

MIT Press • Open Encyclopedia of Cognitive Science

Organizing

Robert J. Glushko¹

¹University of California, Berkeley

MIT Press

Published on: Jul 16, 2026

DOI: <https://doi.org/10.21428/e2759450.c31862e7>

License: [Creative Commons Attribution-NonCommercial 4.0 International License \(CC-BY-NC 4.0\)](#).

Organizing is a fundamental and pervasive activity. We organize things, we organize people, and we organize information. We also organize information about things, people, and information itself. As the scale of organizing increases, businesses, governments, and institutions systematize models like these with more precise semantics to support efficient action and interaction. Many rely on software that ingests data, sorts it, and deletes it to facilitate these organizing activities. An abstract, unifying perspective treats organizing as a design problem guided by three common questions across domains: *What* is being organized? *Why* is it being organized? *How* is it being organized? The answers to these questions define an *organizing system*—an intentionally arranged collection of resources and the interactions they support. Organizing systems can be physical or digital, personal or institutional, or rigid or adaptive, but all of them share the goal of enabling purposeful action with the resources they organize.

History

Thinking about organizing in terms of what is being organized is deeply ingrained in language and culture. The English language includes hundreds of collective nouns like library, museum, pantry, dataset, family, army, and swarm. Each of these names a way of fitting individual resources into a structured model. Organizations also develop domain-specific vocabularies to describe their organizing activities: libraries acquire, catalog, and weed books; businesses hire, promote, and fire employees; and sports teams draft, bench, and trade players.

However, this ubiquity and pervasiveness of organizing, and the rich vocabulary used to describe it, come at a cost. Reinforcing domain-specific characteristics and activities and highlighting their differences makes it harder to recognize what all organizing has in common. Yet, there are broadly applicable concepts and methods for designing and implementing organizing systems according to the idea that organizing is a design problem ([Glushko, 2013](#)).

To fully understand organizing as a design problem, it helps to trace how organizing evolved from a pragmatic activity into an intellectual pursuit. In early human societies, organizing was primarily practical—tracking seasons, storing food, arranging shelter, and coordinating group activities. These practices were embedded in ritual, custom, and environment and likely not conceptualized as explicit design decisions. Much of this organization was influenced or constrained by natural or astronomical time units and cycles like days, seasons, and years, and we still organize our lives this way to a great extent today; consider that the time unit of *month* names the astronomical body whose cycle defines it.

The shift from practice to theory began in the classical world. Philosophers like [Plato \(375 BCE/2007, 370 BCE/2009\)](#) and [Aristotle \(350 BCE/1937, 350 BCE/2007\)](#) articulated foundational ideas about categories, hierarchies, and logical structure. Plato sought to explain how the living world was organized into species with his metaphor of “carving nature at the joints,” using clusters or correlations of feature values to discover categories. Modern data science describes this in algorithmic terms as a form of “unsupervised learning” that creates categories that maximize within-category similarity and minimize between-

category similarity ([Hinton & Sejnowski, 1999](#)). Aristotle's taxonomy of animals and his theory of the four causes or explanations for why things exist and behave as they do laid the groundwork for viewing the organization of knowledge as systematic and analyzable.

Some of the earliest examples of organizing as an institutional and intellectual activity can be found in the creation of ancient libraries and museums. The Library of Ashurbanipal (seventh century BCE) and the Library and Museum of Alexandria (third century BCE) organized knowledge for preservation and scholarly use. These institutions demonstrate that even in antiquity, organizing systems emerged with recognizable design features; resource descriptions, physical arrangement, and curatorial purpose can also be found in institutions that still exist today. The Bodleian Library at Oxford, established in 1602, is one of the oldest libraries in continuous operation and has long served as a hub for cataloging and classification practices. The Ashmolean Museum, founded in 1683 as part of Oxford University, is the world's first university museum.

