Philoctetes Center for the Multidisplinary Study of Imagination
ROUNDTABLE ON ARTIFICIAL INTELLIGENCE:
NEUROANALYSIS AND IMAGINATION
November 17, 2004

EN. Let's begin with Edith.

EA: My relationship to the understanding of artificial intelligence is very diagonal, so I will just try to situate my input properly and hope that it doesn't become too tangential. My interest has always been to try to understand the genesis of children's intuition, or ways of making sense of their world and selves. In my early years, I was working at the Piaget Institute, in Geneva, an interdisciplinary research factory where the focus was on t the ways in which children come to understand different areas in the scientific domain, like mathematics, physics. And how their intuition evolves and whether it parallels the way in which this intuition was actually developed in the discipline of science.
Later on, at the MIT Media Lab we had the technology available that allowed children to build and interact with hosts of very simple-minded robotic devices. These devices had sensors that could actually sense things like light, sound, or obstacles in their environment, and effectors, which gave them a directionality, or tropisms -- they had their own personhood.
Because we could build these devices with the involvement of children, we could also study how, when put in different roles, the children actually talked about the workings in the "minds" of those devices. We imagined different scenarios where the children and also adults were placed in an exhibition room, and where we had the robots do certain things. In a first scenario, participants had to interact with the robots and figure out what they were about, what they were doing, how they were functioning. We then put the same people in a different role, the role of the builder of the device. This time, the metaphor was not : you come to a new planet and there are these creatures, you don't understand quite how they work and you try to make sense of their ways of being. Instead, we used the metaphor of the workshop of the crazy builders. And together with the MIT students, participants were invited to build clones of the creature. We became interested in whether the way they related to the creatures and explained their behavior changed from when they were in the role of surgeon, so to speak, as opposed to the role of observers, just interacting with the creatures.
At the time there were debates between people in artificial intelligence and psychology about questions of purpose, intentionality, goals. What we tried to do was to mimic in the domain of cybernetics what Piaget and his team had done in other domains. In other words, we took apart a seminal interdisciplinary document, referred to as the Macy Conferences, an archive of much of the thinking in the cybernetic community, and we tried to understand which topics, at the time, led to heated debates among people like Gregory Bateson, Wiener, McCulloch, and others., we tried to see how children come to build those notions.

As you see, this indeed is a very diagonal way of entering the field of artificial intelligence!

DG: Perhaps a theme that is going to come up at various times during the evening is the question of whether the models, imitations, approximations of imagination are top-level ones. Many of you know that one of the many arguments against AI was that it was very top level – reproducing the phenomenology but not reproducing the mechanisms of true intelligence. So I think the behavior of something that appears to be alive but may not make use of some of the same principles -- one can have for instance a cartoon produced by Warner Brothers that features creatures that act as if they were alive, that react to stimuli, but it is people drawing things on cels -- doesn't represent anything other than a representation. So perhaps one question is how far do we want to go in the discussion, and how close does it have to be to the actual mechanism of imagination, of intelligence, of other things that we want to model?

RS: It's mostly a misunderstanding, I think, that is the problem. Most people when thinking about the subject don't really get it. AI is not really about computers at all, but it's easy to think that it might be. What happens is that if you are an advertising person, you are interested in how people think. And what you have is the means for testing out your hypotheses. For example if you think that people process a sentence in a certain way, then you might build a program to replicate your hypothesis. The computer is really a kind of testing device.

What happens is philosophers get a hold of it and begin to ask: "is the computer thinking?" which the AI person could care less about if they are a good AI person, because it's hardly the issue. It's very hard to imagine that the computer is thinking in any deep sense no matter what it's doing because it is just doing some narrow thing. For years I built programs that read newspaper articles, would answer questions about them, translated them, paraphrased them. To say it was thinking would require you to walk up to it and say, "So what did you think about that article?" You know perfectly well it isn't going to answer. We used to get a lot of money from the Defense Department and we'd run programs for our sponsors who would show up and want to know how things were working. We would run the stories that would read the best; so for example, we ran a story about an earthquake in Iran, which the computer must have read at least 50 times. But it never said, "Gee I'm bored with this article" or "Boy, there sure have been a bunch of earthquakes in Iran.". Which is foolish – because the reason to do this project wasn't to figure out why it wasn't saying that.

So the next question to ask yourself would be: Well, how would it say that? The whole deal is to figure out what processes people might have and then come up with a replication that is so precise that you could put it on a computer. And what always happens is, you can't. So you put a precise explanation on the computer and you realize all the holes that you have, and all the things that you thought needed to be there. Every once in a while someone gets it and announces "now I am going to solve some problem!" Which is silly, because I don't think it can do that, maybe someday it will, but I don't think that I'll live to tell about it. It's always misunderstood and we always have to get into this silly argument about what computers can do, if they can really think. I don't care about that.

What I do care about is -- and this is a good example of what I think about AI -- is getting computers to plan because I'm curious about how people plan. What I've convinced myself of is that people don't plan and are actually incapable of it. And that all the work in AI that's about planning is about writing up a bunch of rules. So you have all these rules, this is decision-making, you can buy all kinds of books on this, if you are optimized for this and that, right? There are many people who can talk about how computers can do planning, but I don't think that people can do it – people can't plan their way out of a paper bag. And I think that people's planning mechanisms roughly look like "retrieve last plan" because it looked a lot like this situation, let's see if it will work now. Which is why I started talking to Ed [Nersessian] and why I'm interested in psychoanalytic issues, because sometimes the last plan that worked was when you were five and you don't remember what it was. Sometimes that plan was something you dreamt and you don't know if it was something that actually happened. So I've become convinced that people's ability to plan is phenomenally bad.

I'll give you one example of this. Rutgers was playing West Virginia. Rutgers had a terrible team. Through some miracle, they were only down by one touchdown by the first half. WVa had the ball on its own twenty-yard line; it was fourth down with only six seconds left. The Rutgers coach called a timeout. Why? Because there's a standard idea that if you have the ball in your own territory, you have to punt. And if you have to punt, then what the other team needs to do is to block the punt because there's no risk – they might get a penalty, but they could also get a touchdown. So the coach called this play, but he misunderstood: it works when there are twenty seconds left, but when there are four seconds left, the quarterback doesn't actually have to punt, he can just run around for four seconds and it will all be over. Realizing this the WVa coach called a pass play, and WVa went on to make a touchdown. Imagine that the Rutgers coach has a giant plan library in his head, and his plan library consists of all the plays that he's ever used and all the conditions under which he's ever called them. But he has this one miscategorized. It's a very standard play, but no one ever said that if the time is very short then you can really get hurt. He hadn't thought about it and he looked like a complete moron. The point is that's what planning always is. I have a plan library, I'm going to pick one and use it, but I may have mischaracterized it. I don't fully understand how it works and why it works. So your plan library gets very confused. As does people's ability to understand what they do. I think the non-conscious is so much more powerful than the conscious and that the conscious can't deal with it.

