



INTERVIEW WITH FREDERICK ADAMS

This interview presents the tracking of formation of Frederick Adams and the interests that guided him. It was made via email on July-November of 2010 by João Antonio de Moraes.

Tell us a little about your background and reasons that led you to attend the Philosophy course.

If I understand you correctly, you want to know what led me to take my first philosophy class and how I got into philosophy as a field of study. If that is correct, then my answer will be on track. If not, correct me and I'll try again.

I grew up in Belleville Illinois (USA). My father worked for Monsanto (now a company that is well known for bio-engineered crops, but when I was young it was mainly a chemical company. My father wanted me to be a chemical engineer. So I started classes taking the usual physics and chemistry courses and math courses one would take to prepare for this career path. I also took logic. To me this was quite new and interesting. I had never had any course quite like it before. I was good at it and enjoyed it and it was taught in the philosophy department at Southern Illinois University at Edwardsville, where I was a student. So I decided to try another philosophy class. The next class I took was an introduction to ethics. This course was taught by a man who was extremely theatrical in his teaching style. Each class day was a performance for him, complete with jumping on top of the desk and raising his voice and waving his arms and so on. We went through absolutism, relativism, subjectivism, emotivism, utilitarianism, prescriptivism, and a few more theories. Each week I thought we finally were on the right theory. Then the next week he would destroy the theory and we'd start a new one. Anyway, I found this all very exciting and I wanted to know if there was an ethical theory that would not be destroyed. So I took an independent study course with him on ethics to follow the other class.

Now this school was unusual in this respect. There was a fledgling M.A. program, no ph.d. in philosophy, despite this there were 20 ph.d. faculty persons in the department. So, when I decided to take more philosophy courses, I had a wealth of people, knowledge, and fields to choose from.

Now, as I've told you, I was excited by logic and ethics because it was so new to me. But when I decided to take philosophy seriously as a subject matter, I fell back into my roots of looking to the philosophy courses that were most like science. So I took a course in the philosophy of science. Then I took a course in the theory of knowledge. And finally I got very interested in the philosophy of mind and started taking independent study courses in that subject. The main question there that interested me was how conscious and phenomenally rich mental conscious states could simply be neuro-chemical states of the brain. I was pretty convinced by the identity theory of Herber Feigl that this theory was true, but I still didn't know how to think about the mental and the physical and the relations between the two...how the one could be or produce the other.

During the time that I was thinking about issues in the philosophy of mind, I started to learn about the philosophy of language and when I saw the connection



between the two subjects, I started reading on my own many things in philosophy of language. Then I took a course in philosophy of language to help me structure my interests and readings on those topics.

In my undergraduate career I was also very interested in anthropology. Again, I was interested because it was so new and different (and my school also had a very good anthropology department). As it turned out, I got my B.A. degree in Anthropology from the university, even though I had many more credits in philosophy than in anthropology.

When it came time to decide on graduate school, my teachers in anthropology very much wanted me to go to graduate school in anthropology. However, I knew that my real interest were in philosophy, but that point in my life and I was never much interested in doing anything else but getting a graduate degree in philosophy.

I eventually went to the University of Wisconsin in Madison Wisconsin and worked with Fred Dretske, Berent Enc, Dennis Stampe, and Elliott Sober and never looked back.

Could you tell us more about your M.A. and Ph.D thesis?

As I told you before, I became very interested in the relation of mind to body as I was finishing up my undergraduate degree. So as I started graduate study, I started reading everything I could find on the topics. I started with Roy Wood Sellars (father of Wilfrid) who was a double-aspect theorist. So I read about this. That topic leads me to the related view called neutral monism. Bertrand Russell flirted with this view for a time, and many others. Both of these were predecessors in one way or another of the identity theory of mind and brain. And so in my M.A. thesis I started with Descartes (a dualist) and Spinoza (a neutral monist) and traced some of the history of the identity theory. I studied Place, Smart, and D.M. Armstrong on the identity theory of mind. My work was largely uncovering the historical record and seeing relationships between various views. I'm not sure if there was much creativity or novelty in what I came to write, but it pleased my M.A. thesis committee and earned me my first philosophy degree...an M.A. in philosophy. I was pretty well hooked on type-type identity theory of mind by the time I was heading to Wisconsin for graduate school. So when I got to Wisconsin and heard about functionalism (which rejected the view I had just adopted), I was very surprised and eager to decide if I was going to give up Identity theory or fight against functionalism. I fought against functionalism for at least two years, as I recall, publishing my first single authored paper defending Identity theory, Adams, F. 1979a: Properties, Functionalism and the Identity Theory, *Eidos*, 1, 153-180.

