

UC Irvine

UC Irvine Previously Published Works

Title

On automorphism criteria for comparing amounts of mathematical structure

Permalink

<https://escholarship.org/uc/item/90s548c0>

Journal

Synthese, 201(6)

ISSN

0039-7857

Authors

Barrett, Thomas William
Manchak, JB
Weatherall, James Owen

Publication Date

2023-06-01

DOI

10.1007/s11229-023-04186-3

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

On Automorphism Criteria for Comparing Amounts of Mathematical Structure*

Thomas William Barrett¹, JB Manchak², and James Owen Weatherall²

¹Department of Philosophy, University of California, Santa Barbara

²Department of Logic and Philosophy of Science, University of California, Irvine

Abstract

Wilhelm (2021) has recently defended a criterion for comparing structure of mathematical objects, which he calls Subgroup. He argues that Subgroup is better than SYM*, another widely adopted criterion. We argue that this is mistaken; Subgroup is strictly worse than SYM*. We then formulate a new criterion that improves on both SYM* and Subgroup, answering Wilhelm’s criticisms of SYM* along the way. We conclude by arguing that no criterion that looks only to the automorphisms of mathematical objects to compare their structure can be fully satisfactory.

1 Introduction

There is a long tradition in the philosophy of physics of arguments that one theory, or formulation of a theory, is superior to another, empirically equivalent theory, on grounds of structural parsimony. The idea is that if two theories have the same empirical content, but one theory’s models have less structure than the other theory’s models, one should infer that the first theory attributes less structure to the world—and therefore should be preferred. Over the past decade, much effort has been devoted to making the comparisons of “amount of structure” involved in such arguments precise.

The most common type of criterion is based on the idea that one can compare amounts of structure by looking to the symmetries, or automorphisms, of the mathematical objects in question. If a mathematical object has more automorphisms, then it intuitively should have less structure than these automorphisms

*Forthcoming in *Synthese*. The authors can be reached at tbarrett@philosophy.ucsb.edu, jmanchak@uci.edu, and weatherj@uci.edu.

are required to preserve. The amount of structure that a mathematical object has is, in some sense, inversely proportional to the size of the object’s automorphism group. Earman (1989, p. 36) puts this basic idea as follows: “As the [...] structure becomes richer, the symmetries become narrower.”

To have a precise criterion, one needs to clarify the sense in which one object may have ‘fewer’ automorphisms than another. Swanson and Halvorson (2012) and Barrett (2015a,b) have proposed the following.

SYM*. *A mathematical object X has at least as much structure as a mathematical object Y if (and only if) $\text{Aut}(X) \subseteq \text{Aut}(Y)$.*

The condition $\text{Aut}(X) \subseteq \text{Aut}(Y)$, i.e. that the automorphism group of X $\text{Aut}(X)$ is a subset of that the automorphism group of Y $\text{Aut}(Y)$, is one way to make precise the idea that $\text{Aut}(X)$ is ‘not larger than’ $\text{Aut}(Y)$.¹ SYM* makes intuitive verdicts in many easy cases of structural comparison. Moreover, while it is not always explicitly mentioned, SYM* is the standard criterion in the literature.

Wilhelm (2021) has recently argued that a different criterion, which he calls Subgroup, is superior to SYM*.² Subgroup, too, has been used in the literature. For instance, Subgroup arises as an application of a category-theoretic criterion of structure comparison used by Weatherall (2016a,b), Rosenstock (2019), and Bradley and Weatherall (2020), among others, based on the property-structure-stuff framework of Baez et al. (2006), to the case of categories consisting of single objects and their automorphisms. Likewise, Barrett (2015a, p. 823) considers—and rejects—a very similar condition, which he calls SYM**.³ But Wilhelm presents the only sustained defense of Subgroup in the literature, and so we focus on his treatment.

We have a few aims in this paper, all of which dovetail off of Wilhelm’s discussion. First, we want to show that while SYM* does have shortcomings, Subgroup is worse. As we will show, SYM* is weak because it cannot rule on pairs of objects with different underlying sets; Subgroup is weak(er) because it gets clear cases wrong. Second, we will explain *why* SYM* succeeds in the cases to which it applies: it captures an important sense of structural comparison, which we call the “implicit definability conception” (IDC). Subgroup does not have any relationship to the IDC.

Nonetheless, Wilhelm’s arguments are suggestive. Drawing on Wilhelm’s intuition, we propose another criterion that improves on both Subgroup and SYM*. This criterion is, in a sense we make precise, the best one can do using only automorphisms. Even so, we will argue, it is not good enough. We will next argue that *no* criterion for comparing amounts of structure that looks solely to

¹We here follow the statement of SYM* given in Wilhelm (2021). Barrett (2015a,b) states SYM* using the relation “more structure than”, rather than Wilhelm’s “at least as much structure as”.

²Wilhelm introduces two criteria, Subgroup and Subgroup₂, where Subgroup₂ is strictly more liberal. We focus on Subgroup in what follows because our intent is to argue that Subgroup is already too liberal. We return to Subgroup₂ in the conclusion.

³See note 5.

automorphisms can be adequate. In brief, automorphisms alone do not encode all of the relevant facts about the structure of a mathematical object. We conclude with a brief discussion of how to think about comparisons of structure in light of the foregoing.

