SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

1: Advances in Consciousness Research | Awards | LibraryThing

*Simulation and Knowledge of Action (Advances in Consciousness Research) [JÃ©rÃ´me Dokic, JoÃ«lle Proust] on www.enganchecubano.com *FREE* shipping on qualifying offers. The current debate between theory theory and simulation theory on the nature of mentalisation has reached no consensus yet.*

Neuroethics and Philosophy in Responsible Research and Innovation: The Case of the Human Brain Project. Coma patients might feel pleasure and pain like the rest of us. Emerging evidence suggests the unconscious can experience many of the things that conscious people do. What does this mean for medical ethics and even how we perceive ourselves as humans? Large-scale Brain Simulation and Disorders of Consciousness. Mapping Technical and Conceptual Issues. Our primary research areas Disorders of Consciousness. Among other things, this line of inquiry aims at reviewing and conceptually analysing the recent development in the scientific explanations and descriptions of consciousness, particularly focusing on disorders of consciousness DOCs and exploring the potential and actual clinical applications of neuro-technology for diagnosing and assessing DOCs. This clinical condition affects psychologically traumatised children undergoing migration and is characterised by failure to respond even to painful stimuli leaving the patients seemingly unconscious. This line of research focuses on the sociocultural and neurophysiological factors underlying this disorder. The goal is to provide an analysis that merges contextual and neurobiological aspects in a brain model of culture-bound syndromes. As neuroscience provides more knowledge of the structures and functions of the nervous system, it is expected that it will further our understanding of what makes us human and will promote the development and application of neuro-technologies to modify the brain. This might have a significant impact on human and personal identity. This research project seeks to provide an analysis of the relevant issues, starting with the notion of human identity, its meaning and value, and its relation to the debate on human nature, and to examine related theoretical and practical concerns. The brain develops in a natural and cultural context that profoundly influences its functional architecture. Lived developmental trajectories, interactions, and social environments impact synaptic connectivity and contribute to the formation of patterns of neural activity. Synaptic epigenesis theories of cultural and social imprinting on our brain architecture suggest that it is thus possible to culturally influence our neural predispositions. This research project examines the relationships between genotype and brain phenotype: This research line focuses on the concept of dual use, with explicit emphasis upon both the definition of dual uses of brain science within the HBP and more broadly, the examination of the neuroethical, legal, and social issues arising in and from such research and its applications. Specific attention is given to dual uses of neuroscientific research for military, national security and warfare operations, and direct-to-consumer and do-it-yourself applications of brain research science for lifestyle optimisation that may pose risks to public safety. Learn more about these topics at:

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

2: Simulation as a Teaching Strategy for Nursing Education and Orientation in Cardiac Surgery

Simulation and Knowledge of Action (Advances in Consciousness Research) First Edition Edition by JÃ©rÃ´me Dokic (Editor), JoÃ«lle Proust (Editor).

This insight, that digital computers can simulate any process of formal reasoning, is known as the Churchâ€"Turing thesis. Herbert Simon predicted, "machines will be capable, within twenty years, of doing any work a man can do". Marvin Minsky agreed, writing, "within a generation Progress slowed and in , in response to the criticism of Sir James Lighthill [37] and ongoing pressure from the US Congress to fund more productive projects, both the U. The next few years would later be called an " AI winter ", [9] a period when obtaining funding for AI projects was difficult. In the early s, AI research was revived by the commercial success of expert systems , [38] a form of AI program that simulated the knowledge and analytical skills of human experts. By , the market for AI had reached over a billion dollars. S and British governments to restore funding for academic research. Clark also presents factual data indicating that error rates in image processing tasks have fallen significantly since Goals can be explicitly defined, or can be induced. If the AI is programmed for " reinforcement learning ", goals can be implicitly induced by rewarding some types of behavior and punishing others. An algorithm is a set of unambiguous instructions that a mechanical computer can execute. A simple example of an algorithm is the following recipe for optimal play at tic-tac-toe: Otherwise, if a move "forks" to create two threats at once, play that move. Otherwise, take the center square if it is free. Otherwise, if your opponent has played in a corner, take the opposite corner. Otherwise, take an empty corner if one exists. Otherwise, take any empty square. Many AI algorithms are capable of learning from data; they can enhance themselves by learning new heuristics strategies, or "rules of thumb", that have worked well in the past , or can themselves write other algorithms. Some of the "learners" described below, including Bayesian networks, decision trees, and nearest-neighbor, could theoretically, if given infinite data, time, and memory, learn to approximate any function , including whatever combination of mathematical functions would best describe the entire world. These learners could therefore, in theory, derive all possible knowledge, by considering every possible hypothesis and matching it against the data. In practice, it is almost never possible to consider every possibility, because of the phenomenon of " combinatorial explosion ", where the amount of time needed to solve a problem grows exponentially. Much of AI research involves figuring out how to identify and avoid considering broad swaths of possibilities that are unlikely to be fruitful. A second, more general, approach is Bayesian inference: The third major approach, extremely popular in routine business AI applications, are analogizers such as SVM and nearest-neighbor: These four main approaches can overlap with each other and with evolutionary systems; for example, neural nets can learn to make inferences, to generalize, and to make analogies. Some systems implicitly or explicitly use multiple of these approaches, alongside many other AI and non-AI algorithms; [61] the best approach is often different depending on the problem. Learning algorithms work on the basis that strategies, algorithms, and inferences that worked well in the past are likely to continue working well in the future. These inferences can be obvious, such as "since the sun rose every morning for the last 10, days, it will probably rise tomorrow morning as well". The simplest theory that explains the data is the likeliest. Therefore, to be successful, a learner must be designed such that it prefers simpler theories to complex theories, except in cases where the complex theory is proven substantially better. Settling on a bad, overly complex theory gerrymandered to fit all the past training data is known as overfitting. Many systems attempt to reduce overfitting by rewarding a theory in accordance with how well it fits the data, but penalizing the theory in accordance with how complex the theory is. A toy example is that an image classifier trained only on pictures of brown horses and black cats might conclude that all brown patches are likely to be horses. Faintly superimposing such a pattern on a legitimate image results in an "adversarial" image that the system misclassifies. This enables even young children to easily make inferences like "If I roll this pen off a table, it will fall on the floor". Humans also have a powerful mechanism of " folk psychology " that helps them to interpret natural-language sentences such as "The city councilmen refused the demonstrators a permit because they advocated violence". A generic AI has difficulty inferring whether the councilmen or the demonstrators are the ones alleged to be advocating violence. For example, existing self-driving cars cannot reason about the location nor the intentions of pedestrians in the exact way that humans do, and instead must use non-human modes of reasoning to avoid accidents. The general problem of simulating or creating intelligence has been broken down into sub-problems. These consist of particular traits or capabilities that researchers expect an intelligent system to display. The traits described below have received the most attention. They solve most of their problems using fast, intuitive judgements. Knowledge representation and Commonsense knowledge Knowledge representation [80] and knowledge engineering [81] are central to classical AI research. Some "expert systems" attempt to gather together explicit knowledge possessed by experts in some narrow domain. In addition, some projects attempt to gather the "commonsense knowledge" known to the average person into a database containing extensive knowledge about the world. Among the things a comprehensive commonsense knowledge base would contain are: A representation of "what exists" is an ontology: The semantics of these are captured as description logic concepts, roles, and individuals, and typically implemented as classes, properties, and individuals in the Web Ontology Language. Such formal knowledge representations can be used in content-based indexing and retrieval, [88] scene interpretation, [89] clinical decision support, [90] knowledge discovery mining "interesting" and actionable inferences from large databases , [91] and other areas. Default reasoning and the qualification problem Many of the things people know take the form of "working assumptions". For example, if a bird comes up in conversation, people typically picture an animal that is fist sized, sings, and flies. None of these things are true about all birds. John McCarthy identified this problem in [93] as the qualification problem: Almost nothing is simply true or false in the way that abstract logic requires. AI research has explored a number of solutions to this problem. Research projects that attempt to build a complete knowledge base of commonsense knowledge e. For example, a chess master will avoid a particular chess position because it "feels too exposed" [96] or an art critic can take one look at a statue and realize that it is a fake. As with the related problem of sub-symbolic reasoning, it is hoped that situated AI , computational intelligence , or statistical AI will provide ways to represent this kind of knowledge. Automated planning and scheduling Intelligent agents must be able to set goals and achieve them. This calls for an agent that can not only assess its environment and make predictions, but also evaluate its predictions and adapt based on its assessment.

