

THE SCIENCE OF LOST WORLDS

Adrian Mitchell Currie

ISBN 978-1-77385-713-8

THIS BOOK IS AN OPEN ACCESS E-BOOK. It is an electronic version of a book that can be purchased in physical form through any bookseller or on-line retailer, or from our distributors. Please support this open access publication by requesting that your university purchase a print copy of this book, or by purchasing a copy yourself. If you have any questions, please contact us at ucpress@ucalgary.ca

Cover Art: The artwork on the cover of this book is not open access and falls under traditional copyright provisions; it cannot be reproduced in any way without written permission of the artists and their agents. The cover can be displayed as a complete cover image for the purposes of publicizing this work, but the artwork cannot be extracted from the context of the cover of this specific work without breaching the artist's copyright.

COPYRIGHT NOTICE: This open-access work is published under a Creative Commons licence. This means that you are free to copy, distribute, display or perform the work as long as you clearly attribute the work to its authors and publisher, that you do not use this work for any commercial gain in any form, and that you in no way alter, transform, or build on the work outside of its use in normal academic scholarship without our express permission. If you want to reuse or distribute the work, you must inform its new audience of the licence terms of this work. For more information, see details of the Creative Commons licence at: <http://creativecommons.org/licenses/by-nc-nd/4.0/>

UNDER THE CREATIVE COMMONS LICENCE YOU MAY:

- read and store this document free of charge;
- distribute it for personal use free of charge;
- print sections of the work for personal use;
- read or perform parts of the work in a context where no financial transactions take place.

UNDER THE CREATIVE COMMONS LICENCE YOU MAY NOT:

- gain financially from the work in any way;
- sell the work or seek monies in relation to the distribution of the work;
- use the work in any commercial activity of any kind;
- profit a third party indirectly via use or distribution of the work;
- distribute in or through a commercial body (with the exception of academic usage within educational institutions such as schools and universities);
- reproduce, distribute, or store the cover image outside of its function as a cover of this work;
- alter or build on the work outside of normal academic scholarship.



Acknowledgement: We acknowledge the wording around open access used by Australian publisher, **re.press**, and thank them for giving us permission to adapt their wording to our policy <http://www.re-press.org>

Artificing

Shuffling on their lower flippers, Prof. Ichthy shifts to stand over the skull on the podium. “Broadfoot is so unusual, so unfamiliar, that many of its properties can’t be straightforwardly inferred from its fossil. Because we’ve only one specimen,”—a gesture to the skull before them—“and because the specimen reveals an extinct animal so unlike today’s animals”—a gesture to the ancestral tree projected behind them—“finding the answers to some questions seems difficult indeed. For instance, how did Broadfoot move? Modern mammals are largely arboreal, built for scurrying through trees, so are hardly a good guide for Broadfoot’s apparently aquatic or semi-aquatic lifestyle. So, with limited remains, and no good analogues, what can we do?”

“How about . . . ” Ichthy pauses dramatically, allowing a little sparkle into their eye, “building a Broadfoot!?”

Prof. Ichthy gives an emphatic nod to a grad student, who steps rather gingerly onto the stage holding a small object, maybe two feet long. As they place it on the ground, it is revealed to be some kind of small, puppet-like, toy-like object, designed and painted to look a fair bit like Ichthy’s reconstruction. Some kind of simulacrum of Broadfoot, then. The grad student looks expectantly at Ichthy, clutching a length of string attached to a small hook at the front of the model.

“Currently living mammals all walk with what we call an ‘erect gait,’ their legs pointing straight downwards.” Ichthy performs a mime with their flippers, “Not so, it seems, for Broadfoot. This ancient lineage seems to walk with its limbs splayed outwards in what we call a ‘sprawling gait,’ quite similar to some of today’s lizards.” Ichthy’s flippers shift from pushing directly in front of their body to a wider position, reminiscent of an ungainly hug. “Partly based on the aspects of Broadfoot anatomy we’ve puzzled together, partly informed by reptilian and mammalian

analogues, we've built this little model. And . . ." (another dramatic pause) "lo and behold . . ."

Ichthy gestures grandly as their graduate student begins to pull the little model along by its string. Emitting a rather sad squeaking noise, it locomotes across the stage in an awkward (and, members of the audience agreed later, cute) gait. It first moves its limbs on the right side of its body and then the limbs on the left, creating a kind of rhythmic twist across its body. There is a small smattering of applause.




ILLUSTRATION 3.1 A Toy Broadfoot

"Excellent work." The professor nods to their student, who takes up the model Broadfoot and vacates the stage with some embarrassment.

In the last chapter, we discussed Liu et al.'s explanation of ivesheadiomorphs. They argued the fossils are not traditional taxa: instead of belonging to a particular clade, ivesheadiomorphs are the remains of a set of different rangeomorph (and other) lineages strangely preserved by the unique Avalonian conditions. Liu et al.'s argument relied on constructing a sequence of fossil forms, taking us from rangeomorphs to pizza disks in a smooth transformation, and describing a mechanism of decay specific to the Avalonian Ediacaran, which could explain that sequence.

You may have wondered, *how did they test that?* And you may also have thought—*oh, I bet I know how*. Perhaps they built an experiment that mimicked Avalonian conditions, introduced deceased organisms, and examined the subsequent decay patterns. As with Prof. Ichthy and their Broadfoot puppet, Liu et al. would have thus *artificed* a lost world, building a simulacrum and examining how that worked. In the face of loss, that is, a lack of Ns that you might use to construct a robust theory, you make some Ns for yourself.

If you thought something like this, you'd be exactly right.

In Alex Liu's doctoral thesis (Liu 2011), he describes a series of experiments. These involved filling a set of small plastic pots with seawater and a substrate of similar composition to Avalonian environments. The pots are populated with microorganisms collected from the seawater. These are placed in tanks blocked from both the sunlight and visible biotic intruders. Some pots are left for a few months, allowing the microbial community time to form something approaching the microbial mat of the Avalonian, while others are left for a few days. Afterward, body tissue from deceased metazoans is introduced: jellyfish, sea anemones, deep-water worms, and seaweed. These were left to decay, with timed photographs recording the process, before being coated in tuff (rock rich in volcanic ash) and left for a year. After this process, they were CT-scanned.

Liu highlights two important features of the results. First, the sea anemones did, indeed, form lobe structures analogous to those of ivesheadiomorphs (the worms and jellyfish decayed too quickly, and the seaweed decayed too slowly: sea anemones were just right). Second, decay occurred much more slowly in the pots that had been allowed a few months to settle; it seems that undisturbed microbial mats indeed affected decay rates.

Liu didn't necessarily need to artifice their proposed decay process to provide a fairly strong argument in favour of their interpretation; demonstrating a transition pattern from rangeomorphs to ivesheadiomorphs, coupled with a plausible mechanism, may have sufficed. Nonetheless, Liu's experiments go some way to both add plausibility to the ivesheadiomorph interpretation and provide a means to explore how that mechanism may have operated. Such experimental approaches are common throughout historical science (Jeffares 2008). My aim in this chapter is to characterize this strategy for tackling epistemic trouble caused by loss. I'll call the strategy "artificing."

"Artificing" involves conceiving of a denizen as a type (historical or ahistorical), and then constructing a token of that type (or, if you want, something analogous or sufficiently similar to that type). Liu conceived of

ivesheadiomorphs as multicellular organisms that decay under Avalonian conditions (a historical type) and then set up those conditions to test various hypotheses concerning them. Thus, he artficed.

Artifcing is easily linked with both philosophical work on modelling and experimentation. There is much familiar about artifcing from the philosophy of science interested in modelling; and no wonder, artifcing is, in a sense, a form of modelling.¹ Specifically, it is modelling with the aim of understanding lost worlds. The ubiquity of modelling and related scientific practices, such as idealization and abstraction, is well noted and well-studied by philosophers (e.g., Potochnik 2017; Frigg and Nguyen 2020) and certainly noted as important in the palaeosciences as well (Turner 2009; Bokulich and Oreskes 2017; Tamborini 2020; Currie 2024b). Despite this, I'm not interested in tackling any of the big topics surrounding scientific models in this chapter; rather, I'm interested in getting a grip on how strategies like artifcing work to reveal lost worlds. So, artifcing might be understood as a form of modelling (or experimentation!), but crucially, we should understand it as modelling or experimentation toward a particular purpose: tackling the challenge of loss. As such, rather than directly engaging with more general philosophical work on modelling and experimentation, I will pull on a different thread in the philosophy of science—*perspectivalism*—to capture crucial features of artifcing, in particular, the importance of *characterizing denizens*.