2 SYM^* , Subgroup, and Their Problems

Wilhelm’s main concern with SYM^* is a problem we will call **sensitivity**.⁴ In brief, SYM^* is too sensitive to the underlying sets of the objects being compared. To make this point, Wilhelm uses the example of two isomorphic groups that have different underlying sets. SYM^* says that these two groups have ‘incomparable’ amounts of structure, in the sense that neither has more nor less structure than the other. This is because no automorphism of the first group is also an automorphism of the second group. The two groups are isomorphic, however, and therefore should have *the same* structure. (Wilhelm, 2021, p. 6361) remarks that “structural comparisons should imply that if two mathematical objects are isomorphic, then those objects have the same amount of structure. SYM^* violates this condition. And that is a reason to reject it.”

We agree with Wilhelm in this case. One can also find similar examples where one object should have *more* structure than another, such as a topological space (A, τ) and a set B , where A does not equal B . Intuitively, one might think that a set with topology should have more structure than a bare set, even if the sets are not the same. But again, SYM^* does not rule on such cases. Examples like these show that there is a sense in which SYM^* is too *strict* a criterion for comparing amounts of structure. There are pairs of objects X and Y such we want to say that X has at least as much structure as Y , but SYM^* does not make this verdict.

It is to address sensitivity that Wilhelm proposes Subgroup.

Subgroup. *A mathematical object X has at least as much structure as a mathematical object Y if (and only if) $\text{Aut}(X)$ is isomorphic to a subgroup of $\text{Aut}(Y)$.*⁵

Despite the names, the important difference between Subgroup and SYM^* is *not* that Subgroup concerns the *subgroup* relation, whereas SYM^* concerns the *subset* relation. If $\text{Aut}(X)$ and $\text{Aut}(Y)$ are both automorphism groups, and $\text{Aut}(X)$ is a subset of $\text{Aut}(Y)$, then $\text{Aut}(X)$ is also a subgroup of $\text{Aut}(Y)$, since the two groups are groups of automorphisms on the same underlying domain, and thereby have the same identity element and multiplication rule (composition). The key difference is that SYM^* compares group structure relative to a preferred embedding of $\text{Aut}(X)$ into $\text{Aut}(Y)$ (the identity map) determined by the fact

⁴This problem is gestured at in the discussion of a criterion called SYM^{**} by Barrett (2015a), and it is mentioned explicitly by Barrett (2015b, p. 3). It is also discussed by Barrett (2021).

⁵Subgroup differs from SYM^{**} , introduced and rejected by Barrett (2015a), in that SYM^{**} says “more” where Subgroup says “at least as much”. Wilhelm’s modification avoids the problem that some objects have more structure than themselves.

that both are automorphisms on the same set. Subgroup, meanwhile, allows one to compare objects relative to *any* injective group homomorphism between their automorphism groups. No particular relationship between X and Y is required or respected. Thus we see a sense in which Subgroup is strictly ‘more liberal’ than SYM^* .⁶

Subgroup does address sensitivity. For instance, in Wilhelm’s example of two isomorphic groups with different underlying sets, it is trivial to verify that they have the same amount of structure according to Subgroup. Likewise, if a topological space (A, τ) is compared to a set B with the same cardinality as A , then Subgroup again makes the correct verdict: (A, τ) has at least as much structure as B .

Wilhelm offers a few other reasons to prefer Subgroup to SYM^* . One is that automorphism groups are groups, not sets, so the ‘subset relation’ that SYM^* employs is the wrong relation. Second, Wilhelm argues that Subgroup is a “strict generalization of SYM^* ” (p. 6365) and that this gives us reason to prefer it. Wilhelm explains why this is supposed to be a mark in favor of Subgroup as follows:

Subgroup expands the range of objects whose structures can be compared. So it supports more of the structural comparisons that mathematicians, physicists, and philosophers make. This, in fact, is the main reason why I prefer Subgroup [...] to SYM^* . (p. 6365)

Finally, Wilhelm claims that it makes intuitive verdicts. It makes many of the same verdicts that SYM^* did. It also deals well with cases like the isomorphic group case we have mentioned.

As we have seen, Wilhelm’s claim that Subgroup is superior because it explicitly involves the group structure of automorphism groups is a red herring. Both criteria involve group structure; they differ in what embeddings they allow. Moreover, while it is true that Subgroup is a “strict generalization” of SYM^* , we contend that that by itself is not a reason to accept Subgroup. A satisfactory generalization of SYM^* needs to make sensible verdicts in the cases where it differs from SYM^* . But Subgroup does not do this. Consider, for example, the following verdicts that Subgroup makes.

- The group \mathbb{Z}_5 (automorphism group \mathbb{Z}_4) vs any set with cardinality 2. Since the automorphism group of the latter can be properly embedded in the automorphism group of the former, according to Subgroup the set has at least as much structure as the group. This is a puzzling verdict, since one obtains a group from a set by *adding* structure to the set. Note that it also follows from Subgroup that a set of higher cardinality always has less structure than a set of lower cardinality, despite the fact that they are both totally unstructured sets.

⁶Note that there are examples where SYM^* and Subgroup agree that X has at least as much structure as Y , but Subgroup *also* rules that Y has more structure than X and SYM^* does not.

