Journal of the American Philosophical Association (2026) 276–293
DOI: 10.1017/apa.2026.10025

How to Speak an Alien Language

ABSTRACT: *Human languages are powerful representational tools, but can they represent every possible kind of entity? This seems unlikely. We can easily imagine languages—God's language, or that of advanced extraterrestrials—that represent features of reality that our actual languages fail to capture. Eklund (2024) calls these alien languages. Yet despite the intuitive pull of this picture, it is unclear what alien languages, so understood, would amount to. I argue that there are no alien languages in this sense; human languages can represent any entity that can be linguistically represented at all. Still, I propose an alternative sense in which a language can be alien. On my cognitive account of alien language, a language is alien when linguistic understanding of it requires cognitive resources not used in understanding human languages. This account better explains the sense in which we can and cannot speak an alien language. We can represent whatever alien languages represent, but understanding alien languages may require cognitive resources that we lack.*

KEYWORDS: metaphysics, language, metametaphysics, representation, realism, linguistics

1. Introduction

Does reality contain alien structure, the kind of structure that no existing human language could possibly represent? Eklund (2024) says: yes. He calls such structure *alien structure*. Metaphysicians have long suggested that reality is radically different from how ordinary language represents it to be, but the existence of alien structure poses a different kind of challenge. In the case of alien structure, the nature of reality does not simply contradict ordinary thought and talk; it cannot be represented in ordinary thought and talk. Alien structure is — by definition — the kind of thing that can only be represented in languages that are structurally different from human languages. Such languages are called *alien languages* (Eklund 2024: 4).

The possibility of alien language and structure involves a fundamental tension. On the one hand, alien structure seems possible. We can imagine beings—literal extraterrestrials—that develop languages that are fundamentally different from ours, languages that can represent aspects of reality that our languages in-principle cannot. On the other hand, it seems that we can perfectly well describe such structure and language using our own language; after all, there is a sense in which I am currently talking about alien entities. If this is the case, however, then alien language is not so alien, after all.

Eklund (2024) attempts to resolve this tension by suggesting that there is an important difference between presentation and description. Alien structure can only be *presented* (accurately or perspicuously represented) by alien languages, but alien structure can nonetheless be *described* by familiar languages. Alien structure is alien in the presentational sense. Either we cannot speak an alien language, or our ability to speak an alien language would require resources that are in-principle unavailable to us in familiar languages.

Here, I argue that there is no plausible linguistic sense in which alien structure cannot be presented by familiar language. Using a variant of a strategy initially described by Hofweber (2017), I argue that we can use our familiar languages to represent alien entities in the same way that alien languages represent those entities. The basic strategy is to introduce a new word in our language that stands for whatever entity is represented in an alien language. While Eklund anticipates and rejects a simple version of this strategy, I modify the strategy to make it more sophisticated. The resulting modified strategy not only leads to true representations of alien structure, but perspicuous representations of alien structure.

While Eklund's attempt to demarcate alien structure fails, we should not abandon the idea of alien structure. Instead, I suggest that a better definition of alien structure makes reference to cognitive, rather than linguistic, capacities. Specifically, I focus on the idea of direct linguistic understanding; a person directly has a linguistic understanding of an utterance just in case they believe or entertain the content of that utterance upon hearing it (Millikan 2004; Mandelbaum 2014; Longworth 2018). I claim that an alien language is any language such that: to have direct linguistic understanding of the entire language, one must use cognitive resources that are not used in directly understanding familiar languages. An alien structure is then defined as any feature of reality that can only be mentally represented using cognitive resources that are not required to directly understand familiar languages. I call this the *cognitive account* of alien language and structure.

Though alien structure cannot be the target of direct linguistic understanding, there is another type of linguistic understanding possible. Linguistic understanding is indirect if it requires one to represent the semantic facts; on this view, understanding the utterance *U*, where *U* means *X*, consists in knowing, believing, or perceiving that *U* means *X* upon hearing *U* uttered (Evans 1982; Heck 1995; Brogaard 2019). I claim that we can have indirect linguistic understanding of alien languages in the sense that we can represent the semantic facts of alien languages. The resulting view is one in which humans can speak and understand an alien language, but they nonetheless crucially lack a certain kind of understanding common to familiar languages.

I begin with Eklund (2024)'s account of alien structure and why it matters (§2). I then turn to the basic strategy that threatens to undermine the possibility of alien structure: the piggybacking strategy (§3). While the simple piggybacking strategy discussed by Eklund fails, I develop a more sophisticated variant of the strategy, one that takes advantage of the idea of a structure-preserving map between familiar and alien expressions; I argue that Eklund does not successfully draw a principled distinction between alien and non-alien languages (§4–5). Afterward, I show that we can shed light on alien structure and language by looking at the literature on

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

linguistic understanding (§6). I end by giving a new, cognitive definition of alien structure, and I explain how the resulting view differs from nearby proposals by Eklund (2024) and Hofweber (2017) (§7).

2. Alien Structure

I start with a characterization of alien structure and why it matters for metaphysics.

> The notion of *alien structure*—alien metaphysical structure—is best explained via language. An ordinary sentence "*a* is F" represents it as being the case that object *a* is F, or has property F. These familiar sentences then present the world as having object-property structure (where, again, the talk of properties can be replaced by, for example, talk of objects being propertied). Alien structure is structure of a kind that only an alien language, if any, can present. (Eklund 2024: 7)

Alien structure is defined by reference to what ordinary (or familiar) languages cannot present, where "present" is a technical term for a certain form of representation. I will unpack this definition in three steps.