My Ph.D. thesis was on Goal-Directed Systems. In my second publication, [Adams, F. 1979b: A Goal-State Theory of Function Attribution, *Canadian Journal of Philosophy*, 9, 493-518] I explored the idea of tying teleological functions to subserving the goals of goal-directed systems. So the idea was that if the function of the heart is to circulate blood, then there must be some function (survival) that the heart's pumping contributes to and does so via a control loop. For natural functions (hearts, kidneys) the control loop can be constrained and sometimes hard to find. For artifacts, (the function of the lock on the door is keep out unwanted intruders) the control loop is much easier to find. Anyway, I sort of liked the theory of functions, but soon realized that I needed a theory of goal-directed system to go with it. This led me to develop a cybernetic or control model of goal-directed system using information theory as the



basis for the account. I needed three basic ingredients: (1) a goal-state representation to set the goal of the system's behavior, (2) the behavior of the system had to be causally dependent upon the goals of the system, (3) the causal dependency had to be via informational control loops. So these were the basic ideas I had to develop in the thesis.

I looked at behavioral indicators of goal-directed behavior. I found a set of such features and then tried to show how they could be explained by a control-theoretic model of goal-directed system. My account was objective, realistic, and scientifically based. It said goal-directedness was an objective property in the world. I had to compare and contrast my account with various projectivist accounts which tried to make goal-directedness a mind-dependent ascription. I had just come from a semester long workshop on philosophy of mind in Bristol England and directed by Andrew Woodfield (attended by myself, Steve Stich and Dan Dennett). Woodfield was a projectionist about goal-direction and teleology (in his book by the same name). So I felt compelled to differentiate his approach from mine and explain why I preferred mine.

I also spent a chapter on the logical connection argument. This argument had crippled action theory prior to Davidson's causal account of action. So I felt it was important to explain what had gone wrong with the argument and why it would not stand in the way of my project. To my surprise, the argument was to crop up again years later in Fodor's defense of narrow content. Luckily, I had this work in my thesis to rely upon in writing my reply to Fodor (in Adams, F. 1993c: Fodor's Modal Argument, *Philosophical Psychology*, 6, 41-56). So it turned out to be a very good thing indeed that I had thought hard about these matters ten years earlier in my thesis. I still think Fodor got things very wrong in his arguments and he had dropped just about all discussion of these ideas by about 1995.

One very important matter that I needed to discuss was why my account needed internal representations of goals. I had discovered that most cybernetic accounts focused on war. If one wanted to shoot down a plane with a guided missile or automatic gun, one would have the system home in on the existing object as target, shoot, compare the difference of the angle of the miss and re-configure and shoot again to diminish the error. That is, all discussion of how goal-directed systems worked presumed an existing goal object. But what about goal-directed behavior where the goal-state to be accomplished did not exist or the goal-object to be created did not yet exist? There were no accounts of how this went. So I knew that my account needed to supply that missing ingredient.

Now how was I to generate the content of the goal-state representation? My idea was to use information theory. I was thinking about concepts and an information based account of concept formation. The concept C is about dogs, if it is information about dogs that is the basis upon which the concept C is formed. So the meaning of C is dog if it is information about dogs that C is created to track, so to speak.

Now this idea worked fine for knowledge (you can't know what isn't true) or concepts (you can't form empirical concepts of dogs, if there are no dogs). But the information based account has difficulty when it comes to representations of things or even states of affairs (types) that do not exist. It has problems because there is no thing or type of state for the goal-state representation to be informed about and from which to derive its content.