During the Enlightenment, organizing took a secular and increasingly empirical turn. Figures like Carl [Linnaeus \(1735\)](#), Denis Diderot ([Diderot & d'Alembert, 1751](#)), and Francis [Bacon \(1620/1902\)](#) attempted to catalog the natural and intellectual world. Knowledge was now being organized for rational understanding, public education, and utility. Linnaeus exemplified this turn with his hierarchical system of biological classification, introducing a binomial nomenclature (e.g., *Homo sapiens*) that remains foundational today.

Meanwhile, Adam Smith introduced the idea of the “invisible hand” ([Smith, 1776](#))—a metaphor for how individual actions in markets could produce coherent systems without central control. This was a revolutionary shift in understanding organizing as something that could emerge spontaneously rather than being imposed top down.

Herbert Simon later reframed organizing as a design problem in his seminal work *The Sciences of the Artificial* ([Simon, 1969](#)). He emphasized bounded rationality and the importance of making choices that transform the current state into a preferred one, defining organizing as a deliberate, iterative process of system shaping. This tradition carried forward into the industrial and information ages, in which organizing was increasingly seen as a sociotechnical and institutional function visible in everything from library classification and bureaucracies to computer science and data systems ([Bowker & Star, 1999](#)). Building on this trajectory, *The Discipline of Organizing* ([Glushko, 2013](#)) synthesized these strands by proposing a transdisciplinary framework that treats organizing as a generalizable, analyzable design method that can be applied to any domain.

Core concepts

The definition of an organizing system implies three design questions, each embodying a core concept:

- *What* is being organized? What are the resources?

- *Why* is it being organized? Every organizing system must enable users to interact with its collection of resources.
- *How* is it being organized? The possible interactions depend on the nature and extent of the descriptions associated with the resources. How do organizing principles use these resource descriptions to intentionally arrange the resources?

A lifecycle perspective about organizing systems reveals that every organizing system goes through five phases or activities during its lifetime to make and implement these three essential design decisions:

- Scoping the resource domain and selecting the resources to be organized
- Identifying requirements for organizing and interacting
- Organizing resources
- Designing interactions with resources
- Maintaining resources and their organization

In simple organizing systems, these activities are informal or implicit, but in complex organizing systems, it is essential to perform them systematically and often iteratively.

Resources

A resource is anything of value that can support goal-oriented activity. In different domains, a resource might be called an object, item, instance, document, artifact, asset, or entity or given more domain-specific names like “tree” in an orchard. A critical question about resources in the design of an organizing system is the scope and scale of what is being organized. Systems with a large scale have many instances of the same kind of resource, which may demand efficient storage, labeling, and access strategies, but their relative uniformity can make both selecting resources and organizing them simpler. Systems with a broad scope that have many different types of resources pose a more complex challenge, as the diversity of resource types requires more precise descriptions to enable organizing principles that distinguish them.

The scope and scale of an organizing problem are primarily determined by decisions about resource granularity and abstraction.

- *Resource granularity*: Many resources can be analyzed as a single unitary thing or a composite with parts, each of which can also be treated as a separate resource. For example, when a car is sold, it is a single resource, but when it is built, it is a complex system composed of engine, frame, tires, seats, and thousands of other parts. Similarly, a book can be one thing or a collection of chapters or web pages. Decisions about granularity determine the number of resources that need to be organized, how they are grouped and compared, and how they are stored, retrieved, and acted upon.
- *Resource abstraction*: Many resources can be treated as a unique instance, as one of many that are treated as equivalent. The size of this equivalence class is determined by the properties or characteristics in consideration. The choice of these properties depends on context and intent, so the same resource can

be identified abstractly in some situations and very concretely in others. For example, all the copies of a book in a library could be treated as equivalent substitutes, or every copy could be treated as unique because they are annotated, signed, or otherwise distinguishable. The library might even treat different books as equivalent physical objects for inventory and storage purposes, even if their contents differ dramatically.