- The vector space \mathbb{R}^2 vs the group \mathbb{Z} (automorphism group \mathbb{Z}_2). \mathbb{Z} has at least as much structure as \mathbb{R}^2 according to Subgroup, another unintuitive verdict, especially since the vector space has underlying group structure *and* additional vector space structure on top of that.
- The vector space \mathbb{R} vs the vector space \mathbb{R}^2 . \mathbb{R} has at least as much structure as \mathbb{R}^2 according to Subgroup, even though \mathbb{R}^2 contains many subspaces isomorphic to \mathbb{R} and further structure relating those subspaces.

We need not multiply examples; the ones given should suffice to make one nervous about Subgroup. In what follows, we will diagnose what is going wrong with that criterion.

3 The Implicit Definability Conception

We begin by discussing why SYM^* works in the cases where it makes judgments. This argument has been offered before: see Barrett (2021) and the references therein. The basic idea is that there is a sense in which X has at least as much structure as Y according to SYM^* just in case X actually has all of the structures that Y has. We can make this idea precise by considering some simple facts about definability in first-order logic and model theory.⁷

A **signature** Σ is a set of predicate symbols. (Our results generalize to the case of function and constant symbols as well.) The Σ -terms, Σ -formulas, and Σ -sentences are recursively defined in the standard way. A **Σ -structure** X is a nonempty set in which the symbols of Σ have been interpreted. One recursively defines when a sequence of elements $a_1, \dots, a_n \in X$ **satisfy** a Σ -formula $\phi(x_1, \dots, x_n)$ in a Σ -structure X , written $X \models \phi[a_1, \dots, a_n]$. We will use the notation ϕ^X to denote the set of tuples from the Σ -structure X that satisfy a Σ -formula ϕ . A **Σ -sentence** is a Σ -formula with no free variables. An **automorphism** of a Σ -structure X is a bijection from X to itself that preserves the extensions of all of the predicates in Σ .

The basic set-up that we will employ in order to discuss definability is the following:

- Let Σ_1 and Σ_2 be signatures. The elements of Σ_1 and Σ_2 represent the ‘basic structures’ on the two objects that we will consider. These can be thought of as the structures that are explicitly appealed to in the notation we use to describe the objects.
- Let X be a Σ_1 -structure and Y a Σ_2 -structure. We will think of X and Y as the two objects whose structures will we be comparing. We temporarily assume that X and Y have the same underlying set.

We need to make precise what it means for X to define all of the basic structures of Y . So let $p \in \Sigma_2$ be one of the basic structures on Y . There are two

⁷See Hodges (2008) for further details.

particularly natural ways to make precise what it means for X to define p . We say that X **explicitly defines** p^Y if there is a Σ_1 -formula ϕ such that $\phi^X = p^Y$. And we say that X **implicitly defines** some subset $I \subset X \times \dots \times X$ (like the structure p^Y) if $h[I] = I$ for every automorphism h of X . We will focus on implicit definition.⁸

Here is the intuition behind these two notions of definability. If X explicitly defines the structure p^Y , then p^Y can be ‘constructed from’ the basic structures in Σ_1 . On the other hand, suppose that X implicitly defines p^Y . When this is the case, one often says that the structure p^Y is ‘invariant under’ or ‘preserved by’ the symmetries of X . It is common to infer from this that X comes equipped with the structure p^Y .⁹ The relation between these two varieties of definability is already well known. If X explicitly defines p^Y , then X implicitly defines p^Y . But the converse does not hold.¹⁰

We have the following simple result.

Proposition 1. *The following are equivalent:*

1. X implicitly defines p^Y for every symbol $p \in \Sigma_2$.
2. $\text{Aut}(X) \subset \text{Aut}(Y)$, i.e. X has at least as much structure as Y according to SYM^* .

Proof. Immediate from definitions. □

There is a natural desideratum about how structural comparisons should work. One wants to say that an object X has at least as much structure as Y if (and only if) X implicitly defines all of the structures of Y . Call this idea the Implicit Definability Conception (IDC). We take this to be a natural (weak) understanding of what it means to compare amounts of structure. It is also the only conception of structural comparison that has been explicitly articulated and defended (e.g. in Barrett, 2018, 2021). Prop. 1 establishes that X has at least as much structure as Y according to SYM^* when X implicitly defines all the structures of Y , and thus shows that SYM^* implements the IDC.

The analogue of Prop. 1 does not hold for Subgroup. Subgroup does not bear the same relationship to definability that SYM^* does.

Example 1. Let $\Sigma_1 = \emptyset$ and $\Sigma_2 = \{p\}$ be signatures, with p a unary predicate symbol. Consider the Σ_1 -structure A whose underlying set is the natural numbers \mathbb{N} , and the Σ_2 -structure B whose underlying set is \mathbb{N} with $p^B = \{0\}$. SYM^* says A has less structure than B . All the automorphisms of B are automorphisms of A but not vice versa. Subgroup says that they have the same amount of structure. Clearly $\text{Aut}(B)$ is isomorphic to a subgroup of $\text{Aut}(A)$.

⁸There are several varieties of implicit definability in the literature. The condition we consider here is one of the weaker ones. See Winnie (1986), Barrett (2018), and references therein for further details.

