# DOWNLOAD PDF COMPUTATION AND COGNITION

## Chapter 1 : Computation and Cognition | MIT CogNet

*In Computation and Cognition, Pylyshyn argues that computation must not be viewed as just a convenient metaphor for mental activity, but as a literal empirical hypothesis. Such a view must face a number of serious challenges.*

This document may be redistributed and reused, subject to certain conditions. This article has been cited by other articles in PMC. Abstract The cognitive concept of representation plays a key role in theories of brain information processing. However, linking neuronal activity to representational content and cognitive theory remains challenging. Recent studies have characterized the representational geometry of neural population codes by means of representational distance matrices, enabling researchers to compare representations across stages of processing and to test cognitive and computational theories. Representational geometry provides a useful intermediate level of description, capturing both the information represented in a neuronal population code and the format in which it is represented. We review recent insights gained with this approach in perception, memory, cognition, and action. Analyses of representational geometry can compare representations between models and the brain, and promise to explain brain computation as transformation of representational similarity structure. The representational geometry of neuronal population codes The concept of representation is central to the cognitive and brain sciences. We interpret neuronal activity as serving the function of representing content, and of transforming representations of content, with the ultimate objective to produce successful behaviors. The content could be a visual image, a sound or odor, a semantic interpretation of sensory input, a proposition, a goal, a planned action, or a motor sequence. The representational interpretation [1] provides a powerful explanatory framework that makes it easier to understand neuronal activity in the context of the overall function of the brain. Representation links cognition to brain activity and enables us to build functional theories of brain information processing [2]. Neurophysiology has long interpreted the selectivity of neurons as serving to represent various kinds of sensory and higher-level information. The population of neurons within an area is thought to jointly represent the content in what is called a neuronal population code [3]. It is the pattern of activity across neurons that represents the content. The many possible combinations of activity states of neurons provide a rich representational space. Motivated by this idea, recent analyses of neuronal recordings and functional imaging data have increasingly focused on patterns of activity across many neurons within a functional region [4]. The dimensions of the space correspond to the neurons, and a point corresponds to an activity pattern i. A visually perceived object, for example, will correspond to a point in the representational space of a given visual area. The set of all possible objects or pieces of mental content corresponds to a vast set of points in the space. It is the geometry of these points that defines the nature of the representation. Mathematical and cognitive psychology have a long history of investigations of representational geometry on the basis of behavioral data [5â€”10]. However, the notion of representational geometry has only more recently been brought into the analysis of brain-activity data [11â€”15]. To characterize the geometry of a representation, we can compare the brain-activity patterns representing a set of stimuli or, more generally, experimental conditions to each other. The dissimilarity of two patterns corresponds to the distance between their points in the representational space. Having measured these distances, we can construct a matrix, the representational dissimilarity matrix RDM , in which we can look up the representational distance or dissimilarity for each pair of stimuli Figure 1. Intuitively, the RDM tells us which distinctions between stimuli the population code honors and which distinctions it disregards.

## Chapter 2 : Computation and Cognition: Toward a Foundation for Cognitive Science by Zenon Pylshyn

*Information processing, computation, and the foundations of cognitive science. Computation and information processing are among the most fundamental notions in cognitive science.*