Selection criteria for resources embody these decisions about granularity and abstraction. An organizing system might be designed to include all the parts of a composite resource and all of the equivalent instances, or it could be more specific and incrementally created by relying on one-at-a-time selection decisions.

Organizing systems with a narrow scope and small scale can use informal and intuitive selection criteria. In contrast, systems with a large scope and scale benefit from and might require formal and computational approaches for selecting resources; this data-intensive selection is often described as predictive analytics ([Siegel, 2013](#)).

Another important consideration about resource selection is whether selection is path dependent, which means that selections are shaped by previous choices, future alternatives, or the choices that others have made ([Page, 2006](#)). The drafting of players in professional sports is a selection activity with extreme path dependence.

Resource selection has critical implications for maintenance activities in an organizing system. More carefully selected resources should be maintained more carefully. This is just common sense, whether the resource is a piece of art, an automobile, a software package, or a star basketball player; if a resource was hard to find or expensive to acquire, it must be taken care of. Furthermore, for many kinds of resources, it is necessary to preserve more than the primary resource for it to have future use. This is especially true for scientific data because it is often created using specialized technology and is managed and analyzed by complex software, all of which must be maintained to preserve access to the data.

Resource description

Inherent in the decision about what counts as a resource is an act of sensemaking or awareness about which properties of resources or aspects of a situation could be used to understand or organize it. For example, physical resources are often organized according to intrinsic physical properties like their size, color, or shape because the human visual system quickly attends to them.

Other resource descriptions require more explicit effort to recognize. Resources that have sortable names or identifiers can be organized alphabetically, resources with sortable dates or times can be organized chronologically, and resources with category labels or content keywords can be classified.

Thinking abstractly about organizing helps with identifying other resource descriptions or properties that are generic enough to apply to many kinds of resources. These include its creator, its current owner,

its current value, its location of origin, its current location, and frequency of use.

A beneficial and generic approach for identifying resource properties is to compute descriptive statistics and histograms for property values. These reveal if there is any useful or meaningful variation that could be the basis of an organizing principle that distinguishes resources in interactions with them. Age, income, and levels of education are properties that usefully distinguish people for many purposes, but there is very little variation in normal body temperature or number of eyes.

The interpretability of descriptive statistics is limited by the levels of measurement embodied in the property value, which determines how much quantitative processing makes sense. Even for nominal-level data that has no intrinsic order, it can be useful to organize resources according to the frequency of specific property values; the mode is the most frequently occurring one. If property values are ordinal-level data, resources can be organized from the lowest to the highest value. The middle value of a list sorted this way is the median, which divides the values into two categories with the same number of members. This idea can be generalized to create any number of equal-size categories, for example, percentiles (100 categories), deciles (10), or quartiles (4). Named categories can be easily created from ordinal properties (low, moderate, high cost; new, almost new, old, very old). Finally, if property values are interval-level data, it is meaningful to calculate the mean or average. This can then create resource categories using the probability of their property values (e.g., as measured in standard deviations) or distance from the mean.

Organizing principles

Because resource descriptions and organizing principles are inherently interleaved, organizing can be analyzed in the reverse direction from the preceding section. Instead of starting with the resource descriptions that are available to determine what organizing principles are possible, choosing the desired organizing principles is a directive to find or create the required descriptions. For example, if usage information is available, resources can be organized by frequency of use. If a design requires organization by frequency of use, usage information must be collected.

This more exploratory approach to finding potential properties that could be used to organize a collection of resources often leads to resource-specific organizing principles that add more value than generic ones. The periodic table exemplifies this approach. It took extensive scientific analysis to determine how to organize chemical elements by their atomic structure, but the value created by this organizing principle is huge; it predicts physical and chemical properties, chemical reactions, and suggests unknown elements that were later discovered.

In contrast, organizing the elements alphabetically or by when they were first identified would be of little value.