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

3: JÃ©rÃ´me Dokic â€” Open MIND

The current debate between theory theory and simulation theory on the nature of mentalisation has reached no consensus yet, although many now think that some hybrid theory is needed.

References and Further Reading 1. Various Concepts of Consciousness The concept of consciousness is notoriously ambiguous. It is important first to make several distinctions and to define related terms. We sometimes speak of an individual mental state, such as a pain or perception, as conscious. However, some kind of state consciousness is often implied by creature consciousness, that is, the organism is having conscious mental states. Most contemporary theories of consciousness are aimed at explaining state consciousness; that is, explaining what makes a mental state a conscious mental state. More common is the belief that we can be aware of external objects in some unconscious sense, for example, during cases of subliminal perception. Finally, it is not clear that consciousness ought to be restricted to attention. An organism, such as a bat, is conscious if it is able to experience the outer world through its echo-locatory senses. There is also something it is like to be a conscious creature whereas there is nothing it is like to be, for example, a table or tree. For example, philosophers sometimes refer to conscious states as phenomenal or qualitative states. There is significant disagreement over the nature, and even the existence, of qualia, but they are perhaps most frequently understood as the felt properties or qualities of conscious states. The former is very much in line with the Nagelian notion described above. Access consciousness is therefore more of a functional notion; that is, concerned with what such states do. Block himself argues that neither sense of consciousness implies the other, while others urge that there is a more intimate connection between the two. Some History on the Topic Interest in the nature of conscious experience has no doubt been around for as long as there have been reflective humans. It would be impossible here to survey the entire history, but a few highlights are in order. In the history of Western philosophy, which is the focus of this entry, important writings on human nature and the soul and mind go back to ancient philosophers, such as Plato. As we shall see, Descartes argued that the mind is a non-physical substance distinct from the body. He also did not believe in the existence of unconscious mental states, a view certainly not widely held today. Our mental states are, according to Descartes, infallibly transparent to introspection. Perhaps the most important philosopher of the period explicitly to endorse the existence of unconscious mental states was G. He also importantly distinguished between perception and apperception, roughly the difference between outer-directed consciousness and self-consciousness see Gennaro for some discussion. The most important detailed theory of mind in the early modern period was developed by Immanuel Kant. Although he owes a great debt to his immediate predecessors, Kant is arguably the most important philosopher since Plato and Aristotle and is highly relevant today. Kant basically thought that an adequate account of phenomenal consciousness involved far more than any of his predecessors had considered. Over the past one hundred years or so, however, research on consciousness has taken off in many important directions. In psychology, with the notable exception of the virtual banishment of consciousness by behaviorist psychologists e. The writings of such figures as Wilhelm Wundt , William James and Alfred Titchener are good examples of this approach. The work of Sigmund Freud was very important, at minimum, in bringing about the near universal acceptance of the existence of unconscious mental states and processes. It must, however, be kept in mind that none of the above had very much scientific knowledge about the detailed workings of the brain. The relatively recent development of neurophysiology is, in part, also responsible for the unprecedented interdisciplinary research interest in consciousness, particularly since the s. There are now several important journals devoted entirely to the study of consciousness: For a small sample of introductory texts and important anthologies, see Kim , Gennaro b, Block et. The Metaphysics of Consciousness: Dualism Metaphysics is the branch of philosophy concerned with the ultimate nature of reality. There are two broad traditional and competing metaphysical views concerning the nature of the mind and conscious mental states: While there are many versions of each, the former generally holds that the conscious mind or a conscious mental state is non-physical in some sense. On the other hand, materialists hold that the mind is the brain, or, more accurately, that conscious mental activity is identical with neural activity. For something to be non-physical, it must literally be outside the realm of physics; that is, not in space at all and undetectable in principle by the instruments of physics. However, something might be physical but not material in this sense, such as an electromagnetic or energy field. Thus, to say that the mind is non-physical is to say something much stronger than that it is non-material. Dualists, then, tend to believe that conscious mental states or minds are radically different from anything in the physical world at all. General Support and Related Issues There are a number of reasons why some version of dualism has been held throughout the centuries. For one thing, especially from the introspective or first-person perspective, our conscious mental states just do not seem like physical things or processes. That is, when we reflect on our conscious perceptions, pains, and desires, they do not seem to be physical in any sense. Consciousness seems to be a unique aspect of the world not to be understood in any physical way. Although materialists will urge that this completely ignores the more scientific third-person perspective on the nature of consciousness and mind, this idea continues to have force for many today. The metaphysical conclusion ultimately drawn is that consciousness cannot be identical with anything physical, partly because there is no essential conceptual connection between the mental and the physical. Arguments such as these go back to Descartes and continue to be used today in various ways Kripke , Chalmers , but it is highly controversial as to whether they succeed in showing that materialism is false. Materialists have replied in various ways to such arguments and the relevant literature has grown dramatically in recent years. Historically, there is also the clear link between dualism and a belief in immortality, and hence a more theistic perspective than one tends to find among