This problem was significant. It was so significant that at first I did not know how to get around it. Worse yet, none of my teachers knew how to get around it either.



This was a particularly difficult time in my life. In fact, I even went to work at a law firm in Madison Wisconsin because I was not sure I would ever find a way around this matter. I worked at this law firm for about five months. Then one day when I came home I was invited to be interviewed for a one year teaching job at Lawrence University in Appleton Wisconsin (about a two hour drive from where I was living with my wife). I took the job and during my time teaching there I came up with a solution that satisfied my committee and got me through the thesis.

The solution was the concept of virtual representation. Here is a simple idea of it. When you set the temperature on your thermostat for the room temperature you want the room to have, this is not informed by the actual temperature. The actual temperature is not the set point, or you would not have changed the setting. So what determines the content (representation, meaning) of the new set point? The meaning is the room temperature that it would take to shut the system off. So it represents what the value would be when the circuit opens. I called this a "virtual" representation. I now think of it as a dedicated property detector. For example, if you have a sensor on your car that turns on a warning light when your oil pressure gets low, the sensor is a dedicated property detector. It detects the property of low oil pressure and then warns you of it. Its content is always what the pressure would be if the light comes on--it would be too low (or below a certain value). This is still an informational account because the content of a virtual representation is the informational content it would convey if it represented an actual state of affairs. So its content is tied to the informational content of such an actual state of affairs. But, and this was the important part, it could carry that content even prior to the existence of the actual state of affairs. So you could build a system that had the goal of bringing about such a state of affairs. And that was just what I needed if I was going to explain how goal-directed systems could have non-actual objects or states of affairs as their targets.

The rest of my dissertation was about the details...the details of intentionality and intensionality (with an "s") and how they are related; the details about causal dependency in a control system; and the details of feedback and feedforward types of control systems. Naturally, all the important matters are in the details, but what I've said so far tells you the basic plan of my dissertation, how I became interested in information theory and control theory, and how I was presented with these topics. My thesis was finished in Fall of 1981, but my degree was not conferred until spring of 1982. These were in the days before I had a computer, so the spring was taken up with time for my typist to prepare my manuscript.

I did a lot of reading on my own. I took a class with Fred Dretske who was writing his book *Knowledge and the Flow of Information*, which came out in 1981. So he and I were working on similar ideas at about the same time. He thanks me in his book and naturally I thank him in my dissertation. He and I have been good friends ever since.

So, it was in your Ph.D that you faced up with information theory. In your article of 2003, named "The Informational Turn in Philosophy", you highlighted some proposals that relate information theory with philosophical problems. But, as you emphasizes in it, the concept of information, itself, only started to be analyzed in 80th. We know that doesn't yet have a consensus between the researchers about the concept of information. Well, to make our talk more spicing, I would like to



make you the following question: what do you conceive by "information"? That is, what is Adam's concept of information?

I think of information as a mind-independent entity. Before there were minds to exploit it, there was information.

Think of an event. For our purposes, let an event be an object's acquiring a property at a time. Now of all of the properties in the ecological space that the object occupies, there will be some constraints on the kinds of properties it can acquire. But there will also be some freedom. When the possibilities for that object reduce to one (we are just looking at one property for now), that reduction constitutes some information coming into being. It can be measured in many different ways and quantified. But for now, put aside the specifics of how it is measured. The fact that something, some event, happened generates information. It is news. And as news, it can be conveyed or transmitted to another location.

Of course, transmitting to another location introduces the matters of a channel of communication and the various mathematical properties of communication channels...mathematical details that were discussed by Shannon and many others. Let's not get bogged down in those details, just now, but focus on the matter of what it takes for information about the occurrence of the event to be communicated to another...or to a mind.