I'll tell you one more story that illustrates this. A few months ago I was in a two star restaurant in England with my wife and a British colleague. Welsh rarebit was on the menu and I ordered it. My wife and colleague were both laughing – my wife saying you can get that from Stouffers any day and my colleague calling it the most pedestrian food; why order it in a two star restaurant? I told him I have a rule – this is my consciousness talking – that if I see something on a menu that I don't often see, then I order it. Later, though, I examined my unconscious and came up with the realization that when I was ten years old, I took a train trip from New York to Miami and ate on the train for the first time. It was really exciting and I ate Welsh rarebit. So now, I really wanted to eat Welsh rarebit. The end of the story: this was the most phenomenal Welsh rarebit I've ever had.

The point is that I see the planning mechanism as a kind of war between the conscious and unconscious. My bets are on the unconscious as being way smarter. So that's what AI is to me, it's not about the computer.

EB: So why use computers?
RS: As a testing device. I want to build a program that does well at planning. I could be a psychologist and say that I'll run experiments; I could be an artist and draw pictures. This is my mechanism; this is my way of looking at it.
HC: Then how can you build a computer with an unconscious?
RS: That's an interesting question. You don't want to ask that question with that hint of disdain. Of course you can build a computer with an unconscious! I like that question, but I like it with a positive spin, with uplift.
DG: Another answer to the question is if you think of a computer as being honest. You can think of a computer as a kind of glorified adding machine, and the adding machine will give you the answer, will tell you whether the model you came up with actually fits the circumstance or not. If you had a person interpreting the words, perhaps he/she would interpret the bias in one direction or another.
I was very pleased to hear Roger say that human beings are terrible at planning because I would argue that the imagination requires the inability to follow a plan or the ability to deviate from whatever the logical plan would be. So perhaps your genial contempt for people's planning ability exists because they have enough imagination to not follow a plan.
EB: I was wondering, maybe you could figure out how to build a little bit of unconscious in there, something that resembles the unconscious, but you were saying that the reason you were interested in working with computers was to use them to test. But you have used them to test computers that don't have an unconscious. So, to what extent can you really test this?
RS: Let me tell you something about this because people get confused. Throw out the computer. Imagine that I am a theoretical philosopher. I am interested in hypotheses that are all about how things work. So with imagination, I built a program to do creative imagination. I'm not going to tell you that it was the best program on the planet; it can do some cool stuff, but if you push me and say – well is this computer creative? I'm going to tell you no. My algorithm was clever, I thought it up, but it isn't creative. Why not? Because no computer has – I'm not saying that it couldn't have –that wealth of experience. How is the computer going to get an unconscious? It would have to live, that's how. All the new students who come in always want to work on learning when they first start in AI, and I say to them, well you better start it when it's your age and gradually have it live a little bit. You can't just stuff all the knowledge into it. Everyone in AI wants to bypass the learning and growing problem and stuff all the knowledge in. You can't stuff in the unconscious. The unconscious gets developed from experience. So the computer is kind of a red herring. When you think of the computer, you miss the point.
EB: I'm not exactly sure what AI is or what it should be or what it is said to be, but certainly the computer is wonderful in that it is a glorified adding machine that helps us in testing our theories about very concrete questions we want to solve. For instance, we

want it to be able to automatically recognize a scene as complex as this one. Or far more complex. We know that with our glorified adding machines we can't do it.
RS: We can do these things. If the Defense Department puts money into it, you can build enough of it so that the computer can do very well at it. Computers are very good at being able to detect bomb sites and being able to recognize where countries put their bomb, because that's where the DoD put its money. You don't want to make the argument at that point that the machine is very smart. But you might want to make the argument that the machine is very useful. There are lots of face recognition programs that the CIA and others have built over the years; they are very good. You can say that's AI. But you wouldn't want to push it and say that the computer is very smart because it can recognize a face – in fact the computer may not even know that it is a face. So it's a little more complicated.
EB: I'd like to pursue that thought about plans and that maybe all we do really is pull up one of the plans – that we don't really know whether it is appropriate or not, but its a plan that we have, maybe it's the plan that you have in your subconscious, like eating Welsh Rarebit on the train. So it feels like it's something like a whole situation that you pull out, a large list of things that may or may not apply to the situation at hand. I think of computers as being very useful. By the way I would take issue with your comment about scene analysis. I work in computer vision, and I think that computers are very, very far from being able to analyze a scene. Yes, they can locate bombs in a certain setting but not what we do effortlessly--recognizing a face here, recognizing chairs, couches and legs. What computers do do is solve chess games. How do they do it? They probably store tremendous numbers of situations and then just look for one of the correct situations or plans to apply.
RS: You can't say what computers do. There are lots of computers, lots of programmers, lots of algorithms. So there isn't one thing that computers do. Chess is a perfectly good example. The chess playing programs that do very well are brute force playing programs. There are also non-brute force chess playing programs that are trying very hard to replicate the human thought process. They actually don't do as well, but that doesn't make them less interesting. Brute force programs are really boring because what they say is that by testing every possible move in the entire world – something that no person could possibly do – I can play better than a person. Yes that is true but I don't care. On the other hand if you say to me this program does exactly what Bobby Fisher does but better, or at least more efficiently because Bobby Fisher is a human and might forget something, then I am interested.
EB: Yes, I agree. But take a game like Go. You can't program a computer to play Go. I think the reason is that the combinatorial is way beyond what a computer is capable of.
RS: It's probably way beyond what a person can do.
EB: and yet people can...
RS: See, That's what I'm interested in. Maybe they are not doing the combinatorials.
EB: So what do they do?
RS: Well, that's the question.
EB: I wonder. It's a plan, what humans do. They pick out a plan, a complete plan from a list of plans, and that's what intrigues me. Maybe we have a way to decompose a plan, or a situation, or a vision, into subparts and establish a hierarchy. Maybe that's what a

