

On the LAM: A Naturalized Model of *Petites Perceptions*

Christopher Stiso

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Master of Arts

In

Philosophy

Lydia Patton

Ted Parent

Matthew Priselac

April 30th, 2013

Blacksburg, VA

Keywords: Philosophy, Perception, Leibniz

On the LAM: A Naturalized Model of *Petites Perceptions*

Chris Stiso

ABSTRACT

I propose a model of perception by which unconscious sensory input can become conscious in a way that is entirely mechanical and 'non-spooky'. Based on the model of perception found in G.W. Leibniz philosophy, primarily the *New Essays on Human Understanding*, the Leibnizian Aggregate Model (LAM) of perception may, when fully matured through future investigation, provide a way for contemporary philosophers of mind to evade certain anti-materialist objections without having to flatly deny the intuitions that underlie them.

Table of Contents

| | |
|---|----|
| Abstract | ii |
| On the LAM: A Naturalized Model of <i>Petites Perceptions</i> | 1 |
| Bibliography | 18 |

On the LAM: A Naturalized Model of *Petites Perceptions*

The task of explaining the mind and its functions is in many ways more difficult for the physicalist as she cannot appeal to any thing or property beyond those observed and tested in physical science where her opponents need only appeal to mental states which need not be manipulable, testable, or locatable. To suppose that an act or faculty of cognition, used here in its broadest possible sense, could be captured by a purely physical system is often contrary to the naive posits of folk psychology that are ingrained in the public consciousness. Even in the present day where physicalism about the world, again in a very broad sense, is widely accepted the mind remains aloof; the last bastion of anti-materialism. While every advance in technology and method shrinks the list of phenomena that cannot be explained by a physical model, certain intractable problems continue to arise.

G. W. Leibniz's representationalism is widely acknowledged: throughout his published work and correspondence he makes it abundantly clear that he thinks that there is nothing more to mental phenomenology than representation. What is less often appreciated is that Leibniz's philosophy of mind allows for the existence of unconscious perceptions, which enables the him and other representationalists to provide a much more plausible account of phenomenological features of consciousness¹. I believe Leibniz's program of modeling perception as an aggregation of unconscious *petites perceptions* remains an unexamined route of inquiry that may prove fruitful for the present day physicalist.

¹ In calling entities "conscious" or "unconscious" I mean to say that they are or are not available to some consciousness, respectively. This is opposed to implying that the entities themselves have or lack consciousness.

While unconscious sensory input is an accepted occurrence in present day psychology and neuroscience it has not been connected to philosophical issues of perception. What I am proposing is a Leibnizian model of perception that draws from Leibniz's theory of unconscious *petites perceptions* to create a viable argumentative strategy for the present day physicalist. This model, which I will be referencing as the Leibnizian Aggregate Model (LAM), involves two commitments. First, a commitment to representation as the function of the mind. This is not a commitment to representationalism as a philosophical position opposed to direct realism, but rather a much weaker assertion that our mental phenomenology in some way captures a world of physical objects. A model that accepts representationalism may be a more natural fit for the LAM but, those that support direct realism need not reject the LAM outright at the first mention of representation; while Leibniz was, in the end, a committed representationalist about perception, the LAM has more leeway. The second, more vital commitment is that our most basic perceptions are not simple entities but composed of an aggregation of unconscious *petites perceptions*.

My project here is to elucidate and draw attention to a model of perception that was used by Leibniz, that of aggregations of unconscious perception, and, using this theoretical framework, construct a naturalized model. Additionally, I have hopes that, with further thought on the matter, this model may be able to ease the tension seen by many between contemporary naturalized models of cognition and certain longstanding intuitions about cognition. I will give brief introductions to how I foresee this method of investigation working with such large, intractable problems as the so-called “Hard

Problem of Consciousness”. Each of these problems will require a great deal of additional thought before my unconscious aggregation model can adequately address all the myriad issues associated with them but I feel there is some evidence that further research into these area will prove fruitful.

