Synthese (2022) 200:354
https://doi.org/10.1007/s11229-022-03792-x

ORIGINAL RESEARCH

Explaining ambiguity in scientific language

Beckett Sterner[1]

Received: 22 October 2021 / Accepted: 23 June 2022

Abstract
The idea that ambiguity can be productive in data science remains controversial. Efforts to make scientific publications and data intelligible to computers generally assume that accommodating multiple meanings for words, known as polysemy, undermines reasoning and communication. This assumption has nonetheless been contested by historians, philosophers, and social scientists, who have applied qualitative research methods to demonstrate the generative and strategic value of polysemy. Recent quantitative results from linguistics have also shown how polysemy can actually improve the efficiency of human communication. I present a new conceptual typology based on a synthesis of prior research about the aims, norms, and circumstances under which polysemy arises and is evaluated. The typology supports a contextual pluralist view of polysemy's value for scientific research practices: polysemy does both substantial positive and negative work in science, but its utility is context-sensitive in ways that are often overlooked by the norms people have formulated to regulate its use, including prior scholars researching polysemy. I also propose that historical patterns in the use of partial synonyms, i.e. terms with overlapping meanings, provide an especially promising phenomenon for integrative research addressing these issues.

Keywords Ambiguity · Polysemy · Lexical semantics · Pragmatics · Computer ontologies · Philosophy of science · Strategic ambiguity

1 Introduction

The legitimacy and value of ambiguity in scientific language is a central issue for data sharing and reuse (Ribes & Bowker, 2009; Leonelli, 2016; Garnett et al., 2020). Advocates for open data, for example, aim to improve the efficiency, transparency, and equitability of science by making data publicly available for reuse by potentially anyone (Ali-Khan et al., 2018). Similarly, leaders of the FAIR data principles

✉ Beckett Sterner
bsterne1@asu.edu

[1] School of Life Sciences, Arizona State University, PO Box 873301, Tempe, AZ 85287, USA

Published online: 19 August 2022

(Findable, Accessible, Interoperable, Reusable) seek to develop new standards that advance the ability of machine agents to "know what the data mean" (Wilkinson et al., 2016; Mons et al., 2019). Understanding what data mean depends on having access to and the ability to interpret careful descriptions of how the data were produced and processed, but given the dynamic nature of scientific knowledge and research, it is difficult to anticipate what information will be needed by possible future uses and users in this respect (Leonelli, 2016; Leonelli & Tempini, 2020). Even basic information about an observation is subject to major transformations, such as in ecology where scientists seeking long-term data have to reconcile contemporary latitude-longitude measurements of location with historical landmarks used in observations a century ago (Shavit & Griesemer, 2011). In response, scientists in many disciplines have adopted the strategy of standardizing and formalizing the language they use to communicate and reason about their research, for example by using computer ontologies (Leonelli, 2012; Poirier, 2019; Sterner et al., 2020b; Lean, 2021).

A common assumption behind these standardization efforts is that allowing technical terms to have multiple, conflicting meanings undermines scientific reasoning and communication (Sterner et al., 2020b). Controlled biomedical vocabularies in the 1990s, for example, regularly equivocated between treating terms as referring to things in the world such as diseases or anatomical parts versus the conceptions researchers had of those diseases and anatomical parts (Ceusters et al., 2005). In biodiversity science, the ongoing flux of taxonomic revisions to lists of accepted biological species is a frequent topic of complaint for researchers or decision-makers that want a single, stable set of observations for each species (Franz & Sterner, 2018). While human experts are frequently able to interpret ambiguous language used by authors in specific publications, this skill remains a major research challenge for natural language processing in computer science (Loureiro et al., 2021).

However, the claim that ambiguity is generally undesirable is contradicted by results from a variety of disciplines. Important examples include: the cognitive value of metaphors and analogies for generating new possible meanings or concepts (Gross, 2006; Higuera, 2018; Perrault & O'Keefe, 2019; Swedberg, 2020); the strategic value of ambiguity for enabling collective action among heterogeneous actors (Eisenberg, 1984; McMahan & Evans, 2018; Altomonte, 2020); and the increased efficiency of using short, simple words to signify multiple meanings relative to maintaining a larger and more complex vocabulary (Piantadosi et al., 2012).

Many scholars also argue that ambiguity has positive value for integrative research and transdisciplinary projects connecting academic and non-academic partners to address societal problems (Ferraro et al., 2015; Winkler, 2015; Neto, 2020). Picking an arbitrary convention is generally an insufficient basis for scientists to converge on a single meaning for each term, and they typically seek stronger justifications based on definitions offering a uniquely natural or epistemically reliable way of describing phenomena (Sterner et al., 2020b). "Wicked problems" such as climate change or biodiversity loss, though, are characterized precisely by their inability to be represented using a single, objective perspective (Rittel & Webber, 1973): by definition, wicked problems preclude a clean, universal separation of factual and normative assumptions, and they don't allow scientists to stabilize the problematic situation for long-term study independent of ongoing efforts to address it (Ferraro et al., 2015; DeFries & Nagendra, 2017). For wicked problems, “one cannot meaningfully search for information without the orientation of a solution concept; one cannot first understand, then solve” (Rittel & Webber, 1973), p. 162. If accommodating ambiguous language is inherent to addressing wicked problems, then approaches to data science based on eliminating ambiguity may fail to deliver on science’s value for society.

The contested value of ambiguity therefore has important implications for how scientists should engineer their data standards and infrastructure (Bowker & Star, 1999; Bowker, 2000). While some initiatives such as the Open Biological Ontology (OBO) Foundry aim for all scientists in a domain to agree on a single set of terms with univocal meanings, there may be circumstances under which it is desirable for data science to facilitate the use of metaphorical or strategically ambiguous language, e.g. "adaptive radiation" and "biodiversity" (Olson et al., 2019; Takacs, 1996), which serve as boundary objects for interdisciplinary research and collective action (Star & Griesemer, 1989). Similarly, pursuing consensus definitions for terms may accelerate or hinder progress in a subject depending on whether their usage is sufficiently stable and precise to avoid the overhead costs required to coordinate and translate between alternative definitions (Sterner et al., 2020a, b).

Looking beyond data science, the existence and frequency of ambiguity also has fundamental significance for the aims and nature of human language (Gross, 2006). Noam Chomsky, for example, famously concluded that the frequent semantic ambiguity of sentences in languages such as English showed that their primary function could not be communication (Piantadosi et al., 2012). Other scholars, such as George Lakoff and Mark Johnson, have argued that ambiguity in the form of metaphor is pervasive and at the root of all human cognition, including scientific reasoning (Lakoff & Johnson, 2008). This stands in sharp contrast to the common assumption that language predominantly has a precise, literal meaning in scientific and legal discourse (Poesio, 2020).

Despite the practical and theoretical importance of ambiguity, we lack a conceptual framework that can guide interdisciplinary synthesis and empirical research on its properties and value for science. I focus on a particular form of ambiguity, known as polysemy or lexical ambiguity, which applies to words with multiple meanings (Poesio, 2020). (The term "ambiguity" itself has been used to describe different types of things, including words, actions, and situations, and with respect to different aspects, such as their uncertainty, vagueness, or multiple possible interpretations (Poesio, 2020; Sennet, 2021). Polysemy is especially relevant to data reuse when it occurs in technical terms important for describing the contents and significance of empirical results, such as with "species" in systematic biology, "function" for genetic sequences, or "interactions" in ecology (Nakazawa, 2020; Linquist et al., 2020; Stankowski & Ravinet, 2021). I should also note that many lexicographers distinguish polysemy from homonymy (Poesio, 2020), which occurs when the same lexical unit (i.e. string of characters) has unrelated meanings. A classic example is "bank," which many dictionaries treat as two distinct words because the meanings of *financial institution* and *sloping land next to a river* are unrelated and can in fact be traced to different etymologies. However, this distinction is not universally followed within linguistics, and the proposed criteria for distinguishing polysemy and homonymy are contested and difficult to standardize even in the context of dictionary making–for more discussion, see (Poesio, 2020). I will therefore use polysemy in the broad sense that includes homonyms.

There is a large, interdisciplinary body of evidence to consider, including research from History and Philosophy of Science (HPS), Science and Technology Studies (STS), linguistics, computer science, and organizational studies. The research in each these fields has generally proceeded in isolation and varies substantially in scale, type of study object, and methods used. Recent research in corpus linguistics, for example, has examined patterns of polysemy across whole languages such as English and German (Piantadosi et al., 2012), while research in organizational studies has considered the fate of a single term disciplinary term such as "organizational effectiveness" over several decades (Hirsch & Levin, 1999). Case studies in HPS have examined the rhetoric of a small number of texts by individual authors (Ceccarelli, 2001) or synthesized results from prior conceptual analyses (Neto, 2020). There is also a rapidly growing community in computer science that uses text corpora to study lexical semantic change (Tahmasebi et al., 2021) as well as algorithms for word sense disambiguation and induction (Li & Joanisse, 2021).

A general framework would therefore be valuable in several ways. First, similar questions about ambiguity have and continue to arise in parallel across scientific fields, and a general way of categorizing types of ambiguity would help reduce fragmentation of researchers and literature in each field. Second, different disciplines studying ambiguity have evaluated it from the perspective of different aims and uses, e.g. focusing on individual communicative accuracy, efficient reuse of words across whole languages, and pragmatic effectiveness at coordinating joint action among diverse actors. Third, a general framework can serve to show how insights can be combined across varying scales of systems and research methods, for example using careful analysis of concepts from a small sample of texts to inform statistical analyses of large corpora representing whole fields or languages. I will argue that historical patterns in the use of terms with overlapping meanings—what I call "partial synonyms"—provide an especially promising phenomenon for integrative research addressing these issues.