A *familiar language*—like the one this article is written in—has certain syntactic categories: singular term, predicate, quantifier, etc. These linguistic categories, insofar as they represent reality at all, represent ontological categories: object, property, etc. The relationship between the syntactic categories and ontological categories is simplest in the case that Eklund describes in the quotation. There, what corresponds to an atomic sentence in first-order logic represents an object having a property. The relationship between syntax and ontology is almost certainly more complex than this. It may be that singular terms can represent properties, or predicates can represent objects. Or there may be a layer of logical form that ultimately structures semantic content. In any case, our language has certain syntactic resources to map onto features of reality.

An alien structure is some feature of reality that cannot be presented using a familiar language. The feature of reality cannot be represented using a singular term, predicate, quantifier, and so on. As a consequence, alien structure does not fall into the familiar ontological categories like properties, objects, tropes, universals, and so on. Alien structure must be represented by an *alien language*, which is a language that is structurally different from familiar languages. By structurally different, I mean "they are different with respect to the *semantic categories* of their expressions or with respect to *modes of composition*" (Eklund 2024: 5). For example, an alien language contains expressions that allow it to refer to (say) objects*, which are not objects but something else entirely that our familiar languages cannot represent. Alien languages can be thoroughly alien in the sense that they contain no familiar semantic categories or modes of composition. Though when I speak of an alien language, I will assume that an alien language can contain some familiar elements.

Alien structure is, in some sense, unrepresentable in familiar languages. But what is the relevant sense of representation? Eklund uses the concept of *presentation* to identify two different sorts of representations: true representation and perspicuous

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

representation. The notion of a true representation will be taken as intuitive; we understand true representations in the myriad of ways developed in the philosophical literature. The notion of a perspicuous representation is more specialized and stems from the metaphysics literature. Eklund gives two examples to explain perspicuous representation.

The first example is from Rayo (2013: 3).

(1) There are no dinosaurs
(1a) The number of dinosaurs is zero

It is plausible that (1) and (1a) are made true by the same facts. Nonetheless, one might argue that one sentence is more metaphysically perspicuous than the other. For example, you may think (1a) is less perspicuous than (1) because it gives the impression that numbers are objects. Eklund uses the following definition of perspicuous representation: "A sentence perspicuously represents reality when the logical form of the sentence corresponds in the right sort of way to the aspect of reality that it represents" (2024: 17). The assumption is that the logical form of (1) and (1a) diverge, where (1a)'s divergence (for instance) fails to best represent the structure of reality.

The second example of perspicuous representation depends less on semantic features like logical form. Suppose "A" means the same as what "Socrates is wise" means. Both expressions have the same meaning, but "Socrates is wise" is a more perspicuous representation of reality because it displays (in Eklund's language) the object-property structure of Socrates' being wise (2024: 17).

Alien structure constitutes features of reality that cannot be presented in familiar languages. This is not to say that alien structure is ineffable (in Hofweber (2017)'s sense). Ineffable structure cannot be presented in *any* language, alien or otherwise. Moreover, it is important to note that alien structure cannot be presented in our language, but it can be described in our language. Eklund (2024: 17) points out that we describe alien structure in our discussion of alien languages. For example, we might say that alien languages have singular terms* rather than singular terms, where these terms* refer to something other than objects. To say this is to describe alien structure using our language, but our language does not present alien structure.

Much of my discussion will consist of a critical interrogation of the presentation/description distinction, so I will set it aside for now. Lastly, I will describe motivations for theorizing about alien structure. There are many motivations, but I want to focus on the potential of alien structure within metaphysics. There are two places where the notion of alien structure might play a useful role.

The first is fundamentality. Suppose Sider (2011) is right in thinking that reality has intrinsic objective structure. This structure can be more or less perspicuously described using human language. Our expressions are better or worse at carving nature at its joints. The possibility of alien structure means that the language that best describes objective metaphysical structure may in fact be an alien language (Eklund 2024: 40–43). To be clear: the notion of structure involved in alien structure is not necessarily the same as Siderean structure. The relevant possibility is one where, in fact, some alien structure is structural in Sider's sense.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

The second place where the notion of alien structure might be helpful is in hardcore metaphysics. Eklund describes a group of metaphysicians as *hardcore metaphysicians* insofar as those metaphysicians want to move away from reading metaphysical theses off of the structure of language. An example of this is Ramsey (1925)'s discussion of universals. Ramsey seeks to challenge the idea that there is a fundamental distinction between objects and properties, seeing it as an illusion generated by the linguistic distinction between subject and predicate. Other metaphysicians also want to criticize metaphysical approaches that draw conclusions about the structure of reality seemingly based on features of language (Simons 1991; Heil 2003; Smith 2005; Lowe 2012). Eklund argues that if one is truly a hardcore metaphysician, one should be open to the possibility that alien structure is relevant to some of our metaphysical questions.

There is more in Eklund's rich discussion of alien languages and alien structure, but I have sought to summarize the basic idea behind alien structure, and briefly give reasons why metaphysicians, specifically, might care about it.

3. The Piggybacking Strategy

Alien structure is, by definition, unpresentable in familiar languages. One strategy for overcoming the limits of familiar languages is by introducing terms that are stipulated to represent the same things represented in alien languages. Eklund (2024: 240) anticipates and rejects a version of this strategy, writing:

> Suppose the angels represent something, X, using expression E, where E stands in some alien semantic relation to X. How could we piggyback on this? Suppose we try to say "let F refer to what E represents" and then try to use F to represent it. Then even if by this means we can refer to what they represent, we do not thereby represent X in the same way as the angels do.

According to Eklund, the piggybacking strategy fails because our language will never represent alien entities in the same way as alien languages. Recall that Eklund claims that, even if a familiar language can describe alien structure, it cannot present, in his sense, alien structure. I will argue that the only clear and plausible understanding of "same way" or "presentation" must involve psychological facts.