Terminology

The field of philosophy of mind has become littered with a great number of conceptually charged terms, many of which lack the level of precision that we would like form them. In order that I might be understood, I feel it is important that I lay out exactly what I mean when I employ certain terms in this paper.

Perception: The most important concept at work here is perception. In its broadest conceptualization, perception is a process that can be carried out by complex and sophisticated mammalian brains but also simple organisms—plants, bacteria, and even, arguably, viruses have the ability to take in environmental information and react—and non-living systems like robots and mechanical sensors. Obviously, this is too large a concept to work with and will need to be narrowed down. What I mean to convey with the concept is the process by which information about the physical world is taken in and utilized by a mammalian brain, specifically a human brain². In other words, how physical information—positions and distances of trees, ambient temperature and illumination, properties of different waves, etc.—is converted to 'brain information' that allow a human subject to interact with the environment.

Conversely, there is an intuition that in order for something to be a *perception* it

² While it is not only possible but, in my opinion, quite likely that the models or frameworks discussed may describe systems other than the human brain, it is irrelevant to the project at hand.

must be available to some consciousness. This concept relates more closely to what Leibniz would call an apperception. Apperceptions are necessarily conscious, higher-order mental acts.

Subjectivity: There is a way in which all perceptions are subjective, insofar as every perception is had by or instantiated in some thing. Whether this is a non-physical mind, a biological brain, or something else entirely does not matter, the information must belong to something in order for it to constitute a perception. This notion is tautological and therefore uninteresting. Beyond this, however, there is a notion of subjectivity that is not obvious and this has to do with what might be called the *Subjective Elements of Perception*. There is a popular intuition that the perceptions I have are, in some important way, *mine* and that they cannot be explained, described, or analyzed separately from me. This has led to the coining of philosophical terms like 'qualia' which are meant to capture this kind of intuition. I wish to avoid these issues in the main discussion as they are not important to the primary project; however, it is often hard to separate the philosophical issues of perception from those of subjective perception and, in light of this, I will give them some brief consideration in the final section.

Perception in Leibniz

Within the *New Essays*, Leibniz means something very specific when he uses the term 'perception', and, if I hope to use Leibniz' theory of perceptions, it is important to lay out as precisely as possible what he meant. Within his psychology Leibniz posits a number of mental entities. Among these entities there are three that seem to be involved in sensory mental acts: impressions, perceptions, and apperceptions. Impressions are

imprints made on sensory organs, necessarily not available to consciousness.

Apperceptions, which can be thought of perceptions of our own perceptions, are non-sensory.

A perception is a representation, an entity that is neither necessarily conscious or necessarily unconscious. Bare monads that have no faculty of consciousness are still capable of perceptions but, at the same time, those privileged monads that have the faculty of consciousness can become aware of their perceptions. In Leibniz's model, *petites perceptions* are a special type of perception described as "the insensible parts of our sensible perceptions" (NE Preface). They are the millions of unnoticed and unnoticeable perceptions that are going on all the time below the level of consciousness, such as parts of the visual field that are not being focused on, background sounds and smells that have become attenuated, and, in Leibniz cases, the innumerable amount of bare monadic perceptions that reflect the entire universe. Shane Duarte (2009) suggests a reading of Leibniz where these *petites perceptions* are elements of our subjective experience³. While themselves unconscious, and thus not subjective in the important sense, Leibniz states that *petite perceptions* may become available to consciousness only in aggregations. Though he is quite vague on the particulars, Leibniz describes how these insensible perceptions 'build up' until they become 'distinct' and are noticed by the faculties of consciousness. Leibniz likens the process to how the sound of the distant ocean, a singular roar, is composed of the aggregation of the sounds of individual waves crashing. Even if we cannot distinguish each individual wave-sound from the others, and

³ Duarte links *petites perceptions* directly to 'qualia', but I would prefer to avoid discussing such overloaded terms in a paper that is not meant to directly address them.

would not have noticed any one of the wave-sounds taken in isolation, the combination of the sounds is noticeable.