In the remainder of the paper, I present a new conceptual typology in Sect. 2 based on a synthesis of prior research about the aims, norms, and circumstances under which polysemy arises. I show in Sect. 3 how the typology supports a contextual pluralist view of polysemy's value for scientific research practices: polysemy does both substantial positive and negative work in science, but its value is context-sensitive in ways that are often overlooked by the norms people have formulated to regulate its use, including prior scholars researching polysemy. I also argue that the typology clarifies how actors' norms about the proper use of polysemy can explain empirical patterns of usage. In Sect. 4, I illustrate how one can use partial synonyms to derive predictions about expected usage patterns given specified conditions under which polysemy has positive value. Section 5 then illustrates how partial synonyms can also serve to integrate several promising computational approaches to analyzing text corpora.

2 Distinguishing types of polysemy by their pragmatics

This section introduces a general framework for distinguishing instances of polysemy into different types reflecting the pragmatics of the speaker's and audience's aims and capacities. I introduce the distinctions using a sequence of example sentences. Note that the pragmatic approach I take treats ambiguity as relative to the context of speaker and audience rather than inherent to the text itself. I build on the framework introduced by Piantadosi et al. (2012) to accommodate additional reasons for valuing polysemy beyond efficiency.

Let's start with some simple example sentences using the polysemic terms "function" and "biodiversity."

Ex. 1: "Mutations with the function of improving lactose tolerance have evolved by natural selection multiple times in humans."

Before addressing how polysemy is resolved in this example, we need to make sure to distinguish the general term "function" from its token instances. For our purposes, let's assume the *term* "function" has two possible meanings: an evolutionary sense of function that requires a trait of organism to have a history of positive natural selection based on its positive effects on fitness, and a causal role sense of function that requires a part of an organism to make a causal contribution to a capacity of the organism as a whole (Garson, 2016). Any instance of the term "function" may therefore be interpreted as having one of these meanings or potentially both if the speaker's intended meaning is uncertain. (If the listener judges that neither meaning applies, then they may judge the word to be used incorrectly or in a novel way that requires constructing a new meaning.)

Assuming that the reader of the sentence above knows about both of these possible meanings, they must consider which if any is intended in this instance. Following Piantadosi et al. (2012), we can represent this situation in information theoretic terms: the listener has received a signal, the sentence shown above, from the speaker, and now must decode its meaning; see also McMahan and Evans (2018). Because there are two possible meanings of the token word "function," the reader faces potential uncertainty in deciding which one is intended (or if any of them are legitimate). The rest of the sentence provides contextual information by reducing the reader's uncertainty. I will refer to contextual information provided by the surrounding text as "linguistic context," in contrast with "social context" that involves information about the speaker and social setting of the text. In general, I take contextual information to refer holistically and inclusively to all forms of information that actors may use for interpretation.

In the example above, uncertainty is eliminated by the words "improving" and "have evolved through natural selection." However, in other cases the context may fail to be wholly decisive or even totally uninformative. Consider, for example:

Ex 2: "Most of the human genome is not functional"

Whether an evolutionary or causal role meaning is intended is impossible to determine from the sentence alone. Indeed, this confusion is at the root of the recent debate over the proper interpretation of the ENCODE project's analysis of molecular activities associated with stretches of the human genome. Some have argued the ENCODE study conflated the two meanings of function in rejecting claims that most of the human genome is "junk DNA" without assessing whether any observed molecular activities have a history of natural selection (Germain et al., 2014).

Another example highlights how differences in background knowledge, i.e. prior knowledge about the subject being discussed, can also figure into uncertainty:

Ex. 3: "The database compiles pathogens and parasites from wild carnivores obtained from the scientific literature"

I adapted this example from the Global Mammal Parasite Database website, where I myself was confused about why their "carnivore" dataset didn't include any rodent species. A colleague pointed out the website meant carnivores in the phylogenetic sense (i.e. the clade Carnivora), not the ecological sense, which would be evident if you already know which species are in this group of mammals and had previously familiarized yourself with the database contents.

Using the ideas presented so far, we can now articulate several reasons for positively or negatively evaluating polysemy.

1. One of the possible meanings is intended but the reader perceives *insufficient* contextual information to identify it.
2. One of the possible meanings is intended *and* the reader perceives *sufficient* contextual information to identify it.

In the first option, failure to identify the intended meaning could be due to either the absence of cues from the author or the reader's inability to use them effectively, for example due to a gap in background knowledge. This option is clearly undesirable if the author and reader share the aim of accurate communication.

The second option is intriguing because it illustrates how polysemy is conditionally consistent with the aim of accurate communication. In fact, we will see how polysemy can actually be more efficient for accurate communication than having a different word for every possible meaning.

The possibility of uncertainty also raises a third scenario:

3. Neither of the possible meanings is specifically intended but the reader perceives sufficient contextual information to select one meaning as intended.

We can understand this third option as a case where the author is being intentionally ambiguous. Many terms at the center of interdisciplinary research efforts have this quality, such as "genetic engineering" or "biodiversity." This is not necessarily in bad faith on the author's part, as it may be a strategic way to appeal to multiple audiences who bring different background knowledge and values. The author's intended pragmatic effect may be to benefit both audiences even as (or precisely because) they interpret it differently. For example:

Ex. 4: "The costs of declining biodiversity for human well-being are frequently overlooked in quantifying the effects of climate change."

This sentence on its own does not provide a basis for disambiguating whether biodiversity is intended in a species, genetic, or phylogenetic sense. Nonetheless, the sentence serves as a call for the reader to think about how biodiversity should be incorporated into their thinking about climate change, and it is consistent with this strategic aim for different readers to interpret the nature of biodiversity in differing ways.

Of course, people don't always use strategic ambiguity for mutual benefit, and there are other cases of polysemy where the author provides misleading cues relative to what they take to be the truth, either as a failure of communication or intentional deception. Part of the controversy about ENCODE, for example, was whether the presentation of their results had used "function" (as well as "junk" in "junk DNA") in a confusing way to exaggerate the project's scientific achievements (Germain et al., 2014).

A fourth option also emerges if we recall that scientists are also in the business of revising their language over time to accommodate new developments in research. Metaphor is an important pathway to semantic innovation in this respect: by using a familiar term in a new context that is not straightforwardly consistent with any of the term's existing meanings, the author engages the reader in creating a new meaning.

4. No existing meaning is intended but the reader perceives sufficient contextual information to suggest a novel meaning intended by the author.

Context is doing extra work in this scenario since the reader must perceive that no existing meanings are warranted and that the author intends a positive, new meaning instead of simply having used the word incorrectly. Take the example of "living fossil," which historically has referred to species or clades that exhibit relatively little evolutionary change since their origins. It has arguably been a fruitful semantic innovation to call a single gene a "living fossil" (Lidgard & Love, 2018), but it makes little sense if we try to apply the term to a soap bubble.

Figure 1 summarizes the key distinctions between these four types of cases at a general level. While I've presented these in terms of analyzing single instances of a term, we can also use these categories to generalize about how terms are being used across contexts. For example, many scientists use Linnaean names for biological species without indicating which taxonomic usage they are following (if they aware at all that species names' frequently change meanings). The same term may therefore occur in many contexts and prove problematically ambiguous (case 1) most of the time (Monckton et al., 2020). Alternatively, a term such as "function" may accumulate a small number of stable meanings that biologists learn to disambiguate reliably (case 2) through linguistic cues, e.g. by adding "evolutionary" or "biochemical" to the term as qualifiers. "Biodiversity" has become an important boundary object between biology and policy in part because it is strategically flexible (case 3), referring to species diversity, genetic diversity, or all living things and accommodating different moral views about why biodiversity is valuable (Takacs, 1996). At the same time, scientists often intend a single technically precise meaning of biodiversity, e.g. number of species present in a region, in their specialized research articles (case 2). The way that I've formulated the idea of novel usages for a term (case 4) also presupposes the term's ongoing use for existing meanings in other contexts. I will discuss below how some terms such as "organizational effectiveness" in the social sciences can go through shifts in how they are used over time as early enthusiasm for a technical term turns into sustained criticism of its ambiguity followed by clarification or abandonment (Hirsch & Levin, 1999; Giroux, 2006).

Nonetheless, there are some important aspects of polysemy I do not address here. I have not tried to characterize the nature of the differences between meanings of a polysemic term, which may alternately be understood as different thresholds applied to a vague concept, contrasting conceptions of a single core meaning, or family

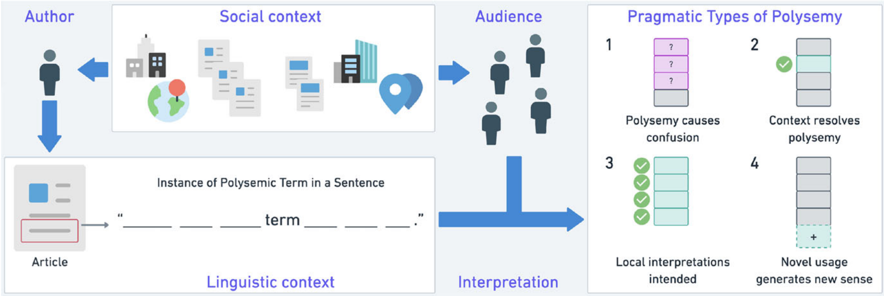

Fig. 1 An author's choice to use a polysemic term—one that has multiple possible meanings—relies on shared context with the audience interpreting the term and can result in one of four distinct pragmatic outcomes. See main text for definitions of linguistic and social context and the four types

resemblance. There are also other uses for polysemy beyond those considered here, especially in humor (Nerlich & Clarke, 2001), and I have not addressed syntactic polysemy, where the same grammatical construction can be construed in multiple ways. (Syntactic polysemy is nicely illustrated by Groucho Marx's joke: "One morning I shot an elephant in my pajamas. How he got into my pajamas I'll never know" (Poesio, 2020)). Nonetheless, having considered four different pragmatic contexts for polysemy, we can turn to consider how they relate to normative prescriptions about the use of polysemic terms.

