

6 INSTRUMENTAL POWER

If we do not know the extent of the power that a particular instrument gives us, we shall not know what to make of our apparent failures, if there should be any, and shall in fact be unable to distinguish failure from a sort of success.

—PATRICK WILSON

Two Kinds of Power: An Essay on Bibliographic Control (1968, 35–38)

Given the described differences between traditional, description-oriented classifications, and composite-based, retrieval-oriented management classifications, one might ask the question, What are the *emergent* properties of these composite systems as they pertain to the biodiversity ecology? As the epigraph by Patrick Wilson indicates, the kinds of power we have over systems must be unpacked if we are to understand the extent of potential control we have over the documentary universe in question. To extrapolate on this potentiality, I outline below the kinds of instrumental power that consensus systems embody in the domain of biodiversity studies. To better understand this power, I expand on two related questions: (1) How do these systems allow us to better understand the extent and limitations of our current biodiversity knowledge? and (2) How does such functionality influence the way biodiversity knowledge is interpreted on a distributed basis, in terms of structuring data entities as well as allowing for the generation of new forms of inquiry?

To answer these questions, I introduce the concept of taxonomic *extensibility*, which I define as the extent to which consensus structures shape

biodiversity knowledge in specific, targeted ways. These kinds of emergent potentials are often, but not always, invoked in domains outside taxonomy proper. Such extension functions both within the Catalogue's own taxonomic space and within external systems. *Internal extension* relates to the Catalogue's ability to identify and compensate for certain gaps in its taxonomic record; that is, by bringing together global data, the Catalogue is able to identify the general contours of the biodiversity data landscape. Taxonomic gaps occur for many reasons, but typically they define areas that contain uncharismatic species (those poor mussels in the Bivalvia class, for example) and underfunded geographies in biodiversity research (such as areas in rural Africa). To bridge these omissions, the Catalogue has introduced the concept of a proto-GSD, which gathers records from many unrelated sources to construct a tentative and intermediary database to temporarily overcome these taxonomic shortfalls.

Management hierarchies are also used to shape external biodiversity databases and to produce new forms of speculative knowledge. These external databases, such as Global Biodiversity Information Facility (GBIF) and the Encyclopedia of Life (EoL), are heavily influenced by the Catalogue's internal and editorial interpretations. In this chapter I outline how one structure, GBIF, integrates the Catalogue into its database structure. The materiality and epistemic influence of these backbones are then explored. Finally, we briefly illustrate some more generative, extensive functions of these hierarchies: their use as knowledge bases intended to predict evolutionary futures based on networked biodiversity data from the past. Such speculative spaces extend the consensus structure's epistemological virtues to its furthest extent, embracing the ever-contingent nature of taxonomic work as a foundation for new modes of hypothesizing work.

This extensibility, I argue, has implications for how we can conceptualize the potential of documentary classification systems in general and how we can reposit how they are positioned as epistemic instruments in information studies.

EXTENSIBILITY

What affordances do consensus structures provide with regard to the production of biodiversity knowledge at the global level? What might we be able to say about the extensibility of the Catalogue versus the extensibility of any other taxonomy? As mentioned in chapter 5, the Catalogue's commitment to access produces an operational and facilitative structure—it helps people get things done. I call what consensus taxonomies can get “done” their *potentialities*, or their *extensive capacities*, and the power they provide to the user (any one user or group of users) an *instrumental power*. In short, consensus taxonomies expand on existing taxonomic knowledge to produce emergent knowledge that exceeds the sum of the taxonomies included in the system. The extensive taxonomic knowledge they produce is understood to be imperfect and is often speculative and temporary.

Looking back to Patrick Wilson's *Two Kinds of Power* (1968), we see that he presents two powers that exist within the bibliographical universe: descriptive power, defined as the ability to describe documents in such a way that we can call up a set of undifferentiated, related documents; and exploitative power, which he describes as the ability to use texts in a manner most relevant to a situational circumstance. A quality of both powers is that they measure what can be performed with documents *within* a system. Further, these powers are predicated on the articulation of a bibliographical instrument (the system) that has particular specifications designed to systematize access to, and maximize, said powers (Wilson 1968, chap. 4). These specifications include the articulation of an instrument's domain and its bibliographical (or documentary) extent, descriptive capacities, logical transparency, and institutional context, among others. The more an instrument meets these ideal instrumental specifications, the more power it affords a user with regard to accessing and using *documents*. The extensibility of a consensus structure, or its instrumental power, similarly depends on these specifications, but they expand on Wilson's original schematic in that this extension focuses on the exploitation of (and access to) the taxonomy system itself, in addition to the documents it contains. Wilson's notion of bibliographical control is concerned with a specific collection of

things (the bibliographical universe), and how we can best situate access to these things. The extensive capacities of a system, on the other hand, tell us how the *system* can be exploited in distinct ways—as a kind of über document—to influence the conceptualization of other biodiversity systems and produce new forms of taxonomic and documentary knowledge.

Seeing consensus spaces more as standardized, repurposable, temporary schema, rather than as composed argument, is an important move for knowledge organizing practice, particularly in the biodiversity sciences. This frame shift redefines the primary assumptions associated with scientific taxonomies—most notably, that they be internally consistent and epistemically coherent. This does not, of course, overshadow the downsides of such a shift (and there are, indeed, many!), which we discuss in chapter 7 and 8. Not only do consensus taxonomies organize classes of things, they also serve as overt interpretive frames that becomes embedded in, and deeply influence, the role and future intentions of global biodiversity work. The flexibility they afford is built into their construction, as part of their design. These instruments provide a structure for other systems that, in practice, actively rearticulates and re-presents the Catalogue's core classification through the lens of its own classificatory expectation and needs—a veritable onion layering of epistemological interpretations.

