

Abstractions

Peter J. Denning

DRAFT 11/6/24 – v1

We do not agree on what our core abstractions mean. They are useful anyway.

Peter J. Denning (pjd@nps.edu) is Distinguished Professor of Computer Science at the Naval Postgraduate School in Monterey, CA, USA, is Editor of ACM Ubiquity, and is a past president of ACM. His most recent book is *Navigating a Restless Sea: Mobilizing Innovation in Your Community* (with Todd Lyons, Waterside Productions, 2024). The author's views expressed here are not necessarily those of his employer or the U.S. federal government.

We claim two things about our profession. Computer science studies information processes, natural and artificial. Computer science is a master of abstraction. To reconcile the two, we say that abstraction is the key that unlocks the complexity of designing and managing information processes.

Where did we get the idea that our field is a master of abstractions? This idea is cosmic background radiation left over from the beginning bangs of our field. For its first four decades, computer science struggled under often blistering criticisms by scientists that the new field was not a science. Science is, they said, a quest to understand nature – and computers are not natural objects. Our field's pioneers maintained that computers were a major new phenomenon that called for a science to understand them. The most obvious benefit of our field was software that controls and manages very complex systems. This benefit accrued from organizing software into hierarchies of modules, each responsible for a small set of operations on a particular type of digital object. Abstraction became a shorthand for that design principle.

Around 2000, the weight of opinion suddenly flipped. Computer science was welcomed at the table of science. The tipping point came because many scientists were collaborating with computer scientists to understand information processes in their fields. Many computer scientists claimed we had earned our seat because of our expertise with abstractions. In fact, we won it because computation became a new way of doing science, and many fields of science discovered they were studying naturally-occurring information processes.

In the remainder of this essay, I will argue that abstraction is a byproduct of our central purpose, understanding information processes. Our core abstraction is information process, not “abstraction”. Every field of science has a core abstraction – the focus of their concerns. In all but a few cases, the core abstractions in science defy precise definitions and scientists in the same field disagree on their meanings. Two lessons follow. First, computing is not unique in believing it is a master of abstraction. Indeed, this claim never sat well with practitioners in other fields. Math, physics, chemistry, astronomy, biology, linguistics, economics, psychology -- they all claim to be masters of abstractions. The second lesson is that all fields have made remarkable advances in technology without clear definition of their core abstraction. They all designed simulations and models to harness the concrete forces behind their abstractions. The profound importance of these lessons was recognized with two 2024 Nobel prizes awarded to computer scientists for protein folding and machine learning.

What is Abstraction?

Abstraction is a verb: to abstract is to identify the basic principles and laws of a process so that it can be studied without regard to physical implementation; the abstraction can then guide many implementations. Abstraction is also a noun: An abstraction is a mental construct that unifies a set of objects. Objects of an abstraction have their own logic of relations with each other that does not require knowledge of lower-level details. (In computer science, we call this information hiding.)

Abstractions are a power of language. In his book *Sapiens*, Yuval Noel Harari discusses how, over the millennia, human beings created stories (“fictions”) which united them into communities and gave them causes they were willing to fight for [YH]. These fictions were abstractions that often endured well beyond their invention. The US constitution, for instance, applies to all its states and has guided billions of people for over 200 years.

The ability of language to let us create new ideas and coordinate around them also empowers language constructs to refer to themselves. After all, we have numerous ideas about our ideas. We build endlessly complex structures of ideas. We can imagine things that do not exist, such as unicorns, or unrealized futures that we can pursue. Self-reference also generates paradoxes. A famous paradox asks: “Does the set of all those sets that do not contain themselves contain itself?” Self-reference is both a blessing and a curse.

One way to avoid paradoxes is to stack up abstractions in hierarchies or connect them in networks. An abstraction can be composed of lower-level abstractions but cannot refer to higher level abstractions. In chemistry, for example, amino acids are composed of atoms, but do not depend on proteins arranging the acids in particular sequences. In computing, operating systems are considered layers of software that manage different abstractions such as processes, virtual memories, and files; each depends on lower levels, but not higher levels.