Why the term “artifcing”? I want to emphasize the active construction involved in the strategy: historical scientists respond to mysterious denizens by attempting to build them, or at least simulacra of them. As we'll see downstream, I also want to avoid overly representationalist approaches to the activities of scientists, and an emphasis on the construction of artifacts suits me—although one needn't agree with me on “artifactualism” (see chapters 6 and 7) to agree with the arguments of this chapter. In case any archaeologists are reading, note that by “artifact” I don't mean artifacts in the sense of traces in our species' material record, but rather artifacts that *scientists construct* to understand the past (the practices described in this chapter can be easily aligned with experimental archaeology; Currie 2022).

I'll first provide two contrasting case studies: one from experimental taphonomy (briefly leaving the Avalonian Ediacaran for the more familiar Cretaceous), and another using geometric modelling. I'll then provide an

1 I take this to be a classic example of the strategy of model-based-science as described by Godfrey-Smith (2006) and Weisberg (2007), although see Currie (2024b).

analysis of the strategy, giving a first pass at what I'll call "strategic perspectivalism." I'll close by discussing an extreme form of artificing: de-extinction.

1. A Dinosaur's Burrow

In his 2011 book *Convergent Forms Most Beautiful*, George McGhee notes that for any large, globally-dominant macrofauna, you'll find examples of lineages adapted to an underground life, that is, burrowers. Current mammal life boasts everything from moles to badgers to meerkats to wombats to aardvarks (and platypuses, of course!). "The number of independent convergences on burrowing morphologies is amazing", McGhee says, "particularly when it would seem that these morphologies are related to a rather restrictive and highly specialized way of life" (McGhee 2011, 26). Beyond current mammals, legless lizards, Carboniferous amphibians, various toads, and an army of insects and annelids are counted among the burrowing menagerie.

On this basis, we should expect burrowing non-avian dinosaurs. But for a long time, we didn't find any, and it's worth noting how surprising this is. Dinosaurs ruled the roost, as it were, for over a hundred million years, and every other successful lineage we know of has taken up a burrowing lifestyle. What gives? Here, we have a case of putative loss. Is there something about dinosaurs as a lineage, or their environment somehow, that meant that burrowers couldn't arise? In considering options, don't focus on dinosaur size. Although some were enormous, the average was around a middle-sized dog in stature (Farlow et al. 1995). Further, recent suggestions that *Lestodon* (one of the Pleistocene's enormous ground sloths that had Richard Owen speculating about a wildly verdant jungled South American past in chapter 1) might have been a burrower imply that size is no insurmountable obstacle to this kind of behaviour (Lopes et al. 2017). So, is there something special about dinosaurs or their environment that explains their apparently unburrowing nature?

This approach to tackling the question of dinosaur burrowing focuses on loss: What properties—which presumably no longer exist—explain the lack of burrowing dinosaurs? This isn't the only way we might explain the phenomenon, however.

A lack of burrowing dinosaurs could be more due to erasure than loss. First, we might think that there were dinosaur burrowers, but there was something about how they burrowed or the conditions of preservation during the Mesozoic that meant that they were not preserved. Here, we have a case of

loss, but loss concerning the status of preservation. This mirrors what we saw concerning ivesheadiomorphs: traces are hard to recognize due to surprising patterns of signal decay.

Second—and much more directly—maybe there were dinosaur burrowers, but we just haven’t found any yet. This is a simple case of erasure, and may well have been the case. In 2007, Varricchio et al. announced the discovery of a new dinosaur genus, one with a hitherto unseen lifestyle; they argued it lived in a burrow and, not only that, it also denned. That is, these animals raised their young in the protection of the earth. They named the genus *Oryctodromeus cubicularis* (the genus name meaning “digging runner” and the species “of the lair”), describing it as a fairly small biped. The find, from Montana’s Blackleaf Formation (dating to around 95 million years ago), “Consists of the skeletal remains of an adult and two juvenile dinosaurs preserved within an infilled burrow” (Varricchio et al. 2007, 1361). Perhaps dinosaurs burrowed after all.

However, Varricchio et al.’s interpretation faces a major objection. As Woodruff and Varricchio (2011) summarize, other ancient burrows were identified “primarily on the presence of articulated or near-complete remains within the burrow.” On the other hand, “the *Oryctodromeus* material differed in being (1) disarticulated and jumbled and (2) elevated in the fill above the burrow floor” (Woodruff and Varricchio 2011, 141). A good, fossilized burrow, then, contains articulated skeletons in a sensible place—on the floor of the burrow where they presumably expired. The jumble of disarticulated bones positioned at the top of the burrow suggests that the animals didn’t expire there, but were washed in, carried with the infill. Perhaps, then, the *O. cubicularis* specimens were not an unfortunate family caught expiring together in their burrow, but some bystanders who were carried in. Both tragic, but if the former is true, we have a much more significant discovery regarding the forms of life dinosaurs took up.

Let’s get clear on how the machinery from the last chapter applies here. There are at least two possible cases of loss on the table. First, non-Avian dinosaurs might be the only successful, global vertebrates that didn’t evolve burrowing; if so, then construed as such, we have a putative loss of that category. That is, the denizens *globally successful, non-burrowing taxa* are lost (recall chapter 2’s brief discussion of absences counting as loss). Second, this particular find might represent an as-yet unseen pattern of burrow preservation; if so, we have putative loss, or at least rarity, of that pattern. On the former interpretation, loss attaches to an ahistorical denizen (non-burrowing globally

successful taxa), on the latter, to the pattern of preservation. Woodruff and Varricchio hope for the latter: in a putative past world, we have a burrowing dinosaur that has left a pattern of burrow decay as-yet unseen.


So, what can we do? Given the topic of this chapter, I imagine you already have a good idea: artifice.

An extremely common mode of artificing asks after the conditions under which preservation occurs. In palaeontology, this often takes the form of “experimental taphonomy” (Briggs 1995). Various interpretations of specimens rely on not-yet-understood or unobserved forms of preservation. So, one approach is to bring about conditions analogous to those postulated and observe them, thus building a model of how those dynamics of information-preservation occur. In short, if you want to understand the processes by which traces form, build an experiment where something like those traces should arise. This is what Liu did in the chapter’s opening.

Woodruff and Varricchio (2011) also take this strategy. They build a half-scale model of the burrow, fill it with rabbit bones in various positions, and use a PVC supply tube to fill the burrow. They do this over and again, varying the angle of the tube and the position of the bones. At base, the apparatus consisted of the tube, filled with a clay-sediment mixture (and, occasionally, rabbit bones), two wooden saw-horses, and the burrow model itself (see Figure 3.1).

The sedimentary composition and structure of the burrow was based on reconstructions of the Blackleaf Formation. Each run of the experiment included a single articulated rabbit skeleton, either “within the burrow chamber or just external to the entire burrow structure within the sediment supply tube or mixed with the sediments” (Woodruff and Varricchio 2011, 143). Variations tested included whether the sediment mixtures were poured through in a single go or in quick incremental pours, at various levels of “in-fill energy” (determined by the angle of the tube), various skeleton positions (starting in the tube or in the burrow) and at various skeleton “densities”—whether the bones were dry or wet. They ran thirteen trials, each time leaving the model for a few days before excavating it, much as they had excavated the fossilized burrow itself. This allowed them to more or less map the relationship between the initial (articulated) skeleton position and various states of infilling. Throughout their discussion, Woodruff and Varricchio are careful to note various points of discrepancy between the original specimens and their artifact, including differences between rabbit and dinosaur bones, variations in temporal and spatial scale, and so forth. In each case, they provide

FIGURE 3.1
Burrow Model
and Supply Tube
(Detail from Fig.
3, Woodruff &
Varricchio 2011).
Reprinted with
Permission from
GeoscienceWorld



a set of reasons to think these discrepancies from the target shouldn't undermine their results.