3 Towards explanation by pragmatic norms

This section introduces contrasting normative principles for scientific terminology and identifies an initial set of empirically accessible conditions for their application. The ideals I introduce may look mutually incompatible as a matter of principle, since one ideal seeks to eliminate ambiguity as undesirable while others embrace it as fruitful in different ways. However, this conclusion would overlook their implementation via methodological prescriptions that conditionalize the costs or benefits of polysemy to specific sets of practical circumstances. The possibility that one ideal may be better suited for a particular situation than the others illustrates the importance of adopting a baseline stance of openness toward the value of ambiguity, i.e. as neither inherently positive or negative for our aims. Critically, each ideal has a characteristic signature that is in principle detectable in scientific discourse. My approach is thus thoroughly pragmatic in the sense of treating normative principles for scientific language as guides for action that are fallible and that may be iteratively revised and improved in light of experience.

Certain purposes for scientific language stand out in the collected literature as especially worthy in light of their broader understanding of its values and aims: communication, reasoning, innovation, and joint action. Philosophers and scientists sometime take these aims to set regulative ideals for scientific language—for instance, that each term should have only a single, literal meaning—and then formulate prescriptive rules to bring these ideals into practice. More broadly, every scientist unavoidably acquires some set of practices about proper language use, even if this amounts to the tacit attitude of "anything goes as long as it doesn't get in my way." Nonetheless, the prescriptive norms that researchers follow provide explanatory reasons for their linguistic behaviors and responses to others' behaviors (Brandom, 2008). An author may use one term instead of another, for example, because the latter option is ambiguous in a way that the audience is likely to misunderstand. Alternatively, an author may select the more ambiguous term precisely because it accommodates divergent understandings without undermining a collaborative task.

In looking for evidence of such norms operating in scientists' linguistic behaviors, I should stress that there are generally also other factors at play. I do not mean to imply, for example, that language use is exceptionally rational or transparent to introspection. I also do not mean to imply that how scientists write or speak is decoupled from broader changes in everyday language and society. Nonetheless, many aspects of any field's terminology and discourse are highly policed by the community, e.g. through mechanisms such as peer-review, graduate training, and comments on conference presentations. The ways that scientists express their personal and social contexts are also often mediated and transformed by the values and histories of their research communities, e.g. so that researchers may respond to skepticism about the public value of science by emphasizing jargon-free writing or applied outcomes. What I am assuming, then, is that scientists form communities with distinctive and enforced practices of language usage that address, among other things, the proper usage of terminology.

Building on the four types of cases in the previous section, I introduce a corresponding set of normative rules. These are intended to formalize and make explicit the often-implicit stances scientists take in their discursive practices. I list them here and then discuss their origins from different domains. I expect these rules will need to be revised and clarified through further conceptual and empirical work, but I believe they provide a useful synthesis of the major reasons for or against polysemy in science documented in the literature.

- Don't use polysemic terms because they lead to errors in communication if enough contextual information doesn't reliably exist.
- Do use polysemic terms because they make communication more efficient if they are linguistically simpler than the alternative and enough contextual information exists.
- Do use polysemic terms because they enable joint action if enough contextual information exists for each listener to select a practically adequate, personal interpretation.
- Do use polysemic terms because they introduce new meanings into discourse if they are used in novel ways and enough contextual information exists for people to generate a personal interpretation similar to the speaker's.

Although the first two rules take contrasting stances on the value of polysemy, they share assumptions in common: they both take for granted that the speaker's aim is to convey a single, determinate, and pre-existing meaning to the audience; both rules also acknowledge the importance of available contextual information. However, the two rules differ over what polysemy has to contribute to communication. In the case of Rule 1, polysemy is merely a source of potential error with no upside. For Rule 2, however, the benefits of reusing comparatively simple terms can outweigh the risks of miscommunication. (Comparatively simple is meant in cognitive linguistic terms, e.g. the number and surprisal of the syllables that form the word; see Piantadosi et al. (2012).) Rule 2 therefore leans heavily on the availability of adequate contextual information for the audience to disambiguate the speaker's intended meaning.

Rule 1's advocacy for univocal rather than polysemic terminology has been endorsed by many scholars. In philosophy, the logical empiricists famously sought to eliminate conceptual confusion and metaphysical nonsense from scientific knowledge in the early twentieth century by reducing the meaning of theoretical concepts and laws to mathematical operations on sets of observation statements. Symbolic logic thus provided a new standard for mathematically precise explications of concepts that promised to eliminate ambiguity from scientific language and justification. This ideal is also central to much of contemporary data science, which uses classificatory theories expressed in description logic in order to integrate and process unprecedented volumes of scientific claims and measurements (Arp et al., 2015). In linguistics, leading theorists such as Paul Grice have promoted the maxim "avoid ambiguity" and view ambiguity as opposed to efficient conversation for the sake of rational cooperation among actors (Grice, 1975). These scholars share the concern that ambiguity is inherently a problem for valid reasoning, whether deductive or inferential, and that the goal of ensuring accurate decoding by the listener is primary.

More recently, cognitive linguists have taken up a new line of research inspired by George Zipf's classical work on scaling laws in human language (Piantadosi et al., 2012). Zipf's best-known law concerns the relationship between the frequency and length of words in a language, but he also proposed another law stating that the number of meanings associated with a word scales with the square root of its frequency. Zipf presented empirical evidence for the first law but derived the second one from theoretical considerations; for a recent analysis, see Ferrer-i Cancho et al. (2020). Using newly available databases such as WordNet, linguists have started to study the second law as an empirical subject and found substantial initial support in multiple languages (Piantadosi et al., 2012; Català et al., 2021). For example, Piantadosi et al. "argue that ambiguity can be understood by the trade-off between two communicative pressures which are inherent to any communicative system: clarity and ease. A clear communication system is one in which the intended meaning can be recovered from the signal with high probability. An easy communication system is one which signals are efficiently produced, communicated, and processed" (Piantadosi et al., 2012, p. 281). They identify two beneficial properties of ambiguity: "first, where context is informative about meaning, unambiguous language is partly redundant with the context and therefore inefficient; and second, ambiguity allows the re-use of words and sounds which are more easily produced or understood." These beneficial properties of ambiguity apply even in the technical settings of scientific communication and reasoning, where collaborative teams can develop a local shorthand and customized understandings for broadly used terms (Sterner & Franz, 2017). Other researchers studying the evolution of folk classifications have proposed that descriptive terms for phenomena such as color or organisms balance a similar tradeoff between informativeness and efficiency (Kemp et al., 2018; Carr et al., 2020).

Accepting a positive value for polysemy therefore does not require one to reject the goal of semantic determinacy, i.e. having only a single intended interpretation of a token term. Recogizing productive ambiguity can simply entail rejecting the ideal of *context-independent* determinacy of meaning.

Turning to consider Rules 2–4 as a set, we can see that they all recognize the importance of adequate contextual information to productive ambiguity, although they identify different ways this positive value comes about. Indeed, the first two rules presuppose a single determinate meaning is intended, while the latter two rules make positive use of indeterminacy relative to the audience. Rules 2–4 are nonetheless consistent with the general norm that the best way to achieve productive ambiguity is by ensuring the reliable availability of sufficient contextual information and shared background knowledge. However, what sorts of context and background knowledge are required will depend on whether a speaker intends to use a term in a semantically multi-valent or novel way. This points brings us to consider how positive value can arise from polysemy when no single meaning can be said to be universally intended by the speaker for the audience.

Indeterminate usage of polysemic terms can be preferable when it facilitates reasoning or communication that would otherwise be obstructed by requiring all actors to identify a single meaning intended by the speaker. Historians and sociologists of science, for example, have shown that ambiguity can scaffold scientific progress by creating and sustaining avenues for communication between scientific communities with only partially overlapping or compatible repertoires of concepts and vocabularies (Star & Griesemer, 1989; Galison, 1996; Gerson, 2008; McMahan & Evans, 2018).

The idea of strategic ambiguity therefore breaks with the common assumption among linguists that the interactions of a speaker and their audience are tightly coupled in a single shared endeavor (Eisenberg, 1984; Davenport & Leitch, 2005; Jarzabkowski et al., 2010; Denis et al., 2011; Johansen, 2018; Altomonte, 2020). Indeed, the strategic value of ambiguity has been formulated primarily in the field of organizational science, where researchers recognized that "people in organizations confront multiple situational requirements, develop multiple and often conflicting goals, and respond with communicative strategies which do not always minimize ambiguity, but may nonetheless be effective" (Eisenberg, 1984, p. 2). It is not necessarily the case, then, that organizations always work better when managers develop ever more precise and comprehensive systems for describing and rationalizing production processes (Shipman & Marshall, 1999). The costs of enforcing univocal, precise meanings are especially visible in health care settings where hospital staff are collectively responsible for competing goals, e.g. ensuring patients are able to care for themselves after discharge and also maximizing insurance payments. In some cases staff reconcile their divergent organizational roles through maintaining alternative interpretations of a shared phrase such "independent aging" (Altomonte, 2020). The value of flexible meaning within practical constraints is also well-established in the study of boundary objects in science (Star & Griesemer, 1989): strategic ambiguity recognizes that actors may share aims only at a certain level of abstraction, and polysemy provides a means for communication at an appropriately coarse-grained level. Strategic ambiguity is also useful when communicating under conditions of high uncertainty, so that leaving the intended meaning open-ended accurately expresses the speaker's current understanding and facilitates future robust action as more evidence comes to light (Currie, 2015; Ferraro et al., 2015).