This kind of extensive activity seems to me appreciably different from the expectations placed on classification systems typically seen in information studies spaces. The Dewey Decimal Classification (DDC) and the Library of Congress (LoC), for example, were built as systems to guide the organization of documents within multiple physical (and digital) locations. The supposition is that these structures are to be centrally controlled to facilitate class consistency, as well as to limit the extent and speed in which the system changes over time. And despite the DDCs “vigorous” policy for change, the rate of change represented in these spaces comes nowhere near that of taxonomies in biodiversity work (Miksa 1998, 25)—and with the rise of genetic taxonomic approaches, we will see the rate of this change speed up dramatically in the future. Further, as systems like the LoC and DDC change, these changes are theoretically expected to trickle down, so to speak, to local repositories, so that they can remain current

with the schedule (we will ignore, for the time being, that public libraries, especially, do not often change book positions on their shelves when such changes occur) (Miksa 1998, 27). On the other hand, while the Catalogue is designed for integration in external systems, its creators understand that those systems may, or may not, ever upgrade their core architecture as the Catalogue evolves, even as the Catalogue continues to make such upgrades easier through the development of persistent identifiers and application programming interfaces. A taxonomic backbone is often seen as a starting point for some external system.

One might argue the position that any classification instrument can be extended. Somebody can, for example, take the DDC classification schedule, download or replicate it, and implement it for their own local purposes, manipulating it as they see fit (aside from whether or not permission was granted or needed for such reuse). This is true, of course; anybody can perform this activity as a kind of imposed power onto the classification instrument. I can take the DDC schedule and do what I wish with it—there is nothing stopping me from doing this. But differentiating this kind of imposed use as categorically different from the use of any document for any other particular purpose would be a fruitless enterprise. An imposed repurposing is not unique to the Catalogue or any other instrument, so it does very little to help us articulate the power of extensibility. The difference lies in the fact that consensus systems like the Catalogue take this repurposing as a core value of their access-oriented commitments, and design the system to facilitate these possibilities. The Catalogue *actively* builds extensibility into its functionality—it is *designed* to be extended into other digital and organizational domains, and so the practices, standards, and policies that arise from within the organization are in support of this facilitative activity. One value that supports this kind of extensive capacity, for example, is that the Catalogue is not trying to control the internal consistency of its taxonomy—conflicts between classes is understood as a possibility, and maybe even an inevitability. Systems like the DDC and LoC depend on classificatory control to maintain consistent descriptive systems over time. A key value of the DDC, for example, is that its classificatory approach must be as consistent as possible over time, lest schedules conflict

with one another. For a local repository to take the DDC classification system and amend it in any number of ways somewhat defeats the original bibliographic and documentary function of the DDC. If a local repository strays too far from the classification standard, updates or revisions to the main classification schedule become irrelevant to the local collection, and thus inhibit the ability to control and exploit the system in the future.

The concept of *extensive power* helps us understand the potential functionalities consensus taxonomic instruments have within digital spaces, as well as their powerful mediative functions. Digging more deeply into the extensive capacities of backbone taxonomies, let us now examine the levels at which new forms of knowledge are produced.

Internal and External Extensibility

The extensibility of consensus taxonomies functions on two distinct levels: internal extension, which focuses on knowledge contained and produced within the consensus taxonomy itself, and external extension, which speaks to the effects on knowledge these structures have on or within external systems (figure 6.1).

First, with regard to internal extension, it is sometimes the case that authoritative global databases are not available for a particular taxon, or

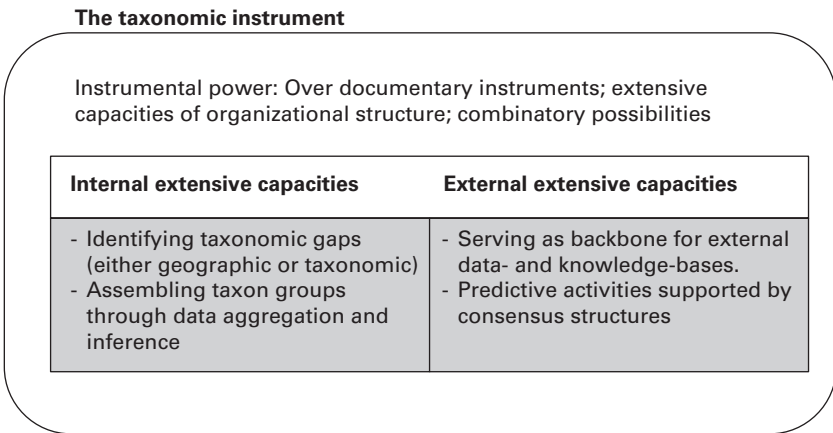


Figure 6.1
The extensive capacities of taxonomic instruments.

data is unavailable for certain global geographies. For example, imagine that a singular authoritative taxonomy did not exist for the genus *Ursus*. Yet, data on *Ursus* may be available in other locations, distributed sporadically in many different contributed taxonomies. In these cases, data is aggregated from these disparate systems and brought together to achieve adequate, or near-adequate, coverage for that particular taxon or area. To best create the most complete taxonomy for *Ursus*, some information might be collected from databases in Mexico, Russia, and Spain to create a semi-complete aggregation of knowledge for that taxon. The resulting internal structure is imperfect at best, to be sure, but yet it still provides a tentative taxonomic view (or, at least, backbone) that allows for the organization and integration of data that exists about *Ursus*. Internal extension, then, is the extension of the knowledge contained *within* the consensus taxonomy itself toward the end of more in-depth and comprehensive taxonomic coverage.

External extension functions in two ways. First, consensus systems provide a data structure (the management taxonomy or taxonomic backbone) and classificatory foundation for organizations and scientists seeking to organize data in online environments. The Global Biodiversity Information Facility, for example, is one global organization that builds on the Catalogue's taxonomy to create a larger data structure outfitted to its purposes. In situations such as these, consensus structures influence the shape and constitution of knowledge in external environments, radiating access-oriented commitments (and the myriad assumptions the edited system represents) throughout the domain of biodiversity work. Second, and perhaps more radically, consensus taxonomies are seen as vital instruments in the emerging discipline of evolutionary informatics, which takes collected information and attempts to answer evolutionary questions and predict ecological and evolutionary trends. External extension is thus the radiant and widespread influence of consensus structures on shaping and modeling the knowledge landscape within the discipline of biodiversity, broadly conceived.

If we are to truly understand the long-term effects of consensus taxonomies on the practice of science, we need to understand how they are implemented to facilitate scientific and knowledge-generating functionalities. I focus our discussion on a select few examples of these uses. Of

course, these structures have many other potential uses—both known and yet to be discovered—but I highlight the most salient uses uncovered in fieldwork conversations. I begin with the internal extensive potentials of consensus structures and, specifically, how taxonomic backbones are used to help survey the geography of biodiversity knowledge and to subsequently fill identified gaps in the taxonomic record. I then broaden my purview, to how these structures influence the articulation of biodiversity knowledge far beyond its systemic borders.