materialists. Indeed, belief in dualism is often explicitly theologically motivated. If the conscious mind is not physical, it seems more plausible to believe in the possibility of life after bodily death. On the other hand, if conscious mental activity is identical with brain activity, then it would seem that when all brain activity ceases, so do all conscious experiences and thus no immortality. After all, what do many people believe continues after bodily death? There is perhaps a similar historical connection to a belief in free will, which is of course a major topic in its own right. To put it another way: Although materialism may not logically rule out immortality or free will, materialists will likely often reply that such traditional, perhaps even outdated or pre-scientific beliefs simply ought to be rejected to the extent that they conflict with materialism. After all, if the weight of the evidence points toward materialism and away from dualism, then so much the worse for those related views. Somewhat related to the issue of immortality, the existence of near death experiences is also used as some evidence for dualism and immortality. In response, materialists will point out that such experiences can be artificially induced in various experimental situations, and that starving the brain of oxygen is known to cause hallucinations. Various paranormal and psychic phenomena, such as clairvoyance, faith healing, and mind-reading, are sometimes also cited as evidence for dualism. However, materialists and even many dualists will first likely wish to be skeptical of the alleged phenomena themselves for numerous reasons. There are many modern day charlatans who should make us seriously question whether there really are such phenomena or mental abilities in the first place. Second, it is not quite clear just how dualism follows from such phenomena even if they are genuine. A materialist, or physicalist at least, might insist that though such phenomena are puzzling and perhaps currently difficult to explain in physical terms, they are nonetheless ultimately physical in nature; for example, having to do with very unusual transfers of energy in the physical world. The dualist advantage is perhaps not as obvious as one might think, and we need not jump to supernatural conclusions so quickly. For example, my desire to drink something cold causes my body to move to the refrigerator and get something to drink and, conversely, kicking me in the shin will cause me to feel a pain and get angry. But a modern day interactionist would certainly wish to treat various areas of the brain as the location of such interactions. Three serious objections are briefly worth noting here. The first is simply the issue of just how does or could such radically different substances causally interact. How something non-physical causally interacts with something physical, such as the brain? No such explanation is forthcoming or is perhaps even possible, according to materialists. Moreover, if causation involves a transfer of energy from cause to effect, then how is that possible if the mind is really non-physical? So any loss of energy in the cause must be passed along as a corresponding gain of energy in the effect, as in standard billiard ball examples. But if interactionism is true, then when mental events cause physical events, energy would literally come into the physical word. On the other hand, when bodily events cause mental events, energy would literally go out of the physical world. At the least, there is a very peculiar and unique notion of energy involved, unless one wished, even more radically, to deny the conservation principle itself. Third, some materialists might also use the well-known fact that brain damage even to very specific areas of the brain causes mental defects as a serious objection to interactionism and thus as support for materialism. This has of course been known for many centuries, but the level of detailed knowledge has increased dramatically in recent years. Now a dualist might reply that such phenomena do not absolutely refute her metaphysical position since it could be replied that damage to the brain simply causes corresponding damage to the mind. However, this raises a host of other questions: Why not opt for the simpler explanation, i. Will the severe amnesic at the end of life on Earth retain such a deficit in the afterlife? If proper mental functioning still depends on proper brain functioning, then is dualism really in no better position to offer hope for immortality? It should be noted that there is also another less popular form of substance dualism called parallelism, which denies the causal interaction between the non-physical mental and physical bodily realms. It seems fair to say that it encounters even more serious objections than interactionism. Other Forms of Dualism While a detailed survey of all varieties of dualism is beyond the scope of this entry, it is at least important to note here that the main and most popular form of dualism today is called property dualism. Substance dualism has largely fallen out of favor at least in most philosophical circles, though there are important exceptions e. Property dualism, on the other hand, is a more modest version of dualism and it holds that there are mental properties that is, characteristics or aspects of things that are neither identical with nor reducible to physical properties. There are actually several different kinds of property dualism, but what they have in common is the idea that conscious properties, such as the color qualia involved in a conscious experience of a visual perception, cannot be explained in purely physical terms and, thus, are not themselves to be identified with any brain state or process. Two other views worth mentioning are epiphenomenalism and panpsychism. The latter is the somewhat eccentric view that all things in physical reality, even down to micro-particles, have some mental properties. All substances have a mental aspect, though it is not always clear exactly how to characterize or test such a claim. Finally, although not a form of dualism, idealism holds that there are only immaterial mental substances, a view more common in the Eastern tradition. The most prominent Western proponent of idealism was 18th century empiricist George Berkeley. The idealist agrees with the substance dualist, however, that minds are non-physical, but then denies the existence of mind-independent physical substances altogether. Such a view faces a number of serious objections, and it also requires a belief in the existence of God.

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

4: Simulation and knowledge of action (eBook,) [www.enganchecubano.com]

Simulation and Knowledge of Action (Advances in Consciousness Research). www.enganchecubano.com is the world's largest independent marketplace for used, rare and out of print books.