To hear the [roar of the sea] as one does, one must hear the parts that compose this whole, that is the noise of each wave, although each of these little noises makes itself known only in the confused collection of all of them together, and would not be noticed if the wave that made it were by itself.
(Leibniz, Preface to the *New Essays*)

This is quite poetic, but there are some immediate problems with it. First and foremost, the simple act of aggregation cannot be all that is required for a perceptions' becoming conscious. It is entirely possible for me to focus on a single voice in a crowded party, despite the fact that the overall stimulus is smaller than the ambient noise. Leibniz claims that memory is important to this process but does not give any speculation to a mechanism of awareness.

Attending to something involves memory. Many of our own present perceptions slip by unconsidered and even unnoticed, but if someone alerts us to them right after they have occurred, e.g. making us take note of some noise that we've just heard, then we remember it and are aware of having had some sense of it.

Leibniz, as a dualist, could be content to leave the workings of the mind mysterious, but this tactic would be frowned upon in a naturalized theory. This problem will be addressed in the next section.

Once *petite perceptions* become distinct, they are noticed by the consciousness and give rise to a sensation. However, while the consciousness is acting on the sensation, the distinct aggregate itself remains, along with its particular phenomenal character. The best way to think of these aggregates within the Leibniz framework is as a *clear but confused perception*. This type of perception is one that allows the perceiver to perceive the complex whole without perceiving the simpler parts that make it up. This distinction

is elucidated by Leibniz the green 'powder example' from the *Meditations on Cognition, Truth, and Ideas*.

when we perceive colors or smells, we certainly have no perception other than that of shapes and of motions, though so very numerous and so very small that our mind cannot distinctly consider each individual one in this, its present state, and thus does not notice that its perception is composed of perceptions of minute shapes and motions alone, just as when we perceive the color green in a mixture of yellow and blue powder, we sense only yellow and blue finely mixed, even though we do not notice this, but rather fashion some new thing for ourselves.

Leibniz would say we have a clear but confused perception of the color green because we can recognize it as an instance of green, but cannot recognize that it is a compound of yellow and blue. In the LAM to follow I intend to posit that perceptions can be thought of in just this manner, as clear but confused cognition of aggregation of unconscious perceptions. What this will translate to in naturalized language will depend on what sorts of entities one is willing to admit.

Naturalized Leibnizian Perception

Leibniz was strongly opposed to the idea that matter could think. It was Leibniz who gave us the famous Mill experiment, which would inspire John Searle's Chinese Room, among other famous critiques of materialism.

17. Moreover, it must be confessed that perception and that which depends upon it are inexplicable on mechanical grounds, that is to say, by means of figures and motions. And supposing there were a machine, so constructed as to think, feel, and have perception, it might be conceived as increased in size, while keeping the same proportions, so that one might go into it as into a mill. That being so, we should, on examining its interior, find only parts which work one upon another, and never anything by which to explain a perception. Thus it is in a simple substance, and not in a compound or in a machine, that perception must be sought for.

(Leibniz, *Monadology*)

Given his stance on the issue, it might seem strange to be appealing to Leibniz as the

savior of the very idea he so vehemently opposed. However, the framework the LAM borrows does not necessitate the additional importation of any of any other Leibnizian commitments. All that naturalized model will need to conserve is two important Leibnizian ideas. First, that conscious perception is *composed of* unconscious perceptions, and, second, that the resulting conscious perceptions are *clear but confused* in the Leibnizian sense. I will first lay out a general structure of the model itself; this will consist of a framework of physical perception that should be acceptable by most current theories. Then I will attempt to test this by applying it to contemporary theories.

