. Our gauge-fixed field is then the package $A_\mu^f := A_\mu + \partial_\mu\lambda^f(A)$. When there are no boundaries, such a gauge-fixing procedure can eliminate all redundancy, allowing access to the “gauge-invariant” degrees of freedom. Equivalently, in the absence of boundaries, a gauge-related configuration, $(A^f)^\lambda$ is mapped back to A^f by a gauge-fixing. That is because

$$\partial^\mu((A_\mu^f)^\lambda) = 0 \quad \text{iff} \quad \lambda = 0. \quad (1.3)$$

Having got what it wanted, the gauge-fixing immediately forgets—it does not distinguish between your or your friend’s original perspective, as long as you both meet its demands in the same way. Geometrically, a gauge-fixing maps two field configurations which differ by a frame (or gauge) transformation to the same representative configuration.

To sum up: when there are no boundaries, gauge-fixings erase the information pertaining to the arbitrary choice of frames, and thus do not leave gauge-variant handles available for describing coupling. The gauge-fixed theory remembers only a gauge-invariant package A_μ^f , and forgets both A_μ and $\lambda^f(A)$. The latter transforms covariantly as $\lambda^f(A') = \lambda^f(A) - \lambda'$. Therefore given two gauge-fixings, $f_1(A) = 0$, $f_2(A) = 0$, the difference

$$F(f_1, f_2)(A) := \lambda^{f_1}(A) - \lambda^{f_2}(A) \quad (1.4)$$

is thus a gauge-covariant,¹⁰ group-valued functional of the perspectives; it relates the two selective perspectives. As we will see in section 2, $\lambda^f(A)$ is an elementary precursor to the object we will use to distinguish changes of frame from physical processes: the *relational connection-form*. Apart from selecting a representation, the connection-form will be used for relating those representations, as in (1.4).

If the main theme of this paper is the role of the connection-form in the treatment of gauge theories in bounded regions, the most prominent sub-theme is that an approach to gauge systems characterized by perspectives can be pried away from an approach to gauge systems characterized by abstract (Lie) groups and algebras. The difference between selecting a perspective and selecting a gauge parameter in bounded regions, which we will now witness, manifests this sub-theme.

Gauge-fixing in field theory with boundaries— This gauge-fixing-induced amnesia causes no trouble if we have access to the Universe as a whole; for example, if Rovelli’s world had a single squadron. Nonetheless, the discussion of the last section, conducted in the absence of boundaries, is also germane when they are present.

Given two regions $M_1, M_2 \subset M$ which share a boundary, ∂M , we have the respective restrictions of the A_μ field onto the regions, A_μ^1, A_μ^2 . For bounded regions, finding the perspectives A_μ^1 (resp. A_μ^2) which belong to the selective class—satisfying (1.2)—also requires us to stipulate boundary conditions on the *gauge-parameter* in that region, λ^1 (resp. λ^2), itself. These are *not* conditions on perspectives. Moreover, the resulting change of frame—e.g.: the solution λ^{1f} —carries

¹⁰Only in this simple Abelian case is it actually gauge-invariant.

a dependence on boundary conditions throughout the region, and not just at the boundary, as we will shortly see.¹¹

As opposed to what is implied by (1.3), in the presence of boundaries two arbitrary perspectives which are related by a change of frame can correspond to selective perspectives which are different everywhere. This is easy to see: Again, we start with A^f satisfying the constraints on perspectives, and ask if a transformed perspective, $(A^f)^\lambda$, would be sent back to the same A^f by the same gauge-fixing constraint which A^f satisfies, e.g. $\partial^\mu A_\mu = 0$. But here equation (1.3) leads to $\nabla^2 \lambda = 0$, whose solution is not $\lambda \equiv 0$ for $\lambda|_{\partial M} \neq 0$,¹² as opposed to the general solution in the absence of boundaries, (1.3), which does vanish.

The possibility of non-zero solutions for λ in this circumstance evinces the strange character of gauge-transformations at the boundary: different choices of λ at the boundary seem to correspond to actually *physically distinct, and yet gauge-related*, configurations. In other words, we face an anxious question: Should we take those perspectives which are gauge-related and yet are not mapped to the same gauge-fixed perspective to represent physically distinct states of affairs?

The appearance of λ at the boundary is therefore far from innocuous; its presence is problematic from two related angles. First, as described above, it may license physical status to gauge degrees of freedom at the boundary, and second, fixing λ on each side incurs the “language translation” issue we saw in the previous section.

To see the second issue more clearly here, even if we choose the same gauge-fixing (e.g. $\partial^\mu A_\mu = 0$) on each side of the boundary, i.e. on M_1 and M_2 , but different λ_1, λ_2 , it follows that A_μ^{1f} and A_μ^{2f} would not smoothly join at ∂M (they would not even be C^0 there, since they differ by $\lambda_1 - \lambda_2$, which is not fixed by the gauge-fixing).¹³ Alas, there are no more gauge transformations available; the gauge-fixing on each side has left no gauge-covariant handles, and we can no longer ‘rotate’ A_μ^{1f} and A_μ^{2f} to match. This can be seen as a field-theoretic analogue of both the metaphor of the two regional languages incapable of translation, and Rovelli’s obstruction to the description of the coupling between the gauge-fixed squadrons.¹⁴

Strangely, upon gluing the regions back together without the prior regional gauge-fixing, the choice of λ at the boundary becomes completely immaterial. If \mathcal{G}_1 is the group of gauge-transformations in M_1 , and \mathcal{G}_1^o are those that reduce to the

¹¹It is also important to note that a Hamiltonian approach would be plagued by precisely the same issues. There, a gauge-fixing is also a function of the canonical variables, $\phi(p, q) = 0$, and requires further, non-physical, determination at the boundaries.

¹²The solution is a harmonic function which depends on the value of λ at the boundary. For example, according to a mean value theorem for $\nabla^2 \lambda = 0$ with a prescribed λ at ∂M_1 , for a domain $M_1 \subset \mathbb{R}^n$, and $\lambda \in C^2(M_1)$, assuming M_1 to be a ball centered at x with radius R , then $\lambda(x) = \frac{1}{n \text{Vol}(B^n) R^{n-1}} \oint_{\partial M} \lambda(y) dy$ where $\text{Vol}(B^n)$ is the volume of the n-dimensional unit ball. This illustrates the ‘percolation’ of boundary gauge transformations into the determination of the frame in the interior of the region.

¹³ Although we are here focusing on the single chart domain, this discrepancy is perhaps best visualized in the case of two intersecting charts. Let $M_1 \cap M_2 = M_{12} \subset M$ be itself an n-dimensional submanifold of M . While it is true that, over M_{12} , one demands $A_\mu^1 = (A_\mu^2)^\lambda$ for some gauge transformation $\lambda : M_{12} \rightarrow \mathbb{R}$, the gauge-fixings of regions M_1 and M_2 could be different, and there are no gauge-transformations to relate the two so selected perspectives. Even more worrisome, due to the non-local nature of gauge-fixing, even if one has chosen the same constraint on the partial perspectives, e.g. $\partial^\mu A_\mu^1 = \partial^\mu A_\mu^2 = 0$, the different domains in which these equations are solved (and in particular the different boundary conditions on λ^1, λ^2), imply that in general $A_\mu^{1f} \neq A_\mu^{2f}$ on M_{12} . So even if they speak the same language they can’t communicate!

¹⁴ In Rovelli’s case, when $\dot{q}_1^1 = \dot{q}_N^2$, the coupling is no longer problematic. In the field-theory case, there might likewise be particular coincidences, $\lambda_1 = \lambda_2$, for which the gauge-fixed configurations match at the boundary, but neither of these cases is generic and so do not resolve the problem.

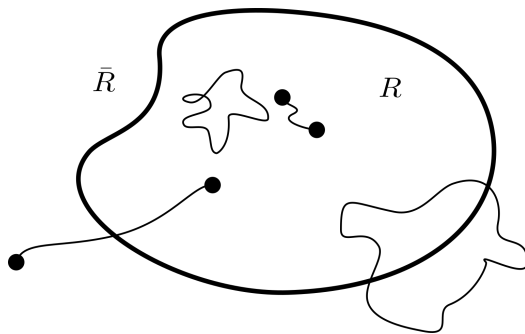


Figure 1: A region and a selection of loops both contained and not contained in it. There exist gauge-invariant functions whose support is not contained in either the inside or the outside of the region. Therefore the space of physical, gauge-invariant functions does not factorize, as the quotient of field space itself does not factorize: $\Phi/\mathcal{G} \neq \Phi_R/\mathcal{G}_R \otimes \Phi_{\bar{R}}/\mathcal{G}_{\bar{R}}$ where Φ is the field space associated to the entire region, and its restrictions are denoted by a subscript (and we are supposing each quotient is inheriting a vector space structure, for which we can form the tensor product).

identity at ∂M , i.e.

$$\mathcal{G}_1^o := \{g \in \mathcal{G}_1 \mid g(x) = \text{Id for } x \in \partial M\} \quad (1.5)$$

we seem to have added physical degrees of freedom in correspondence to $\mathcal{G}_1/\mathcal{G}_1^o$ —for each physical degree of freedom,¹⁵ a different element of $\mathcal{G}_1/\mathcal{G}_1^o$ would correspond to a new physical degree of freedom as well—which must disappear once the regions are glued back together!¹⁶ Unless these boundary degrees of freedom somehow disappear upon gluing, at least naively, gauge-fixing does not commute with gluing—gauge-fixing first, one gets boundary degrees of freedom, whereas gluing first leaves no room for them.

As with Rovelli’s example, the total count of physical degrees of freedom differs from the sum of the regional degrees of freedom. In field-theory, one can imagine Wilson loops which cross the boundaries, and thus support gauge-invariant functions only on the joint region, being gauge-variant when restricted to each region (see figure 1). In this light, it seems reasonable that imbuing the boundary λ with a genuine physical character and appropriate covariance properties could ‘close off’ the Wilson loops, associating extra physical degrees of freedom to each region.¹⁷

New degrees of freedom at the boundary— Here we see that the boundary choices for λ , being non-perspectival and gauge-variant, create puzzles. These puzzles could be to a certain extent ‘resolved’, either by disallowing gauge transformations at the boundary,¹⁸ or, without fixing λ at the boundary, by the following procedure:

¹⁵One could be tempted here to refer to such degrees of freedom as “bulk” degrees of freedom. But I will refrain from doing this, remaining agnostic about localization.

¹⁶This does sound a lot like the job description of ghosts in non-Abelian QFT. We have clarified this analogy in [GR17] and I will comment on it further along the paper.

¹⁷ There are many approaches that endow with physical character the punctures of surfaces by Wilson loops (see e.g. [Rovge]), but usually such surfaces have some prior physical significance: for instance, they may be null surfaces as in [Wie17, Wie18]. But here I will be arguing for what I consider to be a more minimalistic approach.

¹⁸This is related to the algebraic approach to the issue of boundary degrees of freedom, for which a vast literature exists, especially in the context of entanglement entropy, see e.g.: [CHR14, DHM18]. The values of the physical fields at the boundary are also fixed in many of these approaches.

first, promote the boundary gauge parameter required for the gauge-fixing to the same status as the original fields. Now it is part of the perspective, and, if selecting representatives is required, it should be done in a similar fashion as with the standard perspectives. Second, imbue these new degrees of freedom with covariance properties under changes of gauge, so that the coupling of two regions becomes gauge-invariant (thus fulfilling the relational role advocated by Rovelli).

This is essentially what Freidel and Donnelly have accomplished, in a different language, and with different purposes, in [DF16] (see also [Don14, RT74, BK95, IK85, BK60] for related approaches and precursors). As we will discuss in section 1.4, their purpose was to obtain gauge-invariant Noether charges. They introduced new degrees of freedom into the theory which appear only at the boundaries of subsystems. In other words, these degrees of freedom are kept in reserve, to be pressed into battle only when our region is confronted with another, when they can be used for coupling.

But one naturally asks: what is the meaning of these degrees of freedom? And why are they not required if we look at the global physics in the absence of boundaries, i.e. why are they not required if we look at A_μ^f or even A_μ ? Since these ‘edge-modes’ now swell the ranks of the *physical* degrees of freedom—but only to be associated to edges (or boundaries)—what are the implications for conserved charges, that tell-tale sign of symmetry?

1.4 Boundaries and charges: a brief intro.

In the first subsection, 1.1, I have argued that the nature of gauge degrees of freedom is ‘perspectival’. Following Rovelli’s example, I argued that in the presence of subsystems this view underlines the significance of gauge-variant objects, to the detriment of gauge-fixed quantities. Then, in the second subsection, 1.2, I argued contra Rovelli: keeping the flexibility required to couple to different subsystems, and also required to support different gauge-fixings (different choices of gauge-invariant variables), demands a more formal, pluripotent structure than the one Rovelli offers (“a position in the World”). These were the perspectives. The challenge is to extract physical information while keeping a multitude of perspectives in play; this will be the job of the connection-form, as we will see from section 2.3 onward.

Indeed, in the third subsection, 1.3, we saw that the usual procedure of gauge-fixing in bounded regions still suffers from the problem Rovelli identified: coupling between different regions is still hampered. There is again a mismatch in the counting of the physical degrees of freedom belonging to the parts and belonging to the whole. In the field-theoretic context, this mismatch can be associated to an artificial specification of gauge-parameters at the boundary.¹⁹

Its consequences are not merely conceptual, but can be felt “on the ground” in the physics of charges. For even if we grant that the spirit of gauge degrees of freedom is relational, only to be fully seen in the coupling of subsystems, their bite comes from the teeth of the Noether theorems. And the unphysical parameters at the boundaries are implicated in the complicated procedure of obtaining physical charges from *local* gauge symmetries through Noether theorems.

The teeth of gauge: Noether theorems – The first Noether theorem implies that global symmetries result in conservation laws that we observe empirically

¹⁹By “artificial” I mean it is unrelated to conditions $f(A) = 0$ through which perspectives can be selected.