Other important capabilities include the ability to sense e. Tests for confirming human-level AGI [13] [edit] The Turing Test Turing A machine and a human both converse sight unseen with a second human, who must evaluate which of the two is the machine, which passes the test if it can fool the evaluator a significant fraction of the time. Turing does not prescribe what should qualify as intelligence, only that knowing that it is a machine should not disqualify it. The Coffee Test Wozniak A machine is required to enter an average American home and figure out how to make coffee: The Robot College Student Test Goertzel A machine enrolls in a university, taking and passing the same classes that humans would, and obtaining a degree. The Employment Test Nilsson A machine works an economically important job, performing at least as well as humans in the same job. The flat pack furniture test Tony Severyns A machine is required to unpack and assemble an item of flat-packed furniture. It has to read the instructions and assemble the item as described, correctly installing all fixtures. Problems requiring AGI to solve[edit] Main article: AI-complete The most difficult problems for computers are informally known as "AI-complete" or "AI-hard", implying that solving them is equivalent to the general aptitude of human intelligence, or strong AI, beyond the capabilities of a purpose-specific algorithm. History of artificial intelligence Modern AI research began in the mid s. As AI pioneer Herbert A. Simon wrote in Of note is the fact that AI pioneer Marvin Minsky was a consultant [20] on the project of making HAL as realistic as possible according to the consensus predictions of the time; Crevier quotes him as having said on the subject in , "Within a generation Funding agencies became skeptical of AGI and put researchers under increasing pressure to produce useful "applied AI". By the s, AI researchers had gained a reputation for making vain promises. They became reluctant to make predictions at all [26] and to avoid any mention of "human level" artificial intelligence for fear of being labeled "wild-eyed dreamer[s]. Artificial intelligence In the s and early 21st century, mainstream AI has achieved far greater commercial success and academic respectability by focusing on specific sub-problems where they can produce verifiable results and commercial applications, such as artificial neural networks , computer vision or data mining. Most mainstream AI researchers hope that strong AI can be developed by combining the programs that solve various sub-problems using an integrated agent architecture , cognitive architecture or subsumption architecture. Hans Moravec wrote in Fully intelligent machines will result when the metaphorical golden spike is driven uniting the two efforts. If the grounding considerations in this paper are valid, then this expectation is hopelessly modular and there is really only one viable route from sense to symbols: A free-floating symbolic level like the software level of a computer will never be reached by this route or vice versa â€” nor is it clear why we should even try to reach such a level, since it looks as if getting there would just amount to uprooting our symbols from their intrinsic meanings thereby merely reducing ourselves to the functional equivalent of a programmable computer. The term was introduced by Mark Gubrud in [32] in a discussion of the implications of fully automated military production and operations. AGI research activity in was described by Pei Wang and Ben Goertzel [33] as "producing publications and preliminary results". As yet, most AI researchers have devoted little attention to AGI, with some claiming that intelligence is too complex to be completely replicated in the near term. However, a small number of computer scientists are active in AGI research, and many of this group are contributing to a series of AGI conferences. The research is extremely diverse and often pioneering in nature. In the introduction to his book, [31] Goertzel says that estimates of the time needed before a truly flexible AGI is built vary from 10 years to over a century, but the consensus in the AGI research community seems to be that the timeline discussed by Ray Kurzweil in The Singularity is Near [1] i. Finally, projects such as the Human Brain Project [39] have the goal of building a functioning simulation of the human brain. Mind uploading A popular approach discussed to achieving general intelligent action is whole brain emulation. A low-level brain model is built by scanning and mapping a biological brain in detail and copying its state into a computer system or another computational device. The computer runs a simulation model so faithful to the original that it will behave in essentially the same way as the original brain, or for all practical purposes, indistinguishably. It is discussed in artificial intelligence research [34] as an approach to strong AI. Neuroimaging technologies that could deliver the necessary detailed understanding are improving rapidly, and futurist Ray Kurzweil in the book The Singularity Is Near [1] predicts that a map of sufficient quality will become available on a similar timescale to the required computing power. Early estimates[edit] Estimates of how much processing power is needed to emulate a human brain at various levels from Ray Kurzweil, and Anders Sandberg and Nick Bostrom , along with the fastest supercomputer from TOP mapped by year. Note the logarithmic scale and exponential trendline, which assumes the computational capacity doubles every 1. Kurzweil believes that mind uploading will be possible at neural simulation, while the Sandberg, Bostrom report is less certain about where consciousness arises. The human brain has a huge number of synapses. Each of the one hundred billion neurons has on average 7, synaptic connections to other neurons. It has been estimated that the brain of a three-year-old child has about synapses 1 quadrillion. This number declines with age, stabilizing by adulthood. He used this figure to predict the necessary hardware would be available sometime between and , if the exponential growth in computer power at the time of writing continued. Modelling the neurons in more detail[edit] The artificial neuron model assumed by Kurzweil and used in many current artificial neural network implementations is simple compared with biological neurons. A brain simulation would likely have to capture the detailed cellular behaviour of biological neurons , presently only understood in the broadest of outlines. In addition the estimates do not account for glial cells , which are at least as numerous as neurons, and which may outnumber neurons by as much as The Artificial Intelligence System project implemented non-real time simulations of a "brain" with neurons in It took 50 days on a cluster of 27 processors to simulate 1 second of a model. Neuro-silicon interfaces have been proposed as an alternative implementation strategy that may scale better. His results do not depend on the number of glial cells, nor on what kinds of processing neurons perform where. Complications and criticisms of AI approaches based on simulation[edit] A fundamental criticism of the simulated brain approach derives from embodied cognition where human embodiment is taken as an essential aspect of human intelligence. Many researchers believe that embodiment is necessary to ground meaning. Goertzel [34] proposes virtual embodiment like Second Life , but it is not yet known whether this would be sufficient. According to the brain power estimates used by Kurzweil and Moravec , this computer should be capable of supporting a simulation of a bee brain, but despite some interest [52] no such simulation exists[citation needed]. There are at least three reasons for this: Firstly, the neuron model seems to be oversimplified see next section. Thirdly, even if our understanding of cognition advances sufficiently, early simulation programs are likely to be very inefficient and will, therefore, need considerably more hardware. Fourthly, the brain of an organism, while critical, may not be an appropriate boundary for a cognitive model. To simulate a bee brain, it may be necessary to simulate the body, and the environment. The Extended Mind thesis formalizes the philosophical concept, and research into cephalopods has demonstrated clear examples of a decentralized system. One estimate puts the human brain at about billion neurons and trillion synapses. Artificial consciousness research[edit] Main article: In an early effort Igor Aleksander [59] argued that the principles for creating a conscious machine already existed but that it would take forty years to train such a machine to understand language. Relationship to "strong AI"[edit] See also: The word "mind" has a specific meaning for philosophers, as used in "the mind body problem " or "the philosophy of mind ". An artificial intelligence system can only act like it thinks and has a mind. The first one is called "the strong AI hypothesis" and the second is "the weak AI hypothesis" because the first one makes the stronger statement: Searle referred to the "strong AI hypothesis" as "strong AI". This usage is also common in academic AI research and textbooks. Possible explanations for the slow progress of AI research[edit] See also: As William Clocksin wrote in The failed predictions that have been promised by AI researchers and the lack of a complete understanding of human behaviors have helped diminish the primary idea of human-level AI. The intricacy of scientific problems and the need to fully understand the human brain through psychology and neurophysiology have limited many researchers from emulating the function of the human brain into a computer hardware. These explanations are not necessarily guaranteed to be the fundamental causes for the delay in achieving strong AI, but they are widely agreed by numerous researchers. There have been many AI researchers that debate over the idea whether machines should be created with emotions. There are no emotions in typical models of AI and some researchers say programming emotions into machines allows them to have a mind of their own. Consciousness[edit] There are other aspects of the human mind besides intelligence that are relevant to the concept of strong AI which play a major role in science fiction and the ethics of artificial intelligence: The ability to "feel" perceptions or emotions subjectively. The capacity for wisdom. These traits have a moral dimension, because a machine with this form of strong AI may have legal rights, analogous to the rights of non-human animals. Also, Bill Joy , among others, argues a machine with these traits may be a threat to human life or dignity. The role of consciousness is not clear, and currently there is no agreed test for its presence. If a machine is built with a device that simulates the neural correlates of consciousness , would it automatically have self-awareness? It is also possible that some of these properties, such as sentience, naturally emerge from a fully intelligent machine, or that it becomes natural to ascribe these properties to machines once they begin to act in a way that is clearly intelligent. For example, intelligent action may be sufficient for sentience, rather than the other way around. In science fiction, AGI is associated with traits such as consciousness , sentience , sapience , and self-awareness observed in living beings. However, according to philosopher John Searle , it is an open question whether general intelligence is sufficient for consciousness. The strong AI hypothesis is the claim that a computer which behaves as intelligently as a person must also necessarily have a mind and consciousness. AGI refers only to the amount of intelligence that the machine displays, with or without a mind. Controversies and dangers[edit].