The LAM account of perception begins when information about the physical world is captured or received by sensory organs leading, by some wholly physical mechanism, to a number of unconscious perceptions. These perceptions are very similar to Leibnizian *petites perceptions* in that they are unconscious, direct representations of the physical phenomena that activated the sense organs. Unlike *petites perceptions* in Leibniz's original account in the *New Essays*, the unconscious perceptions in the naturalized LAM are entirely physical. It is a widely accepted fact in modern neuroscience that there is a gap between sensory information entering the brain and its becoming available to consciousness, and that an amazing amount of structure is preserved in preconscious neural coding⁴. These structure-preserving unconscious perceptions are not conscious but, just as with *petites perceptions* in Leibniz, they are what makes up our conscious perceptions. In Leibniz, aggregations of unconscious perceptions are *available to* consciousness without necessarily being conscious entities

⁴ For example, the tonotopic arrangement of the human cochlea and its associated nerves. Any modern neuroscience textbook will provide a number of examples. A good resource is *Neuroscience Online*, published by University of Texas Medical School at Houston. Paul Churchland also has an in depth discussion of this in *Plato's Camera* (2012).

due to the fact that perceptual aggregations become available to the awareness—whatever that faculty might be—when they 'gain distinctness'. In the naturalized LAM, the excitation of single sensory neuron is not enough to produce a conscious perceptual state, it requires the excitation of multiple, similar neurons.

As I mentioned previously, awareness is a problem for the Leibnizian model of perception. The mechanisms and requirements of 'gaining distinctness' are not defined in Leibniz's work. While conceptually important, the problem of awareness is one that is better avoided than confronted directly. The LAM does not need a specific theory of awareness in order to function, it merely requires that such a thing exist in whatever parent theory is adopting it. So long as there is some method in the parent theory for unconscious entities to become available to the conscious mind—it should be an uncontroversial fact that such a mechanism must exist, given the existence of such commonplace examples as manual breathing where an unconscious act becomes conscious when attention is drawn to it—then the LAM will simply import that mechanism. For example, in *Consciousness Explained*, Daniel Dennett proposed a model of consciousness known as the multiple drafts theory (MDT). MDT wishes to do away with the 'Cartesian Theater'⁵, claiming that there is not privileged time or place where the boundary between unconscious and conscious is crossed because this boundary is illusory. Dennett's model hypothesizes that the brain does not take in sense data and spit out an interpretation, like a black box, but rather a given event will have attached to it a number of different sensory inputs and also a number of different interpretations that

⁵ The Cartesian Theater Model is the view that, supposedly, derives from the conception of mind found in Descartes' *Meditations*. It states that the mind is something like a screen on which sensory perceptions are projected, allowing the consciousness—which is discontinuous—to sit back and make judgments about them.

exist simultaneously. All of these 'drafts' are available for acting upon at the same time and the intuitions we have about making a single choice is an illusion born from the brains illusory sense of self.

As MDT is a materialist theory the story, applying the LAM would be a very simple matter. It will begin with photosensitive retinal cells absorbing light which the retinal and visual cortex neurons will encode and arrange the unconscious stimuli into aggregates. I can see two ways of proceeding from here. Either these aggregates would, due to neural coding, themselves contain interpretive elements, in which case an aggregate would be a kind of mental draft. Alternatively, aggregates would remain purely sensory and could be coupled with what Dennett calls 'interpretations' and these new aggregate-plus-interpretation clusters would constitute a mental draft. Personally, I feel Dennett would favor the first pathway as, in the second, the attachment of interpretation to sensory aggregate provides room to re-introduce the Cartesian Theater that MDT was meant to eliminate.

The motivation for this is that it would eliminate much of the tension seen in the LAM vis-a-vis the uncertain status of gaining distinctness. Within the context of Dennett's model, gaining distinctness would amount to an aggregations being used by the brain for action. This is a process already accepted by the parent model, thus eliminating that particular source of doubt.

It is important to note that, while this method is entirely physical, the LAM allows for, but need not require, a high degree of subjectivity. While the individual *petites perceptions* bear a direct resemblance to the physical world, the structure of any

individual aggregation need not, allowing for an element of privacy⁶ to be salvaged should one so desire. Leibniz used the term 'clear but confused' to refer to this phenomenon where the complex whole forms a clear representation but the individual parts remain indiscernible, similar to how the phosphorous dots on an old cathode-ray television screen form a picture, even though the dots themselves cannot be distinguished. The conscious perception of aggregations is similarly confused and may not bear any resemblance to thing it is representing. This allows for my unconscious perceptions and your unconscious perceptions to both resemble the same thing in the same way, while still allowing my conscious perception to look nothing at all like yours—a fact that can easily be accounted for by physical differences in the brain. While we hear the same physical sound waves, our ears, cochleas, auditory nerves, and brains are all structurally different and thus, need not have the same qualitative conscious experience.