(see [BB00, BC03] and references therein for a historical account and philosophical discussion). The conservation of electric charge, for example, is a consequence of global $U(1)$ symmetry in non-relativistic quantum mechanics, conservation of energy in special relativity is a consequence of time translation symmetry, etc.

Whereas global symmetries are usually understood to possess empirical significance, for instance in the form of conserved charges, the empirical significance of local gauge symmetries is much looser. For these symmetries are in the domain of Noether’s *second* theorem, whose effect is better construed as constraining the form of the equations of motion, rather than as straightforward conservation of charge.²⁰

Explicitly, the puzzlement about conservation laws comes into focus in the distinction between Noether’s first and second theorem. As two of the foremost experts in the topic characterize the issue: [BB02]:

“[The puzzle] is encountered when one tries to define the charge related to a gauge symmetry “in the usual manner”, by applying Noether’s first theorem on the relation of symmetries and conserved currents. The problem of such an approach is that a Noether current associated to a gauge symmetry necessarily vanishes on-shell (i.e., for every solution of the Euler-Lagrange equations of motion), up to the divergence of an arbitrary superpotential.”

Gauge charges are related to constraints on initial data, at Cauchy surfaces. Therefore, at least in the absence of boundaries of these surfaces, such charges should vanish when the constraints are obeyed [MH92].

Following the standard methods by which one can associate charges to gauge-generators in a covariant setting, these so-called Noether charges indeed vanish in the absence of spatial boundaries.²¹ However, this is not the case for bounded subregions of the Cauchy surface.

Let us take electromagnetism as an example. We define the electric field as a $d-2$ form (for d the spacetime dimension) $E = *F$, where F is the electromagnetic curvature tensor and $* : \Lambda^n(M) \rightarrow \Lambda^{d-n}(M)$ is the Hodge star operator, taking n -forms (i.e. elements of the alternating tensor product of linear functions on TM , denoted by $\Lambda^n(M)$) to the complementary $d-n$ forms, $\Lambda^{d-n}(M)$. Then, let $R \subset M$, bounded by ∂R , be a region of the Cauchy surface. For a gauge-generator $\xi(x)$, the covariant procedure (which we explain in section 3.3) yields a Noether charge for this field content and gauge parameter:

$$Q[\xi] \approx \oint_{\partial R} (\text{Tr}(E(x)\xi(x)) + \alpha[\partial_\mu \xi]) d^{d-2}x, \quad (1.6)$$

where \approx indicates equality up to terms that vanish when the equations of motion (including the constraints) are satisfied, also called “on-shell equality”; and α is an *arbitrary* (field-independent) $\Lambda^{d-2}(M)$ -valued linear functional of $\partial_\mu \xi$; in the language we will introduce in section 3.3, $\alpha \in \Lambda^{d-2}(M) \otimes \Lambda^1(\Phi)$ ($\Lambda^1(\Phi)$ means it is a linear functional on field-space, cf. section 2.2). The point is that such charges are doubly troubled: not only do they seem to be non-zero for those gauge-generators

²⁰For instance, in the case of the internal groups of Yang-Mills, they will yield Bianchi identities for the curvature tensor; while for diffeomorphisms, they will yield generalized Bianchi identities for metric concomitants (see [MH92] chapters 1-3, and [BC03] for a philosophical discussion). In the Hamiltonian language, they give us the local constraints and their closed algebra.

²¹The standard method I refer to here is the covariant symplectic framework, which will be more fully explained in section 3.3. Here I aim to merely illustrate the problem.

which do not vanish on the boundary, but the arbitrariness of α even makes it hard to define a specific charge for any given transformation!

Therefore, although talk of conservation at codimension one spatial surfaces (codimension two in spacetime) might sound a lot like a Gauss law, there are important differences: the conserved quantities one obtains from Noether’s second theorem may be infinite in number, and do not solely depend on the original physical fields and geometrical shape of the d-2 surface ∂R in (1.6). The integrated conserved charges from (1.6) will be smeared by the generators of arbitrary gauge-transformations at the boundaries, and may depend on arbitrary choices of superpotentials, which strains interpretations of such charges in terms of standard physical conservation laws.

Naively calculating charges for continuous symmetries by Noether’s second theorem therefore depends on arbitrary choices at the boundary. According to the arguments from the previous section, this arbitrariness at least partially corresponds to the one we saw in the gauge-fixing procedure.

As Barnich and Brandt assessed the situation in 2003 [BB02],

“The problem of defining and constructing asymptotically conserved currents and charges and of establishing their correspondence with asymptotic symmetries *in a manifestly covariant* way has received of lot of attention for quite some time.” [my italics]

And that interest has only grown in the past years (see [Str18] and references therein).

Gauge-variant ‘edge-modes’ – Taking Regge and Teitelboim’s seminal work on conserved charges at asymptotic infinity [RT74] as a guide,²² Donnelly and Freidel introduced a new special type of gauge-variant degree of freedom [DF16]. It encodes intrinsic boundary degrees of freedom, for boundaries both asymptotic and finite, and either real or imagined.

The manner in which Donnelly and Freidel design the gauge-covariance properties of the boundary degrees of freedom in (1.7) is such that their effect on the charges will end up canceling the first term of (1.6).²³ But if this were the end of the story, Donnelly and Freidel might have had a problem: this procedure would identically cancel all charges. That is where the following equation (1.8) comes in.

Calling the gauge-group-valued new degree of freedom $\kappa(x) \in G$ for $x \in \partial M$, with A for the Yang-Mills gauge potential, ξ the generator of an arbitrary gauge transformation, $D = d + [A, \cdot]$ the gauge-covariant derivative (with the Lie-algebra bracket $[\cdot, \cdot]$ and d the exterior derivative) and δ_ξ indicating the infinitesimal action of the gauge-transformation on the variables, they define:

$$\text{Gauge symmetry:} \quad \delta_\xi \kappa := \kappa \xi, \quad \delta_\xi A := D\kappa \quad (1.7)$$

$$\text{Surface symmetry:} \quad \Delta_\xi \kappa := \xi \kappa, \quad \Delta_\alpha A := 0 \quad (1.8)$$

²²To my knowledge, Regge and Teitelboim were the first to introduce new physical “embedding variables” at boundaries to reinstate a broken symmetry (in their case, Poincaré). But there are many more, very closely related precursors to Donnelly and Freidel. Most notably, [BK95, IK85, BK60], have all introduced degrees of freedom to parametrize gauge choices.

²³In fact, the terms are added at the level of the symplectic potential, and not directly at the level of the action nor of the charge. Moreover, their introduction is predicated on the field-dependence of the gauge-parameters at the boundary. This is less of a stretch of the usual concept than might at first seem, for, as we saw in section 1.3, boundary gauge transformations have a “physical” character. In other words, changing the gauge at the boundary implies a physical change of the field. .

where I have suppressed: (i) indices (notation will be more thoroughly introduced in equation (2.11)) and (ii) the specific form of the left and right action of the Lie algebra on the surface degrees of freedom. The *surface symmetries* are supposed to represent redefinitions solely of the boundary degrees of freedom; redefinitions by the restriction of same gauge-group \mathcal{G} to the boundaries.²⁴

In other words, Donnelly and Freidel are able to precisely exploit the arbitrariness in the definition of α so as to cancel the charges associated to ξ in (1.6). The covariant properties of these new boundary degrees of freedom are designed to render the charges Q completely gauge-invariant. For the action of the gauge parameters given in (1.7), no charges are left. However, these edge modes also bring their own infinite contributions to conserved charges; namely all the charges associated to the action of the gauge-parameters through Δ_ξ , given in the second line, (1.8). In the end, the tally is unchanged: there remain a continuous number of left-over ‘physical’ degrees of freedom associated even to imaginary finite boundaries, and these yield the gauge-invariant charges.

In my view, starting from the field content and the boundaries of a given region, we would like physical charges to be entirely functionally dependent on physical/geometrical features of the boundaries—and not at all dependent on any further arbitrary, unphysical degrees of freedom. As with [BB02], we should then expect charges to be strictly associated to *physically relevant symmetries*, such as regional/asymptotic Killing transformations in the case of general relativity and non-Abelian Yang-Mills.²⁵ But Donnelly and Freidel’s formalism does not solve this issue.²⁶

To sum up, the reason to deem Donnelly and Freidel’s new charges unphysical is simply this: their charges depend on a new field which has no physical interpretation if boundaries are not present. In my personal view, this approach may cause some confusion if boundaries are not material entities, but just figments of the theorist’s imagination. Looking ahead, the point is that both imbuing physical significance to the left-over gauge choice λ at the boundary, and making up new degrees of freedom are unnecessary for solving any of the issues considered above: there is another way!

Up to this point of the paper, we have essentially reviewed some of the vexing issues arising from conjoining gauge and boundaries, and considered one attempt at resolution [DF16]. Such issues are technical and conceptual, and must be embraced and dealt with by all who worry about the nature of gauge.

Now, we arrive at a different attempt at resolution, one which seems to patch up the holes left by the first part of the paper. I will give a conceptual exposition of [GR17, GR18, GHR18], a recent geometric generalization of Donnelly and Freidel’s ‘edge-modes’. Instead of new degrees of freedom held in reserve at the borders, ready to be pressed into duty, we use *relational properties of the original field-content*. These are present everywhere within the regions, and are also able to perform the special translation functions at the boundaries, for their covariance is a consequence of their geometric nature.

²⁴ The motivation for this particular cordoning of redefinitions is not clear to me.

²⁵ Or, more precisely: strictly related to reducibility parameters [BB02].

²⁶ But they also do not hang much on the issue. In the words of Donnelly, “I am comfortable with the point that our charges depend on new degrees of freedom. And indeed, those are not there when the regions are combined [...] And I also agree that one could try to build such degrees of freedom out of fields already in the theory, which is an appealing idea.” Private communication.

2 Horizontal geometry and the connection-form

We now turn the spotlight to the main actor of this paper: the *relational connection-form*. Given our purposes here, the following sections will consist of an extremely abridged account of the work contained in [GR17, GR18, GHR18], without the technical detail that can be found there. The first subsection, 2.1, summarizes the role and construction of the connection-form. The second subsection, 2.2, puts up the technical scaffolding, by establishing some required notation and concepts and giving an overview of the horizontal geometry of field-space. The third subsection, 2.3, then gives an abstract mathematical definition of the connection-form.

2.1 A summary of the construction

The connection-form resolves the problems listed above by a much simpler route—it defines physical and pure gauge transformations intrinsically to each perspective (e.g.: to each A_μ) without having to introduce spurious degrees of freedom [Gom11, GR17, GR18, GHR18]. It embodies the function of ‘Rovelli’s gauge-variant couplings’ by describing the system from a given arbitrary perspective, while keeping the required gauge-variant handles characterizing that perspective.

Let us stress two main differences between a connection-form and a gauge-fixing: as I have repeatedly emphasized, gauge-fixings are amnesiac—they forget information about the frame and therefore do not leave gauge-variant handles around to describe coupling between subsystems.

A second difference is that a gauge-fixing is in the business of *defining a selective class of perspectives*, while a connection-form is in the business of *defining changes in perspective*. That is, given any field configuration φ , a gauge-fixing eats it up and spits out a φ^f , which is a state gauge-equivalent to φ but with a different perspective, viz. one belonging to the selective class.

Here is what the relational connection-form does instead: given the arbitrary state of affairs φ and an infinitesimal transformation of that state, \mathbb{X}_φ (or $\delta\varphi$ in the usual notation), the connection-form decomposes that transformation into a pure gauge translation—the change of frame—and a *physical* change. This change is physical with respect to ϖ and with respect to that particular field-content of φ ; it depends on the perspectives in a covariant manner, as evidenced by equation (2.17), below. The pure gauge part is what provides the handles for coupling at the boundaries as we will see in section 3.2; there are no new degrees of freedom that make curious special appearances there.

Surprisingly, these advantageous properties can be directly inherited from the geometry of field space, as we will see more directly in section 3.1. And, due to its geometrical representation, the connection-form defines a modified, horizontal symplectic geometry in field-space. This sort of horizontal geometry is germane to the distinction between physical and pure gauge (vanishing) Noether charges, more so than standard symplectic geometry, even if the difference is only explicit at corners and boundaries. Indeed, as we will see in section 3.3, the connection-form is much more conservative on the topic of conserved Noether charges than Donnelly and Freidel’s boundary degrees of freedom.²⁷

²⁷Although it has not yet been done explicitly, it is expected that the results for entanglement entropy in gauge theories obtained by the procedure of Agarwal et al [AKN17] are entirely reproduced by the connection-form.

2.2 The geometry of field-space

The more explicit construction relies on geometrical structures of field-space, to which we will now turn.

In the forthcoming, Φ is to be thought of as the space of all possible field configurations—all the possible states of affairs of the field, with no regard to satisfying the equations of motion. This vastly redundant space represents all of the possible ‘perspectives’ on the same state of affairs, as discussed in section 1. Local gauge transformations act locally on the fields in question, relating the perspectives. The action of the group forms “fibers” in this space, partitioning the entire field-space into different equivalence classes.

Such a fibration looks a lot like the standard principal fiber bundles we usually encounter in gauge theories. There, base space is usually just spacetime. However, here, in this infinite-dimensional field-theoretical context, due to a lack of local parametrization of all the local gauge-invariant degrees of freedom, base space can only be characterized as the quotient Φ/\mathcal{G} , “the moduli space of physical field configurations”.