I will start by describing the piggybacking strategy in more detail before describing Eklund's initial objection to it. Imagine that angels speak angelic language, one filled with angelic terms. Suppose that angelic language is an alien language, and our language is not an angelic language.

Here is one way to try to speak an angelic language. Suppose I am listening to angels speak. Eventually, an angel uses the term "flurg." I then introduce the following proposal:

> Let *F* refer to whatever angels refer to using the term "flurg," if it refers at all

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

My use of *F* is part of my language, but the term's meaning is stipulated so as to capture what angels mean in their language. Assuming that "flurg" is meaningful and refers in the alien language, it appears to follow that *F* is meaningful in our language; furthermore, *F*'s meaning is such that it refers to an alien structure (Eklund 2024, 237).

If I can use *F* to represent an alien entity, then it looks like the entities are not so alien, after all. I have successfully represented an alien entity because I have introduced a definite description that piggybacks on the alien use of the term. Piggybacking is a normal way to represent the world in a familiar language. Let "Pete the Robber" represent the man who robbed the grocery store yesterday, if he exists. It seems like the term is meaningful despite the fact that we simply introduced it.

The piggybacking strategy can be leveraged into an argument against alien structure. The argument goes like this.

1. If we can use familiar languages to present alien structure, then it would appear that there is no such thing as alien structure.
2. The piggybacking strategy allows us to use familiar languages to present alien structure.
3. Therefore, there is no such thing as alien structure.

Call this the *piggybacking argument* against alien structure. The argument seems valid. Given the definition of alien structure, it cannot be presented in a familiar language. The second premise is the one that Eklund targets. He says that familiar languages do not represent alien structure in the same way as alien languages.

By "the same way," Eklund mainly means that familiar languages do not present alien structure (Eklund 2024: 241). Remember that presentation means either true representation or perspicuous representation. Piggybacking can plausibly lead to true representation, but is the representation perspicuous? Plausibly not. The stipulated term *F* is not going to perspicuously represent the alien entity, even if it represents it. This is either because alien language better displays alien structure, or the logical form of alien expressions better captures features of reality's structure. In either case, the piggybacking strategy fails to show that familiar languages can present all possible features of reality, and thus fails to show that there is no such thing as alien structure.

The piggybacking strategy should be distinguished from the *adoption strategy*. Instead of using our language to represent alien structures, we might simply adopt aspects of the alien language itself. This would mean directly using alien speech and writing in order to present alien structure. Strictly speaking, this strategy is no threat to alien structure. Eklund's claim is that alien structure can only be presented in alien language; he does not claim that humans cannot speak and learn alien languages. One consequence is that "[an] alien language could be one we in principle could learn and use" (Eklund 2024: 17). This possibility raises different questions, but it does not immediately threaten the claim that familiar languages can present reality in the way that alien languages do.

Here is one possible moral of the story. Piggybacking is an indirect means to represent alien structure, but such a means is not itself a threat to the existence of alien

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

structure. As long as we distinguish between indirect and direct means, or between presenting and representing, we can preserve the sense in which there may exist alien structure.

Instead of giving up on the piggybacking strategy, one might redouble one's efforts. In what follows, I offer a more sophisticated version of the piggybacking strategy, one that can plausibly make sense of how a familiar language might perspicuously represent alien structure.

4. Structure-Preserving Display

In the imagined piggybacking scenario, we use a term to represent alien entities as if they exist. The problem is that our term may not be a perspicuous representation of alien structure. Eklund gives two examples of perspicuous representation. One concerns the surface features of an expression and what it displays. In this section, I argue that (in-principle) familiar languages can display alien structures just as well as alien languages.

Recall the intuitive example of displaying. We have two sentences:

(2) A
(3) Socrates is wise

The intuition was that (3) is more perspicuous than (2), on syntactic grounds alone. Why? Because (3) arguably represents the object-property structure of Socrates' being wise, while (2) does not.

Though whether (2) represents object-property structure depends on how the symbol is being understood. The letter "A" consists of an inverted letter "V" with a crossbar through the middle. Suppose that the object-property structure of Socrates' being wise is represented as follows: the inverted "V" part of the letter represents the property of being wise; the crossbar of the letter represents Socrates. Finally, the fact that Socrates *has* the property of being wise is represented by the fact that the crossbar and inverted "V" are connected as a whole. Thus, (2) does perspicuously represent the object-property structure of Socrates' being wise.

Of course, you may argue that (2) is not the kind of representation I have described. It is unlikely that you could have noticed the object-property structure represented in (2) prior to my drawing attention to a way of understanding the symbol. You could argue that the structural features that I have drawn attention to are not conventionalized, in some sense. This may be true, but the current point is that there is nothing intrinsic to the letter "A" that could stop it from being just as perspicuous as (3).

The informal example I have given is an example of a structure-preserving map, or homomorphism. Homomorphisms are usually defined in the context of mathematical structures. A structure is a set of objects plus some kind of structuring function. Let $\langle F, *^F \rangle$ and $\langle G, *^G \rangle$ be structures. A function h from F to G is a homomorphism if and only if, for any $a, b \in F$, $h(a *^F b) = h(a) *^G h(b)$. This is an abstract way to characterize the idea that a function defined on elements of F and their combination will preserve elements of G and their combination.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

Suppose F contains all the ways of dividing up the letter "A" and that $*^F$ is a structuring function that tells you how to put them together. The combination $a*^F b$ is what you get when you put the inverted "V" and the crossbar together: namely, "A." We define h such that $h(a) =$ "Socrates" and $h(b) =$ "is wise." The structuring function $*^G$ is just the syntactic operation used to combine singular terms and predicates. The example is simplified, but the basic idea should be clear: the same symbols can represent the same structure represented by other symbols, given a structure-preserving map between them. The concrete interpretation of the homomorphism will be some kind of implicit syntactic or semantic structure.