Skip to content The Major in Cognitive Science Cognitive science is the empirical study of intelligent systems, including the human mind. An interdisciplinary science, it combines results from biology, computer science, linguistics, mathematics, neuroscience, philosophy and psychology to the study of language processing, perception, action, learning, concept formation, inference and other activities of the mind, with applications for information technology and the study of artificial intelligence. For general program information, please contact Program Manager Jessica Marcus: Advising hours for Fall will be held on Mondays from  It is offered in the Fall term. Please note that the courses listed are courses that have historically counted towards breadth requirements; we do not have specific course requirements for the breadth requirements. Advanced Placement credit will not be counted toward the major requirements. Cognitive Neuroscience, Computation and Cognition, Language and Mind, or a special Independent Concentration constructed to meet a set of interests not included in one of the other concentrations The Program Director advises students when they are first considering the major and while still fulfilling the breadth requirements; handles administrative duties such as major declaration and certification; and is the final authority in all matters relating to the major requirements. To determine whether a course meets the concentration requirement, please contact Dr. Cognitive Science has become even more interdisciplinary as the field matures. We recognize the importance of specialized skills, especially those honed in the biological, economic, computational and mathematical sciences, in cognitive research, education, and application. At the same time, we strive to ground our program in the empirical studies of cognition in Linguistics, Psychology, and Neuroscience; tools are important, but we also need to know what they are for. For students in the Computation and Cognition concentration, a fifth credit in Artificial Intelligence, or other topics directly related to human cognition, may be allowed upon approval. Those four credits are usually drawn from the list of courses below; for suitability of courses not listed below, please contact the Program Director. We would like our students to maximize their educational experience in the Cognitive Science Program by forming a deeper understanding of some select topics or themes. We especially advise against taking introductory classes that have significant overlapping materials, including similar courses that are offered in different departments. Please contact the Program Director should these concerns arise during your course planning and selection process. Currently, only the Cognitive Neuroscience concentration has a specific required course: For example, at least one course in Statistics, such as STAT , is strongly recommended to students specializing in Cognitive Neuroscience. The list below indicates courses that have historically been approved for Concentration requirements. The courses in Psychology suitable for Cognitive Science generally have an odd course number; they are courses in the area of Brain, Cognitive, and Decision science, following the research program and numbering convention in the Department of Psychology. Courses that are offered in Spring are highlighted in the list below. This is not a list of required courses. To find out whether a course not on this list will be approved for the COGS major, please contact Dr.

## Chapter 3 : The Major in Cognitive Science â€“ Cognitive Science at Penn

*This systematic investigation of computation and mental phenomena by a noted psychologist and computer scientist argues that cognition is a form of computation, that the semantic contents of mental states are encoded in the same general way as computer representations are encoded.*

Rather, a computational system is a symbol manipulator that follows step by step functions to compute input and form output. Alan Turing describes this type of computer in his concept of a Turing Machine. Early proponents[ edit ] One of the earliest proponents of the computational theory of mind was Thomas Hobbes , who said, "by reasoning, I understand computation. And to compute is to collect the sum of many things added together at the same time, or to know the remainder when one thing has been taken from another. To reason therefore is the same as to add or to subtract. Causal picture of thoughts[ edit ] At the heart of the Computational Theory of Mind is the idea that thoughts are a form of computation, and a computation is by definition a systematic set of laws for the relations among representations. This means that a mental state represents something if and only if there is some causal correlation between the mental state and that particular thing. This is sometimes known as Natural Meaning. Conversely, there is another side to the causality of thoughts and that is the non-natural representation of thoughts. If these basic mental states can have a particular meaning just as words in a language do, then this means that more complex mental states thoughts can be created, even if they have never been encountered before. Just as new sentences that are read can be understood even if they have never been encountered before, as long as the basic components are understood, and it is syntactically correct. Criticism[ edit ] A range of arguments have been proposed against physicalist conceptions used in Computational Theories of Mind. An early, though indirect, criticism of the Computational Theory of Mind comes from philosopher John Searle. In his thought experiment known as the Chinese room , Searle attempts to refute the claims that artificially intelligent systems can be said to have intentionality and understanding and that these systems, because they can be said to be minds themselves, are sufficient for the study of the human mind. With the paper, the man is to use a series of provided rule books to return paper containing different symbols. Unknown to the man in the room, these symbols are of a Chinese language, and this process generates a conversation that a Chinese speaker outside of the room can actually understand. Searle contends that the man in the room does not understand the Chinese conversation. This is essentially what the computational theory of mind presents usâ€”a model in which the mind simply decodes symbols and outputs more symbols. Searle argues that this is not real understanding or intentionality. Though originally written as a repudiation of the idea that computers work like minds, it is not a stretch to also argue from this position that minds do not work like computers. Searle has further raised questions about what exactly constitutes a computation: But if the wall is implementing WordStar, if it is a big enough wall it is implementing any program, including any program implemented in the brain. They claim that computational theories of mind fail because computation is insufficient to account for some capacity of the mind. There are also objections which are directly tailored for computational theories of mind. This would mean that a normal Turing complete computer would not be able to ascertain certain mathematical truths that human minds can. Consciousness is the computation, there is no extra step or " Cartesian Theater " in which you become conscious of the computation. Jerry Fodor argues that mental states, such as beliefs and desires, are relations between individuals and mental representations. He maintains that these representations can only be correctly explained in terms of a language of thought LOT in the mind. Further, this language of thought itself is codified in the brain, not just a useful explanatory tool. Fodor adheres to a species of functionalism, maintaining that thinking and other mental processes consist primarily of computations operating on the syntax of the representations that make up the language of thought. David Marr proposed that cognitive processes have three levels of description: Steven Pinker described a "language instinct," an evolved, built-in capacity to learn language if not writing. Hilary Putnam proposed functionalism to describe consciousness, asserting that it is the computation that equates to consciousness, regardless of whether the computation is operating in a brain, in a computer, or in a "brain in a vat.