Finally, polysemy can serve as an instrument for generating creative new ways to anchor language and thought in empirical experience (Hesse, 1988; Ortony, 1993; Geeraerts & Geeraerts, 1997). The philosopher Mary Hesse famously distinguished between positive, negative, and neutral aspects of analogies or metaphors: features known to be shared with the target phenomenon, features that are known not to be shared, and features whose statuses are not yet settled, respectively (Hesse, 1988). The creative use of ambiguity is not limited to the social or biological sciences in this respect, and has been of great value to physics and mathematics in contexts where syntactic rules sometimes permit the use of theoretical constructs outside established semantic frameworks and thereby open up new spaces of conceptual possibilities (Wilson, 2006; Grosholz, 2007). The idea of sensitizing concepts plays a similar role for theorizing in the qualitative social sciences (Bowen, 2006). As "fresh" metaphors or analogies enter common usage, their generative qualities diminish with time as people learn to associate the term with one or more additional, stabilized meanings.

Considered together, these four rules synthesize results from multiple fields demonstrating different ways in which polysemic ambiguity can be productive or harmful. While the list is not exhaustive, developing a more complete and detailed accounting can meaningfully help advance theoretical and empirical research. Based on the analysis I've presented so far, it is important for data-centric science to keep in mind the diversity of circumstances in which scientific terms are used, including the different aims to which polysemy can be used and the different sorts of contextual information and background knowledge required. As computational infrastructure becomes increasingly dominates the means for communicating scientific information and ideas, we should consider what sort of support is deserved by these different uses of polysemic language. If the efficient, strategic, and generative uses for polysemy remain important for the future progress of science, infrastructure that supports only context-independent, determinate uses of scientific language will not meet the full needs of the community.

4 Connecting theory to empirical hypotheses

The four prescriptive rules I formulated are pragmatic in nature: they express how people should use language to achieve a specified aim under certain circumstances. If a community of scientists follows one or more of these rules, this should have consequences for their observable linguistic behaviors. As a result, it is possible to formulate empirically testable generalizations relating three key aspects of a term's usage: how often and in what contexts the term is used, the communicative goal being prioritized, and circumstantial factors of its use, including the specific linguistic and social context available and the background knowledge of participants. To the extent that the rules characterize norms actually at work in scientists' practices, then, it should be possible to predict one or more of these aspects from knowledge of the others. In this section, I introduce some potential theoretical principles we can apply to this purpose and illustrate how they can be used to derive testable predictions.

The key result I will draw on from my analysis so far is that maintaining shared background knowledge and abilities is essential to effective communication when terms have multiple meanings, and hence to any positive use of ambiguity. We can then add the further premise that as ideas spread into new contexts or the participating community changes, e.g. by growing larger or more heterogeneous, additional work is required to prevent communication from breaking down (Gerson, 2008). This work may take multiple forms: for instance, it may involve adjusting how people use existing terminology to accommodate changing contexts, revising terminology by introducing new or altered terms and definitions, or by investing more heavily in social mechanisms to ensure participants have sufficiently similar background knowledge (Dourish, 2001; Gerson, 2008; Dietz, 2012).

We can connect these qualitative principles to observable linguistic patterns using sets of terms that overlap in their possible meanings. For convenience, I will call two or more terms that have overlapping but non-identical possible meanings "partial synonyms." Related ideas have been presented in McMahan and Evans (2018), Hauer and Kondrak (2020), and Alagić and Šnajder (2021) but in different technical contexts that are outside the scope of what I can fully discuss here. Partial synonyms can vary in their degrees of polysemy (i.e. number of meanings), their specific lists of possible meanings, and their textual characteristics such as number of written characters or syllables. For example, two terms may be partial synonyms because they have semantically nested sets of possible meanings, implying that word choice conveys information about meaning even with otherwise identical contextual information. Technical terminology is often structured through hierarchical specification in this way: "evolutionary function" and "biochemical function," for instance, are partial synonyms of the more generic "function," but not of each other. Considered together, though, these three terms are connected by partial synonymy relations and form what I will call a partial synonym network. Edges in a partial synonym network are defined by overlaps in possible meanings and therefore differ from the broader idea of semantic networks in linguistics as they don't represent negations or oppositions among terms. Another example of a partial synonym network from systematic biology would be "subspecies," "species," "race," "form," and "ecotype." There is no restriction on the structure of networks partial synonyms may form, which for example could take the form of a strict semantic hierarchy or of a family resemblance, i.e where no single meaning is shared by all and only the partial synonyms in the network. See Cusimano and Sterner (2019) for more extensive discussion of the types of semantic relationships among partial synonyms that philosophers have formulated in the setting of conceptual analysis.

The token occurrences and contexts of partial synonyms in a text corpus are primary observable features one can study. Partial synonyms are in principle substitutable for each other when the possible meaning they share is the intended one. In practice, however, we may observe substantial variation in partial synonym usage, such as differences in relative frequency of use within or across contexts. Less directly observable but nonetheless significant is how frequently speakers use a partial synonym to convey each of its potential meanings. Indeed, polysemic terms often have one meaning that is substantially more common than the other alternatives (Meyer & Lewis, 2020). The frequency of intended meanings for a term may also be estimated conditional on a single context or as an unconditioned average across all linguistic contexts.

Partial synonyms therefore provide an empirical approach to detecting synchronic or diachronic patterns in polysemic language, enabling investigation into whether these patterns can be explained through differences in the linguistic and semantic properties of terms as well as changes in exogenous variables such as the size or diversity of the population of authors and readers. For example, tracking the relative frequencies of partial synonyms may show significant shifts in whether authors prefer longer but less polysemic terms to shorter terms with more meanings. If it is correct that growth in community size or diversity cause a decline in shared background knowledge, then all else being equal authors will shift to using less polysemic terms because adequate contextual information is less reliable. Conversely, if the community composition stabilizes or increased investment in background knowledge ensures reliable contextual information, then authors will adopt more polysemic terms when these are easier to read and write.

To show how this reasoning can be made more precise and quantitative, I will draw on an information-theoretic view of language as conveying bits of information about the speaker's intended meaning (Piantadosi et al., 2012). Piantadosi et al. originally introduced this approach in an 'equilibrium' or static setting, though, so the application here to partial synonyms in a dynamic environment represents a significant extension. Information theory provides an extremely general and quantitative framework for expressing uncertainty about the mapping between a linguistic form and its possible meanings M. In particular, we can express the polysemic ambiguity of a linguistic form via its entropy, $H[M] = -\sum_{m \epsilon M} P(m) \log P(m)$. That is, the entropy H of a term relative to a set of possible meanings M is the sum of the weighted probabilities that each meaning will occur in use. We will see that the entropy provides a useful way of quantifying ambiguity by measuring the uncertainty of the speaker's audience about the intended meaning.

As noted earlier, Piantadosi et al. argue that an "optimally efficient communication system should look ambiguous, as long as context is informative about meaning" (Piantadosi et al., 2012, p. 282). If we treat the polysemic ambiguity of a single term as simply a function of the number and probability of its possible meanings, this excludes the possibility of disambiguating information available in the environment. An expanded expression would incorporate some contextual information c into the measure of entropy across a set of contexts C:

$$H[M|C] = -\sum_{c \epsilon C} P(c) \sum_{m \epsilon M} P(m|c) \log P(m|c).$$

$H[M|C]$ can thus be interpreted as the expected entropy over meanings, in context (Piantadosi et al., 2012, p. 283). It will be lower than $H[M]$ when taking context into account increases the audience's certainty about the intended meaning (since lower entropy means greater certainty). This contextualized expression for the entropy forms the basis for the simple but very general argument that we should expect polysemy to be a prevalent feature of optimally efficient languages. A maximally efficient communication system will not convey unnecessary information, and since language use always occurs in specific contexts, the amount of information that the word itself must convey is necessarily less than what would be required without context.

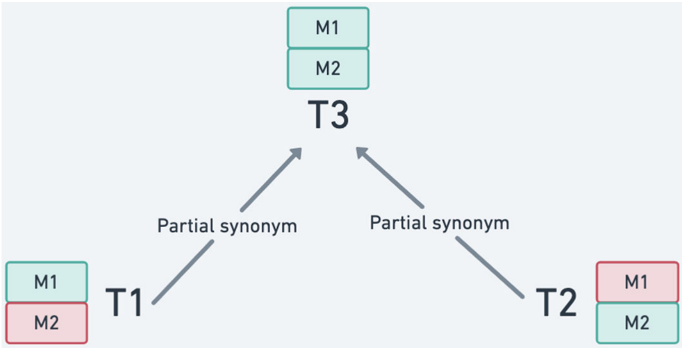

Fig. 2 A simple partial synonym network composed of three terms, T_1, T_2, and T_3. T_3 is polysemic because it has two possible meanings, M_1 and M_2. T_1 and T_2 each have only one possible meaning, M_1 and M_2, respectively. T_1 and T_3 are partial synonyms because they are in principle substitutable for the same intended meaning, and likewise for T_2 and T_3

A major assumption in the expression for this expected entropy across contexts "is that speakers and listeners have the same—or very similar—coding schemes (corresponding to similar probabilistic models of language and the world), and also the same ability to use contextual information to constrain the possible meanings" (Piantadosi et al., 2012, p. 283). In other words, maximally efficient communication depends on background knowledge about the possible meanings of a term and how to use contextual information to identify the intended one(s). Social change is capable of disrupting both of these prerequisites for accurate disambiguation and hence the overall efficiency of the language. If no work is done to repair this loss of common knowledge, changes in the social composition of the community (e.g. through growth or turnover) will tend to cause decay in the prerequisites for efficient communication. Similarly, as usage of a word expands to new contexts, accuracy and ease of communication will tend to decrease unless additional work is done.