In the case of gauge potentials, the moduli would be obtained by setting $A' \sim A \Leftrightarrow A' = R_g A$ for some right action of the gauge group (R_g) on the potential, and we would denote the equivalence classes by $[A] = [A'] \in \Phi/\mathcal{G}$. For diffeomorphisms acting on spatial metrics, elements of Φ/\mathcal{G} would be geometries—i.e. each *indivisible point* of Φ/\mathcal{G} would consist of the complete geometric structure of space, without redundancy. As far as the ‘view from nowhere’ can be implemented, it applies only to Φ/\mathcal{G} .²⁸

Apart from the status of the base space, there is another important disanalogy between standard finite-dimensional principal fiber bundles (PFBs) and Φ seen as a fibered space over Φ/\mathcal{G} . As we will see in section 3.3, it is crucial that the latter is *not* a bona-fide PFB. That is because not all the orbits in Φ are isomorphic—some are effectively of lower dimension than others. This structure emerges from differences in frame which do not amount to differences in perspective. These differences carve out the landscape features of Φ/\mathcal{G} which will give rise to global charges. The fact that such a structure manifests itself directly in the physical moduli space is what we would expect if such charges have physical content.

Nonetheless, as in the case of finite-dimensional PFBs, it is useful to define a notion of ‘sameness of frame’ when moving from one orbit to another on Φ —this is the geometrical role of the connection-form. Note that, unlike a gauge-fixing, this notion of sameness is only defined infinitesimally, and conforms to different perspectives of the same configuration (i.e. is covariant).

Mathematical preliminaries – More specifically, the stage on which we set the pieces is field-space, denoted by $\Phi = \{\varphi^I\}$. In this notation, φ^I stands for a whole field configuration $\{\varphi^I(x)\}_{x \in M}$, where I is a super-index labeling both the field’s type and its various components, and M denotes the underlying manifold (space or spacetime).

²⁸ Each point of Φ/\mathcal{G} characterizes a full gauge-invariant configuration. However, gauge-invariant observables of GR are known to be non-local, e.g. integrals over spacetime such as $\int R^{\mu\nu} R_{\mu\nu} \sqrt{g}$, and therefore each such point in Φ/\mathcal{G} cannot be put into correspondence with a region or point of M . There can be no surprise appearances of Einstein’s ‘hole argument’ in Φ/\mathcal{G} , since each point of it has identified all diffeomorphically-related metrics. However, there are several problems about endowing a differentiable local orbifold structure on the quotient space of Lorentzian-signature spacetime metrics [IM82]. But it is straightforward for the space of Euclidean signature metrics of any dimension.




Figure 2: A pictorial representation of the field-space Φ seen as a principal fiber bundle. I have represented the quotient space of ‘gauge-invariant configurations’ Φ/\mathcal{G} . I have highlighted a configuration φ —it is one perspective on $[\varphi] \in \Phi/\mathcal{G}$, which is intractably non-local,—its (gauge-transformed) image under the action of $R_g : \varphi \mapsto \varphi^g$, and its orbit $\mathcal{O}_\varphi \cong \mathcal{G}$. On the left hand side of the picture, we have ‘zoomed into’ a representation of φ and φ^g as sections of a vector bundle over the spacetime region M .

In the following, a ‘double-struck’ typeface—as in \mathfrak{d} , \mathbb{L} , \mathbb{X} , etc.—will be consistently used for field-space entities. For instance, we introduce on Φ the deRham differential \mathfrak{d} [CW86, Crn87, Crn88]; it should be thought of as the analogue, on Φ , of the standard spacetime differential (or exterior derivative) d . We will also need a notation for field-space vectors, Lie derivatives and the interior product between forms and vectors in field-space, denoted respectively by \mathbb{X} , \mathbb{L} and \mathfrak{i} . For instance, contraction of a vector with a basis element of the differential forms in Φ , denoted by $\mathfrak{i}_{\mathbb{X}}\mathfrak{d}\varphi^I$, is defined by

$$\mathfrak{i}_{\mathbb{X}}\mathfrak{d}\varphi^I = \mathbb{X}^I \quad \text{where} \quad \mathbb{X} = \int \mathbb{X}^I \frac{\mathfrak{d}}{\mathfrak{d}\varphi^I} \quad (2.1)$$

where I omitted dependence on x in the integrand for simpler notation.

Also for simplicity, we will focus on an internal gauge group, but at this level there would be no difference between this action and, say, an action of diffeomorphisms on the space of metrics (it still acts pointwise on field-space [Gom11, GR17]). Given a *charge group*, say $G = SU(N)$, gauge transformations also inherit a group structure, forming the *gauge group*:

$$\mathcal{G} := C^\infty(M, G), \quad (2.2)$$

\mathcal{G} with elements $g(\cdot) \in \mathcal{G}$ and point-wise composition (we are now moving beyond the Abelian case) over M , and (\cdot) denoting the slot for points of M . Similarly, the Lie algebra of the gauge group is given by $\xi(\cdot) \in \text{Lie}(\mathcal{G}) = C^\infty(M, \mathfrak{g})$ where $\mathfrak{g} = \text{Lie}(G)$. There is a right action of the group:

$$R : \quad \mathcal{G} \times \Phi \quad \rightarrow \quad \Phi \\ \left(g(\cdot), \varphi \right) \mapsto R_{g(\cdot)}\varphi =: \varphi^g. \quad (2.3)$$

Given this right action, we can readily define a map from the Lie algebra of the gauge group, $\text{Lie}(\mathcal{G})$, into the vector fields on field-space (denoted by $\mathfrak{X}(\Phi)$),

$$\# : \quad \text{Lie}(\mathcal{G}) \quad \rightarrow \quad \mathfrak{X}(\Phi) \\ \xi \quad \mapsto \quad \xi^\# \quad (2.4)$$

which associates to an infinitesimal gauge transformation $\xi(\cdot) \in \text{Lie}(\mathcal{G})$ the associated flow $\xi^\# \in \mathfrak{X}(\Phi)$ on field-space. We will denote $\xi^\#$ the *fundamental vector fields*, formally defined iteratively from its action on a scalar $f : \Phi \rightarrow \mathbb{R}$:

$$\xi^\# f := \mathbb{L}_{\xi^\#} f := \frac{d}{dt}\bigg|_{t=0} f(R_{\exp(t\xi)}A, R_{\exp(t\xi)}\Psi). \quad (2.5)$$

where we have already used the field-space of Yang-Mills theory, Φ_{YM} , given by the two different sectors—gauge connections, A , and matter fields, Ψ —i.e.:

$$\Phi_{\text{YM}} = \{\varphi = (A, \Psi)\}. \quad (2.6)$$

Gauge connections are \mathfrak{g} -valued 1-forms over the spacetime manifold

$$A = A_\mu^a(x)\tau_a dx^\mu \in \Lambda^1(M, \mathfrak{g}), \quad (2.7)$$

where $\mathfrak{g} = \text{Lie}(G)$ and $\{\tau_a\}_a$ is an orthogonal basis of \mathfrak{g} . We take this basis to be normalized with respect to the trace in the fundamental representation as²⁹ $\text{tr}(\tau_a\tau_b) = -\frac{1}{2}\delta_{ab}$.

We will consider here only matter fields, ψ , which are smooth functions on M valued in W , where W is a vector space carrying the fundamental representation of G , e.g.: for $W = \mathbb{C}^N$ and $G = SU(N)$:³⁰

$$\psi = \psi^m(x)|m\rangle \in \mathcal{C}^\infty(M, W); \quad (2.8)$$

where $\{|m\rangle\}_{m=1, \dots, N}$ is a basis for W .

The transformation properties of A and ψ are given by:

$$A^g = g^{-1}Ag + g^{-1}dg \quad \text{and} \quad \Psi^g = g^{-1}\Psi. \quad (2.9)$$

so we obtain

$$\xi^\# = \int \delta_\xi A \frac{d}{dA} + \int \delta_\xi \Psi \frac{d}{d\Psi} \quad (2.10)$$

where we introduced the standard notation for infinitesimal gauge-transformations (along ξ),

$$\delta_\xi A = D\xi := d\xi + [A, \xi] \quad \text{and} \quad \delta_\xi \Psi = -\xi\Psi, \quad (2.11)$$

with $[\cdot, \cdot]$ the Lie bracket on $\mathfrak{g} = \text{Lie}(G)$, extended pointwise on M to $\text{Lie}(\mathcal{G})$.

Vector fields \mathfrak{X} which are tangent to gauge orbits in Φ will be called ‘vertical’ and their span at a $\varphi \in \Phi$ defines a vertical subspace V_φ of the tangent space $T_\varphi\Phi$. In symbols,

$$T_\varphi\Phi \supset V_\varphi = \text{Span}\{\xi^\#, \xi \in \text{Lie}(\mathcal{G})\}. \quad (2.12)$$

Vertical fields represent infinitesimal gauge transformations—infinitesimal changes of perspective—and they span the gauge orbit through a given point. To briefly illustrate our notation and construction so far, consider figure 2. As shown in the figure, we call an orbit through a point φ , \mathcal{O}_φ .

Now, $[\varphi] \in \Phi/\mathcal{G}$ is defined abstractly by taking the equivalence class, and there is in general no natural parametrization for the gauge-invariant degrees of freedom. φ represents one perspective, or one parametrization of $[\varphi]$. But, since there is no canonical isomorphism between \mathcal{G} and \mathcal{O}_φ , a given $\varphi \in \mathcal{O}_\varphi$ does not itself single out

²⁹For $G = SU(2)$, $\tau_a = -\frac{i}{2}\sigma_a$, with σ_a the Pauli matrices.

³⁰This restriction is only for simplicity in exposition. In [GHR18] we consider also spinorial fields.

a frame. However, for each $T_\varphi\Phi$, we seem to have a canonical isomorphism between $\text{Lie}(\mathcal{G})$ and V_φ , i.e. we can canonically identify a change of a frame (i.e. $\xi \in \text{Lie}(\mathcal{G})$) to a change of perspective (i.e. $\xi^\# = v \in V_\varphi$). In the nomenclature we have been using, this is what (2.4) does.

But in fact, this identification may falter: although for standard principal fiber bundles, V_φ and $\text{Lie}(\mathcal{G})$ are isomorphic vector spaces, this is not always the case for the infinite-dimensional, field-theory cases. There are changes of frame which do *not* bring about changes of perspective. We can say such changes are in the *blind-spot* of the given perspective.

As we shall see in section 3.3, the implicit assumption that the isomorphism above always holds spawns much confusion regarding the distinction between global and local symmetries. As we envisaged in section 1.1, the failure of this isomorphism illustrates the second way in which the standard characterization of gauge systems through their symmetry group detaches from that given by our intuitions on perspective.³¹

2.3 Connection-forms: the formal construction

Now, figure 2 looks like a principal fiber bundle. As with any principal fiber bundle, we should aim to define a connection-form therein, which tells us how to parallel-propagate frames from orbit to orbit.

Indeed, such a definition is equivalent to defining a horizontal complement to the vertical subspaces in the tangent bundle $T\Phi$. Mimicking the finite-dimensional case [KN63], we define ϖ (pronounced VAR-PIE) as a functional 1-form over field-space, valued in the Lie algebra of the gauge group $\text{Lie}(\mathcal{G})$,

$$\varpi \in \Lambda^1(\Phi, \text{Lie}(\mathcal{G})). \quad (2.13)$$

A one-form naturally contracts with vector fields, and thus defines horizontal complements through its kernel (see figure 3):

$$H := \ker(\varpi) = \{\mathbb{X} \in T\Phi \mid \mathfrak{i}_\mathbb{X}\varpi = 0\}. \quad (2.14)$$

Using this notation, we can thereafter identify the vertical projector by $\widehat{V} = \varpi^\#$, i.e.: since $\mathfrak{i}_\mathbb{X}\varpi \in \text{Lie}(\mathcal{G})$, for $\mathbb{X} \in T_\varphi\Phi$, we have $\widehat{V}(\mathbb{X}) = (\mathfrak{i}_\mathbb{X}\varpi)^\# \in V_\varphi$ and therefore, using this notation, the horizontal projection is defined by

$$\widehat{H} = (\text{Id} - \varpi^\#) : \mathbb{X} \mapsto \widehat{H}(\mathbb{X}) = \mathbb{X} - (\mathfrak{i}_\mathbb{X}\varpi)^\# \quad (2.15)$$

The horizontal exterior derivative is obtained by composition with the horizontal projection [KN63]:

$$\mathfrak{d}_H := \mathfrak{d}(\widehat{H}(\cdot)). \quad (2.16)$$

Now, how *does* one define such a ϖ , precisely? First, to be compatible with the gauge structure, a bona-fide connection-form in a principal fiber bundle needs to satisfy the following equivariance conditions:³²

$$\mathfrak{i}_{\xi^\#}\varpi = \xi \quad (2.17a)$$

$$\mathbb{L}_{\xi^\#}\varpi = [\varpi, \xi] + \mathfrak{d}\xi \quad (2.17b)$$

³¹The first concerned the stipulation of the gauge transformations at boundaries, for gauge-fixing.

³²At least in finite dimensions, (2.17a) sums up (2.14) and (2.12). In infinite dimensions some further conditions need to be fulfilled. See e.g.: [Gom11] and references therein.




Figure 3: A pictorial representation of the split of $T_\varphi \Phi$ into a vertical subspace V_φ spanned by $\{\xi_\varphi^\#, \xi \in \text{Lie}(\mathcal{G})\}$ and its horizontal complement H_φ defined as the kernel at φ of a functional connection ϖ . With dotted lines, I have represented a different choice of horizontal complement associated to a different choice of ϖ . Note that each choice of functional connection ϖ defines a distribution of subspaces of the tangent bundle, but this distribution is not necessarily integrable, i.e. not necessarily tangent to ‘horizontal submanifolds’.

The standard way of automatically implementing conditions (2.17) is to have the horizontal spaces be orthogonal to the vertical spaces, with respect to some physically relevant metric on field-space, i.e. with respect to a *gauge-invariant supermetric*. We will see this in section 3.1. In short, each sector of field-space (i.e. a field A or ψ and region in space or spacetime in M) carries an essentially unique choice of ultralocal supermetric; and so practical choices are more constrained than they might at first appear.³³

Therefore, having chosen the sector, and then given the corresponding supermetric, one has a canonical choice of connection. Importantly, as we will see in section 3.1, we can build connection-forms which only employ the original fields. In this way, it measures what is a physical (horizontal) and what is a gauge (vertical) change with respect to the instantaneous configurations of those fields, i.e. with respect to what we have called ‘a perspective’. For these reasons, it is in general called a ‘*relational connection-form*’. It is relational with respect to a perspective, given by the field content and its representation in a frame. According to (2.17), the vertical (pure gauge) component of a process automatically transforms covariantly, fulfilling the role of ‘handles’ required by Rovelli [Rov14].