# DOWNLOAD PDF COMPUTATION AND COGNITION

## Chapter 4 : Computational and Cognitive Musicology | Center for Music Technology (GTCMT)

*Cognition Computation and Development Lab. Home Lab Members Research Publications Methodology and Philosophy Collaborators Join.*

Computation and Cognition Peter Mawhorter pmawhorter gmail. Computation and Cognition Spring  Tuesday and Thursday 1: Assignment Details â€"Details about each kind of assignment. Includes due dates, late policies, and collaboration policies. Documentation and Handouts â€"Various useful documents; other stuff will refer to these. Examples â€"All of the in-class example code in one place. Strategies â€"Concrete programming and debugging strategies in one place. Piazza â€"Ask questions here. You should set up email alerts, because any urgent course business will be announced via Piazza. Debugging Wiki â€"Wiki on Sakai with entries that describe errors and bugs. Overview This course introduces basic programming concepts, while also considering how computation and cognition are related. It also deals with the use of computers to measure cognition. This is both an introductory programming course that assumes no prior experience, and a course that looks a bit at several intersections between cognitive science and computer science. Basic programming skills in Python which are also applicable in other programming languages. Some fundamental models of computation, including the idea of a Turing machine. Some concrete applications of computation in artificial intelligence, including neural networks. Learning Goals Have a strong understanding of the historical context of how the study of cognition has influenced computation and computational models of cognition. Be able to discuss contemporary contributions in the intersection of cognitive and computer sciences. Be able to apply computational thinking, perhaps including programming, to a task in linguistics and cognitive science. Be able to design and implement a solution to a computational problem. Understand the basic models of computation and be able to create and simulate Turing Machines and Finite State Automata. Instructors This course is taught by Peter Mawhorter. His website lists contact information and office hours. Mentoring sessions will be held according to the following schedule the Sunday evening hours and the Tuesday hours from 7â€"8 will each have two TAs:

# DOWNLOAD PDF COMPUTATION AND COGNITION

## Chapter 5 : Computation | Brain and Cognitive Sciences

*Cognitive science is the empirical study of intelligent systems, including the human mind. An interdisciplinary science, it combines results from biology, computer science, linguistics, mathematics, neuroscience, philosophy and psychology to the study of language processing, perception, action, learning, concept formation, inference and other activities of the mind, with applications for.*