Now suppose we have an alien language that uses the symbol "⊙." I use the expression *CO* to represent whatever is referred to by ⊙, assuming that it refers. One complaint is that *CO* may not be as perspicuous as ⊙ with respect to describing alien structure. But there is a structure-preserving map between *CO* and ⊙ that allows *CO* to function so as to perspicuously represent whatever ⊙ represents in an alien language. In the case I have chosen, the symbols involved in both languages are actually topologically equivalent.

Two shapes are topologically equivalent if you can turn one into the other by stretching them, not tearing, gluing, or puncturing. Note that ⊙ is composed of a dot and a circle. Though dots are topologically equivalent to the letter *C* and circles are topologically equivalent to the letter *O*; you can turn them into each other by stretching them. Topological equivalence is just a special sort of homomorphism. The topological equivalence of certain letters of the English alphabet is not a deep result; it is given to students who learn topology for the first time. The purpose of raising the topological equivalence of certain expressions is simply to highlight the fact that the symbols themselves do not "display" reality in more or less perspicuous ways. There must be some additional form that goes beyond the spoken words or written symbols.

The relevant notion of form might be relevant to debates about the ontology of words. It is tempting to identify a word with a particular set of sounds or marks. The immediate problem with this view is that (*a*) some words are homophonic and thus are pronounced exactly like others and (*b*) the same word can be written in different ways. In response to these sorts of problems, some philosophers and linguists posit an abstract (or partially abstract) level of linguistic form used to individuate words (e.g., Nefdt (2019)). Perhaps the relevant display structure is part of the abstract structure that individuates words.

Whatever story we give about words, however, two relevant possibilities present themselves. If display structures are not essential to the individuation of words, then it will be contingent whether a word perspicuously represents some feature of reality better than another. My argument will still stand because we can then make some word have the same display structure as words from an alien language. If display structures are essential to the individuation of words, then things are more complicated. Strictly speaking, the display structure mirroring alien structure will not be attached to the same words (in English, say). In that case, it is possible for something to look and sound like an English word yet actually turns out to be an alien word. This is possible, but it is not plausible insofar as the expression is embedded in the English language in the familiar ways.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

My conclusion is that, in principle, there is nothing stopping our language from having the power to display alien structure. The expression *CO* may not appear to display alien structure, but that is because one may not know where to look. If there is a structure-preserving map between the alien symbols and symbols of our language —which there almost certainly will be, given the topological facts—we can find ways to spot the relevant structure.

5. Structure-Preserving Logical Form

In addition to the syntactic sense of perspicuous representation, there is also the semantic sense. On this approach, it is not the features of the symbols themselves that better correspond to reality, but rather the logical forms expressed by the symbols.

There are various things meant by the idea of a perspicuous representation of reality. Inspired by Rayo (2013) and Russell (2018), Eklund (2024: 16) characterizes the notion as follows: "A sentence perspicuously represents reality when the logical form of the sentence corresponds in the right sort of way to the aspect of reality that it represents." He explains further, writing:

> What sorts of things settle whether it is "Socrates is wise" or "Socrates instantiates wisdom" that corresponds in the right sort of way to what both these sentences represent? If one takes there to be such entities as facts, one may think that it is the number of constituents of the fact that matter. If the relevant fact only has constituents corresponding to "Socrates" and "is wise" (and where the "is" of predication is understood not as standing for a relation but as having a different function), the former sentence perspicuously represents the fact. If it has constituents corresponding to all of "Socrates", "instantiates", and "wisdom", it is the latter sentence that perspicuously represents the fact. (2024: 16)

I will briefly summarize the idea. We have two sentences:

(4) Socrates is wise
(5) Socrates instantiates wisdom.

The logical form of (4) contains two constituents while the logical form of (5) contains three constituents. Both (4) and (5) are true, but (4) is more perspicuous, say, if there are only two constituents of the fact that makes both sentences true.

This notion of perspicuous representation is clear enough. Though I should note that this sense of perspicuous representation is importantly different from Sider (2011)'s notion of a perspicuous representation of reality. The latter notion assumes that reality has an objective intrinsic structure (of a special kind) and that linguistic expressions can do a better or worse job capturing that structure. The current notion of perspicuous representation isn't committed to the existence of metaphysical naturalness.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

Recall that the piggybacking strategy said that we should make the following stipulation.

> Let *F* refer to whatever angels refer to using the term "⊙" if it refers at all

Eklund's claim is that

(6) ⊙

will be a more perspicuous representation of reality than the stipulated use of

(7) F

Why? Because (6) has a logical form that better represents reality than (7), despite both expressions representing the same objects.

I want to challenge this conclusion. Let us start with the simple case of (4) and (5). Suppose that the logical form of (4) is ⟨*Socrates*, *wise*⟩ and the logical form of (5) is ⟨*Socrates*, *instantiation*, *wisdom*⟩. It seems like a contingent fact that (4) is more perspicuous than (5). Presumably, we could use a language in which the logical form of (5) was the same as that of (4). I can simply stipulate that, for the purposes of the current discussion, I take (5) to be the expression whose logical form is the same as that expressed by (4). In such a special context, (5) is just as perspicuous a representation as (4).

This leads to a new strategy for piggybacking on the logical form of alien languages. If we are in the business of stipulation, we might go even further. We say:

> Let *F* have the same logical form as the term "⊙" if it has a logical form at all.

In a purely semantic conception of logical form, the actual surface expressions do not matter. But if they do not matter, we should be able to stipulate that our expressions have the same logical form.