Organizing principles and category hierarchy

The order in which organizing principles are applied creates a category hierarchy with the resource instances as the “leaves of the tree.” With physical resources, this has a big impact on resource interactions. Imagine organizing the shirts in your closet. Visualize an arrangement in which shirts are first sorted by color, then by sleeve length. Contrast this with the reverse application of the organizing principles, in which shirts are first sorted by sleeve length and then by color. This change in order can make finding or comparing shirts easier or more difficult depending on user priorities.

Faceted classifications, commonly used in user interfaces for online shopping, can overcome this inherent constraint on category interactions. Each resource is described using properties from multiple facets, but interactions (most often search) can consider subsets of the properties in any order. This design creates a dynamic organizational structure that makes selection both flexible and efficient ([Broughton, 2006](#)).

Category context: Individual, cultural, and institutional

The meaning and salience of categories often depend on the context in which they are used or created ([Glushko et al., 2008](#)). Categories can reflect the following:

- *Individual context*: Personal experience, goals, or expertise (e.g., a wine lover’s categories for red wines may differ from a sommelier’s)
- *Cultural context*: Shared conventions, norms, or language use (e.g., regional definitions of “breakfast food” or “family”)
- *Institutional context*: Professional standards, regulations, or formal classifications (e.g., legal definitions of disability, medical taxonomies, or library subject headings)

These contexts shape what distinctions are meaningful, what boundaries are enforced, and how flexible or rigid a category system can be.

Category structure: How membership is determined

Categories are defined in different ways depending on the type of resources being organized and the reasons for organizing them [see [Categorization](#)]. This concept is the *category structure*, defined as the way in which membership in a category is determined. The notion of category structure has deep roots in philosophy, psychology, and the information sciences, but for much of its history, it lacked a precise or consistent definition.

In classical philosophy (especially classical philosophy influenced by Plato and Aristotle), categories were treated as ideal or natural kinds—sharply bounded groupings defined by necessary and sufficient conditions. Membership was binary: a thing either belonged or it did not.

This essentialist view was famously challenged by Ludwig Wittgenstein in *Philosophical Investigations* ([Wittgenstein, 1953](#)), in which he argued that many everyday concepts like “game” cannot be defined by a fixed set of shared features. Instead, he proposed that categories are held together by family resemblances: overlapping similarities without a single defining trait. His broader view that “meaning is use” undermined the idea that concepts must have fixed boundaries and opened the door to context-dependent, interaction-based understandings of classification.

Wittgenstein’s insights were empirically validated in studies of color naming and object classification ([Rosch, 1973](#); [Rosch & Mervis, 1975](#)) that demonstrated category membership is often graded and that categories are organized around prototypes—central, highly typical examples. This marked a shift toward understanding category structure in terms of statistical regularities and psychological salience.

The notion of ad hoc categories, groupings like “things to take from a burning building,” which are constructed on the fly to serve specific goals, further reinforced the idea that category structure can be dynamic and context sensitive, not just innate or learned ([Barsalou, 1983, 1985](#)) [see [Ad Hoc Categories](#)].

More recently, Bayesian models of categorization have gained prominence in cognitive science, offering a probabilistic framework that captures how people integrate prior knowledge and observed evidence to make flexible, context-aware category judgments ([Tenenbaum et al., 2011](#)) [see [Bayesian Models of Cognition](#)].

Members in enumerated categories do not share necessary properties or prototypes; membership is determined by inclusion in a fixed list, like a set of zodiac signs, country codes, or approved vendors. Despite being widespread in institutional and information systems, enumerated categories have received little attention in psychological theories.

Making category structure a central analytic tool emphasizes that category structures can be explicit or implicit, fixed or fluid, or formal or informal, and they shape how we organize, access, and interact with resources. Understanding the structure of categories helps us see why category boundaries vary in sharpness, how people learn or challenge categories, and what forms of reasoning a system supports.