3 Deploying connection-forms

We are now ready to provide a couple of explicit examples and put the connection-form to use. We start with the examples in subsection 3.1. In section 3.2, we analyze how the connection-form translates between physical degrees of freedom of different regions. i.e. how it aids “gluing”. Finally, in section 3.3, we show how the horizontal symplectic geometry defined with the aid of the connection-form passes

³³For Yang-Mills there is a single such supermetric, and for gravity a one-parameter family. Ultralocality—i.e. a metric that is a first integral of undifferentiated field-space vectors—is required for the well-posedness of properties of the connection, as we will glimpse in section 3.1 (see section 5 in ??). If any disambiguation is necessary, we can also appeal to the kinematical field-space metric which implicitly appears in the Lagrangian of the theory in the 3+1 case.

the “acid test” of providing physically meaningful charges, without gauge-fixing and without the introduction of extra degrees of freedom.

3.1 Examples of connection-forms

For Yang-Mills with matter, the connection for the A_μ^a -sector is more specifically termed *the Singer-DeWitt* connection, and for the ψ -sector, *the Higgs* connection. This terminology is defended in detail in [GHR18]. Importantly, even for regions with boundary, gauge-transformations are completely unconstrained at the boundary; no conditions apart from the field sector and an appropriate supermetric need to be specified.

Although we will not see here the explicit equations of the connection for gravity, the general connection’s role in selecting preferential frames is best illustrated in examples with non-internal symmetries. Therefore, before we go on to specific examples of connections, I will first present the more palpable example of diffeomorphisms so as to demonstrate this role (see also [Gom11] for more technical details).

Selecting frames: the gravity example– We consider gravity in a 3+1 framework, so that now M is a 3-dimensional manifold representing space, on some definition of simultaneity (or instantaneous slices). We describe our instantaneous metric by $\{g_{ab}(x), x \in M\}$. The covariant tensor $\mathbf{g}(x) = g_{ab}(x)$ (in abstract index notation) can be subject to active spatial diffeomorphisms (which are the ones we consider), and therefore any such description identifies points in M in a rather arbitrary way, i.e. in an arbitrary “frame”.³⁴

The Lie-algebra associated to the group of diffeomorphisms of M is the set of smooth spatial vector-fields $\mathfrak{X}(M)$.³⁵ A ϖ for g_{ab} and a metric velocity, \dot{g}_{ab} in this frame, according to (2.13), we get a spatial vector field $\mathfrak{i}_{\dot{\mathbf{g}}}\varpi^a =: \varpi^a(\dot{\mathbf{g}}) \in \mathfrak{X}(M)$. Suppose $\varpi^a(\dot{\mathbf{g}}) = X^a \neq 0$. This vector field, X^a , tells us that we changed our frame, or rather, we changed our ‘identification of points’ during this transformation of g_{ab} . It pinpoints which part of the field-transformation that we described as \dot{g}_{ab} came from a ‘changing’ identification of spatial points along time. In order to say ‘the identification changed’, we need a preferred way of identifying spatial points during evolution. In the principal fiber bundle language, when integrated over time, ϖ would effect a horizontal lift; it parallel transports a given choice of initial point x along time, providing, in the words of Barbour, an *equilocality relation* [BB82].

According to (2.17b), this equilocality does not really care about what the original frame was; instead, it only evinces the field’s ‘preferred way’ of identifying spatial points amidst an arbitrary change of fields.

The actual physical change is extricated from the total one by rotating the frame back to equilocality (and correspondingly adapting an infinitesimal change in perspective):

$$\widehat{H}(\dot{\mathbf{g}}) = \dot{\mathbf{g}} - (\mathfrak{i}_{\dot{\mathbf{g}}}\varpi)^\sharp = \dot{\mathbf{g}} - \mathcal{L}_{\mathbf{X}}\mathbf{g} \quad (3.1)$$

³⁴We cannot identify ‘coordinate systems’ as such, in either the field space, Φ , or in the gauge group, $\text{Diff}(M)$. Each frame here is better thought of as ‘an identification of spatial points’, which can be shuffled by active diffeomorphisms.

³⁵Although we will not concern ourselves with the order of differentiability here, it does make a difference to the type of manifold structure we endow to both Φ and its group action. See [Gom11] and references therein for more details. In the C^∞ case, the so-called *inverse limit Hilbert* structure is the most appropriate [Ebi70].

where $\mathcal{L}_{\mathbf{X}}$ stands for the spatial Lie derivative along $X^a = \mathfrak{i}_{\dot{\mathbf{g}}}\varpi^a$. The properties of (2.17b) guarantee that upon a time-dependent diffeomorphism, $\widehat{H}(\dot{\mathbf{g}})$ will have the standard covariance properties under time-independent diffeomorphisms [Gom11, GR17].³⁶

It could also be that the transformation involved no physical change at all, and was purely one of changing perspectives. When that is the case, equation (2.17a) demands that ϖ recognizes it to be so, yielding $\dot{\mathbf{g}} = \mathbf{X}^\sharp := \mathcal{L}_{\mathbf{X}}\mathbf{g}$, and therefore, according to (3.1), $\widehat{H}(\dot{\mathbf{g}}) = 0$.

In the case of diffeomorphisms, differently to previous definitions of equilocality relations [BB82], the ones presented here are explicitly gauge-covariant, and are grounded on the field-content of the theory; they can use different fields to inform the choice of frames along time.

The internal gauge example works almost identically to this diffeomorphism in gravity example, the difference being that the gauge-potential value is matched along time at predetermined spatial points.

For matter fields, the situation is similar; here the main difference, as we will see, is one between (3.3), which is non-local for both gauge-fields and gravitational fields, and (3.8), which is local. In other words, matter fields define equilocality relations in a local manner. Maybe we would like the identification of points along time to be grounded on the location of dust particles around ourselves, or maybe we would like points to be identified by that cloud of neutrinos zipping by—no problem, there will be a connection-form associated to each, as long as these fields are not zero anywhere and have no ‘blind-spots’.

‘Blind-spots’ are infinitesimal gauge-transformations, ξ , which are not registered by a given choice of ϖ , i.e. they are $\xi \neq 0$ such that $\varpi(\xi^\sharp) = 0$. They signal the failing of the isomorphism between changes of frame and changes of perspective: thence their name. That is, a blind-spot is a change of frame which is not registered by the perspective.

There would be no such ξ if (2.17a) was completely valid, i.e. if Φ was a bona-fide principal fiber bundle and ϖ a bona-fide connection-form therein. But they are not. And this failure is immensely important in the determination of Noether charges, as we will see in section 3.3.

The Singer-DeWitt connection – Let us take a simple, 3+1 Yang-Mills example to understand how ϖ is obtained through orthogonality: given a Lagrangian for the fields, L , and a choice of field-space sector, say A_μ^a , upon a 3+1 split the Lagrangian itself will yield a kinematical supermetric for A_i^a (where now i, j run over spatial indices). In the Yang-Mills case, the kinematic metric is suggested, not surprisingly, by the kinetic term in the Lagrangian:

$$\mathbb{G}(\dot{A}, \dot{A}) := \int_{\Sigma} d^3x \operatorname{tr} \left(g^{ij} \dot{A}_i^a \dot{A}_j^b \right)$$

This kinetic metric enables us to define those instantaneous field-transformations which are strictly horizontal, since the choice of vertical vectors is canonical (i.e. of the form $\delta_\xi A = D\xi$, also called pure-gauge). The definition of these horizontal vectors at each base point in field-space is tantamount to a definition of ϖ . It splits any ‘change’ into a purely gauge part, $\widehat{V}(\dot{A})$, and a purely physical part, $\widehat{H}(\dot{A})$, with respect to the perspective of A_μ (and what it characterizes as kinetic).

³⁶Indeed, in the ADM formalism [ADM62], the ‘shift’ components of the metric satisfy the same transformation laws, for field-independent diffeomorphisms.

To ascertain the precise form of the Singer-DeWitt (SdW) connection, the procedure is simple: first, we determine the set of horizontal vectors, $\mathbb{X} = \mathfrak{h} \in H_\varphi := (V_\varphi)^\perp \subset T_\varphi\Phi$:

$$G(\xi^\sharp, \mathfrak{h}) := \int d^3x \ (g^{ij} \mathfrak{h}_i^a D_j \xi_a) = 0 \quad \forall \xi \in \text{Lie}(\mathcal{G}) \quad (3.2)$$

using the definition of D in (2.11), and without imposing any a priori restrictions on the content of ξ or \mathfrak{h} at the boundary of the region $\partial\Sigma$. Second, we define the horizontal projection, $\mathbb{X}^h = \mathbb{X} - (\mathfrak{i}_\mathbb{X}\varpi)^\sharp$ such that $\mathbb{X}^h \in H$. In this way, the Singer-deWitt connection is defined as a solution to:

$$D^k D_k \varpi = D^i dA_i \quad \text{and} \quad n^i D_i \varpi|_{\partial\Sigma} = n^i dA_i|_{\partial\Sigma}. \quad (3.3)$$

In other words, the gauge-covariant Poisson equation for ϖ comes automatically equipped with nonzero *gauge-covariant* Neumann boundary conditions. The appearance of the field-content on the right hand side of the boundary problem is essential for covariance and tells us ϖ does not represent an arbitrary new degree of freedom encoded in the boundary, as was the case with the arbitrary stipulation of λ at the boundary for gauge-fixings, explored in section 1.3. The boundary conditions appearing in the definition (3.3) descend directly from the properties of the supermetric in field-space, and therefore are unlike standard boundary conditions, which are restrictions on field content. Lastly, note that although ϖ is non-local— as signalled by inversion of D^2 — these equations are defined intrinsically in each region, from the field content itself, and are for these reasons termed *regional*.

This connection-form is non-local and, for a non-Abelian charge group, also non-flat—it is non-integrable and therefore does not correspond to any gauge-fixing. At this point of the discussion, it is tempting to say a few words about the Gribov problem. The ‘problem’ is more of an obstruction, which tells us that, for non-Abelian gauge theories, Φ does not have a global section (or global gauge-fixing). If Φ were a principal fiber bundle, the lack of a global section would translate to the non-triviality of the bundle (meaning it does not admit a global product structure).

A lack of a global gauge-fixing is certain to spell trouble for non-perturbative quantizations of non-Abelian theories, but let us put that aside for now. The Gribov obstruction tells us that any connection-form which is everywhere defined in Φ does not form an integrable distribution, because, if it did, it would be associated to a global section. Since it is not associated to a global section, parallel transport according to ϖ will be path-dependent. That is, for a non-Abelian gauge theory, ϖ will possess intrinsic curvature, \mathbb{F} , a very interesting concept on its own. For more on this and its relation to the field-space Aharonov-Bohm effect, see [GHR18].

On the other hand, in the Abelian case, i.e. for electromagnetism, the connection is flat. It is given by

$$\varpi = \nabla^{-2}(\partial^i dA_i) \quad \text{with} \quad n^i \partial_i \varpi|_{\partial\Sigma} = n^i dA_i|_{\partial\Sigma}. \quad (3.4)$$

For $n^i dA_i|_{\partial\Sigma} = 0$, and $\Sigma \simeq \mathbb{R}^3$, the unique solution is:

$$\varpi = \int_\Sigma \frac{d^3y}{4\pi} \frac{\partial^i dA_i(y)}{|x-y|} \quad (3.5)$$

For field-space curves $A_i(x, t)$, we can define a horizontal lift in analogy to a Wilson line, even in the general non-Abelian case, through a path-ordered exponential,

i.e. the field-dependent field-space ‘parallel-transport’

$$h_\gamma[A] = \mathbb{P} \exp \int_\gamma \varpi, \quad (3.6)$$

where γ is the history of the field, a path linking the configuration A to the initial configuration $A^\star = A(0)$.

In the Abelian case, in temporal gauge, we identify $\mathfrak{d}A_i = \dot{A}_i$ and we can drop the dependence of h_γ on the path: $h_\gamma \rightarrow h$. If we set out an initial configuration (such as $A(0) = 0$ for electromagnetism), we can construct a unique holonomy associated to each field configuration. The path integration essentially cancels out with the time derivative (see [GR18] for details), and we obtain a ‘dressing’:³⁷

$$h[A] := \int_\Sigma \frac{\mathrm{d}^3y}{4\pi} \frac{\partial^i \mathfrak{d}A_i(y)}{|x-y|} \quad (3.7)$$

In this case, $h[A]$ becomes precisely the gauge-fixing translation $h[A] \equiv \lambda^f(A)$ discussed in section 1.3, associated to temporal Coulomb gauge. Using it, we can ‘dress’ variables with a garment that makes them invisible to gauge-transformations. This procedure precisely recovers the results of Dirac [Dir55]; using the nomenclature of the introduction, he redefined the electron variables as $\psi_h := \lambda^f(A)^{-1}\psi$. These redefined, dressed variables are gauge-invariant.³⁸

Moreover, the garment is versatile: it is equally applicable globally and regionally. It is solely defined by the field-content, including its boundary values; one does not freely fix the boundary value of λ , for here it does not represent a new degree of freedom.

The Higgs connection – The Higgs connection as a concept is perhaps even richer than the Singer-deWitt connection, and, since I will not be able to convey its many facets here, I must redirect the interested reader to chapter 7 of [GHR18].

For now, let’s take a simple example of Yang-Mills theory where the gauge group acting on the matter fields is $G = \mathrm{SU}(2)$, and in this case, $\tau_a = -\frac{i}{2}\sigma_a$.