The central result was that differences in final bone position “generally but not exclusively occurred in association with specific sedimentary deposits” (Woodruff and Varricchio 2011, 145). That is, the initial position of the bones didn't make a difference to the assemblage pattern, rather it was differences in infilling: “The angle of the supply tube and the mode of sediment delivery” (Woodruff and Varricchio 2011, 148). This result puts pressure on the interpretation that non-burrow-dwellers were swept into the burrow or otherwise deposited from outside. That claim, after all, relies on the idea that skeleton position provides important information about initial position

as infilling occurs. This doesn't in itself prove that *O. cubicularis specimens* were burrowers, but it does remove an important objection: "the experiment failed to falsify Hypothesis 1: sediment infilling of a burrow chamber can produce a disarranged and elevated assemblage from bones sitting within the chamber. This lends support to the inference of Varricchio et al. (2007), that the *Oryctodromeus* assemblage resulted from burial of bones located previously within the burrow." (Woodruff and Varricchio 2011, 148).

The results do more than simply leave both options open. The trials where the bones began outside of the burrow generated results that differed importantly from the conditions of the find itself. These runs produced a differently mixed set of bones with, tellingly, many more breakages than seen in the specimens. Woodruff and Varricchio, then, take this result to show that the skeletons' position in the burrow is unlikely to be the result of their being deposited from the outside: "For a debris flow to cover a long distance and account for the *Oryctodromeus* assemblage would require what would seem like a very precise and fortuitous delivery of an associated bone assemblage into the burrow" (Woodruff and Varricchio 2011, 148).

Before drawing some lessons, a few caveats about the case. You might notice that the specimens being in the burrow at the time of infilling doesn't in itself show that *Oryctodromeus* was a burrower. This doesn't distinguish between "burrowing," that is, an organism constructing a burrow, and "denning," that is, the organism living in, and raising its young in, a burrow. Rabbits both burrow and den, while polecats den but do not burrow—they rather find existing burrows (typically rabbit burrows) and take them over. So-called "burrowing owls," which nest in squirrel burrows, are another non-burrowing denner. I suspect this is a case of the kind of productive ambiguity that I've argued is distinctive of some claims in palaeontology (Currie 2015): while current evidence is insufficient to distinguish between burrowing hypotheses and denning hypotheses, palaeontologists adopt a more coarse-grained hypothesis inclusive of both. One way of telling between denning only and burrowing is to point to morphological adaptations of burrowing, and indeed, Varricchio et al. 2007 do highlight some. Their find's morphology "Exhibits features of the snout, shoulder girdle and pelvis consistent with digging habits while retaining cursorial hindlimb proportions" (Varricchio et al. 2007, abstract only). I've abstracted this from the discussion for the purpose of highlighting artifice, but it is nonetheless a nice example of "recontextualization," which I'll discuss soon.

Building on Ben Jeffares' work (2008), I've previously described practices like experimental taphonomy as experiments that aim to establish dependencies between current traces and their upstream causal ancestors (Currie 2018a, chapter 3). Which is to say, these experiments seek to establish a dependency between a current variable (say, the arrangement of a set of fossilized bones) and a past variable (say, the speed of infill). Woodruff and Varricchio's results suggest that there is a dependency relationship between the speed of infill and the arrangement of bones, but that there is a very minor dependency, if any, between the initial position of bones and their resting place. Modelling these causal regularities enables historical scientists to make trace-based inferences from, in this case, a preserved burrow to a burrow-dwelling dinosaur.

Before returning to the Ediacaran, I'll draw out some important features of artificing from this case. These points delve a little deeper into the strategy experimental taphonomists adopt when trying to establish dependencies between now and the past.

First, notice how Woodruff and Varricchio conceive of their target. In these instances of artificing, the denizens are understood in terms of patterns and processes of preservation. *O. cubicularis* specimens were understood as something like "burrowing animals being buried in various forms of sediment within their burrows." More carefully, they were simply understood as "vertebrate bones being preserved in an infilling burrow." This allowed, for instance, for them to take the rabbit bones used in the experiments as the same relevant kind as the specimens. The basic bet is that there is a set of dependencies that are common across these elements, which the experiment captures and allows us to probe.

As I'll explain in the next few chapters, this practice of finding the right way of characterizing denizens in a past world is crucial for how challenges from loss or putative loss are overcome. Plausibly, Varricchio et al. (2007) face loss. There are no currently existing burrowing non-avian dinosaurs—they have been lost through extinction—and until now there have been no traces of dinosaur burrows; either they never existed, we've been unable to recognize them, or they've been erased. The putative burrow itself has idiosyncrasies that make its interpretation difficult: rather than nicely-behaved articulated skeletons familiar from previous fossil burrows, we have a mess. The response is to consider which properties might be isolated for the purpose of empirically testing the relationship between burrow preservation and bone position. As such, Woodruff and Varricchio identify sedimentary

composition and vertebrate bones as properties of the past world that can be brought together in an artifact.

So, a crucial step in artificing is *characterization*: identify a set of properties in the past world that we can reasonably fabricate. As we'll see later in this chapter, processes of characterization—indexing our denizen as a kind under some description—are central features of strategic perspectivalism.

Second, although Woodruff and Varricchio consider properties of the past world in a kind of proxy system, they also *recontextualize* those results, in dialogue with the specimen itself and other aspects of theory, to generate an interpretation and justify aspects of their construction.² For example, their argument against the hypothesis that the bodies came from outside the burrow involved returning to features of the find itself, namely, their largely unbroken morphology and unsorted state, pointing out the differences between them and the results from experiments where bones began outside the burrow. If the find had been generated through that mechanism, they argue, it would likely have involved more breakages and been more carefully sorted, and thus would have been fortuitous indeed.

So, although various properties of past worlds are isolated to make it possible to generate tests, these tests are then brought back into conversation with particular, more fine-grained features of the specimens. I'll next note two further aspects of artificing.

Third, the construction and justification of the artifact are *bespoke*. Decisions made about how to design the experiment and what variations to perform are carefully guided by local, particular features of the specimen at hand. The specific kind of sediment involved, the decision to use rabbit bones, and the size and structure of the burrow (in particular, its “s-bend” shape) were based on the specimen and justified in terms of that. As opposed to the experimental paradigm being used “off-the-shelf,” in this instance, it is constructed for a specific purpose.

Fourth, Woodruff and Varricchio's artifact is what I'll call *pluripotent*. That is, the results it can generate are used for multiple purposes, sometimes at the same time. It explicitly tests a particular interpretation of the specimen while also generating a set of dependency relationships—a kind of causal

2 A notion of “recontextualization” is important for Sabina Leonelli's account of data (2019). Data must be *decontextualized* from the context of its generation to be fruitfully used in, for instance, database practices (to “travel”), but then must be *recontextualized*, particularly using metadata, to be deployed as evidence. I take our uses to be related but different.

model—between skeletal remains and various kinds of burrow preservations and infills. As they put it, “This experiment specifically addresses the recent interpretation of a Cretaceous skeletal assemblage as a dinosaur fill . . . but we hope that results represent a first step in understanding the sedimentology and taphonomy within large vertebrate burrows” (Woodruff and Varricchio 2011,140). They take steps toward this by comparing their results to other fossilized burrows. For instance, a Permian synapsid (the ancient lineage that mammals eventually arose from) was found at the bottom of its burrow with a little disarticulation (Damiani et al., 2003). Woodruff and Varricchio point out that one of their trials, those with dense, water-saturated bones, generated a similar pattern: “This suggests that largely undisturbed skeletons can be preserved provided that bone density is sufficient to resist transport” (Woodruff and Varricchio 2011,149). Thus, their study might have consequences for understanding other burrows, past or present (or future).