Partial synonym networks provide a way to link these high-level theoretical arguments to empirical data through comparing trends in the frequency, properties, and contexts of terms that are in principle semantically substitutable for each other. To illustrate the core approach, consider a simple network of three terms: a polysemic term T_3 with two meanings, M_1 and M_2, and two monosemic terms T_1 and T_2 that are partial synonyms for M_1 and M_2, respectively. See Fig. 2 for a graphical representation. In principle, in any context where the speaker intends to convey M_1, either T_1 or T_3 are adequate choices, and similarly with T_2 and T_3 for M_2. However, several other factors may enter into the speaker's choice of term. The terms may vary in number of characters or other measures of linguistic difficulty such as surprisal (i.e. using unusual syllables). T_3 also has two possible meanings and hence would require contextual information for disambiguation, whereas T_1 and T_2 have only one meaning and therefore don't. Furthermore, background knowledge may enter into consideration if the audience is less likely to know one of the terms, for example because it is new or requires specialized training to learn.

We can observe the effective substitutability of these terms within a context by comparing their actual frequency of usage when that context occurs in a text corpus or other body of linguistic data. When a speaker aims for semantic determinacy, they should use polysemic terms only when adequate contextual information is available and switch to use alternative terms when it is not. If we hypothetically assume T_1 and T_3 to be practically equivalent except for their difference in polysemy, then any difference in their actual frequency of use when M_1 is the intended meaning must be due to the availability of sufficient contextual information for T_3. In particular, T_1 should be used equally when contextual information is low or absent, while T_1 and T_3 should be preferred when contextual information is high. If T_3 is also substantially shorter or simpler than T_1, then T_3 should be more common in contexts with sufficient cues for disambiguation.

A different application of partial synonym networks allows one in principle to estimate relative changes in background knowledge required to use contextual information. In this scenario, assume all features of T_1, T_2, and T_3 are constant except the frequency of their usage. Given some context where an "expert" is able to reliably disambiguate the meaning of T_3 to be M_2, for example, then we expect speakers to use T_3 and T_2 at some stable relative frequency based on the costs associated with their other features. If we then observe a change in this relative frequency over time, assuming all other features of the partial synonyms are fixed, this represents a change in the audience's ability to make use of the contextual information available. If the polysemic T_3 becomes more common relative to monosemic T_2, for example, the audience is using existing contextual information more effectively, and vice versa if T_3 declines relative to T_2.

If we bring social context into the picture, partial synonym networks can also be applied to distinguish between semantically determinate and strategically indeterminate uses of polysemic terms. Consider a simple version of the "trading zone" scenario that Peter Galison introduced for understanding how members of different scientific specialties can interact without unification under a shared vocabulary (Galison, 1996). Assume we have a community of discourse formed of two subgroups and three types of social contexts: a homogeneous social context where the speaker and audience are only from group 1, a similarly homogeneous social context for group 2, and a heterogeneous context where members of both groups are present. Now we can look at how speakers use the same set of three partial synonyms in different social and linguistic contexts. Within each homogeneous social context, we assume speakers aim for efficient semantic determinacy, so that the polysemic term T_3 will only be used where adequate social-linguistic contextual information is available, and speakers will prefer the terms T_1 and T_2 where this is missing. Assuming that members of both groups have less contextual information available to them in the heterogeneous social context, then a strategic use of ambiguity predicts that speakers will use the polysemic term T_3 more often in the heterogeneous social context compared to the same linguistic context in their homogeneous group.

Alternatively, semantic indeterminacy can arise through generative uses of ambiguity, such as novel metaphors, where the speaker uses a term in a creative way that suggests a meaning to the audience beyond any of the already established senses. The interactional theory of metaphor suggests one simple scenario where partial synonyms can be applied, although the full range of theoretical possibilities has yet to be explored. Let's assume we can represent a novel usage for the term T_1 as the assertion of a new metaphor, "T_1 is T_4", where T_4 is a term that previously was outside T_1's partial synonym network. Because people have not previously treated T_1 and T_4 as semantically substitutable, we would expect T_4 generally occurs in a different set of linguistic contexts than T_1 and has its own set of partial synonyms distinct from T_1, T_2, or T_3. If people start regularly using T_1 in this generative way, it should therefore appear as a pattern of novel linguistic (and possibly social) contexts in which T_1 behaves as a partial synonym with T_4, and potentially other members of T_4's partial synonym network. However, we also expect people will experiment with the metaphor at first and that there will be an eventual ossification of how it gets used, which may appear as a peak and then decline in the number of linguistic contexts and partial synonyms for T_1. These predictions are distinct from what we expect for efficient or strategic uses of polysemy, where no fluctuations in the partial synonym networks are expected.

The hypothetical scenarios I've introduced show the utility of partial synonym networks to formulate predictions about polysemic language in light of different normative ideals. I've also indicated in principle how these predictions can be compared to observed patterns in the usage of partial synonyms relative to each other and across linguistic or social contexts. Substantial work remains to turn these scenarios into realistic quantitative models, and in the next section I turn to survey relevant empirical methods and results.

5 Integrating emerging computational methods and results

The largest-scale empirical evidence demonstrating the positive value of polysemy comes from corpus linguistics, where researchers have investigated aggregate trends in multiple languages. Piantadosi et al.'s work represents an important landmark in this respect(Piantadosi et al. 2012), and has inspired multiple further lines of research. Their study used datasets for three languages—English, German, and Dutch—to test whether easier-to-use linguistic units also show more ambiguity, controlling for confounding variables such as frequency of use. They found robust support for their hypothesis in homophones (where the same spoken sound can designate different words), polysemic words (where the same word can designate multiple meanings), and syllables (where the same syllable taken individually may indicate different complete words). Hence polysemy both reduces redundancy in communication by utilizing shared context and lowering the number and difficulty of linguistic units needed to communicate. Later studies have buttressed these conclusions (Gibson et al., 2019), and have found, for example, a positive correlation between the degree of ambiguity and amount of contextual information provided by speakers (Pimentel et al., 2020), and that more polysemous words have higher rates of semantic change, consistent with their innovative reuse (Hamilton et al., 2016).

The typology I've introduced here, though, suggests a much broader scope for investigating the use and value of polysemy. In this section, I highlight empirical methods and results that are relevant to studying the value of ambiguity but that have developed largely in isolation. I also suggest some ways in which these methods and results can be better theoretically and practically integrated through the study of partial synonyms (Grantham, 2004). Corpora of written texts provide an immediate and large source of data for detecting patterns of polysemy. Considering resources for written discourse, there are now rich repositories for scientific articles and monographs such as JSTOR, PubMed, Hathitrust, and Arxiv, and extensive sources for semantic information about terms, such as WordNet (Miller, 1995), Wikipedia, and terminology projects such as Leon-Arauz et al. (2018). Natural language processing (NLP) in computer science has clear potential for developing methods and tools that enable the scaling of qualitative analyses of texts to large text corpora. Computational methods for analyzing semantic relationships between terms in scientific texts in particular are advancing rapidly, most notably with the development of powerful new neural network models such as Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019).

As we saw with the partial synonyms approach, however, explaining patterns in the use of polysemic terms depends on disentangling their semantic relationships, the immediate linguistic and social contexts of their use, and the background factors influencing people's interpretive skills and accuracy. Returning to Fig. 1, we can also read it as a guide for the types of information needed to detect and explain patterns of polysemy. In particular, knowing the distinct meanings of scientific terms is key to applying the partial synonym approach, but the limited availability of this qualitative knowledge is often an obstacle to computational studies using large text corpora. Many computational studies instead have relied on proxies for the number or nature of a term's meanings. A common approach assumes that distinct meanings tend to occur in distinct contexts, so that the number and clustering of a term's linguistic contexts should correlate with its number and distinct meanings, respectively. Of course, this proxy strategy is unable to detect cases of ambiguity in a recurring linguistic context, and it tends to conflate how frequently a term is used across a corpus with the number of meanings it has. Testing explanations empirically will therefore require both qualitative and quantitative methods from multiple disciplines.

Creating datasets focused specifically on polysemic terminology in science therefore represents a potential opportunity for interdisciplinary collaboration among researchers using qualitative and quantitative methods. Linguists have already developed several large databases, such as WordNet, but these typically focus on common words or phrases used in everyday human language. Similarly, NLP problems such as word sense induction and synset induction are active areas of research in computer science (Panchenko et al., 2017; Ustalov et al., 2018), but the need for specialized training datasets are a key obstacle to applying NLP methods to ambiguity in scientific texts. The most popular deep learning algorithms in NLP over the past few years have large gaps in coverage for scientific terminology: SciBERT, for example, has training data from computer science and biomedicine but not ecology or evolution (Beltagy et al., 2019).

Lexical semantics, computer ontologies, and philosophy have the potential to help fill these gaps by contributing formal definitions of senses for polysemic terms and characterizing their linguistic contexts. Determining when peoples' understandings of the terms' meanings are sufficiently distinct to merit formal differentiation into senses is also an important topic for study. Prior work in lexical semantics has been centrally concerned with the challenges of individuating senses for a word (Geeraerts & Geeraerts, 1997; Kilgarriff, 1997) and methods for translating technical terminology into different human languages (Leon-Arauz et al., 2018; L' Homme et al., 2020). In computer science, the ontology alignment problem aims to specify semantic relationships between terms defined in different classification systems, and multiple algorithms have been developed for aligning biomedical or anatomical trait ontologies (Bertone et al., 2013; Dragisic et al., 2017; Oliveira & Pesquita, 2018). A related problem also arises for computer ontologies when they persist long enough to be substantially revised, creating the possibility of unnoticed drift in the meaning of annotated datasets (Fokkens et al., 2016; Duncan, 2020). These apparently narrow technical problems in data classification also have broader interest for the topic of theory integration in philosophy of science (Laubichler et al., 2018). Formalized classification systems could provide a rich source of cases for linguists, philosophers, and ontology designers to analyze partial synonym networks among scientific terms while also contributing to the improvement of data integration and computational studies of text corpora.