INTERNAL EXTENSION

Understanding the breadth and depth of current biodiversity knowledge is central to the recent articulation of consensus structures. Aggregated structures allow biodiversity scientists a bird's-eye view of the biodiversity and taxonomic knowledge known to science. Such a view affords the ability to find gaps, errors, and overlaps in global data. As previously noted, Mesibov (2010) illustrated the extent to which database errors proliferate throughout the data ecology. Such identifications help scientists fix these errors across a broad swath of infrastructure—a process that would be far more onerous (if not impossible) if attended to at the local level many times over at each database site. Much more will be said of error proliferation—which certainly presents a significant data management problem in its own right—but it should at least be noted that identifying any problem is the first step toward managing its impacts on integrated systems.

The identification of taxonomic gaps is another significant potential of consensus structures. As contributed taxonomies are brought together, what emerges is a partial representation of the natural world: that part of the natural world that has been prioritized by systematic scientific investigation. The gaps that inevitably present themselves arise from two different, but related, problems: understudied taxa, and underrepresented geographies. Regarding the first problem, more-charismatic species—those species that get more attention from scientists and policy makers (Bowker 2008, 146)—have a tendency to get described and classified more exhaustively. Entwistle and Dunstone (2000) show how 70 percent of conservation funding

applications to Fauna & Flora International, a global conservation organization, were for large mammals and birds—taxa that have historically received a great deal of attention. These species then come to symbolize organizational efforts (such as those of the World Wildlife Fund and Save the Tiger Fund), which, in turn, funnel more resources to these species over time (Entwistle and Dunstone 2000). Certain other species groups, on the other hand, such as worms (Annelida) and mollusks (Mollusca), are only sporadically described in biodiversity literature (Kunze, Didžiulis, and Roskov 2013), meaning that they are underrepresented by scientific examination, broadly conceived. In rarer cases, intellectual property disagreements regarding user agreements, data reuse practices, and intellectual control cause taxonomies covering large taxonomic groups to be excluded from the Catalogue. Once a taxonomy is incorporated into the Catalogue, the creator of that database has no control over how the data will be used, which some find (understandably) problematic. Because databases are the result of extensive and long-term individual labor, and because one person (or group) often controls the database, an entire segment of the tree of life is excluded from the Catalogue’s global taxonomy if minute disagreements arise from the Catalogue’s open-source approach to data management. The result of this kind of omission has the same operational outcome: these species do not get their fair share of representation.

Underrepresented geographies pose still another nuanced problem for systems like the Catalogue. Such underrepresented areas often coincide with areas with unstable or very little local scientific infrastructure; such is the case with certain African geographies that have increasingly become the focus of Catalogue data collection efforts. And although select areas of the southern hemisphere, such as Brazil and New Zealand, have well-funded data collection and animal monitoring programs, the data is still sparse in comparison with more robust species and ecological documentation programs in the northern latitudes (Chambers et al. 2017). Additionally, international conservation efforts tend to favor areas of the globe with large numbers of endemic species that are acutely threatened, also known as “diversity hotspots,” which tends to overemphasize these areas in research and data (Conservation International 2017). These hotspots, as noted by

Conservation International, represent over 40 percent of bird, mammal, and reptilian species around the globe, and represent significant portions of Central America, South America, the Mediterranean, and Southeast Asia (Critical Ecosystem Partnership Fund 2017), which means upward of 60 percent of species live outside these zones.

On top of these disparities, the Catalogue must also contend with the species and geographic biases of the Global Species Databases, in that they tend to overrepresent Western geographies. For example, GSDs, which typically focus on one taxon worldwide and contain a taxonomic checklist of all species within it (Species 2000 2015c), are focused primarily in European and North American countries and, as such, tend to amplify species in areas they have easy access to. One reason for this Western focus is that taxonomies and species checklists arising from these areas are bolstered by an integrated and historically strong network of biodiversity institutions and natural history museums—institutions that hold the type specimens and literature central to species identification and descriptive processes (Roskov 2016c).

Historically, countries such as China, Russia, and Brazil have taken a different approach to species data collection, focused on national checklists and taxonomies rather than information that strives for a global reach. The result of these more insular collection efforts is a Regional Species Database (RSD), which poses a special series of problems for consensus systems like the Catalogue.

To reiterate, the lion's share of data ingested into the Catalogue of Life comes from GSDs—these databases provide in-depth and far-reaching global coverage of specific taxa (even if a particular taxa's range is not very large). A genus's range might be relatively small, for example, so a global database in this case does not mean that the species are found throughout the world. These databases are professionally vetted and relatively complete, given their descriptive and analytic focus. But not all species are covered to such depth in GSD environments, nor do GSDs have access to each and every part of the globe. RSDs, then, can serve as potential data sets to fill the gaps left after the collocation of GSD data. The Catalogue produces an intermediate taxonomic space composed of RSD data, which they

call proto-GSDs. These proto-GSDs account for these taxonomic gaps, but two major problems present themselves when integrating RSD data into spaces dominated by GSDs: there is often overlap with GSD sources, meaning reconciliations must be debated, and RSDs cannot ensure full and equal global coverage for all species represented in the database, meaning that gaps, to a certain extent, often still persist.

By conglomerating RSDs, proto-GSDs attempt to manufacture taxon lists to the best extent possible given available information. To accomplish this, editors will often compose taxa through both automatic and manual efforts. Such activities, however, present some major complications, such as conflicting taxonomic orientations, duplicated taxon names, and the inevitable conflicts between species concepts of the same name. As covered in chapter 4, two exactly similar names can reference an entirely different set of literatures and type specimens, meaning their circumscriptions can be in conflict with each other (Matthias 2013). This concept differentiation problem is evidenced in the following excerpt from the Catalogue of Life blog:

Take for example the Family Gentianaceae in the Plant Kingdom. This family of plants has an estimated 1650+ species worldwide. We have two current suppliers of Gentianaceae to the Catalogue of Life—[the] ITIS Regional database and Catalogue of Life China. Together they supply the Catalogue of Life with 552 species in addition to 82 infraspecific taxa. The species *Gentianella acuta* (Michx.) Hultén appears in both checklists where it is a synonym in ITIS Regional and an accepted name in the Catalogue of Life China. This is because some of the species of Gentianaceae are cosmopolitan (i.e. present in North America and China) and the taxonomic concept (i.e. accepted name or synonym) is different. To combine the datasets the Catalogue of Life editors had to resolve these issues before publishing it in the Catalogue of Life. Gentianaceae is now part of a “proto-GSD.” (Matthias 2013).