Here is where the notion of a communication channel and the conditions that this channel must have for the information about the event to be transmitted become important. So suppose a child has a fever. Then the child has acquired a complex property and that even is news. Now suppose the child's mother wants to know whether the child has a fever (or is faking so he won't have to go to school today). She wants to take his temperature. So she uses a thermometer. For the thermometer to "tell" her something, it has to be a conduit for information about the child's temperature. It has to not create or introduce new information of its own into the system (about temperature). So it has to be an information channel...something that can detect and convey information to the mother about the child's temperature without itself telling the mother too much about itself. She wants to use it to learn about the child, not about the thermometer. Of course, if she already knows the child's temperature, then she could use this fact to test the thermometer to see if it is working properly. But she can't do both simultaneously, use the thermometer to tell her about the child's temperature and test the thermometer at the same time.

Now if the thermometer is working properly (no cracks or flaws), then it is a perfect conduit for information about the child's temperature to flow to the mom. Of course, it has to be calibrated and she has to know how to read it, but if all that goes well, then the rise in the thermometer to say 101F tells the mom that the child does have a fever (and should be allowed to stay home from school--is not faking).

This only works because there are law-like correlations between temperature and the rise of liquid in a sealed chamber (in old fashioned thermometer). This lawful correlation of events (child's temperature and thermometer reading) allows information to flow. And if the mom knows how to read the thermometer, she learns this from it. Of course, even if she does not know how



to read it, the thermometer still contains this information. It just won't be of use to the mom, if she doesn't know how to read it.

Following the last question, what about meaningful information? For you the meaning, as the information, is: i.i) a mind-independent property? i.ii) a relational property? i.iii) a thing that is only in the subject's mind?

I think information is a mind-independent property. If a tree falls in the forest and there is no one there to hear it, still information has been generated by its fall. When the "big bang" of cosmology happened, it set events in motion that scientist are still able to detect...they can detect reverberating effects of the big bang organ of the universe, and can detect these even today. There was no one present when the universe began, but there was information.

So my answer to this question is (i).

However, (i) and (ii) are not independent. Suppose that Holmes and Watson are detectives. They are trying to find out who killed Jones. They know the murder walked with a limp and had a moustache. Suppose, they both see Brown seated. They both know (acquire the information visually) that he has a moustache. But they can't see whether he limps because he is seated. But suppose earlier in the day Holmes, not Watson, saw that Brown walks with a limp. Then Holmes, not Watson, may come to know (by now seeing that Brown has a moustache) that he is the murderer. So information (how much one learns from an event) may be relative to what else one already knows.

Information itself has no semantic aspect. Semantic = meaning. Information is not meaning. It is a source of meaning. It is something that can generate meaning. But it is not meaning. Why do I say that? For this reason. If a signal carries the information that p, then p (is true). If I water freezes, it expands. Its expansion carries the information that it has frozen. So if something carries the information that the body of water has expanded, then it carries the information that the body of water has frozen. It cannot be otherwise. Information is tied to the facts. This is why people (Dretske, Floridi, me) say that information cannot be false...it cannot send a signal that is capable of telling one something other than the truth about the world.

Meaning, something with semantic content, is not like that. I can tell you "the water has frozen" when it has not. I can lie. I can be mistaken in my beliefs and in my statements. So the important difference between information and something with semantic content (meaning) is that the latter, but not the former, can be false (can misrepresent the facts).

Now how are information and meaning related? That is a longer story than I can tell her, but I tell it in my paper "The Informational Turn in Philosophy" so readers can see the details there.

Since information is mind-independent, and meaning has information as its source, I think meaning is mind-independent in this sense. I believe in something like a language of thought (LOT)...a set of symbols with which we think about the world. One of my concepts (the concept of what a dog is, say) is about dogs. It is formed by obtaining information about dogs. Once formed, I can think about dogs and my concept means dog. Is this "up to me"? In a sense--no. I can't freely make my concept of a dog become a concept of something else just by choice or by training. Yes, it is a concept that arises "in my mind." So it depends on my having a mind for me to be able to form



the concept of a dog. But that is a kind of dependence that is not subjective and not able to be bent by my will. It is not like a convention. We could create an artificial language and we could make up a new word for dogs. That would be subjective and by convention and subject to our wills and completely mind-dependent. But our concept of dogs in our minds (not in an artificial language we make up) is not like that at all.