We also have a unique type of bilinear term for the fields (see equation 7.3 in [GHR18]), and we can pursue the same orthogonality strategy as in the SdW case, obtaining for the connection-form:

$$\varpi_{\mathrm{SU}(2)} = \frac{i}{\overline{\Psi}\Psi} \left((\mathfrak{d}\overline{\Psi})\sigma^a\Psi - \overline{\Psi}\sigma^a(\mathfrak{d}\Psi) \right) \tau_a. \quad (3.8)$$

Unlike what is the case for SdW, we have found a local expression for the connection-form. The general shape of (3.8) can be generalized to include more general charge groups, spinors, and also vielbeins transforming under Lorentz symmetries.

³⁷The obstruction for such dressings in non-Abelian theories is that they cannot be path (e.g. history)-independent, and are thus not intrinsically defined by instantaneous perspective (the state of the field) [GHR18].

³⁸In Dirac’s construction it is also important that the electric field created by $h(A)$ is the Coulomb field of the electron. In classical terms, this means that the Poisson bracket of E with $\hat{\psi}$ is $\{E(x), \hat{\psi}(y)\} = \frac{-1}{4\pi(x-y)^2} \hat{\psi}(y)$. Notice that this is a Poisson bracket between gauge-invariant quantities. It is most easily computed in temporal gauge; in Coulomb gauge, in which $h_{\mathrm{Dirac}} \equiv \mathrm{id}$, but one needs to first introduce non-local Dirac brackets. The supermetric appearing in the kinetic term of the Lagrangian governs both the definition of canonical momenta and also the structure of the Gauss constraint. It is not difficult to convince oneself that the relationship between the SdW connection and the supermetric is what guarantees that the Poisson bracket $\{E, \hat{\psi}\}$ gives the expected, Dirac result.

Using an actual matter field allows us to define the frame point by point. Importantly, this description is only possible for those states of Ψ which are everywhere non-zero, i.e. $\Psi(x)\Psi^\dagger(x) \neq 0, \forall x \in M$. This caveat has profound implications, and can be related to certain aspects of the Higgs mechanism for symmetry breaking.

A Higgs connection-form requires a spontaneously broken phase, using the condensate as a gauge-frame, hence the namesake [GHR18]. Coordinatizing the field ψ as

$$\Psi = \Psi(h, \rho) = \rho h v_o \quad \text{where } h \in C^\infty(\Sigma, G) \text{ and } \rho \in C^\infty(\Sigma, \mathbb{R}), \quad (3.9)$$

and $v_o \in V$ is some non-vanishing reference vector in $W \cong \mathbb{C}^N$, the fundamental representation of $G = \text{SU}(N)$. Note that the reference vector v_o is chosen once and for all, i.e. it is not a field-space coordinate, hence

$$\varpi = -\text{d}h h^{-1}. \quad (3.10)$$

This connection-form recovers Donnelly and Freidel’s edge-modes, with a simple translation $h|_{\partial\Sigma} \equiv \kappa$ (the edge-modes κ appear explicitly in (1.7),(1.8)). The main difference is that here the fields have physical significance, and not just on the boundary. Interestingly, h ’s physical status brings a caveat: the matter field ψ cannot vanish at any point within a region for the regional connection-form to exist.

Barbour-Bertotti’s best-matching – To fully explain the philosophical background to Barbour and Bertotti’s best-matching procedure would take us too far afield, as would delving into its consequences for the relationalist vs substantivalist debate (see chapter 4 [BP13]). The gist of it is depicted in figure 4 (the idea originated in [BB82]; see [Mer17], section 6.2, and [Bar10] for reviews). This will be a ‘bare-bones’ exposition, based on the concepts we have here developed.

In the language being employed here, a ‘best-matched’ frame can be summarized as a horizontally projected one, when configuration space is the one of N - particles in Newtonian mechanics, $\Phi = \mathbb{R}^{3N}$ and the gauge group is the 3-dimensional Euclidean one, $\text{Eucl}(3) := \text{SO}(3) \times \mathbb{R}^3$, acting as:³⁹

$$q_i^a \mapsto R_b^a q_i^b + t^a \quad \forall i$$

where $\mathbf{q}_i \in \mathbb{R}^3$ is the position vector of the i -th particle, and $(\mathbf{R}, \mathbf{t}) \in \text{SO}(3) \times \mathbb{R}^3$. Finally, to determine the connection-form, we need to stipulate the field-space inner product, which is derived from:

$$d(\mathbf{q}, \mathbf{q}') = \left(\sum_i \|\mathbf{q}_i - \mathbf{q}'_i\|^2 \right)^{1/2}$$

where we use the standard Euclidean norm. In this case, the group action is field-independent, and the connection-form would yield the necessary action of the group to attain the “best-matched” frame (in the infinitesimal case, i.e. as applied to $\dot{\mathbf{q}}$).

³⁹Barbour also considers the so-called ‘similarity group’, composed of rigid rotations, translations and dilations (an overall change of scale).




Figure 4: A stacking of three-body configurations: the arbitrarily chosen stacking above is best matched (blue arrows) by translations so as to bring the barycenters to coincidence, after which rotational best matching eliminates residual arbitrary relative rotation (taken with permission from arxiv [Mer17])

3.2 Gluing

Establishing communication: gluing different regions – Now consider a region $M = M_1 \cup M_2$, with $M_{1,2}$ embedded manifolds sharing a portion of their boundary, $S = \partial M_1 \cap \partial M_2 \neq \emptyset$. Let $\iota_1 : M_1 \rightarrow M$, $\iota_2 : M_2 \rightarrow M$, $\iota_S : S \rightarrow M$ be the embedding maps.

To each of these regions, we assign an own field-space, Lie-algebra of gauge transformations, supermetric, and the respective ϖ connection. I will denote a restriction to one of the regions by the same subscripts; that gives us a map $\cdot_I : \Phi \rightarrow \Phi_I$ and so on, with $I = 1, 2$.

Since the Higgs matter connection is local, the decomposition into vertical and horizontal follows suit, and there is not much complication at the boundary. Therefore, we will follow our examples using electromagnetism of section 1, here exploring the consequences solely of the SdW connection-form for gluing. I.e. we take a connection-form based on (3.3), reserving a few comments on the Higgs-type connection form to the end.

We can then define the vertical and horizontal projectors $\widehat{V} = \varpi^\sharp$ and $\widehat{H} = (\text{id} - \varpi^\sharp)$, and similarly for $\widehat{V}_{1,2}$ and $\widehat{H}_{1,2}$. Each of these operators acts on field-space vectors intrinsic to each region. More concretely, given a vector \mathbb{X} supported on M , it can be decomposed into $\mathbb{X} = \mathbb{X}_1 + \mathbb{X}_2$, where $\mathbb{X}_{1,2}$ live respectively on $M_{1,2}$ understood as intrinsic manifolds with boundary.

While it is true that the relational physical (horizontal) component of the restrictions $\mathbb{X}_1, \mathbb{X}_2$ may not match at their boundary, the difference at the boundary is always purely vertical. But by itself verticality does not suffice to characterize a degree of freedom intrinsic to the boundary; the difference must also be a vertical element *intrinsic to the boundary*. And that it is:

$$\widehat{H}_1(\mathbb{X}_1)|_S - \widehat{H}_2(\mathbb{X}_2)|_S = -(\widehat{V}_1(\mathbb{X}_1)|_S - \widehat{V}_2(\mathbb{X}_2)|_S) = \xi_S^\sharp \quad (3.11)$$

where $\mathbb{X}|_S$ means we restrict $\mathbb{X}(x)$ to $x = y \in S$, and where \sharp_S is defined for the field-space and gauge-transformations over S , Φ_S and $\xi_S \in \text{Lie}(\mathcal{G}_S)$, respectively.⁴⁰ This is essentially a more elaborate, infinitesimal, non-Abelian version of the more qualitative (1.4). The gauge transformation ξ_S earns its status as an “edge-mode” because it is intrinsic to the boundary.⁴¹

Here we see Rovelli’s handles at work: we may be given the physical transformations of the subsystems, $\widehat{H}_{1,2}(\mathbb{X}_{1,2})$, but we still have the “embedding” information

⁴⁰ This is shown in equation 5.17 in [GHR18].

⁴¹ There are different definitions of edge-modes in the literature, but being gauge-group-valued and intrinsic to the boundary is a common requirement.

in the full field-space which tells us, through ξ_S , what sorts of “rotations” we need to perform at the boundary for the coupling, or “language translation”, between subsystems. This is only possible because we are projecting onto the physical (or horizontal) subspaces but not discarding the ambient space and their plurality of perspectives.⁴²

There are many interesting questions we can ask about the properties of these projections viz a viz the composition of regions. For instance: what is the relation between the horizontal projections and the restrictions to the given regions? As is shown in [GHR18], the two operations do not commute: given a globally horizontal $\mathbb{X} = \widehat{H}(\mathbb{X})$, we may still find that its restriction is not regionally horizontal:

$$\mathbb{X}_1 := (\widehat{H}(\mathbb{X}))_1 \neq \widehat{H}_1(\mathbb{X}_1)$$

All this means so far is that regionally one might still have to change frames so as to uncover the physical component of \mathbb{X}_1 . That is, it means \mathbb{X}_1 has a non-trivial vertical component, but it can be further removed to reveal the physical part of \mathbb{X}_1 relative to region and field-content. In fact, it could even be that \mathbb{X}_1 is purely vertical:

$$\mathbb{X}_1 := (\widehat{H}(\mathbb{X}))_1 = \widehat{V}_1(\mathbb{X}_1) \quad \text{and therefore} \quad \widehat{H}_1(\mathbb{X}_1) = 0$$

in which case that perspective (i.e. regional configuration) *experienced at most* a change of frames.⁴³ However, it is always true that if the global \mathbb{X} has a physical component (i.e. if $\widehat{H}(\mathbb{X}) \neq 0$), then at least one region also has registered a physical component to the transformation, e.g. $\widehat{H}_2(\mathbb{X}_2) \neq 0$. This fact lends objectivity to physical properties as measured by \widehat{H} .

Due to the possible lack of commutativity between restriction and projection, the global physical transformation does *not* always decompose nicely into the local, purely physical transformations, i.e. i.e. the following can fail to hold

$$\widehat{H}(\mathbb{X}) = \widehat{H}_1(\mathbb{X}_1) + \widehat{H}_2(\mathbb{X}_2) \tag{3.12}$$

In fact, as shown in [GHR18] (see eq. 5.17 there), the necessary and sufficient condition for the validity of the composition in (3.12) is precisely tracked by the discrepancy in rotations at the boundary, discussed in sections 1.1 and 1.2. Namely, (3.12) holds if and only if:

$$i_S^*(\widehat{V}_1(\mathbb{X}_1)) := i_S^*(\widehat{V}_2(\mathbb{X}_2)) \tag{3.13}$$

where i_S^* is the pull-back to the boundary, acting on the spacetime indices of A_μ^I (or \mathbb{X}_μ^I) as usual, namely, for $Y^i \in \mathfrak{X}(S)$,

$$(i_S^* \mathbf{A}^a)(\mathbf{Y}) = A_\mu^a (i_{S^*}(Y^i))^\mu = A_\mu^a \frac{\partial i_S^\mu}{\partial y^i} Y^i$$

where y^i are intrinsic coordinates of S and x^μ are those of M . Therefore, for generic configurations, the physical decomposition (3.12) holds if and only if the RHS of (3.11) vanishes. That is: a physical (or horizontal) process precisely decomposes

⁴² Note that essentially the same procedure would work at the intersection of charts, as in footnote 13.

⁴³We don't mean ‘experienced’ in a conventional, physical sense. We mean solely that the region underwent a change of frames and nothing else.

into two physical processes if and only if the regional physical components match at the common boundary.

In the formalism of Donnelly and Freidel, a “fusion product” is defined on phase space, precisely by quotienting out matching gauge transformations on each side of the boundary (see eq. 2.46 there). In the SdW case, it is unclear how to take such a quotient in any precise mathematical form. But Donnelly and Freidel’s edge-modes are not analogous to the SdW connection; they are analogous to the Higgs connection. The flat, algebraic form of the Higgs connection (3.10) allows closed expressions for the quotients. Again, it is important to stress that here this procedure only involves the existing physical fields. In fact, it is essentially identical to the quotienting out of certain subgroups of gauge symmetries in some modern accounts of the spontaneous symmetry breaking mechanism (see [Fra14] and references therein).

By using ϖ , the present framework keeps tracks of rotations of frames at the boundaries. It is therefore able to connect the two descriptions and diagnose the effects of a mismatch.

Quantum considerations – For a classical system, it is indeed possible that the chosen perspectives give $\widehat{V}_{1,2}(\mathbb{X}_{1,2}) = 0$ on each side, and therefore the horizontal components just match on the nose, as in choosing coordinates such that $\dot{q}_1^1 = \dot{q}_N^2$ in Rovelli’s example in section 1.1, or such that (1.4) vanishes. However, once we pass on to the superpositions of the quantum domain, we will have to account for \mathbb{X} such that $\widehat{V}_{1,2}(\mathbb{X}_{1,2}) \neq 0$.

When performing the path integral using a geometric decomposition (as in [Mot95]), integrations over regions sharing a boundary will not be completely independent. For instance, the second equation in (3.3) requires the projections of the vertical component normal to the boundary to be equal, and since the difference between horizontal projections is identified with minus the difference in vertical projections, as in (3.11), these produce tractable correlations between the two regional path integrals. My expectation (and of others that have worked on this topic) is that such intertwining between the two path integrals (one for each region) lead to the same results for entanglement entropy as was achieved in [AKN17], without the use of additional edge-modes (or “replica tricks”).

In their role as projectors of general degrees of freedom onto the physical ones, the connection-forms closely resemble BRST ghosts [MH92]. There are two points of physical similarity: i) Feynman originally introduced ghosts from cutting a process into parts and comparing it with the glued version (e.g. enforcing the Cutkosky rules for a non-Abelian theory) [Fey63]. ii) In the Kugo-Ojima quartet mechanism [KO78], the job of the ghosts is precisely to avoid counting as physical some degrees of freedom which are fundamentally gauge. Both of these points are familiar to us: the connection comes alive at boundaries, where it is used to identify and subtract the unphysical component of processes. Indeed, this work is motivated by the confusion that boundaries engender in the splitting of physical and gauge degrees of freedom.