Design patterns for organizing systems

Treating organizing as a design problem rather than as the task of applying an existing organizing solution enables more innovative and robust thinking. Familiar organizing systems like libraries and museums can be viewed as common design patterns that are starting points in a design space.

Because they embody typical configurations of the what, why, and how design choices, these starting points are surrounded in the design space by less populated or empty regions where organizing systems that make more innovative or less typical design choices fit. Analyzing these variations is a valuable way to explore the design space of organizing.

For example, a university library is a canonical organizing system for managing books and information; it emphasizes formal cataloging, persistent storage, public access, and circulation of resources. By contrast, a bookmobile, which serves mobile, temporary collections, or a tool-lending library, which lends tools and equipment rather than information resources, are designed to satisfy different requirements. These edge cases stretch assumptions and reveal which organizing principles are essential in a domain and which are context dependent.

Another analytic dimension in the design space of organizing is the degree of intentionality and systematic structure involved. Organizing systems can range from informal, loosely structured arrangements to highly designed, rule-governed systems. At one end of the spectrum are ad hoc or minimally organized settings. For example, a personal notebook, a dinner party seating plan, or a “smash-and-grab” crime spree involves little formal structure or enduring organization. At the other end are systems like tax codes, hospital record systems, or organized crime families, systems that are rule based, role based, and often institutionalized.

Documents illustrate this spectrum clearly; a novel is a narrative document in which meaning emerges through interpretation, with implicit structure, and an invoice is a transactional document in which semantics and structure are explicit and standardized for predictable interaction, often by computational agents ([Glushko & McGrath, 2005](#)).

Positioning an organizing system along this spectrum helps to assess its flexibility, reproducibility, and the expectations it creates for users and participants. It complements the typical versus edge case framework by focusing not on what is being organized but on how deliberately it is organized.

A third dimension in which organizing systems in a particular design domain can differ is the extent to which standards and formal specifications play a critical role. These standards can govern both the description of resources and the interactions with them, promoting consistency, interoperability, and trust.

- *Standards for resource description* include controlled vocabularies, metadata schemas, ontologies, and classification systems. These standards ensure consistency and interoperability when organizing and exchanging information about resources. For example, libraries commonly use machine-readable cataloging ([Library of Congress, n.d.](#)) and Dublin Core ([Weibel et al., 1998](#)) for bibliographic metadata. In health care, the International Classification of Diseases ([WHO, 2019](#)) and the Systematized Nomenclature of Medicine—Clinical Terms ([SNOMED International, n.d.](#)) are widely adopted for classifying diagnoses and clinical findings.
- *Standards for interaction* define how systems exchange, validate, or act upon organized resources. These standards facilitate machine-to-machine communication and integration across platforms. For instance, Representational State Transfer Application Programming Interfaces (REST APIs) ([Fielding, 2000](#)) provide a widely used architectural style for web-based software interactions. In health care, Health Level 7 protocols ([HL7 International, n.d.](#)) standardize electronic data exchange. For digital

repositories, protocols such as the Open Archives Initiative Protocol for Metadata Harvesting ([Lagoze et al., 2002](#)) enable interoperability by supporting structured metadata retrieval across systems.

Adhering to standards can increase efficiency, facilitate integration, and support governance, but it can also introduce rigidity or favor dominant stakeholders. Choosing or designing standards is itself a form of organizing that encodes values, assumptions, and power dynamics.

Questions, controversies, and new developments

Bias in organizing systems

It is often said that an organizing system or the categories it creates are biased, and this is almost always meant as a criticism. This can be indisputably true; no one could defend the race categories used to implement segregation in the Apartheid system. However, too often this criticism misses the point that all organization is biased. Organizing decisions are never neutral; they reflect context-specific goals, constraints, and values. Organizing systems created for personal use will naturally embody individual goals and preferences. In contrast, organizing systems used by more than one person or for more than one interaction must prioritize the users and the interactions. This means that the organizing system will be more efficient or create more value for some users than for others.