Do you think that the information-based concepts provide us a better understanding of philosophical problems?

Yes. I hope that was clear from my presentation in Marília on October 1st. I spoke about three traditional problem areas in philosophy where I find information-based concepts to be particularly helpful in understanding and potentially solving long-standing philosophical problems. The three areas I spoke of were: knowledge of empirical propositions (such as that the Brazilian elections were recently held), intentional action (such as that Adams intentionally came to visit Marília), and gaining a naturalistic understanding of the mind (both the phenomenological qualities of experience and the cognitive concepts we use in thinking). I believe the concept of information is enormously useful in understanding and trying to solve problems in these areas. I offered some solutions in my talk there and used informational concepts to try to explain.

Let me add a few more remarks about this. Knowledge is true belief, yes, but it is more than true belief. It is even more than justified true belief (as Edmund Gettier taught us, 1963). One who has knowledge that a proposition is true has eliminated the possibility of accident. The person who knows has a true belief and it is no accident that his or her belief is true. Information cannot be false. So if one's belief is based upon information, then one will have knowledge because having information eliminates the possibility that one's belief is true by accident. Hence, the concept of information and its connection with being true, is very helpful in solving the traditional philosophical problem of what it takes for a person to know that a proposition is true.

A similar place where this same sort of issue arises is intentional action. A person who intentionally does an action A not only intends to do the action A and does the action A, but his or her doing A is not an accident. Surprisingly, Donald Davidson and others showed us that the same kind of accident that Gettier discovered with respect to knowledge can happen when we consider actions. A person can intend to do A (to break the precious vase) and can do A (can drop it) and can drop it because one intended to drop it, but can also drop it because the intention to drop it so unnerves one that one begins to tremble uncontrollably and drops it accidentally (even though one intended to drop it). One loses control. The notion of control is key here and the notion of control makes use of concepts from information and cybernetics. So an action that is done intentionally is one where S intends to do A, one does A, and one's doing A is under one's control. For me this means the bodily activity of doing A is informationally sensitive to the contents of one's intention and depends upon that content via an information feedback loop (or other informational mechanisms). Again, the key is that information theory (concepts) provide the crucial understanding of what it takes to eliminate the kinds of accident that can arise between an intending to A and a doing of A—the kinds of accidents that can render the doing unintentional. When one does an action intentionally, these sources of accident are eliminated and are eliminated by informational connections. For more of the details about how this goes, one may



consult my “Tertiary Waywardness Tamed” (1989) and “The Role of Intention in Intentional Action” (1988), which I wrote with Al Mele.

Readers can find more information in papers listed on my cv. or on my website about these topics. I did not speak of the usefulness of the concept of information for various problems in the philosophy of language, but I think some of these too are understood better against the backdrop of talking about information and how it can be used to understand meaning and reference. See especially my papers on the semantics of empty names. There are papers on my c.v. [<http://udel.edu/~fa/cv.html>] and website about this too, if students are interested to look further into how I think about these matters.

What do you think about the research in Undergraduate Level?

I am strongly in favor of this and would encourage it for any university. Undergraduate research is not for every student, but for those who want to do it, it is a wonderful opportunity. We have a long tradition of undergraduate research here at the University of Delaware. In my time here (I came here in 1997), I have worked with five different undergraduate students on research projects. Students come to see professors in the early Spring semester (about February or March). When students are interested in working with me, I tell them what projects I am currently working on. I usually have ideas for papers in philosophy of language, philosophy of mind, philosophy of action, or epistemology. So I tell them what projects I am thinking of working on during the summer and ask them if any of those possibilities sound interesting to them. They usually find one topic more interesting than the others. So we agree that if they get the undergraduate research grant, we will work together during the summer on that topic. Then the student writes an application to be funded and submits it to the research office. I write a corresponding letter of recommendation explaining what I'll be doing and how the student's research will intersect with my research. When the student gets the award, then we meet and get a list of specific things we will read and in what order. When summer comes, I tell them that we will meet in my office every other day for 2 hours and discuss the readings for that meeting. This is a very intense schedule. Many other faculty meet with their students only once every two weeks or so.