Let's consider three examples illustrating how different fields use their abstractions.

Computer Science. An "abstract data type" represents a class of digital objects and the operations that can be performed on them. This reduces complexity because one algorithm can apply to large classes of objects. The expressions of these abstractions can be compiled into executable code: thus, abstractions can also be active computing utilities and not just descriptions.

Physics. For physicists, abstraction simplifies complex phenomena and enables models to help understand and predict the behavior of complex systems. Many physics models take the form of differential equations that can be solved on grids by highly parallel computers. For example, the Stokes Equation in computational fluid dynamics specifies air flows around flying aircraft. Other models are simulations that evaluate the interactions between entities over long periods of time. For example, astronomers have successfully simulated galactic collisions by treating galaxies as masses of particles representing stars. Because models make simplifications there is always a tradeoff between model complexity and accuracy. The classical core abstraction of physics has been any natural process; in recent decades it expanded to include information processes and computers.

Mathematics. Abstraction is the business of mathematics. Mathematicians are constantly seeking to identify concepts and structures that transcend physical objects. They seek to express the essential relationships among objects by eliminating irrelevant details. Mathematics is seen as supportive of all scientific fields. In 1931 Bertrand Russell wrote: "Ordinary language is totally unsuited for expressing what physics really asserts, since the words of everyday life are not sufficiently abstract. Only mathematics and mathematical logic can say as little as the physicist means to say."

Anywhere you see a classification you are looking at a hierarchy of abstractions. Anywhere you see a theory you are looking at an explanation of how a set of abstractions interacts. Highly abstract concepts can be very difficult to learn because they depend on understanding much past history, represented in lower-level abstractions.

Differing Interpretations of the same Abstractions

It is no surprise that different people have different interpretations about abstractions and thus get into arguments over them. After all, abstractions are mental constructs learned individually. Few abstractions have clear logical definitions as in Mathematics or in Object-Oriented languages. Here are some additional examples showing how different fields approach differences of interpretation of their core abstractions.

Biology. This is the science studying life. There is, however, no clear definition of life. How do biologists decide if some newly discovered organism is alive? They have agreed on a list of seven criteria for assessing whether an entity is living:

- Responding to stimuli
- Growing and developing
- Reproducing
- Metabolizing substances into energy
- Maintaining a stable structure (homeostasis)
- Structured from cells
- Adaptability in changing environments

The more of these criteria hold for an organism, the more likely is a biologist to say that life is present.

Artificial intelligence. Its central abstraction – intelligence – defies precise definition. Various authors have cited one of more of these indicators as signs of intelligence:

- Passes IQ tests
- Passes Turing test
- Pinnacle of a hierarchy of abilities determined by psychologists
- Speed of learning to adapt to new situations
- Ability to set and pursue goals
- Ability to solve problems

However, there is no agreement on whether these are sufficient to cover all situations of apparent intelligence. Julian Togelius has an excellent summary of the many notions of “intelligence” (and “artificial”) currently in play [JT]. This has not handicapped AI, which has produced a series of amazing technological advances.

Computer science. Its central concept – information process – defies a precise definition. Among the indicators frequently mentioned are

- Dynamically evolving strings of symbols satisfying a grammar
- Assessment that strings of symbols mean something
- Mapping symbol patterns to meanings
- Insights gained from data
- Fundamental force in the universe
- Process of encoding a description of an event or idea
- Process of recovering encrypted data
- Inverse log of the probability of an event (Shannon)

There is no consensus whether these are sufficient to cover all situations where information is present.

Neuroscience. Consciousness is a core abstraction. Neuroscientists and medical professionals in general have agreed on a few, imprecise indicators of when someone is conscious [CK]. Some conscious people may fail all the indicators, and some unconscious people may satisfy some of the indicators. It may be impossible to ever know for sure if someone is conscious or not.