Artificial intelligence There are two main purposes for the productions of artificial intelligence: Until s, economist Herbert Simon and Allen Newell attempted to formalize human problem-solving skills by using the results of psychological studies to develop programs that implement the same problem-solving techniques as people would. Their works laid the foundation for symbolic AI and computational cognition, and even some advancements for cognitive science and cognitive psychology. Researchers started to believe that artificial intelligence may never be able to imitate some intricate processes of human cognition like perception or learning. A chief failing of AI is not being able to achieve a complete likeness to human cognition due to the lack of emotion and the impossibility of implementing it into an AI. This movement led to the emerging discipline of computational modeling , connectionism , and computational intelligence. Computational model As it contributes more to the understanding of human cognition than artificial intelligence, computational cognitive modeling emerged from the need to define various cognition functionalities like motivation, emotion, or perception by representing them in computational models of mechanisms and processes. The results help experimenters make predictions about what would happen in the real system if those similar changes were to occur. Consider a model of memory built by Atkinson and Shiffrin in , it showed how rehearsal leads to long-term memory, where the information being rehearsed would be stored. Despite the advancement it made in revealing the function of memory, this model fails to provide answers to crucial questions like: How long does it take for information to transfer from rehearsal to long-term memory? Similarly, other computational models raise more questions about cognition than they answer, making their contributions much less significant for the understanding of human cognition than other cognitive approaches. Adaptive Control of Thought-Rational is based on the theory that the brain consists of several modules which perform specialized functions separate of each other. Connectionism Another approach which deals more with the semantic content of cognitive science is connectionism or neural network modeling. Connectionism relies on the idea that the brain consists of simple units or nodes and the behavioral response comes primarily from the layers of connections between the nodes and not from the environmental stimulus itself. Neural back-propagation is a method utilized by connectionist network to show evidence of learning. After a connectionist network produce a response, the stimulated results are compared to real-life situational results. The present of parallel distributed processing has been shown in psychological demonstrations like the Stroop effect , where the brain seems to be analyzing the perception of color and meaning of language at the same time. Therefore, the results can be used as evidence for broad theory of cognition without explaining the particular process happening within the cognitive function. Other disadvantages of connectionism lie in the research methods it employs or hypothesis it tests, which has been proven inaccurate or ineffective often, taking connectionist models further from an accurate representation of how the brain functions. These issues cause neural network models to be ineffective on studying higher forms of information-processing, and hinder connectionism from advancing the general understanding of human cognition. Machines Who Think 2 ed.

## Chapter 6 : Information processing, computation, and cognition

*Note: Citations are based on reference standards. However, formatting rules can vary widely between applications and fields of interest or study. The specific requirements or preferences of your reviewing publisher, classroom teacher, institution or organization should be applied.*