At the mathematical level, the properties of BRST ghosts (its anticommuting nature and transformation properties) can be entirely reproduced by a flat ϖ — i.e. one associated to a gauge-fixing section. Moreover, according to geometrical interpretations of BRST [BCR83, TM85, TM80], ghosts can be associated to vertical one-forms and BRST transformations can be associated to vertical derivatives. If c are the BRST ghosts, and s the BRST transformations, the matching $s \leftrightarrow \mathfrak{d}_V$ and

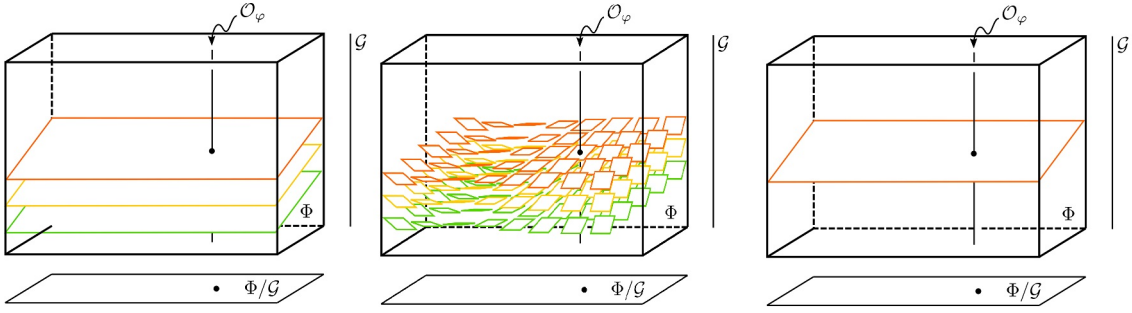


Figure 5: A visual illustration of the different structures provided respectively by: the rigid vertical transformations compatible with BRST symmetry (but see footnote 44), the horizontal distribution given by the connection-form, and gauge-fixing.

$c \leftrightarrow \varpi$ comes about as follows:

$$d_V \varpi = -\frac{1}{2}[\varpi, \varpi] \quad (\text{BRST} : s c = -\frac{1}{2}[c, c]) \quad \text{and} \quad d_V A = D_{\varpi} A \quad (\text{BRST} : s A = D_c A)$$

In the integrable ϖ case, both are ‘rigid’ transformations which may be interpreted as globally shifting the gauge-fixing section vertically (and therefore affording the required flexibility for matching of regions, see figure 5).⁴⁴ These discussions are included in [GR17].

3.3 Horizontal (or relational) Noether charges

Having set up the stage, I will now illustrate the use of a connection-form explicitly in the computation of the Noether charges. We first focus on the general properties of the formalism, and then move on to a particular use of the SdW connection.

Symplectic geometry – Let $L = \mathcal{L}(\varphi)d^d x$ be a Lagrangian (spacetime-) *density*. We then define the presymplectic potential θ implicitly through the field-space derivative of this density:

$$\boxed{dL = EL_I(\varphi)d\varphi^I + d\theta(\varphi),} \quad (3.14)$$

where d is the spacetime exterior derivative, and $EL_I(\varphi)$ are the (densitized) Euler-Lagrange equations for φ^I .

Let us for simplicity assume the Lagrangian *density* and the symplectic potential are strictly invariant under (infinitesimal) gauge transformations, and not only up to boundary terms (at least for ξ ’s which are field-independent)

$$\mathbb{L}_{\xi^\sharp} L = 0, \quad \text{and} \quad \mathbb{L}_{\xi^\sharp} \theta = 0 \quad (d\xi = 0). \quad (3.15)$$

Where \mathbb{L} is the field-space Lie-derivative, and along ξ^\sharp , i.e. along vertical directions. The demands (3.15) are quite restrictive, but do apply to the standard Lagrangian and symplectic potential of Yang–Mills theory.

⁴⁴ Having said that, since connection-forms come from taking derivatives in field-space, we do not currently know how to encode them in an action functional. For this reason, although we can fully reproduce the properties of BRST transformations (d_V) and BRST ghosts (ϖ), recovering the full *BRST mechanism* is still vexing.

Now, if (3.15) holds, then

$$0 = \mathbb{L}_{\xi^\#} L = \mathbb{E}L_I \delta_\xi \varphi^I + d\mathbb{i}_{\xi^\#} \theta, \quad (3.16)$$

and one is led to define the Noether current density j_ξ as (e.g. [IW94])

$$j_\xi := \mathbb{i}_{\xi^\#} \theta \equiv \theta(\varphi, \delta_\xi \varphi). \quad (3.17)$$

Then, one has

$$dj_\xi \approx 0. \quad (3.18)$$

where \approx assumes the equations of motion hold. i.e. the subspace of Φ_{YM} defined by $\mathbb{E}L_I = 0$.

Using the arbitrariness of $\xi \in \text{Lie}(\mathcal{G})$, one concludes that the Noether current must be of the form

$$j_\xi = C_a \xi^a + dQ_\xi \quad \text{where} \quad C_a \approx 0. \quad (3.19)$$

This equation defines the charge density Q_ξ whose relation to the total charge, $Q[\xi]$, introduced in (1.6) is, in shorthand

$$\int j_\xi \equiv Q[\xi] \approx \oint Q_\xi \quad (3.20)$$

Equation (3.19) is an instantiation of Noether's second theorem, which also shows the association between gauge symmetries and canonical constraints [LW90], the latter being the canonical generator of the relevant gauge symmetries. Explicitly, defining $\Omega := d\theta$, we have the Hamiltonian flow equation

$$\mathbb{i}_{\xi^\#} \Omega = -dj_\xi \quad (d\xi = 0). \quad (3.21)$$

Three remarks are in order: i) as can be seen in this case, the current is conserved automatically. If there is a conserved charge, it should be taken to be on a corner d-2 surface (usually taken to be a 2-sphere at infinity). ii) There is one charge per each choice of ξ . iii) We used $d\xi = 0$ (gauge-transformations were taken to be field-independent) but, as we know, this can be quite restrictive.

Horizontal symplectic geometry – Why is considering $d\xi \neq 0$ at the boundary relevant? Boundary, or asymptotic conditions on the gauge-parameters have an indelible reference to the physical field, e.g.: they are the ones that preserve certain properties of the fields; the two—gauge parameters and physical fields—are thus wedded at the boundaries.

This can be confirmed in many ways. Firstly, when dealing with gauge-fixings, one must allow for $d\xi \neq 0$, e.g.: the orbit translation $\lambda^f(A)$ depends on the field. For instance, a one-parameter family of $\lambda(t)|_{\partial M}$ is equivalent to a one-parameter family of gauge-fixings, and since a change of gauge-fixings is field-dependent (see e.g. equation (1.4)), one should also count a change of gauge at the boundary as being field-dependent (a change of gauge in the interior does not incur a change of gauge-fixing section, as seen in (1.3)). Secondly, as we saw in section 1.3, the transformations of the gauge-parameter on the boundary are usually of an ambiguous character—changing them changes the physical field (see footnote 23).

Indeed, in the Hamiltonian formalism, it is clear that the difference between $d\xi = 0$ and $d\xi \neq 0$ only crops up in the presence of boundary terms. To see this in

a simplified setting, consider a set of first-class constraints $\{\phi_\alpha(p, q)\}$ parametrized by α , which can be either a continuous or enumerable set. Again denoting equalities that hold when $\phi_\alpha(p, q) = 0$ by \approx , we have that in the absence of boundaries, even for field-dependent gauge parameters $\xi^\alpha(p, q)$,

$$\{\xi_\alpha \phi^\alpha, q\} \approx \xi_\alpha \{\phi^\alpha, q\}$$

and similarly for p . But once the constraint has a boundary contribution which does not vanish on-shell (i.e. a charge, as in equation (3.20)), then the variation of ξ_α at the boundary clearly appears explicitly.

While it is true that such a requirement—of $d\xi \neq 0$ —is no longer needed once the two bounded regions are glued and the boundary disappears, the same fate awaits the physical degrees of freedom taken to reside at boundaries (i.e. as $\mathcal{G}/\mathcal{G}^o$ for \mathcal{G}_o given in (1.5)), which disappear once gluing is completed. A unified treatment—one that seamlessly treats bounded and unbounded regions—requires the consideration of $d\xi \neq 0$ at boundaries, and, for unification, everywhere else too. In fact, just as dealing with spacetime dependent gauge transformation required the original finite-dimensional connection, dealing with field-dependent gauge transformation requires the field-space connection-form. As we will shortly see, appropriately dealing with $d\xi \neq 0$ suffices for finding a precise definition of horizontal symplectic geometry.

What happens when (iii) is dropped? Many of the previous computations start failing. For instance, we no longer have the relation (3.21). The obstructions can be traced back to a violation of the conditions (3.15) for $d\xi \neq 0$,

$$\mathbb{L}_{\xi^\sharp} \theta = \mathbb{i}_{\xi^\sharp} d\theta + d\mathbb{i}_{\xi^\sharp} \theta \stackrel{(3.15)}{=} \mathbb{i}_{(d\xi)^\sharp} \theta = j_{d\xi} \approx dQ_{d\xi}. \quad (3.22)$$

It was by eliminating this obstruction that Donnelly and Freidel got rid of the arbitrary charges related to the local gauge symmetries (only to replace them with another infinite set of charges).

In our case, following a more geometrical approach to covariant symplectic geometry leads us to the introduction of the horizontal symplectic current [GR17]:

$$\theta_H := \Pi_I d_H \varphi^I = \theta - \mathbb{i}_{\varpi^\sharp} \theta = \theta - j_\varpi. \quad (3.23)$$

If θ satisfies the conditions of equation (3.15), then—even for a field-dependent ξ ($d\xi \neq 0$)—one has automatically $\mathbb{L}_{\xi^\sharp} \theta_H = 0$. Thus the horizontal presymplectic two-form is automatically d -exact, and differs from the standard symplectic form by a spacetime exact term (α is a field-space two-form and spacetime scalar):

$$\Omega_H := d_H \theta_H = d\theta_H = \Omega + d\alpha, \quad (3.24)$$

In other words, Ω_H is d -closed and therefore a viable presymplectic form.

From the above

$$j_\xi^H := \mathbb{i}_{\xi^\sharp} \theta_H = \mathbb{i}_{\xi^\sharp} \theta - j_{\mathbb{i}_{\xi^\sharp} \varpi} = j_\xi - j_\xi = 0 \quad (3.25)$$

where we employed (2.17a). Note moreover that horizontality, but no gauge-fixing, was used to annihilate the local gauge charges.

It is similarly easy to see that $\mathbb{i}_{\xi^\sharp} \Omega_H = 0$ from $\mathbb{L}_{\xi^\sharp} \theta_H = 0$ and the Cartan magic formula. These formulas are valid at the density level, and hold for field-dependent gauge transformations as well. The message they convey is that such gauge transformations carry *no physical charge* with respect to ϖ .

Indeed, local gauge transformations, unlike global gauge transformations, should not produce meaningful physical charges. Whereas the standard symplectic potential does not make this distinction—naively assigning non-trivial currents to both types of transformation through Noether’s theorems—here we seem to have the opposite problem: our modified symplectic potential apparently leaves no charge at all!

Conserved global charges – For transformations which respect (2.17), ϖ correctly identifies the pure gauge ones and (3.25) eliminates the respective charge. This procedure is encapsulated by the validity of $j_\xi^H = 0$ for such ξ . And (3.25) holds for all bona-fide connection-forms in principal fiber bundles. However, Φ is not in general a bona-fide principal fiber bundle, and the relational connection-form ϖ is not in general a bona-fide connection-form.

Let us take a look again at (3.25). It is easy to see that if there existed $\xi_o \neq 0 \in \text{Lie}(\mathcal{G})$ such that $(\varpi(\xi_o^\sharp))^\sharp = 0$ (and therefore (2.17a) did not hold completely), then we could conceivably be left with a non-trivial charge for some such ξ_o .⁴⁵ If ξ_o is in the blind-spot of ϖ , and the standard charge for it does not vanish on-shell, $j_{\xi_o} \not\approx 0$, then

$$j_{\xi_o}^H = j_{\xi_o} \not\approx 0 \quad (3.26)$$

The point of confusion – Usually, there is no essential difference between changes of frame and changes of perspective, i.e. between ξ and ξ^\sharp —there is a bijection between V_φ and $\text{Lie}(\mathcal{G})$. One essential point which has confused discussions in both the philosophy and physics communities is that this bijection does not always hold: the group of *effective* gauge transformations can differ in important ways from the group of gauge transformations.

And indeed, if global gauge transformations are merely specific cases of the local gauge-transformations, how are they to escape extermination by (3.25)? The answer: by exploring the aforementioned blind-spot of perspectives.

The global group is defined from the local one—usually in a field-dependent manner—through ϖ by a projector onto the kernel of a differential operator. For instance, in electromagnetism taking $\xi = \text{const}$ we have $\xi \neq 0$, $\varpi(\xi)^\sharp = 0$, for all base points A_μ . There, the projection is field-independent, but this is not the case in general, since in general the group action is itself field-dependent (e.g. $\xi^\sharp = D\xi$ for Yang-Mills, with D containing A_μ^I). For such group-actions, finding all the ‘blind-spots’ becomes less trivial. Nonetheless, we can safely say that blind-spots always contain *the stabilizer subgroups* of given configurations, i.e. elements $g \in \mathcal{G}$ such that $R_g\varphi = \varphi$. But it is not every φ which has a non-trivial stabilizer. Much to the contrary: In the space of non-Abelian potentials and gravitational fields, generic configurations have no stabilizer subgroups.