As the meetings progress, I usually start to formulate an idea for a paper to write. I share my idea with the student and, if the student is interested, our sessions turn from just reading and learning what others have to say to formulating our own ideas on the topic and reading to see what else we need to consider to present our own ideas.

The way things work here at the University of Delaware, the student takes an Independent Study course with the faculty mentor during the Fall Semester following the summer undergraduate research experience. So, I've found that this afford the opportunity for my student researcher and I to meet weekly and actually write the paper together at my computer in my office.

I have now done this successfully many times. If you look at my c.v., the papers I co-authored with Dietrich (2), Steadman (3), Beighley (1), and Figurelli (1) were all products of undergraduate research (Dietrich and Steadman did research with me two summers in a row). So it was good for the students to have publications before they graduated as undergraduates and it was good for me, because I always learn a great deal in the process of working with them too. Two of these students have gone on to get law degrees (Dietrich and Figurelli...who is now a lawyer working at UD), and the other two



are in graduate school in philosophy (Steadman is now finishing her ph.d. at Arizona and Beighley is completing his M.A. at Georgia State and applying to doctoral programs this Spring). In addition, Steadman and I presented our papers at national and international conferences. Beighley presented joint work at regional conferences and won a second place monetary prize for best presentation at a UD Summer Undergraduate Research Symposium at UD, where all students present the results of their summer research (there are about 150 presenters...so he was 2nd out of 150...not bad!).

So, yes, I strongly recommend undergraduate research programs and encourage undergraduates who have the interest to pursue these opportunities. They can be life-changing.

What are your interests in Human Sciences?

I have started to learn more about Deep Brain Stimulation. It is a technique used to help patients with Parkinsons, Essential Tremor, Dystonia and many more conditions. I have begun a paper on this topic with a graduate student in my department, Evan Bradley, and a colleague at Brown University, Michael Frank. We have a good start and Evan and I have made one presentation on the topic and will make others.

The procedure involves placing two electrodes deep into the brain (targeting different areas for different conditions) and then sending electrical pulses into the brain. A tiny computer is placed in the patient's collarbone and connected to the wires under the skin, so the patient can be completely mobile after the surgery.

My concerns include that there are many potentially significant bad side effects that are largely downplayed by practitioners, the procedure is being rapidly expanded to treat many more conditions, where it seems contra-indicated. And in many cases, the patient leaves the very high tech medical center where the device is inserted and then goes home to a very rural area where few doctors or hospitals specialize in use of such devices. So, the patient may not get the best care at first, if something goes wrong and needs fast attention.

Other than this, I am currently teaching and researching the topic of embodied cognition. The interest in this new approach to cognition is sweeping the planet. As I read the empirical research, I see why scientists and philosophers are interested in embodied cognition. But I think some of the claims that are being made for this new approach are going beyond what the empirical data will support, and I've been expressing my concerns in presentations and in publications about this fairly recently. I won't try to give you all the details, but basically embodied cognitivists claim that the human body conditions the kinds of concepts we have and can have (so that, for example, a computer without a body, could not have our concepts, even if it could think). The view also argues that all of our concepts are symbols that have their origin in the sensory system and in the motor system of the brain, and never lose some of the properties acquired due to their origin. These are called "modal" symbols--referring to the sensory modalities. The more standard view of cognition prior to embodied cognition is that there is a region of the brain that takes in information from the sensory systems and then forms concepts. Once we use these concepts to think about the world, then the brain sends signals to the motor system getting the muscles of the body to do our bidding in making the world the way we want it to be (or at least trying). On the



new embodied view, there is no central system. So thinking literally goes on in the perceptual and motor systems. This, if true, would indeed be revolutionary. That's a strong reason why I'm interested in trying to determine whether it is true or not.