Instead of describing this inevitable property of organizing systems as bias, it is better to characterize it as embodying design tradeoffs required by different contexts. For example, organizing systems used by experts typically have finer resource granularity than those for nonexperts or casual users. Granularity increases flexibility and reuse but may require more complex descriptions and storage.

A similar tradeoff occurs with resource abstraction. Treating resources as interchangeable simplifies sorting and scaling, but overly broad categories may hide unique attributes that matter for interaction or valuation. Treating each resource as unique allows tailored organization but increases cognitive and computational load.

Finally, tradeoffs can be made about resource description. More explicit and detailed descriptions enable richer organizing principles and more precise interactions, but they require effort to create, maintain, and standardize. Minimal or implicit descriptions are easier to use but limit organizing options and reduce interoperability.

The domain of medicine illustrates these tradeoffs. High granularity in patient symptoms supports personalized care and research analytics but also requires high data-entry overhead, interoperability standards, and privacy safeguards. Symptoms are categorized differently for insurance purposes, and communication about symptoms between physicians and patients is typically of low granularity but is also of low abstraction because every patient needs to be treated as unique.

Computational organization?

Software has played an increasingly central role in organizing resources since the early days of information retrieval. In the 1970s, the SMART system (System for the Mechanical Analysis and Retrieval of Text) introduced vector space models and term-weighting schemes like TF-IDF (term frequency—inverse document frequency) to automatically extract index terms and rank document relevance ([Salton, 1971](#)) [see [Distributional Semantics](#); [Natural Language Processing](#)]. These ideas became foundational to search engines and digital libraries. Over the following decades, classification algorithms, clustering techniques, and collaborative filtering emerged as key tools for organizing large-scale digital content. The rise of recommender systems in the 1990s and early 2000s, such as those used by Amazon and Netflix, demonstrated the value of *computational classification* for predicting preferences and personalizing information access ([Linden et al., 2003](#); [Resnick et al., 1994](#)). These techniques typically relied on structured metadata, behavioral traces, or manually engineered features.

However, in the past few years, breakthroughs in artificial intelligence (AI), particularly in deep learning and large language models, have expanded the possibilities for what we might now call *computational organizing* [see [Large Language Models](#)]. As the design space for organizing systems becomes more complex, AI assistants can be powerful tools for exploring alternative configurations. An AI assistant can help organizers by

- Generating multiple design scenarios that vary in granularity, abstraction, and description strategies based on a given context or goal.
- Comparing tradeoffs across these scenarios, highlighting potential consequences for usability, cost, or scalability.
- Recommending metadata schemas or classification models suited to the domain and intended user interactions.
- Learning from prior examples of successful organizing systems to suggest patterns or design heuristics.

For example, a museum curator designing a new digital catalog might prompt an AI assistant to show different ways to organize exhibits by artist, period, material, or theme, depending on user goals. A business analyst might use AI to simulate how product catalog changes affect inventory management or search relevance.

Rather than replacing human judgment, these assistants amplify it, helping organizers surface overlooked options, evaluate tradeoffs more systematically, and make more informed, context-aware design decisions. Of course, it is important to note that AI models also contain bias/design tradeoffs and can hallucinate plausible but incorrect information, so care needs to be taken when assessing the suggestions.

Broader connections

Not all structure in the world is intentionally designed organization. Going beyond the organizing systems created by individuals and institutions, there are many examples in which structure emerges spontaneously from the uncoordinated collective behavior of many agents, like traffic patterns, flocks of birds, or online reputation scores ([Goldstone & Gureckis, 2009](#)) [see [Collective Intelligence](#)].

Another type of structure is imposed automatically by our brains when they implement a simplicity principle to interpret visual and auditory input ([Chater & Vitányi, 2003](#)); this is called Gestalt perception in vision ([Wagemans et al., 2012](#)) and statistical learning when perceiving spoken language ([Saffran et al., 1996](#)) [see [Gestalt Perception](#); [Statistical Learning](#)].