Such taxon negotiations only serve to complicate an already complex consensus space. First, how does one decide if *Gentianella acuta* (Michx.) Hultén is a synonym or an accepted name? This partly depends on the currency of the source database and its reputation. But this process cannot be automated, since each conflict is idiosyncratic and requires expert

taxonomic opinion. Second, moving taxa from one area of the taxonomic tree to another is not a simple editorial task. Recall the process of establishing nomenclature and how, to a certain extent, these names reflect taxonomic *positions* (the first part of a name indicates a generic position, while the second part a specific). If a species is moved into another genus, the first part of the binomial name needs to be changed to reflect this modification. If the Catalogue were to attempt to do so, the new name would qualify as a “new combination” (GNA 2020a) and, to be accepted by the taxonomic community as a valid nomenclatural act, it would need to be published in accordance with the particular code governing that group. Aside from causing undue confusion in the taxonomic community (changing a name for merely consensus purposes rather than for descriptive purposes would be antithetical to normal nomenclatural practice), the process would be incredibly time-consuming and expensive. Not practical and outside the operational bounds of the Catalogue even if the resources were available.

The results exhibited by proto-GSDs are promising, even if the aggregation is fairly primitive at this point and not yet scrutinized by expert communities (Roskov 2016c). At minimum, these proto-GSDs tentatively allow the collocation of taxon data that would otherwise be completely excluded from global systems. These mechanisms effectively increase the Catalogue’s instantiative power (as in, it has the ability to instantiate species concepts that would have otherwise been silenced in the documentary realm). In some cases, taxa once represented at only a 10 percent level in the Catalogue have been increased to about 50 percent coverage with proto-GSD merging. These statistics are not insignificant, especially for a taxon that is understudied and perhaps, even ecologically threatened. Such inclusion could promote more conservation efforts or more research on a particular taxon. Of course, editors are limited in terms of their own taxonomic expertise, and given that RSDs straddle many different taxa, proto-GSDs are far less authoritative than the GSD segments of the hierarchy. This leads to a great deal of understandable skepticism for biodiversity taxonomists who consider this kind of analysis both misleading and partially researched. And within the bounds of the Catalogue, the limitations of these proto-GSD spaces may not be fully visible or understandable.

Proto-GSDs are difficult to build. Algorithmic software designed to flag taxonomic conflicts and recommend master classifications between overlapping structures are promising and have the potential of releasing this burden from individual taxonomic editors. Taxonomy alignment software such as the Euler Project are also gaining in credibility and use (Chen et al. 2017). Larger backbone systems, such as those built by GBIF, enable algorithmic reconciliation of both names and taxonomic ranks (Döring 2015). Yet, despite early attempts at building this automated infrastructure, such software mechanisms never materialized as part of the Catalogue's workflow, partly because of the initial commitment to a more manually curated taxonomic structure. However, as of 2021, no specialized software or interface has been created with which the Catalogue can perform these aggregating tasks. All the editing work is being done directly within the database, using the relational database itself (Roskov 2016c). Often, doctoral students or postdoctoral academics are performing this proto-GSD work, but this approach is costly, and such temporary work is difficult to sustain in the long run, given its dependence on soft money allocations (Schalk 2016b). More often than not, manually building proto-GSDs is undertaken by the Catalogue's executive editor.

Despite these proto-GSD efforts, there is no aggregative practice generally accepted by all taxonomic specialists. Many "professional taxonomists do not appreciate any kind of technical exercises," such as automatic taxonomic building, Yuri Roskov indicated (2017). "For example, a plant list . . . is being built using software by [the] Missouri Botanical Gardens . . . where they are trying to merge regional floras. I spoke to professional taxonomists . . . and they are very much skeptical about this work. It is a very political process" (2017). For the Catalogue, such interventions are better than nothing, however, since its structure is aimed at temporary snapshots of extant knowledge that will inevitably integrate input and get refined over time. It boils down to the Catalogue's commitment to access and the belief that temporary taxa are better than a structure with gaps.

If the goal is to get an up-to-date snapshot of the world's biodiversity in one coherent structure, the Catalogue will do all that it can to achieve this coverage, even if the mechanisms are temporary and imperfect. A major difference between *taxonomic opinion* and *management hierarchies* is this ability

to accept flexibility and contingency as part of an epistemic foundation. The Catalogue's taxonomy is a practical tool that understands and embraces its limitations. As I was told on numerous occasions, it is better to embrace taxonomic disagreement than to wait for a professional taxonomic consensus that will never materialize. The extensibility of the taxonomic space, by way of filling gaps in the taxonomic record to increase the access-oriented powers of a consensus space, has the potential of building structures far more robust than they otherwise would be. This is especially important for external organizations that use the Catalogue as a core element of their data hierarchy. The external functions of the Catalogue are far more generative and provide the maximum amount of extensivity in the biodiversity infrastructure environment.

EXTERNAL EXTENSIBILITY

While the internal extension of the Catalogue increases the breadth of taxon coverage within the system, the external capacities of these backbones radiate outward far more impressively. An essential access-oriented aspect of consensus structures is that they serve as a modeling mechanism for data all over the world. One might think of this kind of use as *taxonomic amplification*, defined as the use of a composite taxonomy to collocate disparate data from multiple sources and to produce taxonomic knowledge by the appropriation and recombination of a taxonomic structure, harkening back to Gerald Guala's concept of synonymic amplification (2016). The Catalogue is intended to support the shift from local practices (hypothesis driven, internally consistent) to global spaces that rely on the embedded standardization for full-scale integration within global infrastructure to support "a quantum increase in the coherence of the world's biodiversity data and analyses" (Species 2000 2015a). And this use of classifications—and the way these systems change because of it—is vitally important to understand.