⁹See Dasgupta (2016) or Barrett (2018) and the references therein for elaboration.

¹⁰The converse *would* hold if we restricted attention to complete theories or if we strengthened our definition of implicit definability.

But conversely, the fact that A is isomorphic to the set $B - p^B$ implies that $\text{Aut}(A)$ is isomorphic to $\text{Aut}(B)$, since every automorphism of B is determined by its action on $B - p^B$. \lrcorner

In this example, A and B have the same amount of structure according to Subgroup, but A does not implicitly define all of the structure of B . Indeed, we generate B by adding structure to A . It is hard to imagine a coherent understanding of structure according to which B does not have more structure than A . At the very least Subgroup does not implement the IDC.

The example also shows that the automorphism group of an object (up to isomorphism) does not fully encode how much structure an object has. This means that a satisfactory criterion for comparing amounts of structure will need to appeal to *more* than merely an object's automorphism group up to isomorphism. In the next section, we consider what else one might need.

4 A Different Generalization of SYM*

As we have observed, the principal difference between SYM* and Subgroup is that SYM* compares the automorphism groups of X and Y only relative to a particular embedding, generated by a specific relationship between X and Y : namely, the identity map on their respective domains. If that map generates a group homomorphism, then $\text{Aut}(X)$ will be a subgroup of $\text{Aut}(Y)$. But it is only the group homomorphism (possibly) generated by this particular relationship between X and Y that matters.

These remarks suggest a different way to generalize SYM* to solve sensitivity. The proposal makes use of the following Proposition. In what follows, for any Σ -structure Z , $\text{dom}(Z)$ will refer to the domain of Z , i.e., its underlying point set. As before, let X be a Σ_1 -structure and Y a Σ_2 -structure.

Proposition 2. *Suppose $f : \text{dom}(Y) \rightarrow \text{dom}(X)$ is injective. Then the following are equivalent:*

1. *there is a group homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(Y)$ that commutes with f , in the sense that $s \circ f = f \circ Fs$ for every automorphism s of X .*
2. *X implicitly defines $f[I]$ for every subset I that Y implicitly defines.*

Moreover, if an F as described in 1. exists, it is unique.

Note that condition 1 does not quite imply that F is injective, and so F does not establish that $\text{Aut}(X)$ is isomorphic to a subgroup of $\text{Aut}(Y)$. But it does imply the following, which is 'close' to F being injective (and becomes 'closer' the 'closer' to surjective f is): If $Fs = Fs'$ for automorphisms s and s' of X , then $s|_{f[Y]} = s'|_{f[Y]}$. This follows immediately from the fact that f is injective. For suppose that $f(x) \in f[Y]$. We know that

$$s \circ f(x) = f \circ Fs(x) = f \circ Fs'(x) = s' \circ f(x)$$

The first and third equalities follow from condition 1, while the second follows from our assumption that $Fs = Fs'$. Since $f(x)$ was an arbitrary element of $f[Y]$, we have that s and s' are equal when restricted to $f[Y]$.

Proof. The proof that 1 implies 2 is exactly as in the proof of Proposition 4, which we will provide later. So we will show here that 2 implies 1. Suppose 2. If $s \in \text{Aut}(X)$, we define $Fs = f^{-1} \circ s \circ f$. First note that Fs is indeed a function from Y to itself, since 2 implies that $s[f[Y]] = f[Y]$. We show that Fs is an automorphism of Y . Since f^{-1} , s , and f are all injective, Fs is injective. And since $Fs(f^{-1} \circ s^{-1} \circ f(y)) = y$ for every $y \in Y$, Fs is surjective.

Let $p \in \Sigma_2$ be a predicate symbol. We show that $Fs[p^Y] = p^Y$. Let $(Fs(y_1), \dots, Fs(y_n)) \in Fs[p^Y]$. This means that $(y_1, \dots, y_n) \in p^Y$, and hence $(s \circ f(y_1), \dots, s \circ f(y_n)) \in s \circ f[p^Y]$. Now since Y implicitly defines p^Y , 2 implies that X implicitly defines $f[p^Y]$. So it must be that $s \circ f[p^Y] = f[p^Y]$, and therefore $(s \circ f(y_1), \dots, s \circ f(y_n)) \in f[p^Y]$. Since $s \circ f = f \circ Fs$, $(f \circ Fs(y_1), \dots, f \circ Fs(y_n)) \in f[p^Y]$. This immediately implies that $(Fs(y_1), \dots, Fs(y_n)) \in p^Y$, so $Fs[p^Y]$ is a subset of p^Y . Now let $(y_1, \dots, y_n) \in p^Y$, and let y'_1, \dots, y'_n be elements of Y such that $Fs(y'_i) = y_i$ for each i . This in conjunction with the definition of Fs then imply that

$$(s \circ f(y'_1), \dots, s \circ f(y'_n)) = (f \circ Fs(y'_1), \dots, f \circ Fs(y'_n)) \in f[p^Y]$$

As before we know that $s[f[p^Y]] = f[p^Y]$, so $(f(y'_1), \dots, f(y'_n)) \in f[p^Y]$. Hence $(y'_1, \dots, y'_n) \in p^Y$, and therefore $(y_1, \dots, y_n) \in Fs[p^Y]$. Altogether, this means that $Fs[p^Y] = p^Y$, and hence Fs is an automorphism of Y . One easily confirms that $F(s \circ s') = f^{-1} \circ s \circ s' \circ f = Fs \circ Fs'$, so F preserves composition and is therefore a homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(Y)$. By construction it commutes with f .