Still other forms of structure—like the arrangement of planets ([Newton, 1687](#)) or the atomic structure of elements ([Mendeleev, 1869](#))—exist independently of human intention, waiting to be discovered, not designed.

This broader perspective shows that intentional design, emergent structure, and discovered order are not separate categories of organizing; they form a continuum ([Table 1](#)). This continuum helps designers make principled choices across different kinds of organizing systems.

Table 1.

A Continuum of Organizing

Dimension	Intentional organizing	Emergent organizing	Discovered structure
Definition	Deliberate structure created for known goals	Structure arises from uncoordinated local interactions	Structure exists independently, revealed through observation
Source of order	Human intent and design	Local interactions, feedback loops	Natural laws, physical forces
Agency	Centralized, explicit	Distributed, implicit	No human agency
Examples	Libraries, closets, and constitutions	Bird flocks, markets, and language	Planetary orbits, DNA, and atomic structure

When an AI system or other application summarizes an article, outlines an argument, categorizes data, or performs some other computational organizing task, this continuum of organizing should be used to evaluate the output. Is the AI reproducing existing structure that it extracted from the massive amount of text used to train it? Did it use the domain-specific organizing principles that would have been incorporated in authoritative sources? Or did the AI create a new structure for organizing its answers that

embodied information patterns extracted from its training text? This new structure might be a novel way to make sense of a domain, but it might be wrong because patterns assembled for coherence might be “hallucinations” that apply generic organizing principles that do not capture the essential properties of a domain.

Further reading

- Glushko, R. J., & Petras, V. (Eds.). (2026). *The discipline of organizing* (5th ed.). Berlin Universities Publishing.
- Morgan, G. (2006). *Images of organization*. Sage Publications.
- Organizing Systems Case Studies Repository. Students in university courses using *The Discipline of Organizing* book often write case studies using the design questions as a final project. Over 100 of these are freely available in a repository maintained by ischools.org (<https://www.ischools.org/tdo-case-studies>).

References

- Aristotle. (1937). *Parts of animals* (A. L. Peck & E. S. Forster, Trans.). Harvard University Press. (Original work published 350 BCE)

↩

- Aristotle. (2007). *Categories* (E. M. Edghill, Trans.). University of Adelaide. (Original work published 350 BCE)

↩

- Bacon, F. (1902). *Novum organum* (J. Devey, Ed.). P. F. Collier. (Original work published 1620)

↩

- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, 11(3), 211–227. <https://doi.org/10.3758/BF03196968>

↩

- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(4), 629–654. <https://doi.org/10.1037//0278-7393.11.1-4.629>

↩

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

↩

- Broughton, V. (2006). The need for a faceted classification as the basis of all methods of information retrieval. *Aslib Proceedings*, 58(1/2), 49–72. <https://doi.org/10.1108/00012530610648671>

↩

- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7(1), 19–22. [https://doi.org/10.1016/S1364-6613\(02\)00005-0](https://doi.org/10.1016/S1364-6613(02)00005-0)

↩

- Diderot, D., & d'Alembert, J. L. R. (1751). *Encyclopédie, ou dictionnaire raisonné des sciences, des arts et des métiers*. Chez Briasson.

↩

- Fielding, R. T. (2000). *Architectural styles and the design of network-based software architectures* [Doctoral dissertation, University of California, Irvine]. Donald Bren School of Information & Computer Sciences. <https://ics.uci.edu/~fielding/pubs/dissertation/top.htm>

↩

- Glushko, R. J. (2013). *The discipline of organizing*. MIT Press.

↩

- Glushko, R. J., & McGrath, T. (2005). *Document engineering*. MIT Press.