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

5: Artificial intelligence - Wikipedia

This volume is part of a series that provides a forum for scholars from different scientific disciplines who study consciousness in its multifaceted aspects. Topics covered include: Simulation theory and mental concepts and neurophysiological evidence for simulation of action.

Introductory Readings Morsella, E. Homing in on consciousness in the nervous system: Behavioral and Brain Sciences [Target Article], 39, The follow-up to the article in Psychological Review below , this is the theoretical foundation of research in our lab. It focuses on the primary function of conscious processing in the nervous system. The evidence for different aspects of the theory can be found in five sections following Books and Special Volumes. These sections are organized by theme. Behavioral and Brain Sciences, 39, The function of phenomenal states: Psychological Review, , The prospective nature of voluntary action: Insights from the reflexive imagery task. Review of General Psychology, 20, This review focuses on the prospective aspects of voluntary action and covers the findings from the first handful of experiments using the reflexive imagery task, which is the primary paradigm used by the lab. The mechanisms of human action: Gollwitzer, Oxford handbook of human action pp. What is an output? Psychological Inquiry, 21, Perspectives on Psychological Science, 3, Cognitive and neural components of the phenomenology of agency. Involuntary polymodal imagery involving olfaction, audition, touch, taste, and vision. Consciousness and Cognition, 62, In the field of the scientific study of consciousness, Consciousness and Cognition is the premier journal. The reflexive imagery task: An experimental paradigm for neuroimaging. AIMS Neuroscience, 5, Involuntary entry into consciousness from the activation of sets: Object counting and color naming. Oxford handbook of human action. Link Consciousness and action control, a special issue 18 articles of Frontiers in Psychology: Communication, cognition, language, and identity. A Festschrift in the honor of Robert M. See next section for consciousness research based on the Reflexive Imagery Task, the new paradigm developed by the lab. Number 90, total Morsella, E. Action-oriented understanding of consciousness and the structure of experience. Toward action-oriented views on cognitive science pp. Subjective aspects of cognitive control at different stages of processing 7-experiment article. The essence of conscious conflict: Subjective effects of sustaining incompatible intentions. Neural correlates of the essence of conscious conflict: Experimental Brain Research, , Unconscious and conscious component processes. What makes us conscious is not what makes us human. Animal Sentience, , Link. Click here for a blog about this article. Competition between cognitive control and encapsulated, unconscious inferences: Frontiers in Psychology, 7: The interdependence between conscious and unconscious processes. In search of the fundamentals of mind pp. The difference between conscious and unconscious brain circuits. The subjective effort of everyday mental tasks: Attending, assessing, and choosing. Motivation and Emotion, 38, The olfactory system as the gateway to the neural correlates of consciousness. Frontiers in Psychology, 4, The conscious control of behavior: The subjective aspects of self-control: Theory and experimental paradigms. Sense of agency as a function of intra-psychic conflict. Journal of Mind and Behavior, 36, Voluntary action and the illusion of conscious will. Conscious versus unconscious processes in action control. Homing in on the brain mechanisms linked to consciousness: Buffer of the perception-and-action interface. Current perspectives on a science of consciousness pp. The three pillars of volition: Phenomenal states, ideomotor processing, and the skeletal muscle system. Voluntary action and the three forms of binding in the brain. The subjective aspects of agency, cognitive control, and self-regulation: Findings from the action and consciousness laboratory. Volume contains a contribution from Dr. The phenomenology of quitting: Effects from repetition and cognitive effort. Korean Journal of Cognitive Science, 23, Cognitive conflict and consciousness pp. A unifying concept in social psychology. Did I read or did I name? Conscious states are a crosstalk mechanism for only a subset of brain processes. Journal of Cosmology, 14, Special issue edited by Sir Roger Penrose. Sources of avoidance motivation: Valence effects from physical effort and mental rotation. Motivation and Emotion, 35, Creating illusory intentions through a phony brain-computer interface. Consciousness and Cognition, 19, The spontaneous thoughts of the night: How future tasks breed intrusive cognitions. Social Cognition, 28, Preliminary evidence from Stroop interference. Journal of Communications Research, 3, Controlled-reflective processes arise from integrative action-goal selection in the ventral pathway. European Journal of Personality, 24, Minimal neuroanatomy for a conscious brain: Homing in on the networks constituting consciousness. Neural Networks, 23, The function of consciousness: Psychological Reports, 93, Theory and practice pp. This review focuses on the prospective aspects of voluntary action and covers the findings from the first handful of RIT experiments.