We now show uniqueness of F . Suppose that $F \neq F'$ are both homomorphisms $\text{Aut}(X) \rightarrow \text{Aut}(Y)$ that commute with f . So there must be $s \in \text{Aut}(X)$ and $y \in Y$ such that $Fs(y) \neq F's(y)$. Since f is injective, $f \circ Fs(y) \neq f \circ F's(y)$. Since both F and F' commute with f , this means that $s \circ f(y) \neq s' \circ f(y)$, a contradiction, and hence $F = F'$. \square

This proposition suggests a different criterion for comparing structure.

SYM⁺. *A mathematical object X has at least as much structure as a mathematical object Y , relative to an injective function $f : Y \rightarrow X$, if (and only if) there exists a group homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(Y)$ that commutes with f .*

SYM⁺, like Subgroup, is a strict generalization of SYM*. SYM* is just the special case of SYM⁺ where f is the identity on Y . SYM⁺ is also compatible with the IDC, as Prop. 2 shows. And SYM⁺ also solves sensitivity, since we can compare structures with different domains using SYM⁺. But according to SYM⁺, this can happen *only* relative to a particular injective map.

It is worth taking a moment to see how SYM⁺ handles example 1. For the object A to have at least as much structure as B , it would have to be the case

that there is an injective function $f : \text{dom}(B) \rightarrow \text{dom}(A)$ and a group homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(Y)$ that commutes with f . The most natural f to choose to compare A and B is the identity. In this case, A does not implicitly define the subset $f[\{0\}] = \{0\}$ since every element of A is mapped to some other element of A by some automorphism. But B does implicitly define $\{0\}$, since it is just p^Y , so Proposition 2 implies that there is no group homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(Y)$ that commutes with the identity map. (Note that this argument actually goes through no matter what map f we pick.) This means that according to SYM^+ , A does not have at least as much structure as B relative to the identity map. And on the other hand, it is easy to see that B has at least as much structure as A relative to the identity map $\text{dom}(A) \rightarrow \text{dom}(B)$, since B implicitly defines every subset that A does.

In many cases, the map f is fixed by context. For instance, isomorphic groups have the same amount of structure, according to SYM^+ , relative to any map f that realizes their isomorphism. A topological space (A, τ) has at least as much structure as a set B of the same cardinality, according to SYM^+ , relative to any bijection between A and B . Still, one might worry that structural comparisons should not require a map between the objects. How much structure something has does not depend on maps to other things; thus, comparing the structure of two different things should not depend on maps either. But this intuition fails for criteria that compare automorphism groups.

The reason is that an automorphism group is not just a group. It is a particular kind of *group representation*: specifically, a representation of a group as the automorphisms on a given object. The injective map f is important because it allows us to compare $\text{Aut}(X)$ and $\text{Aut}(Y)$ as *automorphism groups*. Results like Props. 1 and 2 show that to capture information about implicit definability using automorphism groups, one must keep track of how the abstract group structure is represented as maps on the domain of the object. In order to compare what structures are implicitly definable for different objects using their automorphism groups, one needs to know not only how the groups are related, but how their representations as automorphism groups are related.

Another worry about SYM^+ is that, since it considers only *injective* maps f , it will never rule that an object X has at least as much structure as an object Y if Y has greater cardinality than X . This is because if Y has more elements than X , there will be no injective maps from Y to X . This may seem puzzling, but we set this aside until the end of the next section.

We have emphasized how SYM^+ relates to SYM^* . But it is also related to Subgroup. Consider the following proposition.

Proposition 3. *Let $f : \text{dom}(Y) \rightarrow \text{dom}(X)$ be a bijection. Then the following are equivalent:*

1. *there is an injective group homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(Y)$ that commutes with f , in the sense that $f \circ s = F s \circ f$ for every automorphism s of X .*
2. *X implicitly defines $f[I]$ for every subset I that Y implicitly defines.*

Moreover, if an F as described in 1 exists, it is unique.

Proof. This follows immediately from Proposition 2 and the remark following it. \square

Prop. 3 tells us that SYM^+ looks like Subgroup in the special case where X and Y have equinumerous domains. Relative to a bijection $f : Y \rightarrow X$, SYM^+ says X has at least as much structure as Y only if $\text{Aut}(X)$ is isomorphic to a subgroup of $\text{Aut}(Y)$. Thus, Subgroup provides a necessary condition. But it is not sufficient. The subgroup of $\text{Aut}(Y)$ must be the image of a homomorphism from $\text{Aut}(X)$ that commutes with f . Some bijection must be specified. Different choices of bijection may produce different verdicts.

5 Triviality

We have now seen that SYM^+ is an automorphism-based criterion of structural comparison that is compatible with the ICD and solves sensitivity. But even so, SYM^+ is not fully adequate. The reason stems from another problem, which we will **triviality**. Triviality is a problem for all automorphism criteria of structure comparison.¹¹ Wilhelm mentions this problem but does not fully appreciate it (p. 6365). Here we present a new example and some new arguments that clarify why triviality matters.