Derek Langridge (1992) described some of the roles and applications of classifications, including facilitating searching and arranging documents. Extensibility is far more invasive in terms of how it intersects with the user-interface environment, mostly because taxonomies are the foundation

for all data management. Langridge makes the point that computers do not eliminate the necessity for classification, they merely provide a more sophisticated and effective means to create more nuanced (and obfuscated) classifications (1992, 70). Indeed, this statement still holds true, for the importance of classification, especially with regard to databases and computational modeling environments, continues to rise even more than twenty-five years after Langridge's publication. If databases have done anything, they have forced us to represent and classify *everything* to the minutest detail so as to facilitate the transfer of knowledge within computational spaces. The existence and subsequent widespread use of biodiversity consensus structures is a testament to such a fact, as is the extent to which they structure the understanding and representation of global knowledge. If, as Jerome McGann (2001, xi) states, "our minds think in textual codes," we might also extend this statement to declare that the public and social conception of biodiversity knowledge "thinks in taxonomic codes" through the lenses of these networked structures. And, more often than not, these taxonomic codes are consensus-based. If this is the case—and my contention is that it is increasingly becoming the standard—we need to think more about what this means for the agents and organizations that put these standards to use, and how we might theorize the kinds of changes that are taking place as part of this integration and rearticulation.

A good place to start this line of inquiry is with Joseph Tennis (2015), who postulates the existence of three essential theoretical approaches to the study of classification that define the kinds of work being performed in this domain. Tennis's categories are foundational classification theory, first-order classification theory, and second-order classification theory. First, foundational classification theory, as defined by Tennis, is "concerned with philosophical and definitional aspects of classification" (2015, 246). A. Broadfield's (1946) work is an example of such an approach, as would be the categorical and philosophical expositions of Bliss, Langridge, Richardson, and Wilson. Larger questions in this arena include ontological and epistemological questions about the a priori assumptions structuring organizing systems. One might ask, as part of this level, What are the epistemic differences between traditional and consensus biodiversity structures?

First-order classification, Tennis's second category, is "solely concerned with the methods of classification scheme construction and use" (Tennis 2015, 245). This definition helps us understand the nuts and bolts of system construction—how we build classifications and articulate the processes we manufacture to produce them. Previous chapters expanded on such first-order and also foundational approaches: how the Catalogue conceives of and constructs evidence for concepts, how such concepts can be connected within a historicized nomenclatural system, and how management classifications build their schematics in contradistinction to traditional taxonomic forms.

Third, and most pertinent to our current discussion, is second-order classification theory, which is concerned with the use, reuse, and manipulation of systems once they have been completed (Tennis 2015, 246). With regard to the Catalogue, we are interested in how its management hierarchy is used in various contexts and how that intended use fundamentally changes the Catalogue's composition—and impacts the production and representation of knowledge as it becomes a vehicle for data in both anticipated and unanticipated spaces. Second-order classification is then broken down into three subcategories: (1) how schemes change over time and how we update them, (2) how installed schemes interoperate, and (3) how systems change when they "change context (reapplied or reengineered)" (Tennis 2015, 246). Each of these elements is pertinent to biodiversity databases, but the third is most relevant in relation to the extensive qualities of the Catalogue.

Understanding how the Catalogue is repurposed, reengineered, and changed is absolutely critical if we are to apprehend the impact of the Catalogue on the shape of public knowledge. Modifying a classification enhances the system in some attributes and limits it in other ways (Tennis 2015, 246). Many of the database systems in the biodiversity world, for example, implement the Catalogue's taxonomy to some extent or another in order to organize and present its content. Does this matter? My contention is that it should, for anyone who works in database systems and classification. In the case of the Encyclopedia of Life, the Catalogue's consensus taxonomy is displayed as one option among an array of taxonomic

approaches, alongside taxonomies contributed by National Center for Biotechnology Information, Barcode of Life Data Systems, and Wikipedia. GBIF's case is particularly interesting, given that their taxonomic backbone builds on the Catalogue's hierarchy to organize all the data compiled from sources that may or may not fit within any given consensus hierarchy. Let's expand on the case of GBIF.

The GBIF Backbone Taxonomy builds on the Catalogue's management hierarchy to organize all data uploaded into its management system. GBIF collects a vast array of data types, including occurrence records and nomenclatural data, as well as genetic data from sources such as the NCIB. This is a functionally different emphasis than that of the Catalogue. The Catalogue's primary function is to provide valid and accepted name tokens (as a nomenclature) and the subsequent connection of these tokens into valid taxon groups. GBIF's main function, on the other hand, is to collect globally produced occurrence data and provide a point of access for data points that can potentially differ in scale and that may or may not point to an accepted species name (GBIF 2020). Validated, curated names and taxonomic relationships are the primary focus for the Catalogue, whereas in the GBIF environment, the occurrence data takes priority, which is appended to the best-fit taxon level to aid in data access and use. This isn't to say that GBIF isn't concerned with validated names and most-current classifications (it, of course, ultimately sees this as a vitally important part of the process); this is just to say that the main emphases of each organization influence how they approach the process of building classifications and what concessions they are willing to make in the process of doing so. The Catalogue initially pushed against the use of algorithmic methods, for example, whereas GBIF embraced them—and to generally positive ends as conveyed by many biodiversity scientists. These approaches have changed in recent years, as collaboration between GBIF and the Catalogue has increased—the Catalogue of Life Plus emerged from a partnership between the two.

The Catalogue represents a significant percentage of the name and taxonomic information for GBIF's Backbone Taxonomy, comprising some 3,175,925 names, or approximately 54 percent of the total GBIF namespace ("GBIF Backbone Taxonomy—Constituents" 2017). However,

even with the Catalogue's impressive 2 million species, the information falls short for GBIF's purposes (GBIF.org 2016). As a GBIF staff member conveyed, no global taxonomy can organize the more than 750-million-plus occurrence records that exist within the system. Some data is submitted with names not yet validated, and other data is formulated as the output of genetic analysis that may or may not be yet described at the species level. Information of this nature may or may not find a location in an inherited hierarchy, which means that GBIF needs to articulate methods to bridge these systemic differences. And so, they build on top of the Catalogue management hierarchy. In addition to the Catalogue, other trustworthy checklists are collected, including, for example, the International Plant Names Index (IPNI) and Index Fungorum (2021), to make an even more encompassing taxonomic whole. To accomplish this, GBIF has produced fine-tuned algorithms to offset the labor of this kind of work, which overlay hierarchies over hierarchies that are perfected across time and freely available on GitHub (Robertson 2016; GBIF 2017a, 2017b). And because GBIF gets a high degree of use, feedback is ongoing and iterated into new versions of the algorithms code (Robertson 2016).