Another key element is middle-range modeling, both qualitative and quantitative, to help bridge theory to data by specifying concrete predictions and explanations for empirically observed patterns of polysemy. Historical, rhetorical, and social studies of science, as well as scientometrics, are critical in this respect to provide both qualitative and quantitative insights into the broader social contexts in which partial synonyms vary in usage across communities and time. They are also key to illuminating how communities respond to these shifting circumstances to preserve existing linguistic habits, for example through investing in increased training or by revising their habits to place lower demands on background knowledge. A couple studies, for example, have used text corpora to test claims about the effects of polysemy specifically on the diffusion or usage dynamics of scientific terms (Volanschi & Kübler, 2011; McMahan & Evans, 2018).

In the sociology of science, for example, McMahan and Evans (2018) have proposed and tested several hypotheses about ambiguity in academic scholarship, including the sciences, humanities, and arts. They suggest first that "more ambiguity in science and scholarship will, on average, lead to greater integration and less fragmentation of subsequent work that references it," and second, that "ambiguity yields disagreement and drives productive engagement in science and scholarship," with interaction effects expected between the level of engagement with a work of scholarship and the diversity of fields citing it (McMahan & Evans 2018, p. 870). Using a corpus of about two million article abstracts and bibliographies from Web of Science, they found a small but consistent and statistically significant effect linking increased ambiguity with increased engagement across fields. They also found clear trends in the average ambiguity of words across fields, measured in terms of how frequently synonyms were substituted for each other in the same grammatical context, similar to the theoretical scenarios for partial synonyms I discussed above. These effects were detected both as the average ambiguity of words within each field and the degree of ambiguity of words across different disciplines.

Engagement through ambiguity and creative interpretation are often closely coupled, as demonstrated by several case studies in the social sciences (Hirsch & Levin, 1999; Giroux, 2006; Keuchenius et al., 2021). For example, Mark Granovetter's ground-breaking publication on the strength of weak ties proposes that weak ties are more likely to be effective bridges between communities than strong ties and are therefore crucial for the flow of information among groups (Granovetter, 1973). Keuchenius et al. trace the diffusion of Granovetter's idea through a citation network and use computational community detection methods to identify three distinct citation clusters that form relatively distinct modules of the network. They then analyze leading texts in these communities to show how they develop distinct local interpretations of the strength of weak ties concept: (1) as a feature of social organizations such as businesses, (2) as a property of individuals in a social network, and (3) as a general characteristic of efficient networks not restricted to social interactions. Keuchenius et al. conclude that "the spread of scientific ideas entails a complex process of translation in which scholarly communities emerge as meso-level mediators, cultivating divergent interpretations of the diffusing idea in line with the different research projects in which they are engaged" (Keuchenius et al., 2021, p. 11). Greenhalgh et al. have also explored this theme in the context of conducting systematic reviews of how ideas diffuse across disciplines (Greenhalgh et al., 2005).

Independently, Hirsch and Levin (1999) proposed a general model for the temporal evolution of "umbrella constructs" that at first serve to connect and integrate an emerging research subject but eventually face challenges to their validity and may ultimately abandoned. They suggest that what survives in the long term are component concepts used in the umbrella construct that emerged during the initial enthusiasm but that succeeded in resisting critical challenges, or which can be re-appropriated for use in new trends. Hirsch and Levin illustrate this model using the concept of organizational effectiveness and document four life-cycle stages: "emerging excitement, the validity challenge, 'tidying up with typologies,' and construct collapse" (Hirsch & Levin, 1999, p. 199). The strategic value of polysemy is central to the initial excitement phase and validity challenge, as many researchers seek to link their existing methods and research subjects to the emerging trend but end up producing incompatible or disconnected studies. Two surveys of the organizational effectiveness literature in the 1970s, for example, found that published articles used 19 different measures of effectiveness and showed little overlap across studies (Hirsch & Levin, 1999, p. 201). If advocates for an umbrella construct cannot find a way to clarify and synthesize diverse approaches in the tidying up with typologies phase, collapse is likely to follow. A later study by Hèléne Giroux (2006) analyzes lexical ambiguity in "quality management" in a similar way.

The typology, principles, and partial synonym approach I have introduced provide a promising basis for formalizing and generalizing these models within a common theoretical framework, and the framework also suggests new opportunities for theorizing. In particular, the generative use of polysemy diverges from its strategic use specifically where novel usage drives new interpretations rather than re-using existing meanings. To my knowledge, there are no corpus-based models that clearly distinguish between these circumstances, i.e. that identify cases where a polysemic term is multiply interpreted by an audience according to pre-existing meanings rather than having the audience generate novel interpretations. Nonetheless, substantial work already addresses how to document trends in metaphors in different scientific fields (Gentner &

Grudin, 1985; Del Tredici et al., 2016), detect metaphors in text corpora (Schlechtweg et al., 2017), and use currently living metaphors more thoughtfully (Higuera, 2018; Olson et al., 2019; Perrault & O'Keefe, 2019; Swedberg, 2020). However, we lack models that disentangle the overall usage frequency of polysemic terms by specific contexts in order to document context-specific changes in meaning frequencies. Detecting semantic change is an important step in this direction (Tahmasebi et al., 2021), as is accounting for how broader changes in the popularity of discourse topics can mediate whether partial synonyms (such as "airplane" and "aeroplane") compete or coexist over time (Karjus et al., 2020).

6 Conclusion

The rise of data-centric science poses new opportunities for explaining and evaluating ambiguity in scientific language, especially through the computational analysis of text corpora. We are all collectively creating a new computational "space" for human knowledge and language, and it is far from settled how we should conceptualize and justify the role of ambiguity in this arena. Nonetheless, current practices in data science frequently embrace the idea that technical terms shouldn't have multiple and conflicting meanings. While the norm of devaluing polysemy is common across controlled vocabularies and ontologies in science and the Semantic Web more broadly, it can prove ill-suited to the aims of efficient communication, collective action, and creative thinking. Research dispersed across multiple disciplines has identified and supported several ways in which polysemy can have positive value for communication and coordinated action. Extending an analysis of the core elements of polysemy presented in Piantadosi et al. (2012), I've formulated a typology of prescriptive rules that can explain scientists' linguistic practices of using or avoiding polysemy. I then introduced the idea of partial synonyms to show how we can formulate new hypotheses linking the value of polysemy to empirically observable patterns of usage and external factors such as community size and diversity. I also illustrated how the study of partial synonyms can serve to integrate emerging results and methods across a wide range of fields in the humanities and sciences.

For those interested in understanding how science works and the knowledge it creates, the value and nature of ambiguity holds inherent interest and significance. The four prescriptive rules I formulated express pragmatic generalizations about how scientists should behave for the sake of achieving a general goal such as efficient or innovative communication in light of the circumstances in which they are working. These rules are meant to capture explicit or implicit reasons for why scientists may use or avoid polysemic language in practice. However, the rules I presented are provisional in and need further clarification and testing: they serve here to demonstrate a domain of phenomena for conceptual and empirical investigation whose full scope and importance has been overlooked because so much of the relevant research has happened in disconnected fields.

Beyond the intrinsic interest of ambiguity's use in science, what can explaining ambiguity contribute to the design of computational infrastructure for science? The simplest but nevertheless critical point is that ambiguity, even in highly computationally intensive settings, is not inherently or self-evidently harmful to communication or action. The prescriptive rule against using polysemic terms is demonstrably less efficient for communication when adequate disambiguating context is available. The common preference for one-word-one-meaning is also insufficient for interdisciplinary communication or collaborative projects where persistent disagreements exist among participants. More broadly, the need to communicate and reason about scientific data is not limited to only the most stable and literal research terminologies—metaphors continue to shape and direct the empirical observations and theories of scientific fields.

I've also argued that the study of partially synonymous scientific terms provides a pathway toward deeper insights into the dynamic and context-sensitive utility of polysemy. The key insights I drew from cognitive pragmatics are that words are not merely strings of characters or verbal sound-patterns and that the interpretation of meaning happens in pragmatic situations with constrained resources. Moreover, the relevant resources span several scales: the proximate linguistic context of a term's use, the background knowledge of the speaker and audience, and the institutional environment that scaffolded their common knowledge and skills. We can use partial synonyms to make these implicit features of scientific language visible to empirical investigation because partial synonyms are in principle semantically substitutable but vary systematically in other features such as their length, frequency, possible meanings, and typical contexts of use. Defining and measuring the characteristics and usage of partial synonyms presents an opportunity for novel interdisciplinary collaborations to using large collections of scientific texts and enrich the quality of computer ontologies and other formalized classification systems for scientific data and concepts.

Acknowledgements This project was funded by NSF Science and Technology Studies Grant STS-1827993. My thanks to Joeri Witteveen, Elizabeth Lerman, Nico Franz, and Manfred Laubichler for their conversations and feedback about the ideas presented here. My special thanks also to the reviewers whose constructive comments helped improve the manuscript significantly. All mistakes are entirely my own.

Author Contributions Not applicable.

Funding Funding was provided by NSF STS-1827993.

Data availability Not applicable.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Code availability Not applicable.

References

Alagić, D., & Šnajder, J. (2021). Representing word meaning in context via lexical substitutes. *Automatika*. https://doi.org/10.1080/00051144.2021.1928437.

Ali-Khan, S. E., Jean, A., MacDonald, E., et al. (2018). Defining Success in Open Science. *Mni Open Research*. https://doi.org/10.12688/mniopenres.12780.1.