Of course, having the same logical form might be too strong. You might argue that this constitutes using alien resources entirely as opposed to resources of familiar languages. Luckily, we can use a version of the structure preservation strategy for logical form.

Start with the least perspicuous representation (5). We know it is less perspicuous than (4). To fully understand why this is the case, I must give a more detailed explanation. Whether or not the logical form of an expression corresponds in the right way to reality depends on what the different aspects of logical form are supposed to correspond to. In the wisdom case, (4) "Socrates is wise" has a logical form that is perspicuous with respect to constituents. The idea is simple: the number of constituents of the logical form represents the number of constituents of the fact that makes the sentence true. Let us call this representational structure—connecting logical constituents to factual constituents—the *dimension* of perspicuity relevant to the expression.

It is important to note that the dimension of perspicuity matters when saying whether one expression is more perspicuous than another. Suppose the dimension of perspicuity is as follows. Logical constituents represent everything needed to bind the truthmaker of the sentence together. Suppose we have a view in which the instantiation relation somehow bonds with an object and a property to make

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

something a fact. On that dimension, then (5) is actually a more perspicuous representation than (4).

I raise the issue of dimensions of perspicuity to make the following point. At any given time, two logical forms can be perspicuous along different dimensions. However, there will only be one relevant dimension of perspicuity in typical assessments of perspicuity. With respect to alien logical forms, we need only capture the bits that make a difference to perspicuity.

I cannot create a structure-preserving map in the alien case, because it would presume that the relevant structure can be cast in the terms I am using. However, I can say:

> Let *F* have a logical form that preserves the relevant dimension of perspicuity of the logical form of the alien term"⊙," if it has a logical form at all.

This is an ambitious stipulation, but if it works, then we can piggyback on resources in our language to present (not just describe or truly represent) alien structure.

At this point, there is an obvious objection: these stipulations are too bold, too ambitious. While it is easy to use the term "Bob" to refer to someone you have heard others talk about, it is not possible to simply stipulate terms to have the structure of the logical forms of alien languages.

I accept that it may be hard or impossible in practice to successfully make these kinds of linguistic stipulations. The current question involves whether it is metaphysically possible for us to do such things. These stipulations may be difficult, but it is unclear what makes them impossible. Given the evidence of their possibility, it appears that there can be no such thing as alien languages, as Eklund understands them.[1]

Let me summarize. In this section, I have argued that: even if we understand perspicuous representation in a semantic sense, a version of the piggybacking strategy can be used to undermine the existence of alien structure. Here, we piggyback on the logical form of alien languages in order to perspicuously represent alien structure. I have now surveyed two accounts of perspicuous representation—syntactic and semantic—and argued that these accounts cannot be used to vindicate the existence of alien structure.

6. Linguistic Understanding

Recall the piggybacking argument.

> 1. If we can use familiar languages to present alien structure, then it would appear that there is no such thing as alien structure.

[1] One might object that what really matters is nomological possibility rather than metaphysical possibility. On this view, alien structure exists because it is not nomologically possible for human language to represent the same things that angelic terms represent. In response, I make two notes. First, this view constitutes a concession to my argument rather than an objection. Second, this view of alien language would be substantially different from Eklund's. Thanks to an anonymous reviewer for suggesting this objection.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

2. The piggybacking strategy allows us to use familiar languages to present alien structure.
3. Therefore, there is no such thing as alien structure.

Against Eklund, I have argued that Premise 2 is correct. It would appear that the conclusion follows: there is no such thing as alien structure.

Of course, the conclusion only follows if we also accept Premise 1. This premise is initially plausible because it is the way that Eklund defines alien structure. I instead claim that Eklund's definition of alien structure is a useful provisional definition that nonetheless fails to capture alien structure. I will now suggest that the difference between familiar language and alien language is not that familiar languages cannot present in the same way as alien languages; the difference is that the proper understanding of alien language requires cognitive abilities that humans lack.

To make this case for a cognitive approach to alien language, it is useful to recall the adoption strategy. Instead of using familiar linguistic resources to present alien structure, the adoption strategy requires us to simply use alien language. The adoption strategy is not, strictly speaking, a threat to the possibility of alien structure, because the alien hypothesis does not rule out the possibility that humans can speak an alien language; it only rules out the possibility that familiar languages can present reality in the same way as alien languages. Still, it is independently useful to imagine whether we can actually use an alien language.

Imagine we are given the alphabet of an alien language and its syntax. We are told how alien language works and what it refers to by metaphorical and analogical means. It refers to objects*, which are not exactly objects but are in the vicinity of objects. We can even imagine conversing with aliens who interpret us as making coherent claims (or claims*). In realistic versions of the scenario in which we have adopted an alien language, users of familiar languages—humans—will intuitively be missing something. It is questionable whether humans can understand alien languages, at least to the degree that aliens can.

Eklund (2024: 240) anticipates the possibility of the adoption strategy, thinking of it as the possibility that humans can merely parrot or echo alien languages. He compares it to a situation in which a person could have no conception of interrogatives, have no language that contains interrogatives, yet nonetheless parrot interrogative expressions. Eklund does not follow up on this possibility, but he suggests that the missing ingredient is some form of understanding; the speaker of the alien language has an understanding or at least a better understanding of what alien languages represent (Eklund 2024: 240).

My suggestion is that (*a*) with respect to understanding, the adoption strategy is equivalent to the piggybacking strategy, and (*b*) the difference in understanding may be all we need to distinguish between familiar and alien languages. The adoption strategy is equivalent to the piggybacking strategy in the sense that someone who uses familiar language to present alien structure would still be in a state where they do not understand in the same way or degree as aliens do. My conclusion, which Eklund clearly rejects, is that this difference in understanding is actually the key issue separating familiar and alien languages.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

But what is the relevant sense of understanding? The most natural place to start is with *linguistic understanding*, the state, ability, or process involved in knowing a language (as opposed to knowing substantial features of the world). Linguistic understanding is a notion that philosophers of language and linguists have theorized about quite a bit. For this reason, it is a good place to start.