smarter algorithm would do. But that's not just taking a plan from a list and applying it to a situation.
RS: Now you are getting closer to having a hypothesis.
DG: I think that Go works because humans are doing some kind of pattern recognition. Roger just agreed that some aspects of computer vision are not well suited to analysis of natural scenes. And I think that looking at the Go board and looking at the pattern of white and black pieces is not the combinatorial strategy of moves that works so well for chess. It's "This pattern works well for this response." So I think it is visual pattern recognition and I don't think that computers do so well at that.
EA: Marvin Minsky says that there are two overall attitudes that are brought to AI. One is to use the computer to try to simulate the way in which the human mind works -- and you can see the difference between the ways that a mind gets at a problem and the ways that a computer gets at it. From there you draw lessons...
The second is that the mind is so complicated that any program captures only partial aspects of how the mind works. And this is also true for psychological theories. The mind is so complex that each of these models, say behaviourist or constructivist or associationist, well captures some aspects, while leaving out others. within conclusion, to Minsky the purpose of AI is not , to simulate the way the human mind works but to create alternative ways of perceiving, of thinking, of learning , maybe of designing. Lastly, I want to say is that AI started as a way to recreate high level cognitive functioning. It was about planning; it was about problem solving. For the longest time in cognitive science people were comparing how poorly people were thinking because they'd keep forgetting stuff compared with the computer that doesn't. Or how limited people are because they don't plan. What interests me, instead, is the question of what kind of intelligence allows a person to understand that most of the time planning is not such a good idea, or, likewise, that forgetting may, under certain circumstances, be highly adaptive. ? AI is important because it generates novel forms of intelligences, artificial ones, that you put in dialogue with human ways of achieving the same type of tasks. Then you can start comparing in other ways . I agree with Roger that when we get caught up in whether the computer is smarter or not smarter, or whether it is really alive or not alive, we arrive at a dead end.

EB: From what you say, I'm not completely sure if I understand whether you agree or disagree that a computer is a glorified adding machine. This is a crucial – I hope you agree that it isn't.
RS: I'll make another statement: that the brain is just another glorified adding machine.
EB: How can you say that? The brain is connected to all kinds of sense organs, it is connected to a body, is connected to...
RS: You never want to make these flat statements about a computer. A computer is an evolving thing. How about asking what a computer can be? Could it be connected to sense organs? Of course. You've got to get over this 'computer is' stuff. There are all kinds of things a computer can be. What a computer is today will be very different than what it will be tomorrow. And you can easily imagine a computer with sense organs, so if that's your argument then it's not a good argument.
DG: In fact if we argue about the epistemology, about all the sense organs, all the sensations that the brain has ever gotten come in the form of action potentials, in cranial

nerves or peripheral nerves. And action potentials are a binary code, like Morse code with only dots, no dashes. At the time that you see an action potential, the nerve could be represented by a 1; at the time you don't, it could be represented by a string of 0's. So epistemologically speaking, the input that comes to our brain is understood, and we understand fairly well how the biophysical mechanisms work that process this information. So the inputs are easily explainable in computational terms. But how they work in concert with one another in very large numbers is another question. I think that we should get away from talking about computers and start talking about the brain. That doesn't mean that we can't occasionally start using the analogy of the computer when it is found to be appropriate.
RS: Actually, I find the brain totally uninteresting in the same way that I find computers uninteresting.
DG: The brain hesitates, it thinks, it perceives, it acts, it makes decisions...
RS: No the mind does that. I don't believe that the brain does any of that.
DG: Oh, are we in dualism?
RS: Yes, we are in dualism.
DG: I'd like to know what arguments there are for dualism.
RS: Well if you open up a computer, all you see is a bunch of lights flashing, you aren't going to learn anything about the program that's running. And it's very hard to look at a brain and see what program is running. I haven't found anything in neuroscience that is as fascinating and therefore is telling me something about the mind that I want to know about. Sorry.
DG: You haven't been listening to the right people. Maybe someone can give Roger an idea of some of the programs that the brain is running.
EB: Well, going back to pattern recognition for instance, we are learning things about the programs the brain is running. And the question really is: is there something fundamentally brainy or fleshy, if you wish, about these programs? Well of course there is something about them that computers could do, but they are light years away from doing that. They are hardly any further than these cognitive machines. I don't know what other metaphors people are using.
DG: Every age has a metaphor of what the current technology is. The brain was originally a great shuttling loom when looms were hot things. Then it became a telegraph system, then it became a telephone switchboard, and at the moment it is a computer.
EB: What I'm trying to say is that we are learning by sticking electrodes in the brain, or by putting the whole brain in an MRI machine. We are learning a few things about possible the strategies that the mind uses. And maybe within a few years or a few decades, we are going to learn the way the brain works. Because each neuron is connected to 10,000 other neurons that sit two to three orders of magnitude away from where it is, it has an immediate, instantaneous line that allows it to listen to and to reply to these specified "friends." Because of this scheme, this wiring -- unparalleled to anything else in nature -- maybe there are strategies that the brain uses. Mind is a big word.
RS: I'm sorry, I just don't find that compelling. When we do all that, what are you going to tell me about unconscious planning?
EB: I'm not going to tell you anything about it.

RS: So how am I to think about higher-level consciousness? Essentially I shouldn't listen to people talking about brains because they don't tell you anything. I'd rather think about unconscious planning without having a neuroscientist telling me, yeah, this wonderful meat machine is doing really cool stuff, which I'm sure is true.
DG: Well, one of the wonderful things it does is called learning.
RS: Then tell me about learning.
DG: I will tell you about learning. This argument eerily reproduces one that Donald Hebb tried to address more than 50 years about when dualism was rearing its ugly head, and people said nothing that happens mechanistically in the brain.....
RS: Just tell me something about learning. That's my only subject. I study learning all day and all night, so tell me something about learning that I don't know and will be useful to me.
DG: I am, but I'm also speaking to everyone else in the room. Hebb said, I'm tired of dualism, so what I'm going to do is start putting forth a number of hypothetical mechanisms that could be used by brains, but I'm making them up. Because they are mechanistic they could have some resonance with the properties of mind that people care about. One of the things that Hebb thought about was learning. What could strengthen one of those 10^{15} connections in the brain? He said, suppose we had two neurons, and if one neuron always fires just before the other, than it appears to have a useful influence over the other. And if that connection is strengthened--a difference in timing between the neuron that innervates and the one that is innervated -- perhaps this is a mechanism for learning. And in fact in the last fifteen years, people have found that this mechanism exists very similarly to the way that Hebb postulated. The difference in timing between when a neuron fires and when another neuron fires indicates whether that connection will be strengthened over time or reduced over time. People have also found analogous mechanisms that have caused the strength of that mechanism to be maintained for weeks at a time and people have correlated the size of that strength with the response that an animal makes to a stimulus, which constitutes a primitive form of learning. There is no reason not to think that this form of learning won't apply in other cases because the cellular mechanisms apply throughout the brain. So there's the analogy for you.