Altomonte, G. (2020). Exploiting ambiguity: A moral polysemy approach to variation in economic practices. *American Sociological Review, 85*(1), 76–105. https://doi.org/10.1177/0003122419895986.

Arp, R., Smith, B., & Spear, A. D. (2015). *Building ontologies with basic formal ontology*. MIT Press.

Beltagy, I., Lo, K., & Cohan, A. (2019). SciBERT: A pretrained language model for scientific text. http://arxiv.org/abs/1903.10676 [cs]

Bertone, M. A., Miko, I., Yoder, M. J., et al. (2013). Matching arthropod anatomy ontologies to the Hymenoptera Anatomy Ontology: Results from a manual alignment. *Database, 2013,* bas057–bas057. https://doi.org/10.1093/database/bas057.

Bowen, G. A. (2006). Grounded theory and sensitizing concepts. *International Journal of Qualitative Methods, 5*(3), 12–23. https://doi.org/10.1177/160940690600500304.

Bowker, G. C. (2000). Biodiversity datadiversity. *Social Studies of Science, 30*(5), 643–683. https://doi.org/10.1177/030631200030005001.

Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.

Brandom, R. B. (2008). *Between saying and doing: Towards an analytic pragmatism*. Oxford University Press.

Carr, J. W., Smith, K., Culbertson, J., et al. (2020). Simplicity and informativeness in semantic category systems. *Cognition, 202*(104), 289. https://doi.org/10.1016/j.cognition.2020.104289.

Català, N., Baixeries, J., Ferrer-Cancho, R., et al. (2021). Zipf's laws of meaning in Catalan. http://arxiv.org/abs/2107.00042

Ceccarelli, L. (2001). *Shaping science with rhetoric: The cases of Dobzhansky, Schrödinger, and Wilson*. University of Chicago Press.

Ceusters, W., Smith, B., & Goldberg, L. (2005). A terminological and ontological analysis of the NCI thesaurus. *Methods of Information in Medicine, 44*(04), 498–507. https://doi.org/10.1055/s-0038-1634000.

Currie, A. (2015). Marsupial lions and methodological omnivory: Function, success and reconstruction in paleobiology. *Biology & Philosophy, 30*(2), 187–209. https://doi.org/10.1007/s10539-014-9470-y.

Cusimano, S., & Sterner, B. (2019). Integrative pluralism for biological function. *Biology & Philosophy, 34*(6), 55. https://doi.org/10.1007/s10539-019-9717-8.

Davenport, S., & Leitch, S. (2005). Circuits of power in practice: Strategic ambiguity as delegation of authority. *Organization Studies, 26*(11), 1603–1623. https://doi.org/10.1177/0170840605054627.

DeFries, R., & Nagendra, H. (2017). Ecosystem management as a wicked problem. *Science, 356*(6335), 265–270. https://doi.org/10.1126/science.aal1950.

Del Tredici, M., Nissim, M., & Zaninello, A. (2016). Tracing metaphors in time through self-distance in vector spaces. http://arxiv.org/abs/1611.03279 [cs]

Denis, J. L., Dompierre, G., Langley, A., et al. (2011). Escalating indecision: Between reification and strategic ambiguity. *Organization Science, 22*(1), 225–244. https://doi.org/10.1287/orsc.1090.0501.

Devlin, J., Chang, M.W., Lee, K., et al. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. http://arxiv.org/abs/1810.04805

Dietz, B. (2012). Contribution and co-production: The collaborative culture of Linnaean botany. *Annals of Science, 69*(4), 551–569. https://doi.org/10.1080/00033790.2012.680982.

Dourish, P. (2001). Process descriptions as organisational accounting devices: the dual use of workflow technologies. In: Proceedings of the 2001 International ACM SIGGROUP conference on supporting group work. Association for Computing Machinery, New York, NY, USA, GROUP '01 (pp 52–60), https://doi.org/10.1145/500286.500297

Dragisic, Z., Ivanova, V., Li, H., et al. (2017). Experiences from the anatomy track in the ontology alignment evaluation initiative. *Journal of Biomedical Semantics, 8*(1), 56. https://doi.org/10.1186/s13326-017-0166-5.

Duncan, M. (2020). Terminology version control discussion paper. http://mrtablet.co.uk/chocolate_teapot_lite.htm

Eisenberg, E. M. (1984). Ambiguity as strategy in organizational communication. *Communication Monographs, 51*(3), 227–242. https://doi.org/10.1080/03637758409390197.

Ferraro, F., Etzion, D., & Gehman, J. (2015). Tackling grand challenges pragmatically: Robust action revisited. *Organization Studies, 36*(3), 363–390. https://doi.org/10.1177/0170840614563742.
Ferrer-i Cancho, R., Bentz, C., & Seguin, C. (2020). Optimal coding and the origins of Zipfian laws. *Journal of Quantitative Linguistics*. https://doi.org/10.1080/09296174.2020.1778387.
Fokkens, A., Ter Braake, S., Maks, I., et al. (2016). On the semantics of concept drift: Towards formal definitions of semantic change. Proceedings of Drift-a-LOD (2016): 247–265.
Franz, N. M., & Sterner, B. W. (2018). To increase trust, change the social design behind aggregated biodiversity data. *Database*. https://doi.org/10.1093/database/bax100.
Galison, P. (1996). Computer simulations and the trading zone. In P. Galison & D. J. Stump (Eds.), *The disunity of science: Boundaries, contexts, and power* (pp. 118–57). Stanford University Press.
Garnett, S. T., Christidis, L., Conix, S., et al. (2020). Principles for creating a single authoritative list of the world's species. *PLoS Biology, 18*(7), e3000736. https://doi.org/10.1371/journal.pbio.3000736.
Garson, J. (2016). *A critical overview of biological functions*. Springer.
Geeraerts, D. (1997). *Diachronic prototype semantics: A contribution to historical lexicology*. Clarendon Press.
Gentner, D., & Grudin, J. (1985). The evolution of mental metaphors in psychology: A 90-year retrospective. *American Psychologist, 40*(2), 181–192. https://doi.org/10.1037/0003-066X.40.2.181.
Germain, P. L., Ratti, E., & Boem, F. (2014). Junk or functional DNA? ENCODE and the function controversy. *Biology & Philosophy, 29*(6), 807–831. https://doi.org/10.1007/s10539-014-9441-3.
Gerson, E. M. (2008). Reach, bracket, and the limits of rationalized coordination: Some challenges for CSCW. *Resources, co-evolution and artifacts* (pp. 193–220). Springer. https://doi.org/10.1007/978-1-84628-901-9_8.
Gibson, E., Futrell, R., Piantadosi, S. P., et al. (2019). How efficiency shapes human language. *Trends in Cognitive Sciences, 23*(5), 389–407. https://doi.org/10.1016/j.tics.2019.02.003.
Giroux, H. (2006). 'It was such a handy term': Management fashions and pragmatic ambiguity. *Journal of Management Studies, 43*(6), 1227–1260. https://doi.org/10.1111/j.1467-6486.2006.00623.x.
Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology, 78*(6), 1360–1380. https://doi.org/10.1086/225469.
Grantham, T. A. (2004). Conceptualizing the (dis)unity of science. *Philosophy of Science, 71*(2), 133–155. https://doi.org/10.1086/383008.
Greenhalgh, T., Robert, G., Macfarlane, F., et al. (2005). Storylines of research in diffusion of innovation: A meta-narrative approach to systematic review. *Social Science & Medicine, 61*(2), 417–430. https://doi.org/10.1016/j.socscimed.2004.12.001.
Grice, H. P. (1975). Logic and conversation, syntax and semantics. *Speech Acts, 3,* 41–58.
Grosholz, E. (2007). *Representation and productive ambiguity in mathematics and the sciences*. Oxford University Press.
Gross, A. G. (2006). *Starring the text: The place of rhetoric in science studies*. Southern Illinois University Press.
Hamilton, W.L., Leskovec, J., & Jurafsky, D. (2016). Diachronic word embeddings reveal statistical laws of semantic change. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics, Berlin, Germany (pp 1489–1501). https://doi.org/10.18653/v1/P16-1141
Hauer, B., & Kondrak, G. (2020). Synonymy = Translational Equivalence. http://arxiv.org/abs/2004.13886[cs]
Hesse, M. (1988). The cognitive claims of metaphor. *The Journal of Speculative Philosophy, 2*(1), 1–16.
Higuera, C. R. (2018). Productive perils: On metaphor as a theory-building device. *Linguistic Frontiers, 1*(2), 102–111.
Hirsch, P. M., & Levin, D. Z. (1999). Umbrella advocates versus validity police: A life-cycle model. *Organization Science, 10*(2), 199–212. https://doi.org/10.1287/orsc.10.2.199.
Jarzabkowski, P., Sillince, J. A., & Shaw, D. (2010). Strategic ambiguity as a rhetorical resource for enabling multiple interests. *Human Relations, 63*(2), 219–248. https://doi.org/10.1177/0018726709337040.
Johansen, Winni. (2018). Strategic Ambiguity. In: The International Encyclopedia of Strategic Communication, edited by Robert L Heath, Winni Johansen, et al., 1st ed. Wiley. https://doi.org/10.1002/9781119010722.iesc0170.
Karjus, A., Blythe, R.A., Kirby, S., et al. (2020). Communicative need modulates competition in language change. http://arxiv.org/abs/2006.09277 [cs]