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

6: Neuroethics and Philosophy

Buy Simulation and Knowledge of Action (Advances in Consciousness Research) First Edition by Jerome Dokic, Joelle Proust (ISBN:) from Amazon's Book Store. Everyday low prices and free delivery on eligible orders.

Baars, James Newman, J. Taylor Neuronal mechanisms of consciousness: A Relational Global Workspace framework. Pages in S. The second Tucson discussions and debates. That convergence has now become so persuasive that we believe we are working within substantially the same broad framework. This architecture is relational, in the sense that it continuously mediates the interaction of input with memory. While our approaches overlap in a number of ways, each of us tends to focus on different areas of detail. What is most striking, and we believe significant, is the extent of consensus, which we believe to be consistent with other contemporary approaches by Weiskrantz, Gray, Crick and Koch, Edelman, Gazzaniga, Newell and colleagues, Posner, Baddeley, and a number of others. We suggest that cognitive neuroscience is moving toward a shared understanding of consciousness in the brain. Taylor Constructing the Relational Mind PSYCHE, 4 10 , June "The "relational mind" approach to the inner content of consciousness is developed in terms of various control structures and processing strategies and their possible neurobiological identifications in brain sites. This leads naturally to a division of consciousness into a passive and an active part. A global control structure for the "single strand" aspect of consciousness is proposed as the thalamo-nucleus reticularis thalami-cortex coupled system, which is related to experimental data on the electrical stimulation of awareness. Local control, in terms of excitatory transfer from pre-processing sites to posterior working memory regions, is supported by data on subliminal perception timing and disambiguation of poorly defined percepts. The inner content of consciousness is understood as arising from the resulting relational features between inputs and stored pre-processing and episodic memories. Strong analogies are drawn between emergent properties of the model and suggested properties of "raw feels", supporting the thesis that working memories are the initial sites for the emergence of phenomenal awareness, and the frontal lobes for its further adumbration in terms of higher cognitive processing, including the creation of self. How conscious experience and working memory interact. Trends Cogn Sci Apr;7 4: Global Workspace theory suggests that consciousness is needed to recruit unconscious specialized networks that carry out detailed working memory functions. The IDA model provides a fine-grained analysis of this process, specifically of two classical working-memory tasks, verbal rehearsal and the utilization of a visual image. In the process, new light is shed on the interactions between conscious and unconscious aspects of working memory. Conscious Cogn Jun;4 2: The simplest hypothesis is that both components are necessary and sufficient to support conscious perceptual experience. Neural Netw Oct 1;10 7: Integrating experimental data and models from cognitive psychology, AI and neuroscience, we present a neurocognitive model in which consciousness is defined as a global integration and dissemination system - nested in a large-scale, distributed array of specialized bioprocessors - which controls the allocation of the processing resources of the central nervous system. The basic circuitry of this neural system is reasonably well understood, and can be modeled, to a first approximation, employing neural network principles. Journal of Consciousness Studies, 4, No. Conscious experience seems to create access to many independent knowledge sources in the brain, most of them quite unconscious. When it comes to sensory consciousness, however, the brain shows little difference between humans and many other mammals. How does a serial, integrated and very limited stream of consciousness emerge from a nervous system that is mostly unconscious, distributed, parallel and of enormous capacity? Ciba Found Symp ; Conscious experience on the other hand is traditionally viewed as a serial stream that integrates different sources of information but is limited to only one internally consistent content at any given moment. Global Workspace theory suggests that conscious experience emerges from a nervous system in which multiple input processors compete for access to a broadcasting capability; the winning processor can disseminate its information globally throughout the brain. Global workspace architectures have been widely employed in computer systems to integrate separate modules when they must work together to solve a novel problem or to control a coherent new response. The theory articulates a series of increasingly complex models, able to account for more and more evidence about conscious functioning, from perceptual consciousness to conscious problem-solving, voluntary control of action, and directed attention. Global Workspace theory is consistent with, but not reducible to, other theories of limited-capacity mechanisms. Global workspace architectures must show competition for input to a neural global workspace and global distribution of its output. Brain structures that are demonstrably required for normal conscious experience can carry out these two functions. The theory makes testable predictions, especially for newly emerging, high-speed brain imaging technology. The conscious access hypothesis: Trends Cogn Sci Jan 1;6 1: The theory had little impact at first, for three reasons: Recent neuroimaging evidence appears broadly to support the hypothesis, which has implications for perception, learning, working memory, voluntary control, attention and self systems in the brain. Towards a cognitive neuroscience of consciousness: Cognition Apr;79 We isolate three major empirical observations that any theory of consciousness should incorporate, namely 1 a considerable amount of processing is possible without consciousness, 2 attention is a prerequisite of consciousness, and 3 consciousness is required for some specific cognitive tasks, including those that require durable information maintenance, novel combinations of operations, or the spontaneous generation of intentional behavior. We then propose a theoretical framework that synthesizes those facts: This framework postulates that, at any given time, many modular cerebral networks are active in parallel and process information in an unconscious manner. An information becomes conscious, however, if the neural population that represents it is mobilized by top-down attentional amplification into a brain-scale state of coherent activity that involves