RS: Will you tell me something I can do with that information?
EB: It's more than an analogy, because now the game is to try to take these Hebbian specificity rules and put them in a computational framework with a hierarchy, several levels of representation, starting with a very simple thing, hardly more complicated than an action – retract if there is a noxious stimulus, or in terms of vision, react to an elemental stimulus. Bring them together into more complex elements, more complex features, all the way to a representation of a face or couch, or a lamp. And create an algorithm for the brain that explains pattern recognition. No, we can't do pattern recognition. We know that we can't do it in the way that humans do it.
RS: So did either of you answer my question? Tell me what I should do with this. Tell me something about learning. I work only on learning and only on school reform and learning.
EA: EA: One form of intelligence is to understand and represent the world. Another form of intelligence is to be able to envision alternative ways of doing things. One way of defining intelligence is an organism's ability to start establishing a kind of dialogue

between what is and what could be, between the actual and the possible. And I think there it becomes possible to try to define what imagination is. I have become interested in the kinds of techniques that allow people, young and old, to get to think in ways they normally would not. Or to start focusing on issues they normally would not focus on. The capacity to generate alternatives is an important part of the work of the imagination. It is a huge and difficult question. In this regard, The Grammar of Fantasy by Gianni Rodari is relevant. It proposes very precise techniques, or language games, to engage children to think and use words in ways that they usually would not. Rodari's technique allows for rule substitution or displacement. Obviously, hose rules may not be directly useful in the domain of science. Instead, they are the materials that make for the poetic usages of language and thought.

Some of these mechanisms by which the imagination works, however, are utterly relevant to science. They have to do with being able to slightly reframe a problem in order to look at it from a different viewpoint, a point of view that you were not able to take before. Rodari gives a very nice example of this. If you as a person can project yourself into a particular situation you will be able to use empathy to understand instead of keeping your analytical distance. In doing so, you can resolve many "science" problems , like relative motion, in ways that you cannot if you do not throw a little piece of yourself into that situation, and see the situation from within. And then you can almost consciously play these games with yourself where you begin to displace yourself within a situation and to actually see it from this different perspective, using self-projection as a technique to actually understand a problem from within and to be able to change the point of view.

A Japanese researcher by the name of Saeki has pushed this idea even further in a paper entitled "steps toward an anthropomorphic epistemology". He says that in many types of "science" problems it is very useful to throw oneself into a situation. It is also useful, alternatively, to adopt a god's eye view, and to move back and forth between "getting embedded" and "emerging from embeddedness".

So there is a sense in which in our imaginations we project ourselves into a situation. Now even that ability has a genesis. What are interesting are the kinds of reconstructions that people make. For example if you tell me to imagine a walk that I took on a beach. Many people almost project themselves into that situation. Like in those paintings of Friedrich where you have always a small person, seen from the back, and looking at the scene he paint.. That's what we do: We miniaturize our environment, we project ourselves into it, but we don't do it in a way as if we were babies, because the way we felt an experience, as babies, was sensory motor. In this case,. It's a reconstruction: we throw our min-me into the scene as a way of evoking it: this clearly is a mental activity worth studying in greater detail.

Another technique that the imagination uses related to the ability to project has to do with being able to scale in and scale out. It involves miniaturizing the environment in which you put yourself, and then enlarging it. These are just a few of the possible techniques -- becoming smart at playing these "what if" games as opposed to just the "if then" games.

DG: Imagination requires some measure of selectivity. I think that it is very easy to say "what if" and then imagine an enormous number of possibilities which could leave one overwhelmed by the combinatorial expansion of all the many possible choices. The trick

to imagination is the selectivity of the creativity, to say I'm going to go beyond the plan, but I'm not going to do so in a uniform way or in a combinatorial way perhaps as a computer may, but I'm going to be selective. And there is a trick to selecting that gives us our imagination. But until we know what that algorithm is, it's going to be difficult to program it into one of those things that begins with a 'c'.
EB: What Damasio called a 'somatic marker.' What allows us to not fall into endless recursions of investigating each one of the plans? Well, after all, yes, we are made of flesh. Without it we wouldn't go very far – we'd be computers.
DG: But this is an algorithm. Evolution came up with the algorithm – so it wasn't a top down algorithm. But the existence of the algorithm made us possible and gave us a reproductive advantage. One of the tricks of studying the brain is studying it beyond the level of the immediate circuit, to see how it is that those circuits combine in a particular way. One thing we do know from neuroscience is that much of learning and development has to do with pruning those connections. So that the ability to pick out things in a visual scene depends upon seeing those things in the critical period of visual development. The synapses are tuned so that those aspects of the visual scene of what we tend to see over and over again when we are infants get reinforced and our brain develops these synapses. Synapses that represent other aspects fall away. One can look at this at the synaptic level and find some basis for that tuning. In this case, it is for the tuning of perception but it may also be related to the basis of the tuning for planning in ways that we do not yet understand.
RS: You could have done all of that without having had a look into the synapses. The whole point about the mind-brain distinction is that you can examine the mind to see what people are doing and come to the conclusion that people are good at the things they are familiar with from their childhood and bad at the things they are unfamiliar with and therefore tend to interpret things in terms of what they are comfortable with.
DG: That is very gratifying. And now we have a neurological explanation for it.
RS: Oh, I'm not suggesting that you shouldn't do your work, but what I'm saying is that I'm not likely to learn from you, because every time you come up with an explanation, it reflects something that thinking people since the time of Plato have already noticed
EN: It seems to me that what you were saying about the imagination is that there has to be selectivity. There has to be a selection that you make. In essence, Roger is saying the same thing: that he selected the Welsh rarebit and didn't know why. He had some rational thought of being in England. But when he really looked, he found that he didn't really know why he selected it. So the issue is: what promotes, what leads, what directs that selection?
DG: Well, I'd suggest that Welsh rarebit is never an imaginative choice, but perhaps I'm wrong.
EN: Well, it was for him. I'd like to turn to Bill and ask him two things: first, what could he say about the issue of selecting. And second, how he would respond to Roger's comment that he doesn't care about the brain, he cares about the mind.
BG: Ok. I'll be very practical. It seems to me that Roger really described several types of imagination. Your examples illustrate what your points are. I think that we are talking about totally different levels of these problems, understandably, by virtue of our own interests. Roger started off by saying that there were no plans, that we are horrible at planning. But then he described how planning works and gave a few examples. One was

that there is a role for memory and there are a number of templates, and you can look back on your old plans and choose. I once heard interviewed an extremely intelligent autistic woman named Temple Grandin, who had a lot of interesting thing to say. Apropos of selection what she said was that we should add categorization. She said, if you think of a building you have a general idea of a building. When I think of a building it's as if I have a library of CD's and I now have to look at all the buildings before I have an idea of a building.