↩

- Glushko, R. J., Maglio, P. P., Matlock, T., & Barsalou, L. W. (2008). Categorization in the wild. *Trends in Cognitive Sciences*, 12(4), 129–135. <https://doi.org/10.1016/j.tics.2008.01.007>

↩

- Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, 1(3), 412–438. <https://doi.org/10.1111/j.1756-8765.2009.01038.x>

↩

- Hinton, G., & Sejnowski, T. J. (Eds.). (1999). *Unsupervised learning: Foundations of neural computation*. MIT Press.

↩

- HL7 International. (n.d.). *Introduction to HL7 standards*. Health Level Seven International. Retrieved July 7, 2026, from <https://www.hl7.org/implement/standards/index.cfm?ref=nav>

↩

- Lagoze, C., Van De Sompel, H., Nelson, M., & Warner, S. (2002). *The Open Archives Initiative Protocol for Metadata Harvesting*. Open Archives Initiative.

<https://www.openarchives.org/OAI/openarchivesprotocol.html>

↵

- Library of Congress. (n.d.). *MARC standards*. Network Development and MARC Standards Office. Retrieved July 7, 2026, from <https://www.loc.gov/marc/>

↵

- Linden, G., Smith, B., & York, J. (2003). [Amazon.com](https://doi.org/10.1109/MIC.2003.1167344) recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1), 76–80. <https://doi.org/10.1109/MIC.2003.1167344>

↵

- Linnaeus, C. (1735). *Systema naturae*. Nieuwkoop B. de Graaf.

↵

- Mendeleev, D. (1869). On the relationship of the properties of the elements to their atomic weights. *Zeitschrift für Chemie*, 12, 405–406.

↵

- Newton, I. (1687). *Philosophiæ naturalis principia mathematica*. Samuel Pepys and the Royal Society.

↵

- Page, S. E. (2006). Path dependence. *Quarterly Journal of Political Science*, 1(1), 87–115. <https://doi.org/10.1561/100.00000006>

↵

- Plato. (2007). *The republic* (D. Lee, Trans.). Penguin Classics. (Original work published 375 BCE)

↵

- Plato. (2009). *Phaedrus* (R. Waterfield, Trans.). Oxford University Press. (Original work published 370 BCE)

↵

- Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). GroupLens: An open architecture for collaborative filtering of netnews. *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work*, 175–186. <https://doi.org/10.1145/192844.192905>

↵

- Rosch, E. (1973). Natural categories. *Cognitive Psychology*, 4(3), 328–350. [https://doi.org/10.1016/0010-0285\(73\)90017-0](https://doi.org/10.1016/0010-0285(73)90017-0)

↵

- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4), 573–605. [https://doi.org/10.1016/0010-0285\(75\)90024-9](https://doi.org/10.1016/0010-0285(75)90024-9)

↩

- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928. <https://doi.org/10.1126/science.274.5294.1926>

↩

- Salton, G. (1971). *The SMART retrieval system: Experiments in automatic document processing*. Prentice-Hall.

↩

- Siegel, E. (2013). *Predictive analytics: The power to predict who will click, buy, lie, or die*. Wiley.

↩

- Simon, H. A. (1969). *The sciences of the artificial*. MIT Press.

↩

- Smith, A. (1776). *The wealth of nations*. W. Strahan and T. Cadell.

↩

- SNOMED International. (n.d.). *SNOMED CT*. Systematized Nomenclature of Medicine. Retrieved July 7, 2026, from <https://www.snomed.org/>

↩

- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285. <https://doi.org/10.1126/science.1192788>

↩

- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological Bulletin*, 138(6), 1172–1217. <https://doi.org/10.1037/a0029333>

↩

- Weibel, S., Kunze, J., Lagoze, C., & Wolf, M. (1998). *RFC 2413: Dublin Core Metadata for resource discovery*. RFC Editor. <https://doi.org/10.17487/RFC2413>

↩

- WHO. (2019). *International statistical classification of diseases and related health problems* (11th ed.). World Health Organization.

↩

- Wittgenstein, L. (1953). *Philosophical investigations*. Blackwell.

↩