So, although model construction is closely aligned with the particular specimen of interest (it is bespoke), its use and results are nonetheless extended across a set of other cases (it is pluripotent). Woodruff and Varricchio generate *both* an experimental test of a particular interpretation *and* a more general model of the dependencies between infill conditions and fossil position. Pluripotency, it should be noted, differs from omnipotency: scientists cannot build whatever they want, and their artifacts cannot do anything. In chapter 5, we’ll touch on the possible dangers arising from the path-dependency of this kind of study. As we’ll see, the artifact being at the same time bespoke but also extending beyond its target is a distinctive feature of strategic perspectivalism.

Artificing can be useful for understanding more than conditions of preservation—that is, circumstances where denizens include trace-forming processes—they can also be used to probe many other properties of past denizens. To see this, I’ll examine some computational models of rangeomorph anatomy and growth. So, let’s return to the Ediacaran to further explore artificing.

2. Modelling Rangeomorphs

In discussing experimental taphonomy, I identified a set of properties relevant to artificing strategies. In the face of loss or putative loss, Woodruff and Varricchio *characterized* properties of the past world in a way that allowed them to construct their experimental system. How the system was


constructed and justified was *bespoke*, targeting specific features of the specimens. The results of their experimental work were *recontextualized* with features of the case itself, particularly specimen-relevant ones, thus generating further interpretations. Finally, the experimental system was *pluripotent*: testing hypotheses about a particular burrow, providing a system of dependencies between bone and infill properties, and allowing further recontextualized interpretations of further burrows.

Rangeomorphs are the most common multicellular denizens of the dark, benthic communities of the Avalonian Ediacaran. I've thus far focused on the apparently lost ecological and preservational status of those communities, a status partly generated by the lack of active movement. But that is only one aspect of this past world that seems significantly lost. Another is the morphology of the rangeomorphs themselves (e.g., Bamford and Narbonne 2009; Dunn et al. 2019; Taylor et al. 2023). The Ediacaran fauna contains many kinds of “frond” taxa with leaf-looking morphology. A “frond” is a structure most familiar from plants—their leaves—but is also found in cnidarians and algae. Fronds are formed of (to stick with botanical phraseology) a *root* (which attaches the organism to the ground), a *stem* (which either raises the frond up or increases its “reach” across the ground, while also acting as a hub for transporting nutrients), and, finally, *leaves* that sit along or atop the stem, often functioning in competition for resources (sunlight most familiarly, but also food particles in many non-plant cases).

So, when considered a member of the ahistorical type *organism with frond morphology*, rangeomorphs are certainly not lost, but members (presumably via convergent evolution) of a widespread brethren of all living things that are vaguely leafy. However, the details matter, and if we look closely, we find that rangeomorph morphology is considerably more mysterious. As Cuthill and Conway Morris have put it: “The branching morphology of Ediacaran rangeomorph fronds has no exact counterpart in other complex macroorganisms. As such, these fossils pose major questions as to growth patterns, functional morphology, modes of feeding and adaptive optimality” (Cuthill and Conway Morris 2014, 13122).

Rangeomorph morphology, like many fronds, follows a fractal pattern. A fractal is formed by repeating the same simple pattern over and over again, resulting in a pattern that is “self-similar,” that is, it looks the same across various scales. In rangeomorphs, this consists of at least three levels of branching self-similarity (see Figure 3.2). Thanks to some rather remarkable

FIGURE 3.2
Rangeomorph Fossils
Showing Three Levels of
Branching Self-Similarity
(Detail from Narbonne
2004, Fig. 2). Reprinted
with Permission from
the American Association
for the Advancement of
Science



three-dimensional fossils (Narbonne 2004), it seems clear that these develop more or less along two dimensions—they're flat.

This particular morphological arrangement is unique to the Ediacaran. As Laflamme, Xiao, and Kowalaweski (2009) put it, “Rangeomorphs are constructionally unique, adopting macroscopic morphologies that are uncommon or even unknown from modern or fossil metazoans” (Laflamme, Xiao, and Kowalaweski 2009, 14441). Given the litany of other uncertainties regarding rangeomorphs (no obvious feeding apparatus, unclear phylogenetic position, obscure developmental process, etc.), it is left quite unclear what to make of these strange critters. In other words, this morphology is lost (though not erased!) and, as such, understanding it requires overcoming some tricky epistemic challenges. We have, then, a classic challenge from loss. What to do? Given the topic of this chapter, you can guess: artifice. That is, characterize the past world in a way that affords the construction of a bespoke model and thus gains epistemic traction.




FIGURE 3.3
 Explanation of
 Fractal Model (Left)
 and One Example of
 Final Result (Right)
 (Details from Cuthill
 and Conway Morris
 2014, Figs. 1 and
 2). Reprinted with
 Permission from the
*Proceedings of the
 National Academy of
 Sciences*

Cuthill and Conway Morris (2014) present a fairly simple geometric model that constructs a rangeomorph-like fractal object by starting with a single “axiom” (the stem), adding a branch on the left, and then another on the right slightly further up, and then repeating this over a series of steps, where previous branches themselves become axes for further additions (see the left of Figure 3.3). In this way, they are able to generate shapes very similar to familiar fossil forms (compare the right image of Figure 3.3 with the fossil in Figure 3.2). They take the model to—at some level of approximation—fit the anatomy and indeed the development of rangeomorphs: “Reconstruction of overall morphologies that closely match the anatomy of the fossils . . . validates this model of branching and growth” (Cuthill and Conway Morris 2014, 13122).

So, we have our artifact: a geometric model of the three-branching fractal pattern distinctive of rangeomorph morphology. Cuthill and Conway Morris have characterized rangeomorphs as a historical kind: a lineage that grows via a three-branching fractal process. And they have thus constructed something

like an instance of that kind. The model is bespoke, guided in its construction by particular rangeomorph features. As we'll see, it is also recontextualized in generating interpretations and, perhaps more importantly, is pluripotent in providing a basis for more general exploration of fractally-growing features, both in rangeomorphs and beyond.

What is the use of such a model? For Cuthill and Conway Morris, the model provides an inroad to the anatomy, phylogeny, development, functional morphology, and even feeding strategies of rangeomorphs. I'll now cover phylogeny, functional morphology, and feeding.

Let's start with phylogeny: at base, should we think that rangeomorphs are a distinct taxonomic group, separate from other Ediacaran denizens? That is, are all rangeomorphs cousins? Despite their similarly fractal morphology, it is not immediately obvious that this would be the case, as Laflamme and Narbonne have pointed out: "Many different morphological constructions are present [in Avalonian deposits], and this further substantiates the proposition that Ediacaran fronds represent an ecological adaptation rather than a taxonomic group" (Laflamme and Narbonne 2008, 175). The suggestion here is that fronding evolved convergently across various taxa, rather than being the domain of a particular ancestral grouping.

Cuthill and Conway Morris take themselves to have provided a "unified" model that captures the variety of rangeomorph anatomies via variation on a set of fairly simple rules. They take this to be evidence of rangeomorph ancestral relationships: "Shared possession of an approximately self-similar and fractal, axial, apical, and alternate body plan supports a rangeomorph clade" (Cuthill and Conway Morris 2014, 13123). The thought is that the diversity of rangeomorph anatomies in the Ediacaran can be explained by their following a particular developmental regime. However, this regime cannot be extended across the Ediacaran fauna, suggesting that "the Ediacaran biota was far from homogenous but instead included diverse phylogenetic lineages and body plans" (Cuthill and Conway Morris 2014, 13124). That is, we should think of Ediacaran multicellular organisms as complex, evolutionarily and ecologically divided, and—as we'll now see—well-adapted to the particular environments they found themselves within. So, in addition to providing information about rangeomorph-as-fractal-growers, the model also informs ancestral interpretations. Which is to say, it is pluripotent.

To explore how Cuthill and Conway Morris' model applies to functional morphology, we'll take a tangent into a fairly radical, but increasingly supported, model of rangeomorph feeding. Here we see an example of

recontextualization, bringing the model back into conversation with local rangeomorph features, and a further feature of the pluripotent nature of artificing: being combined with other artifacts.