Roger was describing this as part of the planning of the football coach. But he also said something else that I think was important. The selection [of possible plays] could be limited; it could be an error because there is a faulty analysis of the nature of the context in which this template is being used. So it could be inappropriate, it could be making mistakes. And one could then have another template or another method which has shown its value and that would solve that problem. So you've described a whole set of steps that would go into the production of what we would think of as an imaginative outcome. I think that's an example of imagination at work. I think that the fact that you came up with the analogy is example of another level of imagination at work.

And then Roger described a different problem exemplified in the choice of the Welsh rarebit. He correctly point out that if you are looking around for templates, it pays to go back very far, especially if it has to do with appetites because there's likely to be a long history. That's another feature. It could be unconscious, you could have a reason for it, which might be a good reason that's conscious, but you would also have a reason that is unconscious.

Now it gets complicated because the fact is that there are several kinds of unconsciouses. There can be unconscious reasons that actually have to do with this process of going through a set of templates. What we may call motivational factors might have to do with the choice of the template. The fact that the coach picks a lousy template, or he doesn't pay attention to the whole picture may be based on an early template of failure which he is driven to reproduce unconsciously.

It seems to me that if you do a top down assessment, you'd want to start with this kind of breakdown of what people actually do, which is what Edith is studying with children. So you can delineate these features of the planning process, and there are probably others. Once you delineate this process, there are a number of other systems of templates. Basically that's the crudest form of ancient formulaic systems: there were two systems of templates called primary and secondary process. People don't really like that today. It's more complicated, but it's an important differentiation.

RS: Let me tell you about a computer program that I wrote some fifteen years ago. It was based upon my walking into my class one day and saying, "Swale died." Swale was the racehorse of the year 1984. He had won all types of honors and at the age of three, he dropped dead in his stall. So I said, '"Swale died. Why?" I was talking to a bunch of computer scientists who surely didn't know who Swale was, but because I was interested in the explanation process at the time, I said, "Listen: I'm going to sit here until you come up with some explanation." Someone said, "I'm reminded of Jim Fixx –the guy who died while running? – and maybe he just ran himself to death. Running is really bad for your heart." So I said, "Well, sounds reasonable, but other horses haven't died." Another person volunteered that maybe he was killed for the insurance money. So I said, "Well, yeah, maybe that's true." Someone else said, this reminds me of Janice Joplin – she was a superstar and lived life in the fast lane." So I said, "Ok, let's go with that one for a while: Swale was also a superstar, but how would that cause him to die?" They said he might be taking drugs. I said, "Well, he doesn't have any hands." They said that maybe his trainer was giving him drugs. So I asked why his trainer should give him drugs. To speed him up. And suddenly we had a hypothesis: that this horse had died because he had been pumped with drugs by his trainer in a way that had been bad for his heart. So we built a program to do what we just did. We were being creative and we were just trying to imitate that process.

Trying to look at imagination and creativity is similar to that kind of process, taking a set of past explanations that work, and trying to apply them in a new context and tweaking the explanations in the changed context where it does not work, to see if it could then work. I think a lot of what creativity is about is that process.

BG: There's another side to it which has to do with the kind of recall that you are talking about with your childhood experience. It has to do with the capacity for imagery and the fact that one of the useful hypotheses is that imagery is one of the earliest and most reliable systems. If you can tap into it by whatever means, for instance by free association which is what you've' been describing roughly, and you can get imagery, then you have a certain set of templates of a different kind.

Apropos of the relationship between the arts and sciences, there was a fascinating article in *Science* a number of years ago about imagery and mathematicians. One of the interesting things about the mathematicians who were studied — these were bigwig mathematicians — was that they all had visual imagery when thinking about the problems they were considering. So although the problems and the solutions were abstract ones, the process involved was based on imagery. There's a related line of thought that goes with what Einstein wrote when considering how one gets a theory to work on. He said that there were really two parts to the process. One is that you have to come up with a mathematical model of some kind and apply it to the real world. Teaching such a model is easy. But coming up with the right model, that can't be taught. We can't explain it. That's the creative part.

The questions about creativity and imagination is where do they come from? How do they occur? I think that you have broken down the process to some extent at that level to talk about how it happens. That's what people are interested in these days. For instance the physicist Peter Galison, Professor of Physics and Philosophy at Harvard, writes about what he calls the disunity of science. It's how scientists are all doing what we were doing tonight earlier – they're all doing their own number. The real problem is to bring them together, to cross the boundaries and to get something creative out of it.

EN: That's our purpose.

BG: So you're in good company. Another interesting thinker is Bruno Latour who has written a lot on what he calls science studies. What he does is to show through many examples how the thought process in the arts and sciences is similar. Other people have made a big point of the fact that even though what is going on in arts and sciences at the same time is mutually stimulating, the final tests or goals are different even though there is a lot that is similar as creative process.

A point about the brain: I've now heard of at least two researchers who are coming up with complex models of functional organization. For instance, they investigate how drugs work, where in the brain they have their effect, and what the relationships are between those sites of action and their interaction with other systems. The important part is not about the specifics, but that more and more complex models are being built. It's not just that at the single cell level certain things may be said, and at the learning process and protein synthesis level other things may be stated. The researchers are working at the point where there are certain groupings of functions in the brain and in the mind that are influenced by and influence one another. That begins to look like the process of how people actually think by using different aspects of their thinking process, by using a dynamic model that is responsive to changes in any part of that part of the system.

DG: Steven Kosslyn from Harvard and others have made a career of correlating the fact that when people imagine images, the same areas of the brain light up in brain imaging techniques as when as are seen when these are observed. But there are interesting variations. In some cases the primary visual cortex is activated and in some cases not, and it depends on whether they are interested in the gestalt or the concept that is being visualized or asked to focus on a small detail in the scene. I find that fascinating, because that's a great correlation between a process that we understand is central to imagination yet at a higher level phenomenon of the brain and yet we can measure it in a fairly quantitative way.

BG: And it has the potential for uncovering new functions that we wouldn't ordinarily think of from the grossest level of observation. It may be that part of the answer to Roger's question – what can you tell me about learning – may come from the possibility of adding to the collection of templates and contingencies that would lead to a better understanding of what it is that we observe. This is my argument against analysts who ask what good neuroscience is. What I see is that if listen to this kind of research, you can listen to people differently because you learn something different about the way people are organized. And when you begin doing that, you begin to hear that there are differences at the level of clinical observation. You can listen in the clinical situation for some of those differences in the way that patients are expressing approach some of their material. It offers the possibility of breaking down the system, or enriching the system of templates and contingencies.