Contemporary Applications

While a plausible naturalized account of perception has been built, the LAM is not a complete theory. There are holes in it that must be filled by applying it to full-fledged physicalist theories. Van Gulick (2011) identifies five general types of modern physicalist theories of consciousness: Higher-Order, Representational, Cognitive, Neural, and Quantum. With some tweaking, I am confident the LAM could be appended to a theory from any of these types. For the sake of being concise I will not go through each of these, but choose a few to demonstrate how the process is meant to work and hope the reader can extrapolate from there. Of the five, I find neural and cognitive theories the most promising, in that I find them more likely to produce a complete theory that is supported

⁶ Privacy: the notion that comparison between individuals is impossible. From Dennett, 1988

the findings of contemporary science, so I will choose a representative of each of these categories and apply the LAM. Daniel Dennett's multiple drafts theory, an example of a cognitive theory, was discussed earlier, so I will focus in this section on a neural theory; P.M. Churchland's eliminative materialism.

For Churchlandian materialism (P.S. 2011, P.M. 2012), traditionally conceived “mental states” are causally irrelevant, and should be eliminated from our talk in favor of a strictly physical language given to us by modern neuroscience. When an eliminative materialist is asked what it is to perceive x —or, to believe x or know x or have a concept x or a whole host of other mental states pertaining to x —she will respond that it is just for a brain to be in a specific state that represents⁷ x in the right way. In most cases, this involves embodying a vector through a theoretical brain-state-space. This kind of radical reduction creates very efficacious models for a many mental processes and abilities; for example, color processing, facial recognition, pattern completion, etc. However, according a number of the Churchlands' contemporaries⁸ it has yet to account for so called 'higher level' processes and has trouble integrating many subjective elements of perception and cognition (I have hopes that the application of the LAM will make some of these problems less problematic).

Contemporary neuroscience has made it a fairly uncontroversial fact that underneath the level of conscious representation there is a huge amount of sensory data

⁷ Churchland's system is representational in a way that closely resembles Leibniz's own system. While there is a direct resemblance to physical features, retinal cells respond in such a way that directly correlates to physical wavelengths and cochlear cells are arrayed to map actual physical differences in compression waves, the brain processes aggregations of these that are indirect. An activation pattern of neurons is a representation of physical phenomena in very much the same way an aggregation of Leibnizian petites perceptions is a representation.

⁸ Notably Chalmers and Searle.

entering the brain of which we are not conscious; a fact that parallels nicely with Leibniz's account of *petites perceptions*. Our brains are constantly receiving thermoceptive, proprioceptive, nociceptive, equilibrioceptive,⁹ etc. data, but, similar to *petites perceptions* we only become aware of said inputs if they become unbalanced, abnormal, or distinct in some way. What's more, all sensory data is believed to work this way. Vision, for instance. The human retina is composed of approximately 100 million retinal cells, each of which transmits numerous action potentials every second. The firing of any one of these retinal cells alone would not be recognized by the brain—in fact, there are specific neuronal processing mechanisms designed to ignore such events. But, just like the sound of the distant ocean, these millions of unconscious inputs aggregate together to create a distinct visual perception. This aggregation of neuronal activation potentials represents Leibnizian module as a vector through a constructed neural state-space, which is what the conscious takes notice of.

Further Investigations

The main project of this paper has been to construct from G.W. Leibniz's theory of *petites perceptions* a naturalized model of perception suitable for contemporary philosophy of mind. The goal of this model is to give a plausible, scientifically consistent mechanism by which unconscious neural activity can become conscious. I feel this has been accomplished; however, the full motivations for this project reach slightly further. When fully mature, the LAM, or a similar model, should be able to help the physicalist theories it is attached to deal with hard conceptual problems posed by anti-materialist

⁹ Thermoception=temperature sensations; proprioception=position of one's own body parts; nociception=pain; equilibrioception=sense of balance.

philosophers. In this section I will give brief consideration to one of these, the “Hard Problem of Consciousness”.