Laflamme and Narbonne (2008) interpret the various frond-like rangeomorph morphologies of the Avalonian as evolutionary convergences adapted to drawing nutrients from the water column. This involves seeing rangeomorphs as obligate osmotrophs. Osmotrophy involves feeding by absorbing carbon from the local environment. In single-celled critters, this basically involves interaction between the cellular wall and ambient carbon. Under those conditions, the size-to-surface-area ratio is critical. There is a trade-off between the energy demands of getting bigger and the relative amount of “wall” (compared to total mass) available to absorb that energy. Further, as there is no mechanism to distribute carbonic energy through the body, single-celled osmotrophy relies on passive diffusion, so the larger the body, the less efficiency on that front. Thus, there are size limits to obligate osmotrophy. We do see osmotrophy in some multicellular organisms—it is the basic metabolic driver of most fungi (and why they are decomposers *par excellence*)—but this is supplemented by other ways of gaining nutrients. Obligate osmotrophs as big as rangeomorphs would indeed be remarkable, and unlike anything we’ve seen.

In another case of artificing, Laflamme, Xiao, and Kowalaweski (2009) construct a theoretical model aimed at supporting an osmotrophic model of feeding. They focused on comparing the body-size-to-volume ratios of rangeomorphs and the largest (“megabacteria”) obligate osmotroph bacteria (their use of megabacteria as comparisons is an instance of “anchoring,” which we’ll meet in chapter 5). As should be becoming familiar, Laflamme et al. characterize a particular set of traits to compare past-world denizens with better-known contemporary analogues. At base, because rangeomorphs (and another Ediacaran denizen, the Erniettomorphs, but let’s not worry about them) have that modular, fractal structure, it is fairly easy to model the trade-offs between overall size and surface area to size ratio: you represent the length, breadth, and overall size of the modules (in other words, the fronds). Although the results don’t demonstrate bacteria-like surface-area-to-volume ratio efficiency (how could they?), they *do* underwrite Laflamme et al. exploring morphological adaptations in rangeomorphs that could facilitate this.

Drawing on older work by Gould (1966), Laflamme et al. (2009) point to two morphological strategies that reduce the costs involved in size versus surface-area-to-volume trade-offs (here characterizing the denizens as

morphologies facing that kind of trade-off). First, *constructional flattening*. If you want to increase size, but minimize cost, then increase size in one dimension: worms get longer without getting thicker, and leaves are typically two-dimensional. And indeed, we see this in rangeomorphs: fossils that preserve the third dimension don't have all that much to preserve. The second strategy involves going fractal. Because rangeomorphs are constructed from modular, self-similar fronds of varying sizes, the use of tiny, complex fronds keeps a reasonable surface ratio despite increases in size. Laflamme et al.'s (2009) results suggest that "fractal branching in rangeomorphs results in significant increases in SA/V ratios that are comparable to some giant bacteria" (Laflamme et al. 2009, 14442). In other words, rangeomorph anatomy achieves impressively efficient surface-area-to-volume ratios despite its size. Third, fill out your increasing body size with stuff that doesn't have metabolic costs, such as non-organic sediment or dirt. Although fossil evidence for this is limited—but not ruled out!—if Laflamme et al.'s modelling results are correct, this suggests that to be efficient at those sizes, rangeomorphs likely did incorporate non-organic, non-metabolic material into their growth.

So, by characterizing some of the geometric features of rangeomorphs and comparing them to megabacteria, Laflamme et al. (2009) are able to explore the constraints individual rangeomorphs would face if they did turn out to be obligate osmotrophs. Where Laflamme et al. focus on the efficiency of particular rangeomorphotypes, Cuthill and Conway Morris' capacity to generate a set of varying morphotypes allows them to compare them to each other, as they are represented in Figure 3.3. If we follow Laflamme et al. in imagining rangeomorphs as osmotrophic, what ecological and evolutionary information might we gain by comparing their morphologies as generated by Cuthill and Conway Morris' model?

Rangeomorph body types can be organized along the usual spatial dimensions (width, depth, and height) to compare them and hunt for patterns, as in Figure 3.4 Cluster analysis of this information reveals three strategies: go for height, such as *Trepassia wardae*, find a trade-off between height and volume, such as *Bradgatia linfordensis*, and maximize width, as in rangeomorphs spread across the ocean floor, such as *Hapsidophyllas flexibilis*.

This information is then brought into discussion with yet another model (the details of which we'll spare ourselves!), which examines the effects of dense rangeomorph canopies on water flow. The model suggests, in a way somewhat reminiscent of plants competing for sunlight, that a dense rangeomorph canopy would stymie water flow, so that there could be selective




FIGURE 3.4
Comparative
Rangeomorph Space-
Filling (Fig. 3, Cuthill
and Conway Morris
2014). Reprinted with
Permission from the
*Proceedings of the
National Academy of
Sciences*

pressure for increased height. Some rangeomorphs reached above the canopy where less-restricted water flow would increase the efficiency of osmotrophy. Cuthill and Conway Morris weave all of this into an evolutionary narrative:

Among the in situ communities of the earliest Avalon Assemblage, rangeomorphs are by far the most diverse and abundant macroorganisms. Furthermore, establishment of the major rangeomorph space-filling strategies preceded the appearance of a much wider range of macroorganisms . . . This suggests that early diversification of rangeomorph branching patterns effected a radiation in the key characteristics of body-size and

microhabitat . . . and was driven by ecological competition between rangeomorph taxa. (Cuthill and Conway Morris 2014, 13125)

In other words, early metazoan ecosystems were dominated and shaped by an adaptive radiation of rangeomorphs, a radiation largely due to that lineage exploiting the various space-filling strategies available to its body plan. In Cuthill and Conway Morris' and in Laflamme et al.'s (2009) models, we see a similar strategy as in the burrowing case: characterize a set of properties, build a bespoke model that instantiates those properties, generate results, recontextualize them with other information, and use these as a basis for further models. Let's think about this strategy yet more abstractly.

3. Strategic Perspectivalism

"The real secret to my success . . ." Prof. Ichthy glances at their students, "Well, our success—relates to how we *look* at Broadfoot. We don't just think of it as a fossil specimen, nor a mammal. Instead, we think about it from many, many different perspectives."

Ichthy begins to gesture a little erratically, their flippers tracing abstract geometries in the air.

"I like to imagine the animal as a kind of puzzle, a puzzle consisting of multiple complex dimensions. To solve the puzzle, you need to focus on just a few dimensions at a time, thus generating partial, small-scale solutions. But then, with those partial solutions, you get hints for how to solve aspects of the puzzle across other dimensions. Imagine a crossword puzzle, but instead of just a horizontal and vertical axis, puzzles work across three, four, many axes. And imagine that questions and solutions aren't simply linguistic, but are sometimes pictograms, sometimes sudoku, and at times purely geometric. Understanding such a puzzle would require looking at it in many different ways and considering how these ways of looking and the partial solutions they provide relate to each other. Just like that, each way we look at Broadfoot gives us crucial information, and that information bears on the other ways we look at Broadfoot." They bring their flippers around in a final circular motion.

"All this perspective-switching can be dizzying sometimes, but you get used to it."

How do historical scientists uncover lost worlds? How do they solve challenges from loss? One strategy is artificing, and we'll meet more in the following chapters. Each of these strategies belong to a larger category: the creative and opportunistic use of what I'll call *strategic perspectivalism*. That is, the adopting of a perspective that allows characterizing apparently lost denizens in ways that enable strategies like artificing to gain traction. In this section, I'll characterize this via a discussion of perspectivalism more generally.

We've examined two main cases of artificing. The first involves recreating something like the conditions of preservation—the processes of trace formation—the second involved recreating a past world itself. These two don't neatly divide: modelling the patterns created by collapsing burrows involves some understanding of the structure of those burrows and burrowing critters, and recreating past Avalonian conditions involves an intimate understanding of the conditions of its preservation. Regardless, these examples provide some clues as to how the artificing strategy works, and some features which, as we'll see, are common across the challenges of understanding lost worlds.