Business. Innovation is a core abstraction. Business leaders want more innovation. Definitions vary from inventing new ideas, prototyping new ideas,

transitioning prototypes into user communities, diffusing into user communities, and adoption of new practice in user communities. Each definition is accompanied by a theory of how to generate more innovation. The definitions are sufficiently different that the theories conflict. There is considerable debate on which definition and its theory will lead to the most success.

Conclusion

The accompanying table summarizes the examples above. The “criteria” column indicates whether a field has a consensus on criteria for their core abstraction. The “explanatory” column indicates whether a field’s existing definitions adequately explain all the observable instances of their core abstraction. The “utility” column indicates whether they are concerned with finding applications of technologies enabled by their core abstraction.

Table 1. A few fields and their core abstractions

Field	Abstraction	Criteria?	Explanatory?	Utility?
Computing	Information	No	No	Yes
Physics	Natural phenomena	No	Yes	Yes
Mathematics	Math concepts	No	Yes	No
Biology	Life	Yes	Yes	Yes
Artificial Intelligence	Intelligence	No	No	Yes
Neuroscience	Consciousness	No	Yes	Maybe
Business	Innovation	No	No	Yes

Thus, it seems that the core abstractions of many fields are imprecise and, with only a few exceptions, the fields have no consensus on criteria to determine if an observation fits their abstraction. How do they manage a successful science without a clear definition of their core abstraction? The answer is that in practice they design systems and processes based on validated hypotheses. The varying interpretations are a problem only to the extent that disagreements generate misunderstandings and confusions.

A good way to bring out the differences of interpretation is to ask people how they assess if a phenomenon before them is an instance of their core abstraction. Thus you could say “Life is an assessment”, “intelligence is an assessment”, and so on. When you put it this way, you invite a conversation about the grounding that supports the assessment. For example, a biologist would ground an assessment that a new organism is alive by showing that enough of the seven criteria are satisfied. In other fields the request for assessment quickly brings out differences of interpretation. In business, for example, where there is no consensus on the indicators of innovation, a person’s assessments reveal which of the competing core

abstractions they accept. That, in turn, opens the door for conversations about the value of each abstraction.

There is a big controversy over whether technology is dragging us into abstract worlds with fewer close relationships, fear of intimacy, and interaction limited to exchanges across computer screens. This is a particular problem for young people [JH]. Smartphones are intended to improve communication and yet users feel more isolated, unseen, unappreciated. Something is clearly missing in our understanding of communication, but we have not yet put our collective finger on it.

Two books may help sort this out. In *Power and Influence*, Nobel Prize economists Daron Acemoglu and Simon Johnson present a massive trove of data to claim that increasing automation often increases organizational productivity without increasing overall economic progress for everyone. They argue that the abstractions behind automation focus on displacing workers rather than augmenting workers by enabling them to take on more meaningful tasks [DA]. In *How to Know a Person*, David Brooks presents communication practices that help you see and appreciate the everyday concrete concerns of others [DB].

Maybe we need to occasionally descend from the high clouds of our abstractions to the concrete earthy concerns of everyday life.

Bibliography

1. (DA) Acemoglu, Daron, and Simon Johnson. *Power and Progress: Our Thousand Year Struggle over Technology and Prosperity*. Public Affairs. 2023
2. (DB) Brooks, David. *How to Know a Person: The Art of Seeing Others Deeply and Being Deeply Seen*. Random House. 2023.
3. (JH) Haidt, Jonathan. *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. Penguin. 2024.
4. (YH) Harari, Yuval Noah. *Sapiens: A brief History of Humankind*. Harper. 2015.
5. (CK) Koch, Christof. *Then I am Myself the World: What Consciousness Is and How to Expand It*. Basic Books. 2024.
6. (JT) Togelius, Julian. *Artificial General Intelligence*. MIT Press. 2024.