Nonetheless, stabilizers do exist. Therefore the orbits $\mathcal{O}_\varphi \subset \Phi$ are not isomorphic for *all* $\varphi \in \Phi$. Indeed, different stabilizer groups produce qualitatively different orbits. When taking the quotient to obtain the space Φ/\mathcal{G} , such discrepancies show

⁴⁵ In [WG14], a similar criterion is proposed for the definition of ‘empirically significant symmetries’. In the language employed here (in 3.2 in particular), the conditions there would be formulated, in infinitesimal form, as: a transformation \mathbb{X} is an empirically significant symmetry if for two regions, $M_1 \cup M_2 = M$, $\widehat{H}(\mathbb{X}) \neq 0$ but $\widehat{H}_1(\mathbb{X}_1) = \widehat{H}_2(\mathbb{X}_2) = 0$. The necessary property they find for this is that there exists a ξ such that $\xi_{|\partial M}^\sharp = 0$. An analysis of this framework in our language shows that no such empirically significant symmetry exists in finite bounded regions corresponding to local gauge symmetries [GR].

up in the topological structure itself: Φ/\mathcal{G} is not a manifold, but only a stratified manifold, a nested union of submanifolds of different dimensions [Fis70, KR83].

Global charges appear from these irregularities in Φ/\mathcal{G} ; they present themselves as transitions between different regions. More specifically, global charges appear when Φ_φ/\mathcal{G} for a given field φ has such a transition, and another sector does not (we need the standard $j_\xi \not\approx 0$, as we see below, in equation (3.29)).

In essence, the connection-form is finding a projection from the local to the global transformations, and the modified symplectic charge gives us the corresponding physical charge.

‘Slippage’: an illustrative example – Again, in certain respects examples with non-internal symmetries are more palpable than the ones with internal symmetries. Therefore it useful to get the gist of what is to come through a heuristic example of how horizontal charges arise from the qualitatively different actions of the gauge group on different physical fields. I.e. they arise from the relational non-triviality of ‘blind-spot transformations’.

For example, generic spacetime metrics have no non-trivial Killing vector fields, and therefore have no non-trivial conserved charges. For the few orbits that do possess isometries, the group of diffeomorphisms acts in a qualitatively different manner. Namely, there is some ‘slippage’ of the action of the group on the orbit: a finite-dimensional subgroup $\mathcal{D}_o \subset \text{Diff}(M)$ does not move some metrics— those which are said to have stabilizer group \mathcal{D}_o , and which we can call

$$I_o := \{g_{ab}^o \mid f^* g_{ab}^o = g_{ab}^o \text{ for } f = h \circ f' \circ h^{-1}, h \in \text{Diff}(M), f' \in \mathcal{D}_o\}. \quad (3.27)$$

For example, for Minkowski metrics, $g_{ab}^o = \eta_{ab}$, in which case $\mathcal{D}_o = \mathcal{P}$ is the Poincaré group.⁴⁶ Reciprocally, note that for the Poincaré group to be well-defined in this manner one must restrict attention to elements of the corresponding I_o .⁴⁷

Suppose now that we don’t have just Minkowski spacetime, but also a particle within it. Poincaré symmetries leave the Minkowski metric completely invariant but have a non-trivial action on the particle. Therefore, apart from Poincaré transformations—which should produce charges by the mechanism above—generic diffeomorphisms should not produce any charge.⁴⁸

Conserved global charges: an explicit example – Finally, we can see in an explicit example how the horizontal symplectic geometry due to the SdW connection-form filters out the global charges from the local gauge transformations.

The Yang-Mills Lagrangian in form language can be written as

$$L_{\text{YM}} = -\frac{1}{2e^2}(F^a \wedge *F_a) + (i\bar{\psi}\gamma_\mu D\psi) \wedge *dx^\mu \quad (3.28)$$

where $F = dA + \frac{1}{2}[A, A]$ and $*$ is the Hodge dual, γ^μ are Dirac’s gamma-matrices, $D_\mu\psi = \partial_\mu\psi + A_\mu^a\tau_a\psi$, in the τ_a algebra basis, and e is the Yang–Mills coupling

⁴⁶Note that in the definition of I_o must be independent of the representation of the metric, i.e. a metric $h^*\eta$ will have isometries given by $h \circ f' \circ h^{-1}$ where f' are the isometries of η .

⁴⁷This is why I disagree with Greaves and Wallace, when they say that “In any theory that has a ‘local’ symmetry group, the ‘global’ symmetries remain as a subgroup of that local symmetry group. (For example, in general relativity, the ‘global’ translations and boosts form a subgroup of the group of all diffeomorphisms.)” [WG14]. It is only for highly homogeneous metrics that this subgroup is unambiguously defined.

⁴⁸A caveat: we have not yet completed the work on relational symplectic geometry for gravity. These statements are based on an assumption that the gravity case will work as well as the Yang-Mills one.

constant. Also, $\bar{\psi} = \psi^\dagger \gamma^0$, with the understanding that ψ and ψ^\dagger have to be considered as two independent (complex) variables.

One can find the Noether current by the standard method, as:

$$j_\xi := \mathfrak{i}_{\xi^\#} \theta_{\text{YM}} = -e^{-2} d(*F_a \xi^a) = d(\text{Tr}(E\xi)) \quad (3.29)$$

where we identified $*F_a$ pulled-back to the d-1 surface with the electric field (and also used index-free notation and absorbed the coupling constant), and the standard symplectic potential is

$$\theta_{\text{YM}} = -e^{-2} dA^a \wedge *F_a + i\bar{\psi} \gamma_\mu d\psi * dx^\mu. \quad (3.30)$$

Although θ_{YM} in (3.30) is invariant under field-independent gauge transformations (as demanded in the previous section), under field-dependent ones:

$$\mathbb{L}_{\xi^\#} \theta_{\text{YM}} = -e^{-2} d(*F_a d\xi^a) = d(\text{Tr}(E d\xi)). \quad (3.31)$$

Skipping technical details (contained in [GHR18]), let ϖ be associated to one field, the reference-field—say the gauge potential A —and suppose there is another field—say the matter field, ψ —which does not share all its stabilizer subgroups with the reference field. I.e. the field whose perspective we are using to map changes has blind-spots which are not shared by another charged field.

Now the non-zero term in (3.31) can be eliminated by using the horizontal (or physical wrt ϖ) symplectic potential, $\mathbb{L}_{\xi^\#} \theta_{\text{YM},H} = 0$, where

$$\theta_{\text{YM},H} = -e^{-2} d_H A^a \wedge *F_a + i\bar{\psi} \gamma_\mu d_H \psi * dx^\mu = \theta_{\text{YM}} - d(\text{Tr}(E\varpi)) \quad (3.32)$$

$$\Omega_{\text{YM},H} = d_H \theta_{\text{YM},H} = d\theta_{\text{YM},H} = \Omega_{\text{YM}} + dd(\text{Tr}(E\varpi)) \quad (3.33)$$

where ϖ (and therefore the horizontal projection) is defined by the Singer DeWitt connection for A^a . Note that the modification from the Yang-Mills symplectic form is both spacetime and field-space closed (i.e. an element of the form $d\alpha$), which is a guarantee that it defines a (pre)symplectic geometry and that we recover the standard Yang-Mills symplectic geometry in the boundariless case. Moreover, we have $\mathbb{L}_{\xi^\#} \theta_{\text{YM},H} = 0$, and, *generically* $\mathfrak{i}_{\xi^\#} \theta_{\text{YM},H} = j_\xi^H = 0$.

So far in this example, we have obtained all the desirable results of introducing an ‘edge-mode’, in the spirit of Donnelly and Freidel [DF16]. Their symplectic potential is given by

$$\theta^{\text{DF}} = \text{Tr}(E\delta A) - d(\text{Tr}(E\delta\kappa\kappa^{-1})) \quad (3.34)$$

which we fully recover with the Higgs connection, $\varpi = -dh h^{-1}$, given in (3.10). What are the differences?

In their nomenclature, Donnelly and Freidel recover charges from a different transformation of their edge-modes, which they call “symmetry transformations” [DF16]. In general, these are not available to us.⁴⁹ As mentioned before, they recover a plethora of charges at any given boundary, related to the intrinsic boundary degrees of freedom, and I consider this to be an unphysical overshoot. But how do we recover charges using the formalism here developed?

⁴⁹Essentially gauge transformations are given by a left action and symmetry transformations by a right one, as seen in (1.8) and (1.7). The meaning of these transformations on the new degrees of freedom inhabiting the boundary are unclear to me. If they are somehow supposed to correspond to choices of ‘observers’, why constrain such transformations to be related to the gauge group?

In our case, as shown in [GHR18] and described in the previous section, $j_\xi^H = 0$ for ξ which are *not* stabilizers of A^a . However, if there exists a set of Lie algebra elements $\{\chi_n^a\}$ forming a basis for the stabilizer of A^a , such a transformation will not annihilate a non-vanishing matter field, since $\delta_{\chi_n}\psi = -\chi_n\psi \neq 0$. From this and the above remarks, for such a “Killing” transformation χ_n we obtain

$$j_\chi^H = -\chi^a i\bar{\psi}\gamma_\mu\tau_a * dx^\mu = \chi^a J_a \quad (D\chi = 0). \quad (3.35)$$

where we introduced the matter current density:

$$J_a = -i\bar{\psi}\gamma_\mu\tau_a\psi(*dx^\mu) \quad (3.36)$$

Thus, the horizontal current for a global transformation is precisely given by the *matter current density* J_a contracted with the globally defined χ_n^a . In electrodynamics, $\chi = \text{const.}$ and j_χ^H is precisely the total current density of electrons.

Using the Gauss law, $D * F_a + J_a \approx 0$ and the Killing condition $D\chi_n = 0$, the horizontal charge can be written as

$$Q^H[\chi_n] = \int_\Sigma \chi_n^a J_a \approx -e^{-2} \int_{\partial\Sigma} *F_a \chi_n^a \quad (D\chi = 0). \quad (3.37)$$

We thus see that *the SdW connection with boundary picks out the global charges—when they exist—as the only physical ones.*

While the standard j_ξ is non-trivial for any gauge transformation, j_ξ^H is trivial for all, except for a few ‘global’ gauge transformations. Ultimately, horizontal charges appear related to objectively conserved physical quantities, like the total charge in electromagnetism, while the standard Noether charge is non-trivial for any gauge parameter.

The space generated by all the blind-spot directions of a given perspective will have a maximum, finite dimension (like the 10-dimensional group of Poincaré transformations for the metric, or the one-dimensional group of constant potential shifts in electromagnetism). This procedure therefore provides a projection from dizzying infinite-dimensionality of local gauge transformations, to a finite number of gauge transformations—the ‘global ones’. For these, we cannot vary the gauge parameters point by point, and Noether’s first theorem applies in standard fashion. Importantly, when obtained from the connection-form, these charges do not magically appear once one conjures a boundary, but correspond to physical degrees of freedom living in both bulk and boundary. They emerge from symmetries which act qualitatively differently on each of the interacting fields.

4 Conclusions

I will now briefly summarize our trajectory through this paper:

Perspectives – We started by relating the glitz of fiber bundles to a more “down-to-earth” understanding related to perspectives—essentially arbitrary frames in which we describe our system.

Because we have no local description of the physical degrees of freedom of gauge theories,⁵⁰ we use particular (but arbitrary) perspectives. Nonetheless, such

⁵⁰E.g. we could try to use holonomies as basic variables, but these are non-local and cumbersome to couple to sources [Hea07].

a plethora of perspectives is usually seen as superfluous, since the assumption is that we can always gauge-fix to a single class and thereby attain a gauge-invariant description.

But in fact, when our access to the system is only regional, not global, the survival of some perspectival flexibility is compulsory for describing the coupling of subsystems. It seems gauge degrees of freedom really come into their own when, as Rovelli advocates, they are related to subsystems [Rov14]. This argument suggests that describing subsystems in particular gauge-fixings may be premature if we allow for limited access to the Universe. We need to treat field-theoretic systems in the entirety of the field-space framework, and then extract whatever physical information we can from that larger domain.

Gauge-fixing and boundaries – In field-theory, we associate subsystems to different fields and to regions. However, gauge degrees of freedom are awkward in the presence of boundaries; since physical degrees of freedom are non-local, distinguishing them from the pure gauge degrees of freedom in a given region is a hard problem. This is the first point at which perspectives become unstuck from gauge symmetries: fixing a selective class of perspectives in a given region, by $f(\varphi) = 0$, does not suffice to fix the gauge; boundaries may require extra information beyond the simple determination of the regions and their field-content. This extra information is a replacement for Rovelli’s gauge-variant objects, which suggests that fixing it once and for all might be similarly short-sighted.

Noether charges – We can track the physical effects of this extra information by their impact on the computation of conserved charges—the teeth of symmetries found by Emily Noether. When using symplectic methods in finite regions, this extra information at the boundaries can correspond to an infinite number of conserved charges. Such charges depend on unphysical choices—i.e. not associated to perspectives—of gauge generators at the boundaries. Therefore, the physical status of these charges is dubious at best, unlike what is the case for those corresponding to rigid, global symmetries.

It is not easy to go from these Noether charges of local gauge symmetries to ones generally recognized as physically meaningful conserved charges. An attempt by Donnelly and Freidel leaves the boundary information completely free, by introducing ‘edge-modes’ [DF16]. These degrees of freedom eat up the spurious conserved currents associated to the local gauge symmetries, but they generate another infinity of conserved currents which are now dressed by new degrees of freedom—i.e. not in the original field-content—existing only at boundaries (imagined or not). Having the correct covariance properties for coupling regions, they could be seen as an implementation of Rovelli’s handles. However, these handles are being attached to the system as extra baggage.

The relational connection-form – The connection-form, ϖ , follows Rovelli’s idea, finds something similar to Donnelly and Freidel’s construction, and ends up with physically meaningful notions of charges. For ϖ is a device that fills three needs: i) providing a reference to measure *the physical change* of states, $\hat{H}(\mathbb{X})$; ii) keeping track of the perspective from which it does this in a covariant manner (allowing for the gluing of regions), $\hat{V}(\mathbb{X})$; and iii) finding relational physical charges requiring solely the physical fields, $Q^H[\xi_n]$ in (3.37). Let us quickly review these points.