In brief, the problem is that all criteria that look (only) to automorphism groups make implausible verdicts about objects that have trivial automorphism groups, i.e., whose only automorphism is the identity map. Examples of such objects are: any set with one element; any vector space with a fixed basis; the group \mathbb{Z}_2 ; any prime field; and so on. Such objects form a diverse collection; automorphism-based criteria struggle with that diversity.

Consider SYM^+ .¹² According to this criterion, given any two objects X and Y , if X has trivial automorphism group and cardinality as least as great as Y , then X has at least as much structure as Y . This is true no matter what structure X and Y actually carry. The problem is arguably even worse for Subgroup. According to Subgroup, given any mathematical structure Y , *every* object with trivial automorphism group has at least as much structure as Y .

Wilhelm is aware of triviality. But he argues that it “has some independent motivation”:

For suppose X is a spacetime with a trivial automorphism group. Then only the identity transformation preserves all of the structure of X . Only the identity transformation leaves X invariant. In other words, the spatiotemporal structure of X is so rich and complicated that every other transformation fails to preserve it. Therefore, X 's structure is ‘maxed out’. X is as structured as can be. (Wilhelm, 2021, p. 6365)

¹¹Triviality has been discussed by Barrett (2021); see also Weatherall (2021), who raises a related triviality concern for category theoretic approaches.

¹²Barrett (2021) offers an example to make the same point in connection with SYM^* .

Something similar can be said for SYM^+ . According to SYM^+ , if an object has trivial automorphism group, no other object can have more structure, relative to any injective map f . SYM^+ , too, says such objects are “maxed out”.

This is the wrong verdict. Consider the following example.

Example 2. It is well known that there exist spacetimes with trivial isometry groups. David Malament has sketched an elegant way to construct an example: start with Minkowski spacetime and then excise a compact region “shaped like a giraffe” from the manifold. Here, we present a precise variation of this idea. We restrict attention to the giraffe region itself, take its interior, and consider it as a spacetime in its own right. The resulting example is flat and, if the (radically idealized!) giraffe region is suitably chosen, it has an underlying manifold diffeomorphic to \mathbb{R}^n .




Figure 1: The ‘giraffe’ spacetime (M, g_{ab}) with vectors t^a and x^a at p .

Let (\mathbb{R}^2, g_{ab}) be two-dimensional Minkowski spacetime where $g_{ab} = \nabla_a t \nabla_b t - \nabla_a x \nabla_b x$. Let $M_1 = \{(t, x) \in \mathbb{R}^2 : 0 < t < 2, 0 < x < 3\}$. Let $M_2 = \{(t, x) \in M_1 : 1 < t < 2, 1 < x < 3\}$. Let $M = M_1 - M_2$ and consider the spacetime (M, g_{ab}) . We see that since M is a non-empty open star domain in \mathbb{R}^2 , it is diffeomorphic to \mathbb{R}^2 . Let $p = (1/2, 1/2)$ and consider the vectors $t^a = (\partial/\partial t)^a$ and $x^a = (\partial/\partial x)^a$ at p (see Figure 1). In order to preserve the lengths of maximal geodesics through p , one can verify that any isometry $\varphi : M \rightarrow M$ is such that $\varphi(p) = p$, $\varphi_*(t^a) = t^a$, and $\varphi_*(x^a) = x^a$. From this it follows that φ is the identity map given that isometries are rigid to first order at a given point Geroch (1969). \lrcorner

One does not want to say that the structure on this relativistic spacetime is maxed out. For instance, one could add an orientation field ϵ_{ab} to this Minkowski giraffe spacetime. That should generate an object with more structure, since metric+orientation is more structure than metric. But none of the criteria under consideration reflect that. One might also compare this example to a similar construction where one begins with a metric that admits fewer automorphisms. For instance, if Newtonian spacetime has more structure than Minkowski spacetime, then a giraffe modeled on Newtonian spacetime should have more structure than a Minkowski giraffe. None of the criteria we have discussed capture this, either.

The example is particularly striking because, while it has a trivial automorphism group, *locally* the structure on the Minkowski giraffe is hardly maxed out (cf. Manchak and Barrett (2022)). Suitably small regions the Minkowski giraffe with an added orientation will have more structure—according to SYM^* , Subgroup, and SYM^+ —than the corresponding small regions of the Minkowski giraffe. The moral we wish to draw is that no automorphism criterion will fully capture all of the structural comparisons one might wish to make. SYM^+ is useful for some kinds of cases, but not for others—much like SYM^* . On the other hand, one can capture a sense in which the giraffe intuitions just described may be made precise in a way that is compatible with the IDC.

Suppose again that X is a Σ_1 -structure and Y a Σ_2 -structure. If $f : \text{dom}(X) \rightarrow \text{dom}(Y)$ is an injective map, we let $f[X]$ be the Σ_2 -structure obtained from Y by ‘restricting’ Y to the image of f . Then we have the following.