Of course, this says nothing of the critiques one might levy against algorithms as organizing mechanisms, which Safiya Noble has written about extensively in recent years (2018). Many scientists feel that such methods are not as refined as taxonomies assembled through individual mediation, in addition to the more political concerns regarding the substitution of computational methods for hard-earned taxonomic expertise. Taxonomy is seen in many circles as a dwindling career, in part because of computational methods. The extent to which these taxonomy-building algorithms can remain effective is ultimately dependent on the quality of individual data sources. As Roderic Page, former science director for GBIF and professor of taxonomy at Glasgow University, has asserted on Twitter, "some data [GBIF does] have is poor (e.g., @catalogueoflife has mangled butterfly names)" (2016c). In addition, many of the data sources are "[aggregations] of sources that [may] themselves be aggregations," making it incredibly difficult to fix the data at the source (Page 2016a, 2016d). Minute editorial miscalculations inherent in the Catalogue of Life have to

be debugged and restructured downstream to meet various infrastructure-specific needs.

Within this collocation activity, information conflicts inevitably emerge. The hundreds of millions of circulating records that find their way into GBIF collectively constitute an environment wherein species concepts, data types, data granularity, and taxonomic hierarchy formats may be in mismatch with one another, both semantically and syntactically. Data quality varies across the platform. Certain professional scientists and biodiversity informaticians, including Roderic Page, have indicated that perhaps GBIF should “take more ‘ownership’ of data quality, but that’s politically tricky” (Page 2016b). True enough, indeed, as all data mediators should. But then again, one needs to weigh the benefit of access-oriented systems against that of those that strive for curatorial perfection. After all, to practice science and produce valid results, one needs data as a starting point. All taxonomic work is iterative, so consensus structures are certainly not immune to this fact.

EPISTEMIC INFLUENCE

What can we make of the Catalogue’s impact on the organizing of global data? Even if the GBIF’s nub taxonomy comprises many independent systems, given its overwhelming presence in GBIF, it is undeniable that the Catalogue impresses quite an impact on the constitution of knowledge in this space. Regardless of whether the Catalogue is a visible aspect of a system such as GBIF, it behooves us to unpack how the assumptions that underlie its construction deeply influence the shape of knowledge throughout the biotaxonomy environment and beyond. Web-based media are rife with representational and graphical structures, and biodiversity database expressions are particularly laden with mapping software that expresses data in easily consumable formulations. GBIF is powerful in this respect—the opening webpage displays a marvelous map-based analysis of its data that includes data density by geography, year, and the basis of a record (observation, literature, preserved specimen, and so on) (GBIF 2017c). Yet, as Johanna Drucker (2014b) reminds us, the space between data and the

graphical interface is filled with interpretive acts—acts that begin with the initial construction of data structures that organize how knowledge should and can be displayed for ready access. One of the primary goals of unpacking and theorizing structures like the Catalogue is to narrow this interpretive gap between data structures and the mediated mechanisms by which we access those data in precoordinated graphical environments.

It goes without saying that the usefulness of these sites as tools that facilitate biodiversity work is powerful to certain extents, so the purpose of such a critique is merely to shed light on the machinations that make *all* taxonomies useful, to highlight how our results might be weighted in favor of certain epistemological values—and, in this particular space, the Catalogue stands as a central mechanism by which those values are constructed. This kind of deconstructive taxonomic work is particularly important given the extent to which sites like the Encyclopedia of Life are implemented in educational environments, illustrated by the integration of educational resources aimed at a primary-school level audience. But given the focus of other biodiversity data platforms such as GBIF and the International Barcode of Life, the impact of this consensus work extends beyond the classroom, affecting scientific research in a wide array of subdomains. GBIF has attained such widespread use that the question of how it is made is often overlooked, as data points are downloaded, referenced, and integrated into the practice of analysis, interpolation, and interpretation.

As a case in point, I had a meeting with an evolutionary biologist studying the sleep patterns of biological organisms. Their team was using GBIF data to fill in some ancillary data elements that weren't part of their original collected data sets. In this case, the research group had collected primarily individual, organism-level information (sleep patterns for one organism, for example), and such individual records needed to be placed alongside other species records to create a kind of ad hoc taxonomy depicting sleeping behavior. To create this map, species-level data was needed, to connect individual specimen data with larger species range and distribution trends. This data also helped scientists map their data using GIS software. And although this particular individual knew that taxonomic errors were possible in the case of any platform—and indeed, in this case, data was

confirmed at the data source to confirm its validity—they were intrigued to learn about the various layers of aggregation, consensus work, and inter-polation that were working together to create the robust knowledge structure represented by the Catalogue, and, to a certain extent GBIF. Taxonomies become invisible to users, so making them visible helps one understand their capacities—and, more importantly their limitations—in structuring easy and ethical delivery of information. An undergraduate biology student just beginning a research career would likely not have been as savvy as these scientists—which says nothing of the lack of expertise a politician or journalist might bring to the table.

The materiality of consensus taxonomies, then, becomes a central point of examination as we think about consensus structures and their impact on the production of new forms and domains of knowledge. In Karin Knorr Cetina's work, especially *Epistemic Cultures*, she offers a depiction of scientific knowledge societies that are fundamentally tied “to the machines deployed in knowledge production” (1999, 11) and notes that the scientists who perform the process of knowledge-creation are derivative agents to the instruments and tools that structure and facilitate their examinations. In Cetina's analysis, instruments limit the capacity of what can be done within a laboratory setting, but perhaps more importantly, these central objects create a bounded, ontological reality that communities become embedded within as part of daily practice. How these instruments are deployed both rhetorically and materially has fundamental impacts on how the cultures that employ them arrange themselves—consciously or not—but also on how they construct a narrative for the production of scientific activities and analysis. Scientists relate to machines quite differently in the domain of high energy physics, for example, than within the space of molecular biology. In the former, instruments take on a symbolic quality and become the internalized, closed, and primary ontological system of the lab against which scientists define their intervention. In the latter, however, the focus is not on a symbolized system, but rather on analysis of the external natural world—mice specimens, for example—that scientists then control toward the goal of knowledge through experimentation. The “real” of a laboratory (the space of primary intellectual focus and interest) is created, in part,

from how instruments embody and build a narrative that constructs the way an investigation can take place.