Received Apr 9; Accepted Jul  Abstract Computation and information processing are among the most fundamental notions in cognitive science. They are also among the most imprecisely discussed. Many cognitive scientists take it for granted that cognition involves computation, information processing, or both â€" although others disagree vehemently. In addition, computation and information processing are surrounded by several myths; first and foremost, that they are the same thing. In this paper, we address this unsatisfactory state of affairs by presenting a general and theory-neutral account of computation and information processing. We also apply our framework by analyzing the relations between computation and information processing on one hand and classicism, connectionism, and computational neuroscience on the other. We defend the relevance to cognitive science of both computation, at least in a generic sense, and information processing, in three important senses of the term. Our account advances several foundational debates in cognitive science by untangling some of their conceptual knots in a theory-neutral way. Classicism, Cognitivism, Computation, Computational neuroscience, Computational theory of mind, Computationalism, Connectionism, Information processing, Meaning, Neural computation, Representation Information processing, computation, and the foundations of cognitive science Computation and information processing are among the most fundamental notions in cognitive science. Many cognitive scientists take it for granted that cognition involves computation, information processing, or both. Many others, however, reject theories of cognition based on either computation or information processing [ 1 â€" 7 ]. This debate has continued for over half a century without resolution. An equally long-standing debate pitches classical theories of cognitive architecture [ 8 â€" 13 ] against connectionist and neurocomputational theories [ 14 â€" 21 ]. Classical theories draw a strong analogy between cognitive systems and digital computers. We are interested not so much in the distinction between connectionism and computational neuroscience as in what they have in common: Thus, for present purposes, connectionism and computational neuroscience may be considered together. In recent years, some cognitive scientists have attempted to get around the foundational debates by advocating a pluralism of perspectives [ 22 , 23 ]. According to this kind of pluralism, it is a matter of perspective whether the brain computes, processes information, is a classical system, or is a connectionist system. Different perspectives serve different purposes, and different purposes are legitimate. Hence, all sides of the foundational debates can be retained if appropriately qualified. Although pluralists are correct to point out that different descriptions of the same phenomenon can in principle complement one another, this kind of perspectival pluralism is flawed in one important respect. There is an extent to which different parties in the foundational debates offer alternative explanations of the same phenomenaâ€"they cannot all be right. Nevertheless, these pluralists are responding to something true and important: The way to make progress is therefore not to accept all views at once but to provide a clear and adequate conceptual framework that remains neutral between different theories. Once such a framework is in place, competing explanations can be translated into a shared language and evaluated on empirical grounds. Lack of conceptual housecleaning has led to the emergence of a number of myths that stand in the way of theoretical progress. Not everyone subscribes to all of the following assumptions, but each is widespread and influential: Computation is the same as information processing. Semantic information is necessarily true. The Churchâ€"Turing thesis entails that cognition is computation. Connectionist and classical theories of cognitive architecture are mutually exclusive. We will argue that these assumptions are mistaken and distort our understanding of computation, information processing, and cognitive architecture. Traditional accounts of what it takes for a physical system to perform a computation or process information [ 19 , 24 â€" 26 ] are inadequate because they are based on at least some of assumptions 1â€"6. In lieu of these traditional accounts, we will present a general account of computation and information processing that systematizes, refines, and extends our previous work [ 27 â€" 37 ]. We will then apply our framework by analyzing the relations between computation and information processing on one hand and classicism, connectionism, and computational neuroscience on the other. We will defend the relevance to cognitive science of both computation, at least in a generic sense we will articulate, and information processing, in three important senses of the term. We will also argue that the choice among theories of cognitive architecture is not between classicism and connectionism or computational neuroscience but rather between varieties of neural computation, which may be classical or nonclassical. Our account advances the foundational debates by untangling some of their conceptual knots in a theory-neutral way. Getting rid of some myths The notions of computation and information processing are often used interchangeably. Here is a representative example: This statement presupposes assumption 1 above. Why are the two notions used interchangeably so often, without a second thought? Cyberneticians did not clearly distinguish either between Shannon information and semantic information or between semantic and nonsemantic computation more on these distinctions below. But, at least initially, they were fairly clear that information and computation played distinct roles within their theories. Their idea was that organisms and automata contain control mechanisms: Then the waters got muddier. Many people accepted that computation and information processing belong together in a theory of cognition. After that, many stopped paying attention to the differences between the two. To set the record straight and make some progress, we must get clearer on the independent roles computation and information processing can fulfill in a theory of cognition. The notion of digital computation was imported from computability theory into neuroscience and psychology primarily for two reasons: These reasons are not sufficient to actually establish that cognition is digital computation. Whether cognition is digital computation is a difficult question, which lies outside the scope of this essay. The theory that cognition is computation became so popular that it