Proposition 4. *Let $f : \text{dom}(X) \rightarrow \text{dom}(Y)$ be injective. Then the following are equivalent:*

1. *there is a group homomorphism $F : \text{Aut}(X) \rightarrow \text{Aut}(f[X])$ that commutes with f , in the sense that $f \circ s = Fs \circ f$ for every automorphism s of X .*
2. *X implicitly defines $f^{-1}[I]$ for every subset I that $f[X]$ implicitly defines.*

Moreover, if an F as described in 1 exists, it is unique.

In this case, when condition 1 holds, F is guaranteed to be injective, by a similar argument as given after Proposition 2.

Proof. The proof that 2 implies 1 is exactly as in the proof of Proposition 2. We will show here that 1 implies 2. Assume 1. Suppose I is implicitly defined by $f[X]$. Let s be an automorphism of X . We show that $s[f^{-1}[I]] = f^{-1}[I]$. Let $(y_1, \dots, y_n) \in I$ and consider $(s \circ f^{-1}(y_1), \dots, s \circ f^{-1}(y_n))$. We know that $s \circ f^{-1}(y_i) = f^{-1} \circ Fs(y_i)$ for each i , since F commutes with f . Since $f[X]$ implicitly defines I , $(Fs(y_1), \dots, Fs(y_n)) \in I$, which implies that $(s \circ f^{-1}(y_1), \dots, s \circ f^{-1}(y_n)) \in f^{-1}[I]$. And therefore $s[f^{-1}[I]]$ is a subset of $f^{-1}[I]$. Now let $(x_1, \dots, x_n) \in f^{-1}[I]$, so $(f(x_1), \dots, f(x_n)) \in I$. Since $f[X]$ implicitly defines I and Fs^{-1} is an automorphism of $f[X]$, this implies that $(Fs^{-1} \circ f(x_1), \dots, Fs^{-1} \circ f(x_n)) \in I$. Since F commutes with f , we know that $Fs^{-1} \circ f = f \circ s^{-1}$, so $(f \circ s^{-1}(x_1), \dots, f \circ s^{-1}(x_n)) \in I$. This means that $(s^{-1}(x_1), \dots, s^{-1}(x_n)) \in f^{-1}[I]$ and hence $(x_1, \dots, x_n) \in s[f^{-1}[I]]$. So $s[f^{-1}[I]] = f^{-1}[I]$. One shows that F is unique in precisely the same manner as in Proposition 2. \square

In order to understand how to use this result it is helpful to reconsider Example 1. We can use Proposition 4 to capture the sense in which A has the same structure as a part of B . Consider the map $f : \text{dom}(A) \rightarrow \text{dom}(B)$ defined by $x \mapsto x + 1$. It is easy to verify that the map $F : \text{Aut}(A) \rightarrow \text{Aut}(f[A])$ generated by f is in fact a group isomorphism. So Proposition 4 implies that A implicitly defines all of the structures that $f[A]$ does. (This is easy to verify by hand as well: the Σ_2 -structure $f[A]$ has underlying set $\mathbb{N} - \{0\}$ and $p^{f[A]}$ is

the empty set. So trivially A implicitly defines all of the structures that $f[A]$ does.) Altogether this illustrates that A has the same structure as the part of B corresponding to $f[A]$.

Proposition 4 is helpful in a few ways. First, it captures a sense in which an object X may have at least as much structure as ‘part’ of another object Y (In the proposition the ‘part’ is represented by the structure $f[X]$.) In the giraffe example, suitably chosen regions of two-dimensional Minkowski spacetime have at least as much structure as (regions) of the Minkowski giraffe. This construction allows us to say how one can ‘add’ structure to the Minkowski giraffe even though its automorphism group is trivial.

Prop. 4 is also useful in that it clarifies how an object X may have more structure than an object Y with greater cardinality, relative to some map f (now going from X to Y). The proposition shows that X may have more structure than the *part* of Y (given by $f[X]$) when X defines all of the structure that Y has on the image of X under f .

6 Conclusion

This paper has considered several automorphism-based criteria for comparing the amount of structure of mathematical objects. They all faced difficulties. SYM^* suffers from sensitivity, whereas Subgroup makes non-sensical verdicts. We went on to introduce a new criterion for structure comparison, SYM^+ , which extends SYM^* , solves sensitivity, and conforms with the IDC. We also showed how SYM^+ captures (some of) the intuition behind Subgroup. Even so, we then argued that no automorphism-based criterion is completely satisfactory, because all of them suffer from triviality. Triviality is not new, but we introduced an example that we feel highlights the difficulty. The example shows that objects with trivial automorphism groups need not have “maxed out” structure, since in many cases, one can add additional structure to such objects.