My goal in bringing up Cetina's work is to illustrate that systems (tools, machines) like the Catalogue have the capacity to change the nature of particular scientific work based on the suppositions these systems make about the "real" objects they purport to organize. The existence of consensus structures allows scientists the freedom to perform new kinds of practices based on an entirely different set of expectations about what these taxonomies are trying to say about the natural world. In the case of description-based taxonomies, the "real" biological world is presumed to be mimicked (perhaps idiosyncratically, and certainly hypothetically) within a closed system that functions under a uniform set of suppositions about how to define a taxon and how that taxon is related to another based on a series of metrics. The "real" in the case of retrieval-oriented consensus spaces is markedly different, and far more speculative and processual. Consensus structures mimic an amalgam of opinions about how we can potentially classify the natural world through the process of mediation. The focus becomes not a reflective, empirical model of the natural world, but the processes, practices, and struggles for instantiative power among many competing taxonomic spaces. And so, as individuals, systems, and organizations integrate consensus data into their workflows, these underlying suppositions must be, first, acknowledged, and second, negotiated and rendered transparent through the access systems they employ.

With this new kind of consensus-based epistemic influence, new forms of knowledge can arise that build on these suppositions, most notably new initiatives focused on using repositories like the Catalogue for predictive purposes.

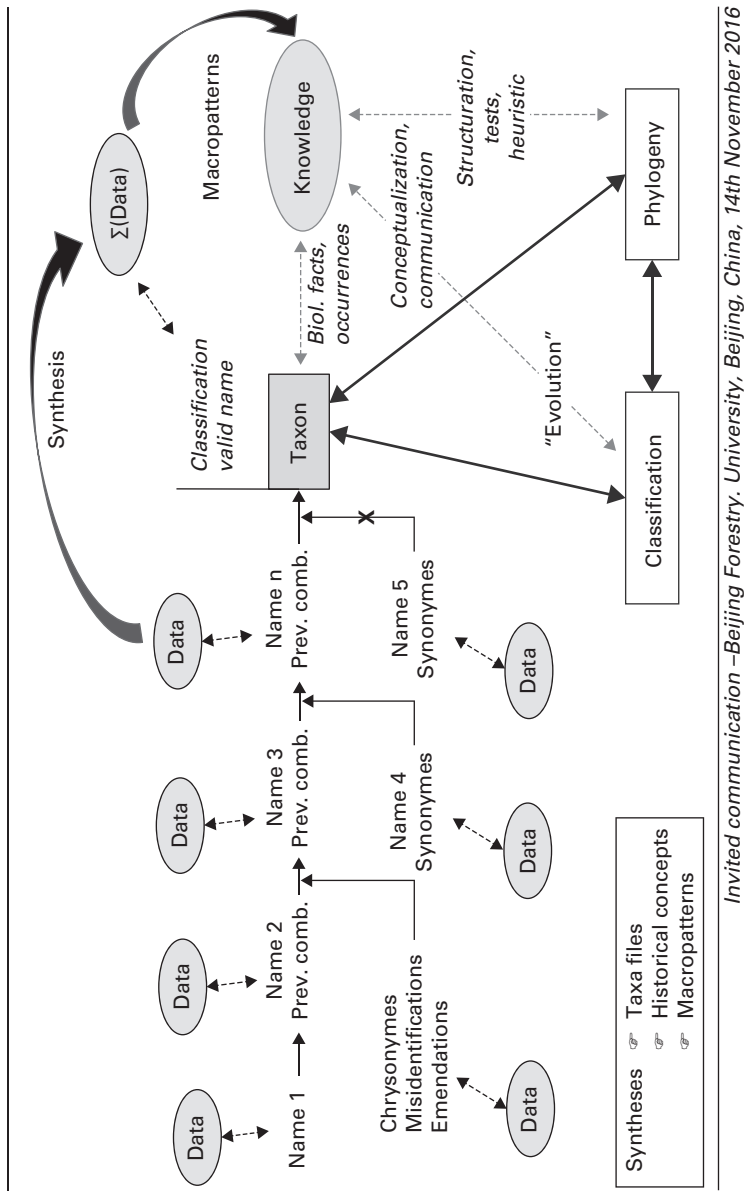
HORIZONS

A more experimental, but nonetheless powerful, extensive ability of consensus systems is their potential use as a springboard from which to assess larger trends in the taxonomic world. Nomenclatural networks, designed as they are to trace the historical connection and development of concepts

via a system of internal taxonomic relationships, provide a rich space of knowledge about our ever-changing understanding of the natural world. As systems like the Catalogue are integrated into backbone systems such as those created by GBIF (which is a proper consensus system in its own right), these consensus structures organize a huge variety of information about species, including geographic distribution, migration patterns, ecological and species relationships, genetic information, and more. Collectively, these data have the potential to help scientists—and especially evolutionary biologists—understand historical trends regarding how scientists have, historically, documented and organized biodiversity. And such historicity can potentially help predict some of these attributes.

As a result, these repositories need to be able to capture the full scale of taxonomic opinion over time to appropriately model evolutionary patterns in a multitude of ways. As taxonomist Nico Franz has noted, “The real-life challenge for these information repositories is to capture more than one authoritative classification; they are built to represent the full spatial and temporal dynamic of the taxonomic process” (2005, 499). Capturing this dynamism is essential, especially since taxonomic production (and the interpretation of the species in the natural world), as a scientific practice, is constantly changing. The infrastructures that we build to support and document these processes should be able to mirror the various transmutations of taxonomic opinion and representation that happen over time.

Building on this, a key extensive capacity of consensus systems is to provide the capabilities to extrapolate “macropatterns” (Bourgoin 2016) using the documentation of historical concepts within the database structure, as well as the associated data that is mapped to these species concepts. In figure 6.2, Thierry Bourgoin, faculty member at the Muséum National d’Histoire Naturelle in Paris, illustrates a possible schematic for how we can think about knowledge bases as they relate to structures such as the Catalogue. Bourgoin makes a distinction among classification, phylogeny, and evolution, as discrete layers needed for the production of knowledge bases. Within this schema, the production of a classification is a preliminary step, which depends on the network of name associations lying at the core of most biodiversity platforms. These names are then synthesized into



Invited communication – Beijing Forestry University, Beijing, China, 14th November 2016

Figure 6.2

From taxonomic databases to knowledge bases: understanding evolution. The goal in a system of this nature is to (a) take name data that is synthesized into validated taxon forms and (b) use these taxa as the building blocks for classifications that can then be used to articulate many possible phylogenetic hypotheses.