The other piece that I think is very important – it's conceivable that you could build it into a computer --is the drive aspect. And one of the important things about human brains is that it already has the motivation mechanism built into it. Gerald Edelman who has built little creatures that do have drive functions wrote about this. He once described how he put together such a creature, and once he got it adapting he and his colleagues were in the position of examining where the activity was going and how things were integrated in the circuits. It turned out that there was some part of it where they couldn't figure out what it was doing – it was very active but it wasn't obvious what it was doing except that maybe it was a transport place for other things. The point is that he builds something that he calls 'value' and that is equivalent to drive into his system. And partly what drives do is they place a certain value on things like Welsh rarebit. It seems to me that it's possible to break down the process if people throw all that they know from the surface into it. And there are correlations that may feed back into it and then help to understand how to further break it down.

DG: We've been talking about machines as if they couldn't have an unconscious; I think machines have an unconscious all the time. What they don't have is a conscious. --

Everything that machines do, they do unconsciously. The interesting thing about brains is that they have both.
FL: But they don't have an unconscious like a human does.
EN: No, but it's an interesting way of looking at it, an obvious way – that computers have no consciousness.
RS: Those are good questions. Why does everyone always want to say about things they don't know how to do that they can't do them? Why don't people just say, "That's an interesting question, I wonder how you would do it?"
EN: That's an interesting question; I wonder how you would do it?
BG: One of the reasons it's a problem to make such comparisons is because we don't know how the brain does it. We know something, we know a lot about imagination of certain kinds and in certain ways. But we don't know what to compare. Even if we wanted to compare it to the computer, we don't know what to compare.
RS: I'd make the same statement about brains. How do I know they have consciousness? For all I know, you are just black boxes staring at me. The only consciousness that you know is your own. The rest could just be a good imitation. That's why I say all people are fleshist. So if you are sitting there and it turns out that you are a computer all this time, I would now discount you as being unconscious. Meanwhile up to this point you had been conscious. So this is all about my fleshist attitude – that anyone sitting there looking like meat is probably thinking, but anything sitting there looking like silicone is not.
EN: I'd like to ask Elie to say something about the role of the visual and of the image in imagination. Imagination is often connected with the image, seeing and the visual system though we know there are other kinds.
EB: I was actually thinking about what we call 'generative models' in models for vision – whether actual or artificial. We don't quite use the word imagination, but it is something that comes close, this notion of a generative model which is something that creates rather than putting out templates or schemes from past experience. It generates as though by playing, by imagining, themes that resemble the outside world. The task is very concrete. It's one of recognizing songs, for example, or something more ambitious, where one can recognize faces or other visual objects. Imagine a system that had some innards, some internal machinery that interfaces with the world, something like a retina which allows information to pass back and forth between itself and the world. What the system does is to generate images on this interface, on this retina. If it was working as well as we would like it to work, then one would see on the interface – just as on a television screen – images that are as believable and as varied as images those emanating from the real world. This is something that we can't do. It's easy to generate templates, snapshots of images, individual objects; but to have both the variability and the accuracy, the faithfulness of the real world is something that we can't do. If we could then we would be very close to having a good algorithm for pattern recognition.
There's an interesting interplay here -- I wonder if imagination is the right word -- and a very concrete path where we can actually measure something in terms of how well the computer is performing something akin to vision and pattern recognition. The way that many people approach this is to have models that have hierarchy, hierarchy structured in a way somewhat analogous to a grammar, a grammar of shape. So we have rules, though not strict rules, usually cast in a probabilistic framework so they are regularities rather than rules: Object tends to be so and so, delimited by boundaries, and within these

boundaries, texture tends to be similar. These rules are nested within each other somewhat like linguistics rules. The question then becomes how to use these rules to generate and synthesize images.
That's the general framework and approach. The category is Bayesian probability, the notion is that there is a prior probability, that is, the probability of the world that we are trying to capture in a generative model based on an observation. Once we observe something, we use this generative or prior probability to generate a posterior probability after the observation. There is an interplay here between creating, generating an internal mapping, and letting the information that comes in from the outside guide this process.
DG: But it's recursive, because the posterior generated at one iteration eventually becomes the prior for further refinement.
EB: That is the learning part, yes, exactly right. I'm wondering whether there is any connection here with imagination...
EA: George Stiny, a colleague at the School of Architecture, developed generative tools for designers which he calls 'shape grammars'. What this tool does is to allow one to play the game of seeing a same shape in different ways, as foreground or background, for example.
DG: I have a comment about the content of applying Bayseian work and graphs of inferences that one could deal with as being compounded by several people as a mechanism for scientific creativity for learning. People are beginning to explore whether machine algorithms can be developed that can make use of Bayseian recursive thinking applied to the question of working towards a machine method for discovery for -- I hesitate to say for creativity or imagination -- a separate approach for reasoning. We've touched on this problem of having a universe of choices and saying that imagination perhaps is a selective one. This idea is expressed in terms of probabilities, not in the classic sense that if you pull a certain number of white balls out of a bin and a certain number of black balls, you will get a sample of the total universe. Not one that is derived from some sampling but one derived from this sort of Bayseian analysis in which the probabilities become estimates of what we know about a situation and what things we can do to extend these into the future. This relates in a nice way to the problems that we have been trying to get at.
EB: Is the question then whether in all creative or imaginative aspects of the human brain we rely on visual image? Or is there a form of the imagination that is not strictly visual?
EA: I would like to suggest that – yes.
EN: Why?
EA: The word "imagination" is very strong, has a very strong visual component. But when you look at the beginnings of symbolic functioning in children, it has a different manifestation. It's the beginning of language. It's the beginning of pretend play. It's the beginning of having imaginary companions. It's the beginning of dreams. And it's also the beginning -- as soon as the children learn to talk -- of the obsession with leaving traces behind. I have the impression that in the field of psychology there has been a big emphasis on the role of language and figurative forms of representation. And there has been an underestimation of what Jerome Bruner calls the enactive modes of representation. With the beginning of the pervasiveness of technologies that allow for dynamic modeling and for simulation, there is a resurgence of interest enactive representation. These are more of the kind that I mentioned earlier. It's more about

reenacting an original scene as in pretend play and being aware that the scene that is being reenacted is not the original, almost with a sense of humor that goes along with it -- the idea that you can replace, substitute, the props that you normally use to enact certain scenes with other props because it's not quite the original. You can also give a different ending to the scenarios that you are reenacting through pretend play, because the world of play offers a space in which it is less risky to try out all kinds of outcomes or to change the original conditions. The great connection between this ability to reenact such scenes and the ability to develop a sense of humor or a sense of incongruity or slips of logic is a very important part of how humans mediate their experience through symbolic activity. This doesn't mean that the visual aspect is not present. Instead, configurations emerge from mental journeys.. So, for example, if you ask a person to imagine a square, and you capture that person's eye movement, you will see that the mind, the brain actually walks the shape that it is representing visually. So my take is that the temporal and the figurative aspects cannot exist without one another.