Those are the two things that are occupying my interest the most at present that are issues in the human sciences.

Once we are still living the impacts of the informational turn in philosophy, what would be the ethical implication of this turn? Well, we are seeing and experiencing changes in our environment with the insertion of informational technologies (cameras, ubiquitous computing, iPhone, notebook, etc.) that are causing the emergence of new patterns of conduct and reshaping older ones. Is possible conceive and built an Information Ethics, as Floridi, Capurro and Gonzalez are proposing?

Let me start with the last part of the question...is an informational ethics possible? I don't know. I think the question I would need answered is "what determines moral standing?" In the literature on abortion and on euthanasia, the main questions are when does someone have moral standing? In the abortion literature, it is when does the foetus become a person? In euthanasia, when does a human being lose the potential for consciousness (permanently)? The questions sometimes turns on when "brain death" occurs. Now why do people ask these sorts of questions? I think it is because they want to know what kinds of things have moral standing? What determines whether something is in the moral realm, such that I need to consider it when deciding what I should do (morally)? I've always believed that moral standing came with significant mental states...the capacity to consciously feel or reason or feel pleasure or pain. So I've never thought of plants or trees as having moral standing because I've never thought of them as having mental states. I know there are people "deep ecologists" who want to extend moral rights to trees and plants and other things that don't have mental states of any kind. I've never found the reasons given for this to be very persuasive.

If information itself or the capacity to carry information generates a moral standing, I need to know why. Now having mental states involves the capacity to carry information, but it also involves carrying information in a special way (producing qualitative states or producing cognitive states). Not just everything that carries information (this computer, the internet) has these mental states. And so in my view, this computer and the internet itself, cannot be morally harmed (though they can be physically harmed).

At any rate, an information based ethics would have to explain the basic principle behind having moral status? What is the principle? What kinds of things have it and don't and why? That's the place to start.

Now as for the other ethical issues that come up due to the informational revolution, they mainly revolve around privacy. Now that people can tell which web sites you visit and where you are on the planet (by tracking your cell phone), we literally have the situation from the novel "1984." It is not that the STATE is watching us, but the CITIZENS. It is hard to be alone or anonymous anymore. Just being observed or the potential for it 24hrs a day 7 days a week is bad enough, but the fact that corporations all over the world know what I like and try to sell me things by pop-up adds on my computer or cell phone (as determined by my past habits), is annoying.



One does not have to constantly think, "If I do this (make a call, visit a website), who will know and do I want them knowing?" But that is what things have come to.

Of course, there are many good things that come with this too. Instant information about what you want (food, entertainment) or need (where is the closest doctor or hospital to my current location?). So it is not all bad, after all. But there are many things one worries might happen that are bad--sharing my medical records with potential employers or insurers, sharing my financial records with them, or businesses selling my personal information to commercial entities who want to advertise to me (even if I don't want the advertisements). The question for information ethics is when the line is crossed from my being merely annoyed by all of this to my being morally harmed by all of this. When that line is crossed, then governments will have to find a way to protect the innocent and regulation will be introduced. There is work for philosophers to do in helping the public understand and see these issues.

Since this is an interview for an Undergraduate Journal of Philosophy, we would like to know your opinion about what should be required in a course of formation of philosophers today?

I think every student should take a basic course in symbolic logic and learn to translate sentences into logical notation, learn to do proofs, and learn about completeness of systems.

Every student should take a course in philosophy of science and learn about what a scientific theory is and the nature of an explanation.

Every student should learn something about philosophy of language and theories of reference and meaning. I would prefer that they learn some naturalized semantics that connects up these topics with the concept of information.

And, of course, every student should learn something in value theory: ethics or aesthetics and theories of value and learn how to evaluate such theories.

This plus some history of philosophy would provide an excellent foundation for any student of philosophy.

We are getting to the end of this interview. Would you like to say anything in especial?

I would like to say thank you João. I think what you are doing with this journal is a terrific idea and I wish you great success at bringing ideas such as those we've discussed to the students of Brasil.