Across our cases, we saw a particular process. First, characterization: denizens are indexed via particular properties such as Avalonian-like conditions of preservation, burrow infill, and rangeomorph growth patterns. Second, these properties are instantiated or generated by constructing some scientific object: a scale model burrow or a geometric model, say. These models are bespoke both in terms of construction and validation. At base, this involves providing reason to think that the procedures and various pragmatic features associated with them will not bias the results as applied to the target past system, and will further make sense of the resultant artifact itself. Third, recontextualization occurs. Here, other features of the find and other aspects of context are brought into dialogue with the results of the study (Turner and Turner 2021). Roughly, an interpretation is generated and defended in various ways by bringing it into dialogue with other features of the past world. These bespoke models are pluripotent: providing information about other instances, as we saw with the burrowing case, and affording the construction of further objects, as we saw with the various models of rangeomorph growth, phylogeny, functional morphology and ecology.

I think we can understand the three-step process of artificing as an instance of a *perspectival strategy*. To understand what that is, we first need to get perspectivalism on the table.

Perspectivism, or perspectivalism, is at base a claim about the nature of knowledge. It claims that all knowledge is *from a perspective*, which is to say,

there is no epistemic “God’s-eye view” to be had. Instead, perspectivalists insist that knowledge is always had from some vantage point. How exactly to characterize “vantage point” is a vexed issue, as is the extent to which perspectives, and therefore knowledge, end up being private or subjective. Perspectivalism is a philosophical view with a grand history, which has fairly recently been a focus in the philosophy of science. These discussions are often in the business of threading needles between constructivist and realist conceptions of knowledge. Is it possible to reconcile the idea that all knowledge is from a perspective with the idea that the knowledge is of a single, concretely existing world? Philosophers of science developing perspectivalism often take their cue from visual imagery or other metaphors from the arts (see, in particular, Giere 2006; Elgin 2017; Massimi 2022). The views developed there bump up against other major traditions in epistemology and the philosophy of science, particularly standpoint theory and various forms of scientific pluralism.

As we’ll see, my interest in perspectivalism is partly instrumental, as committing to some general epistemic view about the relationship between perspectival knowledge and reality is beyond my remit, as is committing to a particular story about what a “perspective” is. Nonetheless, recent discussion, particular Michaela Massimi’s, is insightful regarding the nature and importance of what I’ve been calling “characterization.” Crucially, “strategic perspectivalism” is *not* a philosophical view in the same sense as perspectivalism is: it isn’t a claim about the nature of knowledge generally that a philosopher might adopt. Instead, it is a philosophical account of the strategies scientists adopt to overcome challenges from loss or putative loss.

Michaela Massimi identifies two kinds of perspectivalist representation (Massimi 2022). The distinction she draws is useful for us because it captures both perspectives in the sense of characterization (which denizens and properties we target) and in the sense of recontextualization and pluripotency (what can be done with our artifacts). Drawing from pictorial representation in art, she distinguishes between *perspectival₁*—“representation drawn from a specific vantage point [italics removed]” (Massimi 2022, 32)—and *perspectival₂*; a “representation drawn towards one (or more than one) vanishing point(s) [italics removed]” (Massimi 2022, 32). Often, images are intended to be viewed as if from the perspective of an onlooker. Consider Charlotte Kenchington’s reconstruction of the Avalonian Ediacaran from chapter 2 (Figure 2.3). Their image provides a specific sense of where the viewer is located in space: some rangeomorphs are nearby, and others disappear into the

distant gloom. Images having an implied onlooker who has a specific spatial position are an example of perspectival₁. For Massimi, such perspectival features are relevant when the “content of the representation is affected by the very vantage point from which the representation takes place” (Massimi 2022, 34).

Perspectival₂, by contrast, is less concerned with specific vantage points themselves, and more with how representations open “up a ‘window on reality’ that extends well beyond the boundaries of the representation itself and where the depth, angle, and scale of the representation leave enough room to make inferences about the space and what’s in it” (Massimi 2022, 39). Kenchington’s perspectival effect is achieved by representing the rangeomorphs as if they were illuminated: a specific light source illuminates the closest rangeomorphs, and those farther away get progressively dimmer, providing the illusion of an expansive, somewhat ominous space. The effect generated creates a *vanishing point*, much like carefully placed lines can give the impression of a landscape’s horizon. Enough information, and information directed toward a vanishing point, underpins our ability to make inferences about things *not* explicitly represented in the image. Adopting a perspective sanctions inferences beyond the explicit content of the representation. Kenchington’s image allows us to imagine further rangeomorphs, not represented, beyond our light source. This is perspectival₂.

Let’s briefly run these concepts through our burrowing dinosaur case. We conceive of the specimen as vertebrate bones deposited in an infilled burrow. The experimental apparatus instantiates the vantage point analogously to how the arrangements of lines in a pictorial representation create a perspective. The aspects of the target that are explored turn on the bespoke, particular features of that experiment: its material composition (PVC pipe, dirt, rabbit bones) and what the scientists choose to do in their experiments (varying pipe position, initial bone position, and so on). That is perspectival₁—how the denizen is characterized. However, the experiment also affords inferences beyond what it explicitly instantiated. Through recontextualization, it helps generate an argument that the specimens were preserved inside the burrow. Through pluripotency, it acts as a model for burrow preservation more generally and links the particular specimen to other cases. This is perspectival₂.

Massimi takes there to be a lesson here for philosophers interested in scientific representation: “Philosophical discussions of perspectival representations have to start by acknowledging the role of the human agent not just as a spectator, but first and foremost as an architect in arranging lines of

composition and creating in painstaking detail the perspectival effects of a ‘window on reality’” (Massimi 2022, 32).

The crucial point to note at this stage is that the content of scientific representations is affected by decisions scientists make regarding what to include, how to contextualize, and what to omit. Scientific agency, in deciding how to characterize and with what apparatus, is a critical component of our story. In the language I’ve developed thus far, scientific representations require characterization. And in artificing, this characterization and the construction of an object instantiating those relevant properties allow inferences beyond the object itself. Cuthill and Conway Morris’ geometric model, when recontextualized, enabled inferences about rangeomorph phylogeny, for instance.

In characterization, scientists represent their targets of inquiry *as* particular things. *Representation-as* has a fair bit of currency in current philosophy of science and related discussions (Elgin 2017; Frigg and Nguyen 2017). Kenchington represents the Avalonian Ediacaran *as* a dark, benthic, rangeomorph-haunted ecosystem lit by a highly unlikely light source. They do this by incorporating various techniques—such as the progressively dimmer rangeomorphs. In virtue of this (and their labelling it as such), we take it as an attempted representation *of* the Avalonian Ediacaran. Similarly, Alex Liu conceives of the Avalonian as a set of preservation conditions, Varricchio et al. characterize their specimen as bones deposited in an infilled burrow, and Cuthill and Conway Morris take rangeomorphs to grow via a geometric pattern.

Perspectival₁—the effect of a representation being from a particular place or observer—and perspectival₂, the generation of vanishing points enabling inferences beyond what is explicitly represented, are both achieved using what I’ll call *perspectival tools*: the representational strategies, technologies, and techniques that enable a particular perspective. A perspectival tool is something that enables a scientist to instantiate, represent, or explore the denizens characterized as having such-and-such properties. Liu and Varricchio et al.’s experimental apparatus, and Cuthill and Conway Morris’ geometric models, are all examples of perspectival tools, and we’ll meet many more in the next two chapters. Suffice it to say, “perspectival tool” should be understood in terms of epistemic function: it enables characterization.

I’ve emphasized how, in artificing, historical scientists adopt very particular characterizations of target denizens. In brief, the denizens in question are understood at a level of grain that, all going well, enables them to be enfolded into a particular kind of historical individual. Varricchio et al. characterize

their specimen as *vertebrate bones preserved in a burrow*, Cuthill and Conway Morris characterize their specimen as (a particular set of) *fractal morphologies*. This is perspective₁; they represent them from the perspective of objects with those kinds of properties. This perspective enables further connections and inferences beyond what is explicitly represented in their model; thus, it is perspectival₂. How they do this—how the effect is achieved—is via the application of perspectival tools. For instance, Cuthill and Conway Morris’ fractal growth model is a perspectival tool that enables their taking that perspective, and the fossils they examine are also perspectival tools. In artificing, this is achieved by particular properties being *isolated* and *built*—represented in some way. For example, by building a scale model of a burrow and adding disarticulated rabbit bones, or constructing a geometric model with fractal growth patterns.