The way I formulate this to my students is to say that when you have a configuration a good way to wrap your mind around it is to actually unravel it is by walking around it, through it, so you transform a more visual representation into something more kinesthetic, rhythmic. Conversely, the best way to keep track of a journey is to spatialize it by drawing it. One of the experiments that a student of ours, Bonne Smith, in the School of Architecture is doing is called "Dwelling into the Drawing." The question addressed to the fellow students was: Can you draw the floor plan of the house in which you lived when you were five years old? Bonne videotaped the way in which her subjects were drawing the floor plan of their childhood homes and recorded their spontaneous commentaries. What you get is something like: "Well there was an entrance here and I remember that I went down a staircase..." as if the pencil served as a means of mental teleportation. And then all of a sudden the pace of the drawing changes: "…and then there was grandmother's room [subject draws a square]. There was a bed… I don't quite remember…." So it was possible to identify very precisely just by the pace of the drawing whether the subjects were just trying to jot down a particular shape of a container (the room) or if they were actually inhabiting the space and using the pencil as an extension of themselves, or avatar

DG: This is a motor recall really--this is kinesthetic, which answers the question perhaps to the visual and the representational, by words or concepts.

EN: I wanted to ask Hallie how it works in terms of painting. How much is kinesthetic, tactile?

HC: I'd rather approach it in terms of how I work with students who are drawing something. Are you [EA] also describing what you were talking about in terms of scale change, in terms of moving in to the specific or the general, and therefore moving closer to the experience or further from it?

EA: Yes.

HC: When I teach drawing, I employ many ways – there is not only not one way to draw, but each individual has his/her own experience of creating. So I do many different kinds of exercises with students. One of them has to do with a much more kinesthetic experience. I use their signature as the first template. They write their signature which is a series of familiar and automatic gestures that create shape on a page. Then I ask them to go through a series of exercises with that self, with that name. They can try for example to imagine that they are just learning to write their name and to write it that way. They go through a series of processes where I ask them to re-experience different stages of either a feeling or of a time, chronology. Then I'll ask them to be a very old person signing their name, like their grandmother giving them their birthday check, watch her drawing out her signature... It relates in that they experience the creation of something. That is as you say a very physical experience.

But then I also have very visual exercises where I have them look at something and be able to see it in various relationships. And apropos of what you were saying about putting together and taking apart, I try to explain that one never draws some thing. We're always drawing something in context, so whatever it is can be very hard for some students and extremely easy for others. You are constantly juggling visual stimuli as you are doing one thing, so you are constantly referring to things. But you've got to be careful of the over stimulation of things. You have to constantly circle around, go back to what it is that you are seeing and thinking and reproducing, and then constantly see it in relationship to everything else: How high is it? How far away? When people who are not artists draw, when you ask them to draw, they are still writing -- which means that they are only seeing it from one part of themselves. But drawing is a whole body experience, even though you may only use your hands.

EB: There is a fascinating experiment carried out by Paolo Viviani. A very simple experiment: The subject looks at a computer screen, where a dot describes a circular pattern. It's a perfect circle. The subject is asked to determine if it is an ellipse, and if so if it is elongated in the horizontal direction or in the vertical direction. The subject is tricked in a sense in that the speed of the dot along the circle varies as a function of where it is on the circle: In some cases it will go faster when it is at the bottom or at the top, but slower along the edges; in others it will go faster on the vertical part than on the horizontal part. And it turns out that the interpretation is always consistent with the laws of motion of the wrist or of the finger when the person writes or draws -- that when there is a high curvature, the hand slows down; when the curvature is weak, a straighter line, the speed is faster. So you incorporate in your perception the laws of motion of your hand.

EA: What is interesting about some of the techniques you were describing before is that they have more to do with learning to see. In order to learn to see, you have to learn not to see the way you normally interpret what you see. So it's a peeling off of the habitual meanings that you tend to attribute to a scene. It's the assimilation idea of Piaget. Paintings are shown reversed, or a grid is used. There are other techniques of a similar kind to train the mind to play this pretend game. Learning to see doesn't have to do with getting rid of the interpretation and perception of reality. Instead, by multiplying the ways of interpreting, or putting on different lenses, you actually learn to see through things.

HC: Well you have to "unknow." When I put my students through some of the more rigorous exercises, they have expressed that their brains hurt. I don't know if that's a technical possibility.... Have you ever tried to look at something and not see it the way you see it? That hurts. You have to stop using language, stop identifying that something as you know it. And it's very painful. So you almost have to go through a certain kind of regression to get to a point of some place that is very unfamiliar, very difficult, to be able to capture something in a different way.

FL: Isn't that what a great work of art does for people? It makes them see things as they have never seen them?
HC: I've had that experience: I went to a show of Holbein drawings at the Morgan Library a number of years ago. It was very crowded and people were carefully examining his gorgeous drawings. What I experienced was I saw the drawings but then I looked around at the people and I saw the people for the first time differently because I saw them as Holbein had clearly been seeing them. So I myself experienced his particular perceptual abilities and then could actually see people differently -- I saw the person with the large head, and the prominent nose, and the high forehead, and they all became they way Holbein saw them. I asked my students: So what is the art? Is the art the drawing? We know that would sell for quite a bit of money and most of us know that good art is expensive. But the art is not the object. The art is some kind of experience of seeing and knowing and understanding. It is more than just the visual, because it is an experiential thing.
DG: These are acts of imagination. This is perhaps the best definition of imagination that we've had all evening.
HC: That's sort of what I traffic in most of the time. Because imagination transcends and goes through time and space and people.
DG: I like the use of the word 'unknow' in your first example. You talk about the experience of your students rather than your personal experience. And indeed the ability to go beyond what is obvious may require some form of unlearning.
HC: In an educational environment much depends on the age of the student. I teach from college age on, and much of it is about taking away their perceptions and their knowledge.
EN: It's a very hard thing to do because one tends to want to look and see and think about something the same way over and over. To stop and think differently is not easy.
EA: I love Davis Byrne when he says "Stop Making Sense".
FL: Earlier in the evening we were dealing with the question of whether to understand the human imagination by understanding the computer as aspiring towards human imagination, we would try to understand some aspect of the human imagination. I can't help but think about an Errol Morris film we saw last night, "Fast Cheap and Out of Control." I kept thinking about the ant colony like the multitasking of computers – and that somehow a colony of ants seems to have something in common with the fragmentation that you might have in any kind of attempted facsimile of the human mind. In other words you have the ant colony with all the fragments working as a whole, and producing a consciousness through multiparts. Is there a multitasking element to the computer's trying to facsimilize the imagination? We talked about the visualizations process, the hand movement, the mind. Could you program a maximum number of these faculties in succession that can network the computers that would become like an ant colony?
DG: Again, the numbers involved get very, very large. We have 10^{12} neurons in the brain, which is a very large number, especially because individual neurons can have different parts that are operating independently of one another. And with 10,000 or 60,000 inputs to one of those 10^{12} nerve cells, there's a lot going on. Most computers have got a handful of processing elements so maybe they are doing half a dozen or sixteen things at a time. It's a little hard to talk about computers and humans in the same