many neurons distributed throughout the brain. We postulate that this global availability of information through the workspace is what we subjectively experience as a conscious state. A complete theory of consciousness should explain why some cognitive and cerebral representations can be permanently or temporarily inaccessible to consciousness, what is the range of possible conscious contents, how they map onto specific cerebral circuits, and whether a generic neuronal mechanism underlies all of them. We confront the workspace model with those issues and identify novel experimental predictions. Neurophysiological, anatomical, and brain-imaging data strongly argue for a major role of prefrontal cortex, anterior cingulate, and the areas that connect to them, in creating the postulated brain-scale workspace. It distinguishes two main computational spaces: Workspace neurons are mobilized in effortful tasks for which the specialized processors do not suffice. They selectively mobilize or suppress, through descending connections, the contribution of specific processor neurons. In the course of task performance, workspace neurons become spontaneously coactivated, forming discrete though variable spatio-temporal patterns subject to modulation by vigilance signals and to selection by reward signals. A computer simulation of the Stroop task shows workspace activation to increase during acquisition of a novel task, effortful execution, and after errors. We outline predictions for spatio-temporal activation patterns during brain imaging, particularly about the contribution of dorsolateral prefrontal cortex and anterior cingulate to the workspace. Investigating the biology of consciousness. Nonetheless, by combining cognitive and neurobiological methods, it is possible to approach consciousness, to describe its cognitive nature, its behavioural correlates, its possible evolutionary origin and functional role; last but not least, it is possible to investigate its neuroanatomical and neurophysiological underpinnings. In this brief essay I distinguish between two kinds of consciousness: Core consciousness corresponds to the transient process that is incessantly generated relative to any object with which an organism interacts, and during which a transient core self and transient sense of knowing are automatically generated. Core consciousness requires neither language nor working memory, and needs only a brief short-term memory. Extended consciousness is a more complex process. It depends on the gradual build-up of an autobiographical self, a set of conceptual memories pertaining to both past and anticipated experiences of an individual, and it requires conventional memory. Extended consciousness is enhanced by language. Keenan, Markus Nowak, Troels W. Sackeim, and Sarah H. Such explicit "autonoetic consciousness" is thought to emerge by retrieval of memory of personally experienced events "episodic memory". During episodic retrieval, functional imaging studies consistently show differential activity in medial prefrontal and medial parietal cortices. With positron-emission tomography, we here show that these medial regions are functionally connected and interact with lateral regions that are activated according to the degree of self-reference. During retrieval of previous judgments of Oneself, Best Friend, and the Danish Queen, activation increased in the left lateral temporal cortex and decreased in the right inferior parietal region with decreasing self-reference. Functionally, the former region was preferentially connected to medial prefrontal cortex, the latter to medial parietal. The medial parietal region may, then, be conceived of as a nodal structure in self-representation, functionally connected to both the right parietal and the medial prefrontal cortices. To determine whether medial parietal cortex in this network is essential for episodic memory retrieval with self-representation, we used transcranial magnetic stimulation over the region to transiently disturb neuronal circuitry. There was a decrease in the efficiency of retrieval of previous judgment of mental Self compared with retrieval of judgment of Other with transcranial magnetic stimulation at a latency of ms, confirming the hypothesis. This network is strikingly similar to the network of the resting conscious state, suggesting that self-monitoring is a core function in resting consciousness. Gusnard, Erbil Akbudak, Gordon L. Shulman, and Marcus E. Raichle Medial prefrontal cortex and self-referential mental activity: Relation to a default mode of brain function PNAS This high metabolic rate and this behavior suggest the existence of an organized mode of default brain function, elements of which may be either attenuated or enhanced. Extant data suggest that these MPFC regions may contribute to the neural instantiation of aspects of the multifaceted "self. In this functional magnetic resonance imaging fMRI study, subjects made two judgments, one self-referential, the other not, in response to affectively normed pictures: These increases were accompanied by decreases in both active task conditions in ventral MPFC. These results support the view that dorsal and ventral MPFC are differentially influenced by attentiondemanding tasks and explicitly self-referential tasks. The presence of self-referential mental activity appears to be associated with increases from the baseline in dorsal MPFC. Reductions in ventral MPFC occurred consistent with the fact that attention-demanding tasks attenuate emotional processing. We posit that both self-referential mental activity and emotional processing represent elements of the default state as represented by activity in MPFC. We suggest that a useful way to explore the neurobiology of the self is to explore the nature of default state activity. Functional specialization within the anterior medial prefrontal cortex: Previous studies have shown that the aMPFC is involved in evaluative judgment and self-referential processes. Specifically, different sections of the aMPFC are differentially influenced by attention demanding processes. Whereas the dorsal section is supposed to be involved in self-referential processes, the ventral section is assumed to be attenuated during attention demanding processes. The present study investigates the involvement of the dorsal and ventral aMPFC in evaluative judgment by using functional magnetic resonance imaging with spin-echo echo-planar-imaging. Processes involved in evaluative judgment are attention-demanding, self-referential and activate regions in the dorsal and ventral section of the aMPFC. Attention demanding tasks do not necessarily lead to an attenuation of the ventral section of the aMPFC, a region mainly involved in emotional and affective processing. One brain, two selves.