The primary form of linguistic understanding concerns understanding an utterance. Someone says, "The cat is on the mat" and you understand what the speaker meant. Theorists of linguistic understanding attempt to theorize about utterance understanding. More precisely, attention is largely focused on the idea of utterance understanding as a state, rather than an ability, achievement, or process.

Theories of the state of linguistic understanding can be divided into two broad camps: indirect and direct. In what follows, I will consider both theories of linguistic understanding and argue that understanding an alien language (i.e., alien elements of an alien language) is possible in the indirect sense but not the direct sense.

Suppose an utterance *U* means *X*. According to indirect views of linguistic understanding, to understand an utterance *U* is to know, believe, or perceive that *U* means *X* (Evans 1982; Heck 1995; Hunter 1998; Fricker 2003; Brogaard 2019). Indirect views require that we represent, not just the meaning of an utterance, but the fact that an utterance means what it does. The main disagreement between indirect theorists is whether the relevant representation is perception-like or not. To motivate the perception view (held by Fricker (2003) and Brogaard (2019)), imagine the contrast between what happens when we hear an expression of a language we speak versus an expression of a language we do not speak. In the former case, we immediately and readily understand what is being said; like perception, we take upon the semantic facts non-inferentially and automatically.

Assume that the indirect view is correct. I argue that humans can understand alien languages, given the indirect view. We have linguistic understanding because we know something of the metasemantic relations involved. We know broadly what ⊙ is supposed to refer to. We lack the worldly knowledge that consists of the exact nature of the thing being represented, of course. But there is plausibly a difference in linguistic and worldly knowledge in the non-alien case. You can know the meaning of water without, at least in some sense, knowing anything about hydrogen and oxygen. Aliens may have substantially more worldly knowledge about alien structure, but it does not appear that this is a barrier to humans having a linguistic understanding of alien utterances.

On direct views of linguistic understanding, understanding an utterance *U*, where *U* means *X*, requires directly believing or entertaining the content *X* (Millikan 2004; Mandelbaum 2014; Longworth 2018). The direct view is motivated by the thought that linguistic understanding is not so demanding as to require knowledge or perception of semantic relations. Millikan (2004: 121) writes:

> We do not first understand what is said and then evaluate whether to believe it. Rather, we first believe what is said and then, if we are not under too much cognitive stress, we may think it over critically and reject it.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

For belief versions of the direct view, it may seem counterintuitive to think that every time you hear an utterance, you believe its content. Though Mandelbaum (2014) has argued that reflection and introspection are not good guides to what we believe; he gives empirical evidence to suggest that we immediately believe the contents upon hearing the utterance, even if our beliefs change soon after.

If the direct view is correct, it is less clear that humans can understand alien languages. There is a big difference between representing that *U* means *X* and directly believing or entertaining *X*. To believe *X*, one would have to have the right kind of cognitive abilities, but in the case of alien structure, we cannot guarantee that such abilities exist. Aliens have these abilities, but it is unclear that we do. Moreover, it may be that belief is not the right kind of relation relevant to alien understanding. Instead of believing things, aliens believe* things. We cannot believe* anything, and therefore we cannot have linguistic understanding of alien languages.

When the alien uses a symbol from an alien language, I do not entertain or believe* the same things that they do; alternatively, I do not believe them in the same way. This is the nugget of truth in Eklund's resistance to the idea of piggybacking. While familiar languages can present alien structure just as well as alien languages can, we cannot have a linguistic understanding of alien structure, assuming that linguistic understanding works in a direct way.

My discussion of linguistic understanding has mixed results. If understanding requires indirect representation, we can understand alien languages. If understanding requires direct representation, we likely cannot understand alien languages. Given that it is impossible to settle the direct/indirect debate in this article, what should we conclude?

One conclusion might be to rest our confidence in whatever theory of linguistic understanding is most popular. As far as I can tell, this means taking the indirect view as the correct one. Accordingly, most people believe the indirect view is true. If the indirect view is true, humans can understand alien languages. Therefore, we should believe that humans can understand alien languages.

Another possibility is that indirect and direct accounts of linguistic understanding are better understood as capturing two different aspects, or types, of linguistic understanding (Dänzer 2017; Grodniewicz 2023). This pluralist perspective is the most congenial to the current discussion. Recall the main question that prompted the discussion of linguistic understanding: what is the intuitive difference in understanding between the alien speaker of an alien language and the human speaker of an alien language? At the end of our investigation, the difference appears to be that humans lack direct linguistic understanding as opposed to indirect linguistic understanding.

7. The Cognitive Account of Alien Language and Structure

I have argued that the difference between alien and familiar languages concerns the fact that we (users of familiar languages) cannot directly understand alien languages. I now want to use this difference to give a definition of alien structure.

Recall Eklund's definition of alien structure. He starts by defining an alien language as a language that presents the world in a way that is structurally

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

different from familiar languages. Alien structure is the structure represented by alien languages. The cognitive definition of alien structure will also involve two steps. I start by defining alien language by reference to alien cognition. I then define alien structure by reference to an alien language.