Source: Thierry Bourgoin (Bourgoin 2016). Used by permission.

accepted taxon groups, which can then be iteratively placed in any number of classificatory arrangements. The ultimate, long-view goal is for platforms such as the Catalogue to confront phylogenetic and ecological questions so as to better “understand evolutionary [concerns]” (Bourgoin 2016). Such a broad and extensive approach to speculative work of this nature, however, requires a networked set of coordinated infrastructures, each with their specific data contribution to this broader whole.

Biologists and informaticians have termed this kind of macro-level analysis, among other monikers, *evolutionary informatics*, which “concerns the capturing, storing and integrating of all these data (about biological specimens, images, genomes, etc.), as well as developing the analytical techniques that use them to answer evolutionary questions” (Parr et al. 2012, 94–95). Platforms such as the Catalogue of Life stand at the center (Parr et al. 2012, 100) of this informatics integration, either as part of the taxonomic backbone to platforms like GBIF or as a source of valid names and synonyms by which all the data within a system is collocated. This makes the Catalogue a key, but certainly not the only, player in the constitution of any knowledge base. But the biases engrained in the Catalogue’s construction must be negotiated and acknowledged in these integrative environments. As Parr et al. (2012), indicate, an endeavor such as this collocates many discrete data types, standards, and integrated repositories in a way that allows them to function as a unified whole—but also in a way that has not yet been perfected in biodiversity work. Aside from those organizations that are part of the iLife consortium, the authors identify a broad number of platforms that must integrate for such predictive functionality, including the Map of Life (2018), which provides geographic distributions; GenBank, a repository for gene sequences; and MorphBank, which offers image annotation functionality. This entire ecology of historical documentation provides a sandbox of potential and emergent knowledge products beyond the sum of the database content itself. Such work recombines data in ways that can provide long-term predictions, climate and species modeling, and simulation (Landers 2016).

Large-scale coordination of this nature takes many years to articulate and perfect. For starters, the success of this scale of data integration will

involve the careful documentation of how taxon concepts, and the taxonomies they are embedded in, change over time. Capturing this contingency in some useable, structured, and representationally sound manner is of utmost importance. It has been acknowledged that the long-term prospects and success of the Catalogue depends, in one sense, on the system's transformation from one that organizes a stable taxonomic backbone that facilitates data sharing to a system that can map the changes that are taking place over time within its own database, as well as among the databases it inherits as part of its backbone taxonomy. The complications involved with the historical tracing of this kind of change, however, far exceeds the capacity of the Catalogue alone. The successful prediction of evolutionary trends will rely on the successful evolution of biodiversity practices, in which coordination and collaboration are central tenets of their mission and values. This is precisely why consensus structures are such important elements in biodiversity work today: they push the practice of taxonomic work beyond its traditional epistemological boundaries, and in doing so, support the articulation of new generative spaces of scientific speculation.

Trying to accommodate a complex arrangement of multiple taxonomies, as well as document how those arrangements evolve over time, is no easy task, particularly when such complexity must somehow be fixed in database fields that are relatively limited in their composition. Implementing knowledge bases is, in many ways, pushing the boundaries of what consensus spaces are situated to perform. As Johanna Drucker conveys, “knowledge forms are never stable or self-identical but always situated within conditions of use” (2014b, xiv). If we embrace taxonomy as an iterative process, then the platforms and structures that we use to collocate this knowledge should be seen as in-process. Consensus spaces are tied to the beliefs of a particular space and time—as are any organizing structures. Indeed, even the predictive knowledge produced from these systems is speculative and hypothesis-oriented, not gospel. But neither is any descriptive taxonomic system a final opinion. The motivational core of consensus systems is not that they strive for perfection or epistemic constancy, but that their imperfections are to be completely on display, so that the users

of these systems can use data responsibly and with a keen understanding of what knowledge can be produced from these platforms.

It would be a mistake, I think, to assume that documentary classifications such as the DDC do not also radiate their epistemic and structural capacities in spaces far beyond their boundaries in ways that are not dissimilar to the Catalogue. The previous discussion on scientific warrant unearthed some of these effects—library systems also are built with particular values and have material consequences. Thinking about how taxonomic knowledge is being enacted, negotiated, and compromised within consensus biodiversity work can perhaps help us reformulate the very foundation of what we mean by *control* in documentary classificatory spaces within information science. Should we rethink the potential of instrumental powers in light of lessons learned from the Catalogue? Perhaps our systems need to provide the flexibility to design organizing systems not only for consistency and stability, but also that are understood as emergent, contingent knowledge in their own right. The *epistemology of contingency* has broadened the horizons of scientific taxonomic work—a domain with roots far deeper in history, and more rigidly theorized, than documentary classification in information science. How might we in IS reimagine our systems in light of these developments? Would such an approach increase diversity of opinion within taxonomic spaces? For now, however, let me table this argument, and turn our attention to a critique of consensus systems, which will place us in a much more effective position to judge their possible efficacy for documentary classifications.

The domain of information science has seen a resurgence of literature that critiques organizing systems as biased, culturally situated, and potentially suppressive and oppressive infrastructures (Olson 2002; Adler 2017; Noble 2018; Furner 2009a; Bowker and Star 1999). These critiques have been useful toward conceptualizing how information science can construct classifications that are more robust and attentive to the needs, sensitivities, and realities of a diverse set of users. It is only fitting then, now that we've talked about the virtues of consensus systems, that we assess the limits of these hierarchies. It is certainly true that consensus systems are not universally accepted in practice. In fact, some professionals feel they do more

harm than good to the biodiversity taxonomic world, undermining the processes, assumptions, and use of description-oriented systems, which are crafted through hard scientific work and empirical analysis. For the field of information and library science, the implications for these critiques are broad and deep. Critiques levied toward scientific consensus systems can help information scholars identify similar weaknesses in documentary-based consensus systems. Secondly, these implications can help rearticulate information systems that are specifically built for flexibility and that strive to manage consensus while remaining attuned to the multiple interpretative layers built into their design.