We wish to conclude by tying up three loose ends. First, observe that we did not propose a new criterion for structural comparison modeled on Prop. 4, as SYM^+ was modeled on Prop. 2. One could do so. But we would urge a different perspective. Prop. 4 captures a fine-grained relationship whose interpretation depends on things like the map f and the relationship between $f[X]$ and Y . It would not be the right criterion in all cases. We think the right moral to draw from the arguments in the foregoing is that (1) the IDC is the right way of thinking about structural comparisons; but (2) there is not a single criterion that implements that conception in all cases. Indeed, there is a sense in which the IDC itself *already provides* a fine criterion for comparing amounts of structure: X has at least as much structure as Y if (and only if) X implicitly defines all of the structures of Y . As we have shown, this is the kind of criterion that automorphism-based criteria are implementing when they work well; this is what Prop. 1 and Prop. 2 capture. Of course, this IDC criterion is itself not useful until one makes precise exactly what kind of implicit definability is at play. That is where the interesting further work on these questions will be. It will

involve investigating different varieties of implicit definability and proving more results like those above—capturing how we can tell when one object defines all of the structures of another—and then applying these results in cases of interest. These kinds of considerations naturally leads one to a kind of tool-box view of how to compare amounts of structure. There will be a variety of precise criteria that implement the IDC in different ways, and context and careful consideration of the questions at hand will guide which one should be used to compare the structure of different objects.

Second, note that when we argued that the Minkowski giraffe spacetime was not structurally maxed out, we were looking at other spacetimes — namely, small regions of the Minkowski giraffe — and asking ‘how many’ structure-preserving maps there are from these spacetimes into the Minkowski giraffe. When we add an orientation field there are ‘fewer’ such structure-preserving maps. This suggests that instead of just looking to automorphisms to tell how much structure an object X has, we should take a more ‘holistic’ approach and look to the entire class of structure-preserving maps between X and other objects of the ‘same kind’ as X . It is precisely this idea that is implemented by the Baez et al. (2006) “property-structure-stuff” forgetful functor approach. This approach, in general, has been adopted by several philosophers to compare, for instance, different formulations of Newtonian gravitation and electromagnetism. It is important, however, to recall a remark from the Introduction: Subgroup arises as the specialization of the property-structure-stuff approach to the case where one is comparing individual objects. So our arguments against Subgroup are also objections to this category-theoretic approach. We suggest that something is therefore missing from the category theoretic approach, and that more work is needed on that subject.

Our final remark concerns Subgroup₂, the second criterion of structural comparison proposed by Wilhelm. This criterion adds to Subgroup the possibility that one object has at least as much structure as another if the former’s automorphism group can be generated from that of the latter as the limit of a one-parameter family of representations. This proposal strikes us as an *ad hoc* attempt to save the idea that Galilean spacetime has more structure than Minkowski spacetime, since the Galilean group can be generated as the limit of a one parameter family of representations of the Poincaré group. But whatever its motivations might be, we observe that there is no relationship between group contraction and the IDC. Thus, if group contraction is to be included in a criterion of structure comparison, some other view of what the criteria intend to capture must be articulated.

Acknowledgments

We are especially grateful to Hans Halvorson and David Malament for many thoughtful discussions about topics related to this paper. We are also grateful to Isaac Wilhelm for his comments on the manuscript and to an anonymous referee who described the paper as ‘bizarrely mystical’.

References

- Baez, J., Bartels, T., Dolan, J., and Corfield, D. (2006). Property, structure and stuff. Available at <http://math.ucr.edu/home/baez/qg-spring2004/discussion.html>.
- Barrett, T. W. (2015a). On the structure of classical mechanics. *The British Journal for the Philosophy of Science*, 66(4):801–828.
- Barrett, T. W. (2015b). Spacetime structure. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 51:37–43.
- Barrett, T. W. (2018). What do symmetries tell us about structure? *Philosophy of Science*, 85:617–639.
- Barrett, T. W. (2021). How to count structure. *Forthcoming in Noûs*.
- Bradley, C. and Weatherall, J. O. (2020). On representational redundancy, surplus structure, and the hole argument. *Forthcoming in Foundations of Physics*.
- Dasgupta, S. (2016). Symmetry as an epistemic notion (twice over). *The British Journal for the Philosophy of Science*, 67(3):837–878.
- Earman, J. (1989). *World Enough and Spacetime: Absolute versus Relational Theories of Space and Time*. MIT.
- Geroch, R. (1969). Limits of spacetimes. *Communications in Mathematical Physics*, 13(3):180–193.
- Hodges, W. (2008). *Model Theory*. Cambridge University Press.
- Manchak, J. and Barrett, T. W. (2022). A hierarchy of spacetime symmetries: holes to Heraclitus. *Manuscript*.
- Rosenstock, S. (2019). *A Categorical Consideration of Physical Formalisms*. PhD thesis, UC Irvine.
- Swanson, N. and Halvorson, H. (2012). On North’s ‘The structure of physics’. *Manuscript*.
- Weatherall, J. O. (2016a). Are Newtonian gravitation and geometrized Newtonian gravitation theoretically equivalent? *Erkenntnis*, 81(5):1073–1091.
- Weatherall, J. O. (2016b). Understanding gauge. *Philosophy of Science*, 83(5):1039–1049.
- Weatherall, J. O. (2021). Why not categorical equivalence? In Madarász, J. and Székely, G., editors, *Hajnal Andr eka and Istv an N emeti on Unity of Science: From Computing to Relativity Theory Through Algebraic Logic*, pages 427–451. Springer International Publishing, Cham.

Wilhelm, I. (2021). Comparing the structure of mathematical objects. *Forthcoming in Synthese*.

Winnie, J. (1986). Invariants and objectivity: A theory with applications to relativity and geometry. In Colodny, R. G., editor, *From Quarks to Quasars*, pages 71–180. Pittsburgh: Pittsburgh University Press.