The connection-form does not provide a fixed selective class of perspectives for comparing arbitrary states—as a gauge-fixing would,—it only provides a reference for examining infinitesimal changes in a given state (see figure 5, for an illustration in terms of non-integrable distributions in Φ). ϖ keeps track of the change in perspective by telling us how the fields extend an arbitrary choice of frame infinitesimally into the future. It then pinpoints any unwitting rotation of our arbitrary frame in the description of the process. Therefore it implicitly also tells us how much our system *physically* changed (infinitesimally) with respect to the given field and perspective, and according to some norm on field space (e.g. the supermetric appearing in the kinetic term in the Lagrangian). One of its key properties is that it is able both to select notions of physical change *and* negotiate the gluing of regions.

ϖ is also exclusively dependent on the field-sector and on the region of space occupied by the physical system; no extra degrees of freedom are required. For this reason I call it a “relational connection-form”.

Mathematically, ϖ defines a notion of horizontality in field-space. A horizontal change in the field corresponds to a change that is ‘physical with respect to the perspective’. Although it is not strictly relational in the sense of Rovelli—i.e. “requiring another subsystem to be made sense of”— ϖ also comes into its own in the presence of boundaries coupling to different regions.

With a connection-form, we are also able to define an alternative (pre)symplectic structure in the space of fields. The horizontal symplectic geometry defined with the aid of ϖ only differs from the standard symplectic geometry in the presence of boundaries to the system.⁵¹ Nonetheless, this boundary difference in symplectic structure is enough to annihilate charges associated to local gauge symmetries, irrespective of boundary conditions. Fortunately, the formalism rescues the baby before throwing out the bathwater: it filters out charges associated to global symmetries. These charges are associated to particular irregularities in the ‘physical base space’, Φ/\mathcal{G} , and are due to the second instance of detachment between gauge symmetries and perspectives.

Usually, not much attention is paid to the structure of the reduced space Φ/\mathcal{G} . It is assumed to abstractly represent the view-from-nowhere of the physical content of the system, without any redundancy. But Φ/\mathcal{G} ’s irregularities are physically relevant, representing the blind-spots of perspectives, and without first describing the full configuration space it is impossible to find them; after all, blind-spots emerge from transformations of perspective which are not registered by the field in question. In this respect, the horizontal charges resemble the definition of “empirically significant symmetries” of [WG14] (see footnote 45), but one more relational fact is needed to make conditions ripe for global charges: a given subsystem’s ‘blind-spots’ must not be shared by the coupled subsystem. The emerging charges are therefore of a patently relational character; they crop up when one subsystem slips from the other during a gauge transformation. A trivial example would be constant gauge transformations in electromagnetism in the presence of charged matter.

Open questions – There are many questions posed by boundaries in gauge theory; no approach can skirt them. The formalism described in the second half of this paper is still young; it may successfully deal with the sampling of questions described in the first half of the paper, but multiple questions remain. Here I will only briefly comment on what are for me the most pressing ones.

⁵¹By boundaries here I mean $d-2$ surfaces, boundaries to the spatial hypersurface, sometimes called ‘corners’.

The inherent tension between locality and gauge invariance is usually expressed by the failure to factorize the physical Hilbert space into local tensor products. This failure is relevant in studies of entanglement entropy [Don14]. In fact, one of the main motivations for [DF16] were computations of entanglement entropy in gauge theories [Don14]: when one is tracing over all degrees of freedom on one of the regions; extra information at the boundary—such as e.g.: fixing λ by hand—need to be counted. As mentioned in footnote 27, the expectation is that we can recover the correct computations of entanglement entropy through path integral methods (without the use of replica tricks, as is done in [AKN17]), but this remains to be seen.

Another glaring gap in our approach is the adaption of the treatment to the diffeomorphisms of general relativity. This is work under way and I see no reason to doubt its success.

A third important open question is the adaptation of the formalism to restrictions on field-space, as occurs when stringent boundary conditions are imposed on field-content—a very useful tool to study particular regimes or types of subsystems. Our work shows that boundaries by themselves may generically have no charges associated to them (e.g. for non-Abelian theories). But it does not mean that a) there are no charges for a particular field-content. The charge comes from the field, not from the boundary. Particular field-contents which have charges associated to them include i) base-point field-configurations with stabilizers but no restriction on field-variations, and, hopefully, ii) restrictions on field-space (e.g. boundary conditions). The open problem is to show that (ii) really does come to pass once one considers the appropriate restrictions in the derivation of ϖ .

Finally, we have mostly used a single manifold chart to describe our constructions. This assumption is also not compatible with regional access to the World. The work done in [GR17, GR18, GHR18] needs to be extended to more general topologies of the underlying manifold.

Acknowledgments

I would like to thank Florian Hopfmüller, and especially Aldo Riello, co-authors of the papers on which this analysis is almost entirely based. Any mistakes of interpretations not found in the original papers are of course my own. I would also like to show my immense gratitude to Jeremy Butterfield, for a meticulous reading, offering invaluable comments and advices on a first draft.

References

- [ADM62] R. Arnowitt, S. Deser, and C. Misner. The dynamics of general relativity, pp. 227-264. In *in Gravitation: an introduction to current research*, L. Witten, ed. Wiley, New York, 1962.
- [AKN17] Abhishek Agarwal, Dimitra Karabali, and V. P. Nair. Gauge-invariant Variables and Entanglement Entropy. *Phys. Rev.*, D96(12):125008, 2017.
- [Bar10] Julian Barbour. The Definition of Mach’s Principle. *Found. Phys.*, 40:1263–1284, 2010.
- [BB82] J. B. Barbour and B. Bertotti. Mach’s principle and the structure of dynamical theories. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 382(1783):295–306, 1982.

- [BB00] Katherine Brading and Harvey R. Brown. Noether’s theorems and gauge symmetries. 2000.
- [BB02] Glenn Barnich and Friedemann Brandt. Covariant theory of asymptotic symmetries, conservation laws and central charges. *Nucl. Phys.*, B633:3–82, 2002.
- [BC03] K. Brading and E. Castellani, editors. *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, 2003.
- [BCR83] L. Bonora and P. Cotta-Ramusino. Some remarks on BRS transformations, anomalies and the cohomology of the Lie algebra of the group of gauge transformations. *Communications in Mathematical Physics*, 87(4):589–603, dec 1983.
- [Bel76] John S. Bell. How to teach special relativity. *Progress in Scientific Culture*, 1, 1976.
- [BK60] Peter G. Bergmann and Arthur B. Komar. Poisson brackets between locally defined observables in general relativity. *Phys. Rev. Lett.*, 4:432–433, Apr 1960.
- [BK95] J. David Brown and Karel V. Kuchar. Dust as a standard of space and time in canonical quantum gravity. *Phys. Rev.*, D51:5600–5629, 1995.
- [BP13] Robert Batterman and Oliver Pooley. Substantivalist and relationalist approaches to spacetime, 09 2013.
- [CHR14] Horacio Casini, Marina Huerta, and José Alejandro Rosabal. Remarks on entanglement entropy for gauge fields. *Phys. Rev. D*, 89:085012, Apr 2014.
- [Crn87] Cedomir Crnkovic. Symplectic Geometry and (Super)poincare Algebra in Geometrical Theories. *Nucl. Phys.*, B288:419–430, 1987.
- [Crn88] Cedomir Crnkovic. Symplectic Geometry of the Covariant Phase Space. *Class. Quant. Grav.*, 5:1557–1575, 1988.
- [CW86] Cedomir Crnkovic and Edward Witten. COVARIANT DESCRIPTION OF CANONICAL FORMALISM IN GEOMETRICAL THEORIES. 1986.
- [Dew17] Neil Dewar. Sophistication about symmetries. *The British Journal for the Philosophy of Science*, page axx021, 2017.
- [DF16] William Donnelly and Laurent Freidel. Local subsystems in gauge theory and gravity. *JHEP*, 09:102, 2016.
- [DHM18] Xi Dong, Daniel Harlow, and Donald Marolf. Flat entanglement spectra in fixed-area states of quantum gravity. 2018.
- [Dir55] Paul A. M. Dirac. Gauge invariant formulation of quantum electrodynamics. *Can. J. Phys.*, 33:650, 1955.
- [Don14] William Donnelly. Entanglement entropy and nonabelian gauge symmetry. *Class. Quant. Grav.*, 31(21):214003, 2014.
- [Ear03] John Earman. Tracking down gauge: An ode to the constrained hamiltonian formalism. In Katherine Brading and Elena Castellani, editors, *Symmetries in Physics: Philosophical Reflections*, pages 140–62. Cambridge University Press, 2003.

- [Ebi70] D.G Ebin. The manifold of riemannian metrics. *Symp. Pure Math., AMS, 11*, 11,15, 1970.
- [Fey63] R. P. Feynman. Quantum theory of gravitation. *Acta Phys. Polon.*, 24:697–722, 1963. [,272(1963)].
- [Fis70] Arthur Fischer. The theory of superspace. In *Proceedings of the Relativity Conference held 2-6 June, 1969 in Cincinnati, OH. Edited by Moshe Carmeli, Stuart I. Fickler, and Louis Witten. New York: Plenum Press, 1970., p.303*, 1970.
- [Fra08] Bas Van Fraassen. *Scientific Representation: paradoxes of perspective*. Oxford University Press, 2008.
- [Fra14] Jordan François. *Reduction of gauge symmetries: a new geometrical approach*. PhD thesis, Aix-Marseille University, 2014.
- [GHR18] Henrique Gomes, Florian Hopfmüller, and Aldo Riello. A unified geometric framework for boundary charges and dressings: non-Abelian theory and matter. 2018.
- [Gom11] Henrique Gomes. Gauge theory in riem: classical. *J. Math. Phys.* 52, 082501, 2011.
- [GR] H. Gomes and A. Riello. The lack of empirical status of local gauge symmetries. *Forthcoming*.
- [GR17] Henrique Gomes and Aldo Riello. The observer’s ghost: notes on a field space connection. *JHEP*, 05:017, 2017.
- [GR18] Henrique Gomes and Aldo Riello. Unified geometric framework for boundary charges and particle dressings. *Phys. Rev. D*, 98:025013, Jul 2018.
- [Hea07] Richard Healey. *Gauging What’s Real: The Conceptual Foundations of Gauge Theories*. Oxford University Press, 2007.
- [IK85] C.J. Isham and K.V. Kuchar. Representations of spacetime diffeomorphisms. ii. canonical geometrodynamics. *Annals of Physics*, 164(2):316 – 333, 1985.
- [IM82] James Isenberg and Jerrold E Marsden. A slice theorem for the space of solutions of einstein’s equations. *Physics Reports*, 89(2):179–222, 1982.
- [IW94] Vivek Iyer and Robert M. Wald. Some properties of Noether charge and a proposal for dynamical black hole entropy. *Phys. Rev.*, D50:846–864, 1994.
- [KN63] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry. Vol I*. Interscience Publishers, a division of John Wiley & Sons, New York-London, 1963.
- [KO78] T. Kugo and I. Ojima. Manifestly covariant canonical formulation of yang-mills field theories. *Physics Letters B*, 73, 459-462., 1978.
- [KR83] W. Kondracki and J. Rogulski. *On the Stratification of the Orbit Space for the Action of Automorphisms on Connections. On Conjugacy Classes of Closed Subgroups. On the Notion of Stratification*. Preprint: Instytut Matematyczny. Inst., Acad., 1983.
- [Lan99] Serge Lang. *Fundamentals of Differential Geometry*, volume 191 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 1999.

- [LW90] J. Lee and Robert M. Wald. Local symmetries and constraints. *J. Math. Phys.*, 31:725–743, 1990.
- [Mer17] Flavio Mercati. Shape dynamics: Relativity and relationalism. *Oxford University Press*, 2017.
- [MH92] C. Teitelboim M. Henneaux. *Quantization of Gauge Systems*. Princeton University Press, 1992.
- [Mot95] Emil Mottola. Functional integration over geometries. *J. Math. Phys.*, 36:2470–2511, 1995.
- [Rov96] Carlo Rovelli. Relational quantum mechanics. *Int. J. Theor. Phys.*, 35:1637–1678, 1996.
- [Rov14] Carlo Rovelli. Why Gauge? *Found. Phys.*, 44(1):91–104, 2014.
- [Rovge] Carlo Rovelli. *Quantum Gravity*. Cambridge University Press, 2007, Cambridge.
- [RT74] Tullio Regge and Claudio Teitelboim. Role of Surface Integrals in the Hamiltonian Formulation of General Relativity. *Annals Phys.*, 88:286, 1974.
- [Str18] Andrew Strominger. *Lectures on the infrared structure of gravity and gauge theory*. Princeton University Press, 2018.
- [Teh15] Nicholas J. Teh. A note on roveli’s ‘why gauge?’. *European Journal for Philosophy of Science*, 5(3):339–348, Oct 2015.
- [TM80] Jean Thierry-Mieg. Geometrical reinterpretation of faddeev-popov ghost particles and brs transformations. *Journal of Mathematical Physics*, 21(12):2834–2838, 1980.
- [TM85] J Thierry-Mieg. Classical geometrical interpretation of ghost fields and anomalies in yang-mills theory and quantum gravity. In *Proceedings of a symposium on anomalies geometry topology*, 1985.
- [WG14] D. Wallace and Hilary Greaves. Empirical consequences of symmetries. *British Journal for the Philosophy of Science*, 65(1):59–89, 2014.
- [Wie17] Wolfgang Wieland. New boundary variables for classical and quantum gravity on a null surface. *Class. Quant. Grav.*, 34(21):215008, 2017.
- [Wie18] Wolfgang Wieland. Conformal boundary conditions, loop gravity and the continuum. *JHEP*, 10:089, 2018.
- [YM54] C. N. Yang and R. L. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Phys. Rev.*, 96:191–195, Oct 1954.