breath when we start talking about the degree of massive parallelism involved in the way the brain works.
FL: Someone like Ray Kurzweil thinks that we will be replacing the brain some day. You won't die, you'll never die. You'll have microprocessors, duplicating processes.
EN: Who wants that? So are you saying it's only a question of numbers?
DG: I'm not saying that it's only a question of numbers; I'm saying that the question of numbers is daunting, that it's much more than a question of numbers. The more we discover about a single nerve cell, the more complex we discover their function is. The individual gate inside the computer is very, very simple in comparison to the processing units of the brain. Again, these analogies are highly suspect because although the brain may perform some of the aspects of computation -- and I can be glib about the way we can digitally simulate what goes in -- there's no reason to think that the way a computer does things and the way a brain does things are similar. There's no reason why they should be: one is an artifact constructed according to a series of rules making use of available technology; the other is an evolutionary product that has been honed and refined in many different ways, which seems to work rather well, because nerve cells have essentially the same properties in almost any organisms. Insects and crustaceans tend to have nerve cells that work a little bit differently from the ways ours do, but worms, leeches, and all vertebrates including mammals and humans seem to have nerve cells that function much the same way. I don't want to draw too close an analogy here, and I'm not too sure that the computers are yearning to be conscious. I'm personally looking to know how the brain works. I don't know if computers are going to be anything more than a tool to help discover in the way the microscope, or the microelectrode, or the MRI is a tool.
FL: So the computer isn't ever going to know that it's a computer.
DG: I don't know what that means. Because I don't know what it means to be a computer, and I don't know how I would know if I were. There's a certain solipsism involved in talking about self-knowledge of anyone, from Roger Shank onward.
HC: Science fiction does that a lot.
DG; It does but I'm not sure that it illuminates the questions that are of interest to us. I'd rather hear about art because I think that does illuminate the questions.
EB: There is, though, one interesting somewhat devious theme of the Turing test. When I talk to a person at FedEx, I think that she's a human. Turing was imagining a dialogue which would allow you to tell a computer from a human. Now some nasty programs on the web tend to attack websites and pass themselves off as humans. Some of the devices that have been imagined to fool these or to erect a barrier in front of these programs involve asking the person or the program-- because the program won't be able to do it -- to analyze a visual scene. For instance you see a number of characters with some kind of variability and some kind of background that fools you, and when you as a human look at it you solve the task quickly and without any problem. One of these 'human interactive proofs' is a situation where you have a face broken down into a number of fragments and illuminated from different angles to reproduce a different position of the image. There is only one face that is whole. You see it within a fraction of a second, you point to it. No computer vision algorithm today has been able to do that. So it's a sort of Turing test in reverse.

FL: A basic human property of imagination is that it retains the ability to know mental representation.
EB: Well the question is: is that imagination or is it more pattern recognition? Is it a form of pattern recognition that is so complex that it requires imagination? That is what I would like to think. In other words, in order to solve such a task of pattern recognition, you have to be able to imagine a face in a multitude of different situations and the combinatorial explosion is such that just pulling templates one after another will never in your lifetime, in fact never in evolutionary time be enough to solve it. What is needed is an entirely different approach, a different strategy, a generative approach. This is, of course bias, but I think that it points at the power of imagination.
Audience member: I recently saw an interview with Phillip Roth about his newest novel. He said that while reading a work of history by Schlesinger, he read a sentence that talked about the fact that in 1934 the Republicans were thinking about running Lindbergh for president. Roth said that that sentence immediately jumped out at him and his novel evolved. Is that pattern recognition? Somehow at that moment he was free enough mentally to be riveted by that.
EB: It's not pattern recognition in the sense that one recognizes cognition again; it's pattern cognition. It creates a new pattern.
AM: But is that imagination?
EB: I think it is. I think that at some point it should become indissociable. Algorithmically they both represent the same kind of process.
AM: In my view, that Roth was free enough mentally at that moment for something to jump out at him is a clear example of imagination.
DG: Pattern recognition tells us it would have to be '44 rather than '34 because this was the pattern of presidential elections. That's a limited application of pattern recognition…
EN: What do you think about that Bill when he read this one sentence and he had the whole novel come to him?
BG: Again it's a question of what kind of answer you want. It's clearly what we mean by imagination. That doesn't actually explain anything, but at least it clarifies our vocabulary. But as an analyst, I would ask myself or ask Roth what his interest was in Lindbergh. There could many reasons. He's approximately my age, but I have my own associations, and I'm sure that his are not a lot different in certain respects. I think you are asking into what network of ideas this belongs that it should have an appeal to him now as a possible line of further elaboration. It would be interesting to know if this event occurred to him after he wrote Patrimony or while he was writing Patrimony because that would certainly make a big difference in the kinds of associations he would have. You might say that my thinking of Patrimony at this moment would also be in the same category. So is it pattern recognition? Well, not exactly. But it must depend on pattern recognition because, after all, a story is a pattern too. So we might say that from one point of view it belongs to a set of stories that get evoked by the name Lindbergh and that the question of whether he going to be elected coming up at this time in our current history suggests a lot of possibilities. I think that's the kind of thing we are talking about, and I think a lot of creative imagination works that way. I tend to think that in some ways, artists are not a good place to study creativity simply because their creativity is so dramatically evident. If you want to know what goes into creativity it seems to me that one of the ways to find out is to be interested in the creativity of ordinary people. Because

they are creative in similar ways but a lot of things lead them to identify themselves as ordinary, not creative. So from our point of view, if creativity is a problem, then we have to look at what is ordinary creativity.
EN: Thank you everybody.

C:/my documents/AI Full Edit with Participant Edits 7-05
11/05