The perspective—the characterization—of the denizen, then, is achieved via strategies of isolation, emphasis, distortion, and so on. These strategies themselves play out in the construction and application of specific scientific objects and techniques.

Massimi emphasizes that different perspectives take place against a “horizon of alternatives,” that is, in light of a set of different possible perspectives that could have been (and might still be) taken. As she puts it, “Many models for the same target system represent selected features of the target system *as y*, or *as j*, or *as k*, where *y*, *j*, and *k* are different properties belonging to a ‘horizon of alternatives’ and often enough incompatible or inconsistent with one another” (Massimi 2022, 40). For instance, Laflamme et al. represent rangeomorphs in terms of size and surface-to-area ratio given a fractal geometric pattern (thus enabling comparison with single-celled osmotrophs), while Cuthill and Conway Morris represent rangeomorphs through a basic set of rules from growing fractal geometric patterns, variations of which generate various rangeomorph forms (thus enabling comparison between rangeomorphs). These, and various other options, such as representing rangeomorphs as ecotypes, as Mitchell et al. do, constitute the horizon of alternatives (Massimi draws the idea from van Fraassen 2008).

What sets the horizon of alternatives? Massimi is interested in how perspectives enable particular forms of inference, inferences that are not necessarily tied to that particular perspective. While this is a crucial aspect of my explanation of how strategic perspectivalism is a good strategy in the face of loss or putative loss, she hasn’t much to say about what the properties of these horizons are. I want to emphasize how the horizon of alternatives

emerges from the available perspectival tools. Scientists cannot meaningfully or productively take up just any old characterization of a denizen. They adopt characterizations that they have the tools to underwrite, or plausibly might develop tools to do so. As such, the horizon of alternatives is constrained and enabled by the available perspectival tools. As Christophe Malaterre has said of the characterization (or “granularity”) of traces: “The instruments and background theories that are available to research teams at certain points of time certainly influence the granularity at which evidence and statements are formulated” (Malaterre 2024, 23). I think this holds more generally: how scientists might describe or characterize their targets is more or less a product of the perspectival tools at hand.

What makes this perspectivalism *strategic*? First off, the phenomena I’m interested in have a lot in common with those that animate perspectivalists. Drawing on van Fraassen’s (2008) discussion, Catherine Elgin (2017) identifies the properties of a representation that generate perspectives. For her, representations are perspectival when “They can show how occupants of that [logical or physical] space appear from a certain vantage point” (Elgin 2017, 155). Further, they are *indexical*, that is, “They represent how things look from *here* (for some value of ‘here’)” (Elgin 2017, 153). I’m not exactly certain what this means (see below), but I take the indexical nature of denizens, as well as the selection and construction of perspectival tools, to be very close if not identical to the phenomena Elgin discusses. Perspectival representations are at a level of grain, and their effectiveness often turns on getting the grain right (see, for example, Borg’s recent discussion of “level-switching” in geochronology; Borg 2025). We saw Cuthill and Conway Morris’ and Laflamme et al.’s models switching grains from individual-morphology to community-level as the questions they asked switched from individual optimization to group optimization in an environment. Finally, perspectival models often involve “occlusion”: that is, capturing some element of a system often precludes including others. In characterizing a past world, selecting a bunch of denizens, and bringing a set of perspectival tools to bear on them, historical scientists inevitably occlude some features they might otherwise have included.

So, the phenomena I’m interested in are close enough to be legitimately labelled “perspectival.” However, as mentioned above, my philosophical target diverges from traditional perspectivalists. This is partly due to being interested in different questions: perspectivalists are often in the business of trying to reconcile how the perspectival nature of representation can be married to views on the reality or otherwise of the world represented, and of arguing

that knowledge (or representation) is necessarily perspectival. Depending on mood, I'm either asking a different but potentially complementary question, or I think that debates concerning realism and anti-realism have become an unproductive research direction in the philosophy of science (Currie 2018a, Postscript). Regardless, our job here is to understand the strategies historical scientists adopt to grapple with lost worlds, and an instrumentalist (or strategic!) conception of perspectivalism is the relevant one.

So, in part, putting "strategic" in front of "perspectival" functions to sidestep some substantive issues. For instance, are *all* scientific models perspectival? Massimi appears to think not, as she identifies a particular class of scientific models as being perspectival (thus implying that some are not). Meanwhile, van Fraassen argues that because of the perspectival nature of the data-generating processes which go into building scientific representations, all are perspectival. For another example of an issue hopefully sidestepped, try as I might, I'm never quite sure what perspectivalists mean when they refer to indexicality: that there are non-perspectival models where "One need not locate oneself in the space of representation to understand what and how it represents" (Elgin 2017, 156), whereas other models do involve somehow situating *yourself* in the representation. I suspect the right way to understand this question is whether characterization, that is, identifying a set of denizens, is sufficient on its own to generate a perspective. I suspect not, as other features of scientific interest—the kinds of questions they ask and the kinds of answers they expect—might also be included in a perspective.

So, given my interests, aims, and perhaps philosophical failings,³ I'm hoping to not take a stance on some of these questions here.

I also suspect there are more bona fide, first-order philosophical disagreements between myself and perspectivalists, especially in terms of the nature and importance of representation. In chapter 6, I'll argue for *artifactualism* about perspectival tools, which doesn't take scientific tools as primarily representational, but instead as material objects that scientists build and interact with. This won't involve me denying the legitimacy of representations (see, in

3 There's an apocryphal story that I've been unable to find a source for that goes something like this. Bertrand Russell is giving a research seminar, and during question time an interlocutor has come back several times claiming they just don't *understand* what he is getting at (philosophers should well recognize this form of "misunderstanding as somehow an objection" seminar question), to which Russell eventually replied, "Your philosophical failings are not my problem!" I suspect on this count that perspectivalists perhaps ought respond to me with the same ...

particular, chapter 7's discussion of palaeoart), but it does involve me folding representational tools into a broader category of "artifact."

So, my view is perspectival-adjacent and certainly consistent with perspectivalism (at least thus far), but what makes it *strategic*?

In approaching a new piece, an artist might ask themselves, "What techniques and tools do I have at my disposal? What effects can I achieve with these?" In other words, what kinds of affordances are there to generate different perspectives₁—what are the available horizons of alternatives—and which would be the most effective given my aims? Similarly, historical scientists, when artificing, consider what techniques and technologies are available to enable them to construct a perspective. And they pay particular attention to what other objects these perspectives enable connections between. Laflamme et al.'s perspective enables comparison with single-celled osmotrophs and the incorporation of the geometric/metabolic logic of body size to surface area ratio; Cuthill and Conway Morris' perspective enables comparisons between various rangeomorphs, the incorporation of fluid dynamical representations of waterflow, and models of ecological competition.

Note that I've not yet said anything directly about what makes an instance of artificing good, successful, or justified. But we've had a few hints in my discussions of recontextualization and pluripotency, and Massimi's notion of perspectival₂. How much grip do perspectival tools give us on the case at hand? What links do they provide to other cases, and what further inferences do they enable? As we've seen, historical scientists pay careful attention to how their artifacts and other perspectival tools might diverge from their lost target. Liu's pizza-disk experiments, for instance, made use of contemporary biota, which likely differ in many ways from those composing the Ediacaran mats. Their artificing relies on the idea that similar metabolic processes are held between them despite those differences. In the next chapter, I'll more explicitly discuss the kinds of overlap we might have between perspectives, and in chapter 5, I'll describe the iterative relationships between them. So, if you can, put a pin in your scepticism for now.