Kemp, C., Xu, Y., & Regier, T. (2018). Semantic typology and efficient communication. *Annual Review of Linguistics, 4*(1), 109–128. https://doi.org/10.1146/annurev-linguistics-011817-045406.
Keuchenius, A., Törnberg, P., & Uitermark, J. (2021). Adoption and adaptation: A computational case study of the spread of Granovetter's weak ties hypothesis. *Social Networks, 66,* 10–25. https://doi.org/10.1016/j.socnet.2021.01.001.
Kilgarriff, A. (1997). I don't believe in word senses. *Computers and the Humanities, 31*(2), 91–113. https://doi.org/10.1023/A:1000583911091.
L' Homme, M.C., Robichaud, B., & Subirats, C. (2020). Building multilingual specialized resources based on FrameNet: Application to the field of the environment. In: Proceedings of the International FrameNet Workshop 2020: Towards a Global, Multilingual FrameNet. European Language Resources Association, Marseille, France (pp 85–92) https://www.aclweb.org/anthology/2020.framenet-1.12
Lakoff, G., & Johnson, M. (2008). *Metaphors we live by*. University of Chicago press.
Laubichler, M. D., Prohaska, S. J., & Stadler, P. F. (2018). Toward a mechanistic explanation of phenotypic evolution: The need for a theory of theory integration. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution, 330*(1), 5–14. https://doi.org/10.1002/jez.b.22785.
Lean, O. M. (2021). Are bio-ontologies metaphysical theories? *Synthese*. https://doi.org/10.1007/s11229-021-03303-4.
Leon-Arauz, P., Martin, A. S., & Reimerink, A. (2018). The EcoLexicon English corpus as an open corpus in sketch engine. http://arxiv.org/abs/1807.05797 [cs]
Leonelli, S. (2012). Classificatory theory in data-intensive science: The case of open biomedical ontologies. *International Studies in the Philosophy of Science, 26*(1), 47–65. https://doi.org/10.1080/02698595.2012.653119.
Leonelli, S. (2016). *Data-centric biology: A philosophical study*. University of Chicago Press.
Leonelli, S., & Tempini, N. (2020). Data journeys in the sciences. *Springer*. https://doi.org/10.1007/978-3-030-37177-7.
Lidgard, S., & Love, A. C. (2018). Rethinking living fossils. *BioScience, 68*(10), 760–770. https://doi.org/10.1093/biosci/biy084.
Li, J., & Joanisse, M. F. (2021). Word senses as clusters of meaning modulations: A computational model of polysemy. *Cognitive Science, 45*(4), e12955. https://doi.org/10.1111/cogs.12955.
Linquist, S., Doolittle, W. F., & Palazzo, A. F. (2020). Getting clear about the F-word in genomics. *PLoS Genetics, 16*(4), e1008702. https://doi.org/10.1371/journal.pgen.1008702.
Loureiro, D., Rezaee, K., Pilehvar, M. T., et al. (2021). Analysis and evaluation of language models for word sense disambiguation. Computational Linguistics (pp 1–57). https://doi.org/10.1162/coli_a_00405
McMahan, P., & Evans, J. (2018). Ambiguity and engagement. *American Journal of Sociology, 124*(3), 860–912. https://doi.org/10.1086/701298.
Meyer, F., & Lewis, M. (2020). Modelling lexical ambiguity with density matrices. http://arxiv.org/abs/2010.05670 [cs]
Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM, 38*(11), 39–41. https://doi.org/10.1145/219717.219748.
Monckton, S., Johal, S., Packer, L., et al. (2020). Inadequate treatment of taxonomic information prevents replicability of most zoological research. *Canadian Journal of Zoology, 98*(9), 633–642. https://doi.org/10.1139/cjz-2020-0027.
Mons, B., Schultes, E., Liu, F., et al. (2019). The FAIR principles: First generation implementation choices and challenges. *Data Intelligence, 2*(1–2), 1–9. https://doi.org/10.1162/dint_e_00023.
Nakazawa, T. (2020). Species interaction: Revisiting its terminology and concept. *Ecological Research, 35*(6), 1106–1113. https://doi.org/10.1111/1440-1703.12164.
Nerlich, B., & Clarke, D. D. (2001). Ambiguities we live by: Towards a pragmatics of polysemy. *Journal of Pragmatics, 33*(1), 1–20. https://doi.org/10.1016/S0378-2166(99)00132-0.
Neto, C. (2020). When imprecision is a good thing, or how imprecise concepts facilitate integration in biology. *Biology & Philosophy, 35*(6), 58. https://doi.org/10.1007/s10539-020-09774-y.
Oliveira, D., & Pesquita, C. (2018). Improving the interoperability of biomedical ontologies with compound alignments. *Journal of Biomedical Semantics, 9*(1), 1. https://doi.org/10.1186/s13326-017-0171-8.
Olson, M. E., Arroyo-Santos, A., & Vergara-Silva, F. (2019). A user's guide to metaphors in ecology and evolution. *Trends in Ecology & Evolution, 34*(7), 605–615. https://doi.org/10.1016/j.tree.2019.03.001.
Ortony, A. (1993). *Metaphor and thought* (2nd ed.). Cambridge University Press.
Panchenko, A., Ruppert, E., Faralli, S., et al. (2017). Unsupervised does not mean uninterpretable : the case for word sense induction and disambiguation. In: 15th Conference of the European Chapter of

the Association for Computational Linguistics : proceedings of conference, volume 1: Long Papers. Association for Computational Linguistics, Stroudsburg, PA, pp 86–98.Stroudsburg, PA, pp 86–98. https://ub-madoc.bib.uni-mannheim.de/42007

Perrault, S. T., & O'Keefe, M. (2019). New metaphors for new understandings of genomes. *Perspectives in Biology and Medicine, 62*(1), 1–19.

Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition, 122*(3), 280–291. https://doi.org/10.1016/j.cognition.2011.10.004.

Pimentel, T., Maudslay, R.H., Blasi, D., et al. (2020). Speakers fill lexical semantic gaps with context. http://arxiv.org/abs/2010.02172

Poesio,M. (2020)."Ambiguity".In: The Wiley Blackwell Companion to Semantics, edited by Daniel Gutzmann, Lisa Matthewson, et al., 1st ed., 1–38. Wiley. https://doi.org/10.1002/9781118788516.sem098

Poirier, L. (2019). Classification as catachresis: Double binds of representing difference with semiotic infrastructure. *Canadian Journal of Communication*. https://doi.org/10.22230/cjc.2019v44n3a3455.

Ribes, D., & Bowker, G. C. (2009). Between meaning and machine: Learning to represent the knowledge of communities. *Information and Organization, 19*(4), 199–217. https://doi.org/10.1016/j.infoandorg.2009.04.001.

Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences, 4,* 155–169.

Schlechtweg, D., Eckmann, S., Santus, E., et al. (2017). German in flux: Detecting metaphoric change via word entropy. http://arxiv.org/abs/1706.04971

Sennet, A. (2021). Ambiguity. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (fall 2021). Metaphysics Research Lab, Stanford University.

Shavit, A., & Griesemer, J. (2011). Transforming objects into data: How minute technicalities of recording "species location" entrench a basic challenge for biodiversity. In: M. Carrier & A. Nordmann, Science in the context of application (Vol. 274, pp. 169–193). Springer. Dordrecht: Springer. https://doi.org/10.1007/978-90-481-9051-5_12.

Shipman, F. M., & Marshall, C. C. (1999). Formality considered harmful: Experiences, emerging themes, and directions on the use of formal representations in interactive systems. *Computer Supported Cooperative Work (CSCW), 8*(4), 333–352. https://doi.org/10.1023/A:1008716330212.

Stankowski, S., & Ravinet, M. (2021). Quantifying the use of species concepts. *Current Biology, 31*(9), R428–R429. https://doi.org/10.1016/j.cub.2021.03.060.

Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, "translations" and boundary objects: amateurs and professionals in Berkeley's museum of vertebate zoology, 1907–39. *Social Studies of Science, 19,* 387–420.

Sterner, B. W., & Franz, N. M. (2017). Taxonomy for Humans or Computers? Cognitive Pragmatics for Big Data. *Biological Theory12*(2), 99–111. https://doi.org/10.1007/s13752-017-0259-5

Sterner, B. W., Gilbert, E. E., & Franz, N. M. (2020). Decentralized but globally coordinated biodiversity data. *Frontiers in Big Data, 3*(519), 133. https://doi.org/10.3389/fdata.2020.519133.

Sterner, B. W., Witteveen, J., & Franz, N. M. (2020). Coordinating dissent as an alternative to consensus classification: Insights from systematics for bio-ontologies. *History and Philosophy of the Life Sciences, 42*(1), 8. https://doi.org/10.1007/s40656-020-0300-z.

Swedberg, R. (2020). Using metaphors in sociology: Pitfalls and potentials. *The American Sociologist, 51,* 240–257.

Tahmasebi, N., Borin, L., Jatowt, A., et al. (2021). Computational approaches to semantic change. Language Science Press. https://doi.org/10.5281/zenodo.5040241.

Takacs, D. (1996). *The idea of biodiversity: Philosophies of paradise*. Johns Hopkins University Press.

Ustalov, D., Chernoskutov, M., Biemann, C., et al. (2018). Fighting with the sparsity of synonymy dictionaries for automatic synset induction. Lecture Notes in Computer Science. In W. M. van der Aalst, D. I. Ignatov, M. Khachay, et al. (Eds.), *Analysis of images, social networks and texts* (pp. 94–105). Springer International Publishing.

Volanschi, A., & Kübler, N. (2011). The impact of metaphorical framing on term creation in biology. *Terminology International Journal of Theoretical and Applied Issues in Specialized Communication, 17*(2), 198–223. https://doi.org/10.1075/term.17.2.02vol.

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data, 3*(1), 1–9. https://doi.org/10.1038/sdata.2016.18.

Wilson, M. (2006). *Wandering significance: An essay on conceptual behavior*. Oxford University Press.

Winkler, S. (2015). Exploring ambiguity and the ambiguity model from a transdisciplinary perspective. *In: Winkler, S. Ambiguity. De Gruyter*. https://doi.org/10.1515/9783110403589-002/html.