On the cognitive account of alien language, an alien language is any language such that: to have direct linguistic understanding of the entire language, one must use cognitive resources that either (*a*) are not used in directly understanding familiar languages, or (*b*) are used in fundamentally different ways than they are used in familiar languages. I specify that one must use resources to understand the *entire* language, as opposed to some proper part of the language, because alien languages can contain fragments of familiar languages.[2] By cognitive resources, I mean the different possible mental states, attitudes, and processes used to represent or perceive reality. I will assume that directly understanding familiar languages uses some (but perhaps not all) of the cognitive resources of human beings. A language spoken by angels will be alien insofar as direct understanding of the angelic language requires resources that are not involved in the direct understanding of familiar languages. A case where the relevant language uses familiar cognitive resources in a fundamentally different way might be the case of languages that only use predicates and no singular terms, what Eklund (2024: 146) calls feature-placing languages. Even if humans can engage in predication, feature-placing languages may stretch that ability beyond its normal capacity. I will not further discuss the cases where cognitive resources are used in fundamentally different ways; for the sake of simplicity, such uses will be treated as if they constituted entirely different cognitive resources.

What it takes to directly understand a language will depend on details about the language itself. An extended example will be necessary, here. I will describe an imagistic language, which is a language that essentially requires language users to directly understand it using imagistic cognition. Roughly, imagistic cognition "is 'seeing' in the absence of the appropriate immediate sensory input, auditory mental imagery is 'hearing' in the absence of the immediate sensory input, and so on. Imagery is distinct from perception, which is the registration of physically present stimuli" (Kosslyn et al. 1995: 1335).[3] The idea is that one has internal representations that are image-like as opposed to propositional.

To directly understand an imagistic language, one must have the ability to engage in mental imagery. I take each language to specify what it would take to directly

[2] Given that most languages have infinitely many expressions, what does it mean to understand an *entire* language? I assume that one must have resources to understand the basic types of expressions of the language. For example, I can understand conjunction and atomic sentences, even if I cannot, for performance reasons, understand infinitely long conjunctive sentences. A language can have infinitely many token expressions, but ordinary humans can still understand the language if they understand the basic types of expressions. If there are infinitely many types that cannot somehow be generated by a finite set of fundamental types, then this may count as an alien language. Thanks to a reviewer for noting this complication.

[3] It is important to note that the nature of mental imagery is contested. For example, Pylyshyn (1981) argues that mental imagery is ultimately a form of propositional knowledge. The debate is sometimes understood as between the view that mental images are pictorial (Kosslyn et al. 2006; Nanay 2023) or propositional (Pylyshyn 2003). There is no obvious scientific or philosophical consensus in favor of one view over the other. My example of imagistic cognition presupposes something like the standard pictorial view. Thanks to a reviewer for this note.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

understand it. Thus, it lies in the nature of an imagistic language *L* that: necessarily, if someone directly understands *L*, they directly understand it using mental imagery. An imagistic language is alien if and only if humans are not required to use imagistic cognition in directly understanding familiar languages. Like any substantive thesis about the nature of language, it will be contested whether languages determine what cognitive resources are essential for directly understanding them. Nonetheless, I claim that the notion of language that makes sense of alien language should be understood in this way.

Let us now give a cognitive account of alien structure. Recall that Eklund's definition of alien structure holds that structure is alien if it can only be represented by an alien language. On the cognitive account, the notion of "what a language represents" is ambiguous. It is important to distinguish between what a language (independently of cognition) represents and what cognitive resources are used in understanding the language. The former notion of linguistic representation posits a relation between words and the world. I call this *narrowly linguistic representation*. The paradigm case of narrowly linguistic representation would be something like a radically externalist theory of linguistic representation. On such a view, meaning is constituted by causal relations between words and things, independently of what we think. The causal relations between words and things may go beyond what we can mentally represent. This notion of representation therefore does not straightforwardly correspond to language-like mental representations. Cognitive resources fundamentally involve mental representation. While mental representations and narrowly linguistic representations may be related, they certainly are not the same. The cognitive account of alien structure appeals to the former notion of representation.

Here is the official definition of alien structure. A feature of reality *F* is alien if and only if: *F* can only be mentally represented using cognitive resources that are not required to directly understand familiar languages. This is my proposed definition, but a simpler, equivalent definition can be given. I should first define two different kinds of cognitive resources. Say that a cognitive resource is familiar if it is required to directly linguistically understand familiar languages. Say that a cognitive resource is alien otherwise. We can now give a simpler definition of alien structure. A feature of reality *F* is alien if and only if: *F* can only be mentally represented using alien cognitive resources. There are two fundamental ways in which a feature of reality can be alien, given the cognitive theory.

In the first case, a feature of reality is alien if it is ineffable. For Hofweber (2017: 142), an ineffable fact is a fact that we (humans) cannot, in principle, conceptually represent. Ineffability, in the intended sense, is a structural inability. Ineffability is also a relative inability. There are facts that may be ineffable for us but effable for other creatures. Ineffable entities are necessarily alien. If we cannot mentally represent an entity, we cannot mentally represent that entity using familiar cognitive resources.[4]

[4] Given my earlier discussion, one might wonder if it is possible for me to accept ineffable entities. If the piggybacking strategy prevents us from making sense of language being structurally unable to represent certain facts, then one might think that we would have similar problems making sense of ineffability. I claim that the leap

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

In the second case, a feature of reality is alien if it is effable but cannot be represented using familiar cognitive resources. For example, there may be features of reality that can only be mentally represented using imagistic cognition. Suppose humans possess imagistic cognition, but imagistic cognition does not constitute a familiar cognitive resource because it is not required to directly understand familiar languages. Then it follows that the relevant features of reality are alien.

What about cases where a language is alien but contains familiar fragments? Imagine that a language contains atomic expressions that correspond to simple English sentences. For example, suppose *A* and *B* can be understood using familiar resources, but the language contains a connective "⊓" that cannot be understood using familiar resources. Now consider the expression $(A \sqcap B)$. On my view, $(A \sqcap B)$ represents alien structure even if *A* and *B* do not. For Eklund, whether or not $(A \sqcap B)$ represents alien features of reality is independent of issues of mental representation.