David Chalmers proposed a problem which he dubbed the Hard Problem of Consciousness. Unlike the soft problem of consciousness which can be solved by looking at mechanisms alone, the hard problem will "persist even when the performance of all the relevant functions is explained" (Chalmers 1995). This is obviously concerning to physicalists like the Churchlands and Dennett and responding to the hard problem has become a major preoccupation for philosophers in this school of thought.

The Hard Problem, as it is presented by Chalmers, focuses on a perceived separation between perception and the subjective elements of perception.

It is undeniable that some organisms are subjects of experience. But the question of how it is that these systems are subjects of experience is perplexing. Why is it that when our cognitive systems engage in visual and auditory information-processing, we have visual or auditory experience: the quality of deep blue, the sensation of middle C? How can we explain why there is something it is like to entertain a mental image, or to experience an emotion? It is widely agreed that experience arises from a physical basis, but we have no good explanation of why and how it so arises. Why should physical processing give rise to a rich inner life at all? It seems objectively unreasonable that it should, and yet it does. (Chalmers 1995)

On one reading, Chalmers is simply begging the question. He posits properties that have few definite properties other than being necessarily not physical and then finds fault with anyone who can't explain these properties physically. Of course, there are more interesting readings based on intuitions about inner mental workings. Many of these can be boiled down to appeals to inconceivability. The most charitable reading of this problem that I can come up with is to see Chalmers as highlighting features that appear obvious from introspection.

To better bring out the issues at hand, we can consider the well known thought experiment known as Mary's Room. In this experiment, proposed by Frank Jackson, we imagine a neuroscientist named Mary who has lived in a room devoid of color. Mary knows everything there is to know about neuroscience and color vision as well as everything there is to know about her own brain's structure and processing. Due to this, Mary knows the physical characteristics of her other color perception aggregates and could accurately approximate what sorts of actual physical aggregates would arise if her eyes were bombarded with 475nm light—this is not saying she could predict the phenomenal character of the aggregate, only actual neural impulses that make up the aggregate. If Mary can predict the aggregates, she then knows everything there is to know about how they will physically interact with her brain and produce results and can thus answer any question about perception of blue, 475nm wavelength light that a normal neuroscientist who has seen blue can answer; all she gains upon exiting her room is an experience, which brings with it no extra knowledge to Mary. The experience of other color aggregates is not necessary for Mary to predict her own aggregate upon being exposed to a particular wavelength of light because the unconscious *petites perceptions* are direct captures of the physical world. Mary, who has been granted perfect knowledge of her own physical brain, can take a physical brain event and know exactly what response her brain will have—namely, exactly what kind of aggregate will be produced¹⁰.

Chalmers could accept all of this, while continuing to claim that it is not equivalent to subjective experience. This is true, but vacuously so. LAM shows that there

¹⁰ I would say this is practically impossible for any real human, but Jackson's experiment relies on Mary being super-human in her knowledge about neuroscience.

need not be any subjective experience beyond perceptual aggregates, that Chalmers is looking for something does not exist.

Now, a philosophical dualist might then complain: “You've described how hurting affects your mind — but you still can't express how hurting feels.” This, I maintain, is a huge mistake—that attempt to reify 'feeling' as an independent entity, with an essence that's indescribable. As I see it, feelings are not strange alien things. It is precisely those cognitive changes themselves that constitute what 'hurting' is — and this also includes all those clumsy attempts to represent and summarize those changes. The big mistake comes from looking for some single, simple, 'essence' of hurting, rather than recognizing that this is the word we use for complex rearrangement of our disposition of resources.
(Marvin Minsky in *Edge*)

An experience of yellow, for instance, seems like a single, distinct thing but it is, in truth, not a single thing at all, it is a representation created by a confused complex of unconscious perceptions. Perceptions, as constructed by the LAM, are not simple but aggregates of unconscious *petites perceptions* which have the subjective elements built in. The subjective elements are not a separate thing that emerges or attends to neural activity, they are neural activity. What LAM allows us to see is why our introspective intuitions would look the way they do. It points out that perceptual aggregates are opaque and do not allow us to see their clear constituents upon introspection.