progressively led to a stretching of the operative notion of computation. Of course, the theory is quite well developed by now, as witnessed by the explosion of work in computational and theoretical neuroscience over the last decades [ 20 , 44 , 45 ]. By contrast, the various notions of information processing have distinct roles to play. By and large, they serve to make sense of how organisms keep track of their environments and produce behaviors accordingly. Other notions of informationâ€”specifically, semantic informationâ€”can serve to give specific semantic content to particular states or events. Whether cognitive or neural events fulfill all or any of the job descriptions of computation and information processing is in part an empirical question and in part a conceptual one. Ok, but do these distinctions really matter? Why should a cognitive theorist care about the differences between computation and information processing? The main theoretical advantage of keeping them separate is to appreciate the independent contributions they can make to a theory of cognition. Conversely, the main cost of conflating computation and information processing is that the resulting mongrel concept may be too messy and vague to do all the jobs that are required of it. As a result, it becomes difficult to reach consensus on whether cognition involves either computation or information processing. Debates on computation, information processing, and cognition are further muddied by assumptions 2â€”6. Assumption 2 is that semantic information is necessarily true; there is no such thing as false information. Therefore, we will reject assumption 2 in favor of the view that semantic information may be either true or false. Assumption 3 is that there is no computation without representation. Most accounts of computation rely on this assumption [ 19 , 24 â€” 26 , 38 ]. As one of us has argued extensively elsewhere [ 27 , 28 , 50 ], however, assumption 3 obscures the core notion of computation used in computer science and computability theoryâ€”the same notion that inspired the computational theory of cognitionâ€”as well as some important distinctions between notions of computation. The core notion of computation does not require representation, although it is compatible with it. In other words, computational states in the core sense may or may not be representations. Understanding computation in its own terms, independently of representation, will allow us to sharpen the debates over the computational theory of cognition as well as cognitive architecture. Assumption 4 is that cognition is computation because of the Churchâ€”Turing thesis [ 51 , 52 ]. The Churchâ€”Turing thesis says that any function that is computable in an intuitive sense is recursive or, equivalently, computable by some Turing machine [ 40 , 53 , 54 ]. But recent scholarship has shown this view to be fallacious [ 29 , 55 ]. The Churchâ€”Turing thesis does not establish whether a function is computable. It only says that if a function is computable in a certain intuitive sense, then it is computable by some Turing machine. Furthermore, the intuitive sense in question has to do with what can be computed by following an algorithm a list of explicit instructions defined over sequences of digital entities. Thus, the Churchâ€"Turing thesis applies directly only to algorithmic digital computation. The relationship between algorithmic digital computation and digital computation simpliciter, let alone other kinds of computation, is quite complex, and the Churchâ€"Turing thesis does not settle it. Assumption 5 is pancomputationalism: There are two ways to defend assumption 5. Some authors argue that everything is computational because describing something as computational is just one way of interpreting it, and everything can be interpreted that way [ 19 , 23 ]. We reject this interpretational pancomputationalism because it conflates computational modeling with computational explanation. The computational theory of cognition is not limited to the claim that cognition can be described modeled computationally, as the weather can; it adds that cognitive phenomena have a computational explanation [ 28 , 31 , 34 ]. Other authors defend assumption 5 by arguing that the universe as a whole is at bottom computational [ 56 , 57 ]. It is not a widely accepted notion, and there is no direct evidence for it. The physical form of pancomputationalism is not directly relevant to theories of cognition because theories of cognition attempt to find out what distinguishes cognition from other processesâ€"not what it shares with everything else. Insofar as the theory of cognition uses computation to distinguish cognition from other processes, it needs a notion of computation that excludes at least some other processes as noncomputational cf. Someone may object as follows. Even if pancomputationalismâ€"the thesis that everything is computationalâ€"is true, it does not follow that the claim that cognition involves computation is vacuous. A theory of cognition still has to say which specific computations cognition involves. The job of neuroscience and psychology is precisely to discover the specific computations that distinguish cognition from other processes cf. We agree that if cognition involves computation, then the job of neuroscience and psychology is to discover which specific computations cognition involves. But the if is important. The job of psychology and neuroscience is to find out how cognition works, regardless of whether it involves computation. The claim that brains compute was introduced in neuroscience and psychology as an empirical hypothesis to explain cognition by analogy with digital computers. Much of the empirical import of the computational theory of cognition is already eliminated by stretching the notion of computation from digital to generic see below.

## Chapter 7 : Home - MCCL@LSU

*Cognition, Computation, and Intelligent Systems The University of Georgia has always viewed Cognitive Science and Artificial Intelligence as interdisciplinary fields where computer science meets philosophy, psychology, linguistics, engineering and other disciplines.*

## Chapter 8 : Computational theory of mind - Wikipedia

*Our research combines multiple levels of computational modeling and experimental work to understand the neural mechanisms underlying reinforcement learning, decision making, and cognitive control.*

## Chapter 9 : Computational cognition - Wikipedia

*Computational cognition (sometimes referred to as computational cognitive science or computational psychology) is the study of the computational basis of learning and inference by mathematical modeling, computer simulation, and behavioral experiments.*