In artificing (and, as we'll see, in grafting and anchoring), historical scientists adopt or construct perspectival tools that enable fruitful characterizations of past denizens. It is in the ingenious use of strategic perspectivalism, I'll argue, that historical scientists are able to recover lost worlds.

4. Coda: De-Extinction

“We know a lot about Broadfoot”, Prof. Ichthy continues, a faraway look drifting across their face, “but there are still so many unanswered questions. What were their mating habits? Were they territorial? Were they fully or semi-aquatic? Did they raise their young? Were they solitary or did they live in groups? Without a time machine, there just doesn’t seem to be any way of finding out. Unless . . .” Ichthy’s eyes come into sharp focus as they slam a flipper emphatically on the lectern and lean forwards eagerly “. . . Unless we can *bring them back!*”

Artificing can be understood as the construction of a lost world. This might sound metaphysically dramatic: are you literally bringing a past world back into existence? Is the constructed simulacrum the same thing or kind of thing as the lost denizen? Recall De la Beche’s poking fun at Lyell, illustrating the absurdity of thinking that, were Jurassic conditions to return, so too would Jurassic fauna. We might think this absurd causally (the chances of the same fauna arising just because we have similar environments are tiny!), but also metaphysically: how is the Jurassic fauna and the new fauna *the same*, exactly? Does artificing itself also involve these apparent causal and metaphysical absurdities? This question is a nice opportunity to discuss the notion of “de-extinction.”

As we’ll see, artificing needn’t involve anything so dramatic as implied by the above questions. Because denizens are sensitive to description—in particular, to the kind they are characterized as falling under—and because artificing involves strategic perspectivalism, “bringing back” a lost world is just a way of saying that something of *that kind* instantiating *those properties* has been fabricated. This point, it seems to me, sheds some light on discussions of de-extinction, a set of practices sold as a way of “bringing back” the once-lost denizens of past worlds, specifically, extinct species. My aim in this subsection is to introduce and critique the notion of de-extinction in light of my discussion of artificing. In doing so, I’ll articulate a few important features of the notion of loss, in particular, how it can foreground value (as we’ll return to in chapter 10) by, in a sense, putting the resolution of metaphysical questions partly in our own hands.

I’ll understand de-extinction as an attempt to, in some sense, bring extinct animals back from extinction. Sometimes the term is used in tandem with notions like rewilding, which is to return an ecosystem to something resembling a past state. Further, sometimes animals are considered extinct

if they only live in captivity, and so attempts to return them to the wild are a kind of de-extinction. Here, I'll stick to "de-extinction" as a set of projects aiming to reconstruct (artifice!) a lineage known only from traces.

Calls for de-extinction are often made against highly moralized or romanticized backdrops: if we could bring back the victims of anthropogenic extinction, say, the passenger pigeon or thylacine, would this not be some kind of act of atonement for our ecological sins? Would re-populating the Arctic Circle with mammoths return us to some pre-human balance of nature (Seddon 2017; Odenbaugh 2023)? But on the other hand, don't such efforts risk overplaying our domination of nature or further undermining genuine nature with artificiality (Lindquist 2020; Katz 2022)? Such normative considerations need to be understood against the empirical reality of de-extinction.

Mechanically, there are at least two paths to de-extinction. First, back-breeding, which makes use of artificial selection to walk a lineage back along its genetic and phenotypic history (most typically associated with aurochs, the recently extinct wild European cattle; Stockstad 2015). Second, cloning, which involves artificially inseminating an animal with an artificially-grown zygote on the basis of recovered genetic material (perhaps most well-known for the mammoth; Shapiro 2015). These are extremely limited from both a technological and a practical perspective. Cloning, for instance, depends on the availability of ancient DNA. So, anything more than a few million years old is highly unlikely (despite the important influence such ideas have had over both public perceptions of palaeontology and dinosaur science generally; see Jones 2022), and this is to say nothing of the enormous difficulties of bringing unusual zygotes to term in either a different species or an artificial womb/pouch/egg. Back-breeding has a somewhat less science-fictional tone, but also involves tricky assumptions about the reversibility of genetic pathways, is extremely labour-intensive, and is extremely limited. As Thomas Reydon (2022) has put it:

The scientific and technological reality is that de-extinction in this literal sense is fundamentally impossible. What is possible is "de-extinction" in a scientifically (and technologically) informed sense. But this merely involves the creation of organisms that are phenotypically similar in a limited number of traits to those of an extinct species, and/or that can perform the same ecological role as an extinct species. (Reydon 2022, 114)

Alongside scientific and technical worries, philosophers and biologists have also pointed out conceptual and metaphysical issues. What does it mean to *bring back* a species, especially if it is understood as a historical individual? Some might argue that no matter how similar a cloned mammoth is to a true mammoth, as it forms a clade with (Asian) elephants (its mother is an elephant, after all), it is at best an odd elephant, not a mammoth. Others question whether a mammoth can truly be a mammoth if it is not situated in a mammoth-like environment—is there a mammoth without the Great Steppe? Philosophers have had several field days applying various species concepts to de-extinction (see, for instance, Turner 2017; Siipi and Finkelman 2017; Finkelman 2018; Lean 2020). And the general lesson is, well, it depends on what you mean: what notion of species you have in mind or—to put it more generally—which characterization of the denizen interests you. If species are historical individuals, and that individuality requires spatio-temporal continuity, then you can't bring them back. But that's reliant on the characterization and, as we've seen, others are available.

The conceptual machinery from the last chapter gives us some traction on these kinds of questions. If you want to artifice a denizen, you need to ask yourself how you want to characterize it: what properties are you interested in, and what “kind of kind” (historical individual versus historical type, etc.) do you have in mind? Further, consider strategic perspectivalism. De-extinction efforts and speculation are underwritten by our technological capacities: cloning, back-breeding, and so on. Our characterization (that is, perspective,) is in part generated by what we can successfully build and represent, that is, our perspectival tools. So, we can't just characterize the denizens in any which way we want in this context.

Certain kinds of de-extinction are plainly beyond our technological—and perhaps conceptual—reach. If we used DNA to clone a passenger pigeon via (say) a rock-dove intermediary, we do not thereby generate something of the same lineage as passenger pigeons. But, if successful, we have brought *something* back: there will likely be some properties that were once lost and now are not. It is up to us to decide if we think it counts as a passenger pigeon according to what is morally, or environmentally, or epistemically, or aesthetically significant. Potentially, there is much to learn from de-extinction projects about the reversibility of genetic pathways, the capacities of cloning technologies, as well as extinct animals' morphology, behaviour, and so on. Artificing is a very productive strategy, after all.

But one of the lessons of thinking about lost worlds is that, in many senses, you can't go back: some elements of the past can be rebuilt, but some cannot. And the question is, which aspects are we actually interested in, what gains might we get from pursuing that path, and what costs are there? Given the enormous expense and very real animal suffering (Browning 2018; Odenbaugh 2023) that goes into developing and deploying de-extinction technologies, it is worth our time to ask if we can't learn these things through less financially and ethically costly means, and indeed whether they are worth learning at all. Artificing a lost denizen requires choosing a perspective, taking a stance on its characterization. There is no characterization-independent answer, then, to whether de-extinction is successful or should be attempted.

How, then, should we answer the metaphysical questions underlying this discussion: does artificing literally involve resurrecting a lost world—when successful, do we have a lost world, found? Excitingly, sometimes yes! However, less excitingly, the answer is, well, it depends on how you characterize that world and its denizens.

There's a further epistemic corollary to note. If you attempt some kind of de-extinction, how will you know you've succeeded? Given that we don't know many of the properties of the extinct organism, how can you be sure that your artificed creature has the same properties as the past denizen? The possibility of uniformity failures is one of the major challenges of thinking about loss—not whether loss has occurred, but just that it might have. If we resurrected a mammoth or a passenger pigeon, how would we be sure it behaved as its extinct brethren did?

In artificing, you tackle loss—a lack of “Ns”—by constructing some Ns. But how can you know that they, in fact, are Ns? The answer to this question will emerge as we examine further strategies in the next two chapters.

“Or, maybe not . . .” Ichthy says, deflating with a guilty look at their grad students.