Conclusion

In the *New Essays on Human Understanding*, G.W. Leibniz proposed a model of perception that posited a host of unconscious *petites perceptions*. I feel that the importance of this model has been overlooked in the contemporary field. My goal has been to naturalize this model, creating a model that retains the core of Leibniz's perception and incorporates the best scientific theory of the present day. By envisioning singular neural impulses as a naturalized analog of *petites perceptions* I develop a line of

investigation that, I believe, shows promise as an eventual model of perception. My goal here has not been to provide answer questions about Chalmers-style intuitions or the difficult issues involved in subjective experience, but to construct a mechanism by which non-conscious entities become conscious in a wholly material and non-spooky way In doing so, I lay a groundwork for future model that there is reason to think, when fully matured, may provide a way for reductive physicalists to avoid Hard Problem of Consciousness in a way that the non-reductive physicalist might be willing to accept.

Bibliography and Works Cited

Brockman, John (2012) "CONSCIOUSNESS IS A BIG SUITCASE - A Talk with Marvin Minsky." Edge, 27 Feb. 1998. Web. 05 May 2012.
<http://www.edge.org/3rd_culture/minsky/index.html>.

Carruthers, Peter (2011) "Higher-Order Theories of Consciousness", *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*, Edward N. Zalta (ed.), URL =
<<http://plato.stanford.edu/archives/fall2011/entries/consciousness-higher/>>.

Chalmers, David (1995) "Facing Up to the Problem of Consciousness", *Journal of Consciousness Studies* 2 (3), pp. 200–219.

Churchland, Paul M. (2012). *Plato's Camera: how the physical brain captures a landscape of abstract universals*. Cambridge, Mass.: MIT Press, 2012. Print.

Churchland, Patricia Smith. (2011). *Braintrust: What Neuroscience Tells Us about Morality*. Princeton, NJ: Princeton UP, 2011. Print.

Dennett, Daniel (1991) *Consciousness Explained*. Boston: Little, Brown and, 1991. Print.
Dennett, Daniel (1988) Quining qualia. In Anthony J. Marcel & E. Bisiach (eds.), Oxford University Press.

Duarte, Shane (2009): Ideas and Confusion in Leibniz , *British Journal for the History of Philosophy*, 17:4, 705-733

Jackson, Frank (1982). Epiphenomenal qualia. *Philosophical Quarterly* 32 (April):127-136.

Jackson, Frank (1986). What Mary didn't know. *Journal of Philosophy* 83 (May):291-5.

Leibniz, Gottfried Wilhelm., Jonathan Bennett, and Peter Remnant. *New Essays on Human Understanding*. Cambridge [u.a.: Cambridge Univ., 1997. Print.

Lycan, William G (2006). Consciousness and Qualia Can Be Reduced. In Robert J. Stainton (ed.), *Contemporary Debates in Cognitive Science (Contemporary Debates in Philosophy)*. Blackwell.

Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83 (October):435-50.

Neuroscience Online: An Electronic Textbook for the Neurosciences. Ed. John H. Byrne. University of Texas Medical School at Houston, 1997. Web. 28 Sept. 2012.
<<http://neuroscience.uth.tmc.edu/>>.

Searle, John R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences* 3:417-57.

Searle, John R. (1993). The problem of consciousness. *Social Research* 60 (1):3-16.

Simmons, Alison.(2001). Changing the Cartesian Mind: Leibniz on Sensation, Representation and Consciousness. *The Philosophical Review*, Vol. 110, No. 1 (Jan., 2001), pp. 31-75

Tye, Michael (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.

Van Gulick, Robert, "Consciousness", *The Stanford Encyclopedia of Philosophy* (Summer 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2011/entries/consciousness/>.