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

7: Simulation and Knowledge of Action | Edited by JÃ©rÃ´me Dokic and JoÃ«lle Proust

Contributions to the development of theory and method in the study of consciousness; Series B: Research in Progress. Experimental, descriptive and clinical research in consciousness. Provides a forum for scholars from different scientific disciplines and fields of knowledge who study consciousness in its multifaceted aspects.

Table 1 Options for simulated events Drug administration can be simulated, and with the use of the drug recognition unit, the simulator will respond physiologically. For example, a simulated morphine injection will cause the pupil size of the mannequin to change and the respiratory rate, heart rate, and blood pressure to decrease. The response to any drug depends on the dose of the drug and the weight and clinical condition of the simulated patient at the time. Instructors can pause the simulation to review assessments, detect problems, or discuss treatment. Unlike the situation in a clinical setting, with GUS, mistakes cannot actually harm a human being and are therefore useful opportunities for learning. The subsequent consequences can be witnessed in the simulated scenario. With the use of this simulator, assessment skills; pharmacological, physiological, and pathophysiological concepts; and basic and advanced cardiac life support techniques can all be taught, reinforced, and evaluated. With GUS, the classroom is transformed into a realistic practice environment. Instructors can set up and control many variables in the clinical learning environment. A top-of-the-line high-fidelity human patient simulator is a large investment. They do have the advantages of being portable and easier to use than the HPS V6 is. Previous Section Next Section Simulation as an Essential Component of Clinical Education All healthcare professionals must have a combined knowledge of physical and behavioral science and technical and clinical education. Unfortunately, much of the technical and clinical learning often takes place in the clinical setting, posing risks to the safety of both patients and learners such that close supervision by experienced preceptors and instructors is required to avoid disastrous consequences. Yet in the clinical setting, preceptors often do not have control over the types of experiences a learner will have or the conditions under which skills can be observed, learned, or practiced. A new critical care nurse could potentially complete an entire orientation period and not experience a common or high-acuity event that the nurse must be competent to deal with in order to practice safely in an intensive care unit. In contrast with the real clinical setting, simulated clinical situations involve only a few safety concerns and allow instructors and preceptors to completely control the events. At Georgetown University School of Nursing and Health Studies, all 4 levels of the undergraduate curriculum include classroom, technological, and clinical instruction. GUS is used as an essential teaching tool in clinical nursing courses. Simulation sessions have also been incorporated into the curriculum of all of the graduate programs. The nurse anesthetist students are the most frequent users of the simulator. The laboratory features a hemodynamic monitor and an anesthesia machine with the appropriate gases. The students practice intubation, induction of anesthesia, continuous administration of anesthetic agents, and monitoring of level of consciousness. Developing and demonstrating critical-thinking skills are strongly emphasized during these simulation sessions. In their text Critical Thinking in Nursing, 17 Bandman and Bandman define critical thinking as follows: This examination covers scientific reasoning, includes the nursing process, decision making, and reasoning in controversial issues. The four types of reasoning that comprise critical thinking are deductive, inductive, informal or everyday, and practical. The scenarios are developed to require students to use classroom knowledge, incorporate assessment skills, and create and implement a plan. They are then given an opportunity to witness the outcome and evaluate their plan and make the appropriate changes if necessary. With the use of simulation, they can implement the entire nursing process and are required to think critically. Previous Section Next Section Simulation as an Essential Component of Critical Care Education Simulation is an excellent teaching and evaluation method for critical care and also for enhancing and evaluating critical thinking, problem solving, and team leading for proficient and competent senior staff. With the opening of a new cardiac surgery step-down unit, many nurses required additional training. After completing the traditional courses in cardiac surgery eg, electrocardiographic rhythms, pacemakers and temporary pacing wires, chest tubes, cardiac drugs, discharge teaching , the nursing staff for this new unit attended a simulation session. Three scenarios were developed to allow the nurses the opportunity to integrate and use the theory they were taught in class. High-frequency and high-acuity situations were selected, such as new-onset unstable atrial fibrillation in a patient who had valve replacement; new premature ventricular complexes leading to ventricular tachycardia after diuresis in a postoperative patient; and an inferior wall myocardial infarction, heart block, and papillary muscle rupture in a postoperative patient with history of cardiac disease. The nurses were able to demonstrate many skills through simulation, including interpretation of rhythms, evaluation of hemodynamic stability, and assessment and reporting of data. Although the skill levels and abilities of the nurses varied, evaluations of the simulation sessions were universally positive. The simulation session held at the end of the course for the intensive care unit received an equally positive response. This session was developed around a single scenario: The first part of the session focused on respiratory assessment, intubation, initiation and management of mechanical ventilation, and progression of acute respiratory distress syndrome. The second part focused on insertion of pulmonary and arterial catheters, hemodynamic monitoring, and pharmacological management. The nurses were able to assist with intubation and with insertion of catheters. Complications such as intubation of the right main bronchus and ventricular tachycardia were assessed and treated. Requiring the nurses to assemble the equipment while caring for a simulated patient in a safe and controlled environment was a great learning experience. Through the use of simulation, the nurses were able to think through their actions and the events without jeopardizing care of an actual patient. The sessions were originally scheduled to take 4 hours but invariably ran longer because the nurses requested to try or see a few more things. Previous Section Next Section Review of the Literature on Simulation The success with simulation as an educational tool in critical care education is not unique to George-town University. Articles in research and critical care literature 2â€" 4, 6, 10, 13, 15, 16, 20 indicate that compared with traditional methods, this method of teaching and evaluating learners is more realistic, enhances both acquisition and retention of knowledge, sharpens critical-thinking and psychomotor skills, and is more enjoyable. Issenberg et al 4 maintain that the use of simulation will reduce the pitfalls inherent in skills practice: The students were evaluated before and after the rotations by using a multiple-choice written examination, a skill station test, and an interactive simulation with a high-fidelity simulator. Although the test results before the rotations were similar for all 3 types of evaluations, the results after the rotations differed. The students performed much better on the written examination than on the simulation tests after the rotations, showing that although theory could be applied in a written case study, application of theory was not as easily demonstrated in a clinical simulation. Gordon et al 16 surveyed both students and educators about simulation as a teaching tool.

8: Action and Consciousness Laboratory

Get this from a library! Simulation and knowledge of action. [JÃ©rÃ´me Dokic; JoÃ«lle Proust;] -- The current debate between theory theory and simulation theory on the nature of mentalisation has reached no consensus yet, although many now think that some hybrid theory is needed.

9: Consciousness | Internet Encyclopedia of Philosophy

Consciousness Emerging: The Dynamics of Perception, Imagination, Action, Memory, Thought, and Language (Advances in Consciousness Research) by Renate Bartsch 39 Consciousness Recovered: Psychological Functions and Origins of Conscious Thought (Advances in Consciousness Research) by George Mandler.

SIMULATION AND KNOWLEDGE OF ACTION (ADVANCES IN CONSCIOUSNESS RESEARCH) pdf

The termination phase. Armor Battles of the Waffen SS, 1943-45 A last forlorn chronicle of despair James joyce dubliners A Writers Dream A Cowboys Heart Some Things Better Left Unsaid Haru library vc++ The teddy bear lovers companion. Fundamentals of engineering dennis k lieu The ruins of ambrai Form designer software Cancers and tumor markers Corporate governance in Latin America Alberto Chong and Florencio Lopez-de-Silanes Physician Practice Compliance Resource Manual (Ahlcc Compliance Series) Decision 3 : Whether you have children Bertrand russell science et religion The rational unified process an introduction 3rd edition Insight Guides Smart Guide New York Animal stories from Rudyard Kipling 2005 ford explorer sport trac owners manual Alternative methods of project selection DB2 for Windows for Dummies Lukes way of looking V. 6. Dombey son. Curriculum and Instruction in Nursing Tangipahoa and the forced migration Beholding God in many faces Cases in consumer law Social relations in the Mahdhis city 1st ser. Philip Massinger, by B. T. Spencer. John Ford, by Mary E. Cochnower. Samuel Butler, by Dan Gibso Track, enduro, and motocross-unless you fall over Florida in the Civil War (The Civil War History Series) Politics and power: the movement for woman?s rights V. 4. Studies in philosophical inquiry Max and Rubys Busy Week with Flash Cards Tent city, mosh pit, trampled clover At last there is nothing left to say The city and the battlefield: Coriolanus W.B. Yeats and Irish folklore The changing face of senior housing Jon Pynoos and Christy M. Nishita