The resulting notion of alien language and structure differs from Eklund's account in substantial ways. Most obviously, it does not attempt to draw a distinction between how reality is presented and how it is described. Or at least, it does not draw such a distinction without reference to independently held cognitive capacities. As a result, we can clearly identify what is missing between alien and non-alien speakers: the requirement of familiar cognitive resources.

To end my discussion of alien structure, I will briefly explain the relationship between the cognitive definition of alien structure and the metaphysical motivations for alien structure. Metaphysicians highlight potentially radical divergences between how our language represents the world and what a metaphysically perspicuous representation of the world would be like. The cognitive account not only allows for the possibility of such divergences, but it also clearly explains the nature of such divergences. Reality can be accurately represented by different kinds of cognitive resources. The mistake is to think that reality is represented (or best represented) only by the cognitive resources used in directly understanding familiar languages.

There is more work to be done elucidating the cognitive account of alien language and structure, but my goal here has been to sketch a basic alternative to Eklund's account. The key to my cognitive account of alien language and structure is that it fundamentally appeals to cognition — specifically, it appeals to the cognitive resources required for directly understanding a language.

8. Conclusion

I have proposed a new definition of alien language and alien structure using the idea of direct linguistic understanding. The definition differs from Eklund's in virtue of giving us a clearer picture of the difference between familiar and alien languages. Furthermore, I should note that the cognitive account of alien language and structure

from linguistic to mental representation is non-trivial. Arguably, language is public in a way that thought is not. My stipulations can work because they are taken up by others, or they fit into a network of causal relations with the world. It is unclear whether the same thing is possible in the case of cognition. Thus, it is possible for ineffable entities to exist, even if there are no alien entities in Eklund's sense. Thanks to a reviewer for prompting this.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press

suggests a specific kind of research program. To further understand alien language and structure, we must inquire further into the language-thought interface, specifically with regard to (*a*) what cognitive resources humans have (or could possibly have) and (*b*) which of these abilities are (or could be) used in languages. By studying these questions further, we can begin to gain some insight into an initially opaque area of inquiry.

KEVIN RICHARDSON
PHILOSOPHY, DUKE UNIVERSITY, UNITED STATES
kevin.richardson@duke.edu

References

Brogaard, Berit (2019) 'Seeing and hearing meanings: a non-inferential approach to utterance comprehension'. In Anders Nes and Timothy Hoo Wai Chan (eds.), *Inference and Consciousness* (London: Routledge), 99–124.
Dänzer, Lars (2017) 'Utterance understanding, knowledge, and belief'. *Ergo*, 4, 529–555.
Eklund, Matti (2024) *Alien Structure: Language and Reality*. Oxford: Oxford University Press.
Evans, Gareth (1982) *The Varieties of Reference*. Oxford: Clarendon Press.
Fricker, Elizabeth (2003) 'Understanding and knowledge of what is said'. In Alex Barber (ed.), *Epistemology of Language* (Oxford: Oxford University Press), 325–366.
Grodniewicz, J. P. (2023) 'The representational structure of linguistic understanding'. *Inquiry*, 66, 1–26.
Heck, Richard (1995) 'The sense of communication'. *Mind*, 104, 79–106.
Heil, John (2003) *From an Ontological Point of View*. Oxford: Oxford University Press.
Hofweber, Thomas (2017) 'Are there ineffable aspects of reality?' *Oxford Studies in Metaphysics*, 10, 124–179.
Hunter, David (1998) 'Understanding and belief'. *Philosophy and Phenomenological Research*, 58, 559–580.
Kosslyn, S. M., M. Behrmann, and M. Jeannerod (1995) 'The cognitive neuroscience of mental imagery'. *Neuropsychologia*, 33, 1335–1344.
Kosslyn, Stephen M., William L. Thompson, and Giorgio Ganis (2006) *The Case for Mental Imagery*. Oxford: Oxford University Press.
Longworth, Guy (2018) 'Understanding what was said'. *Synthese*, 195, 815–834.
Lowe, E. J. (2012) 'Categorial predication'. *Ratio*, 25, 369–386.
Mandelbaum, Eric (2014) 'Thinking is believing'. *Inquiry*, 57, 55–96.
Millikan, Ruth Garrett (2004) *Varieties of Meaning*. Cambridge, MA: MIT Press.
Nanay, Bence (2023) *Mental Imagery: Philosophy, Psychology, Neuroscience*. Oxford: Oxford University Press.
Nefdt, Ryan M. (2019) 'The ontology of words: a structural approach'. *Inquiry*, 62, 877–911.
Pylyshyn, Zenon W. (1981) 'The imagery debate: analogue media versus tacit knowledge'. *Psychological Review*, 88, 16–45.
Pylyshyn, Zenon W. (2003) *Seeing and Visualizing: It's Not What You Think*. Cambridge, MA: MIT Press.
Ramsey, Frank P. (1925) 'Universals'. *Mind*, 34, 401–417.
Rayo, Agustín (2013) *The Construction of Logical Space*. Oxford: Oxford University Press.
Russell, Jeffrey Sanford (2018) 'Quality and quantifiers'. *Australasian Journal of Philosophy*, 96, 562–577.
Sider, Theodore (2011) *Writing the Book of the World*. Oxford: Oxford University Press.
Simons, Peter (1991) 'Ramsey, particulars, and universals'. *Theoria*, 57, 150–161.
Smith, Barry (2005) 'Against fantology'. In Johann C. Marek and Maria E. Reicher (eds.), *Experience and Analysis* (Vienna: HPT&ÖBV), 153–170.

https://doi.org/10.1017/apa.2026.10025 Published online by Cambridge University Press