Notes on "The Society of Mind"

Manushaqe Muco

12/13/2015

I. Introduction

"The Society of Mind" is a book by Marvin Minsky, attempting to answer the questions of what the human mind is and how it works. This paper is meant to be an exercise in active reading of this book. I will go through each chapter of the book in the form of notes. The nature of notes will vary, from summaries of what I got from reading each portion of the chapter, to questions and ideas the reading inspired, to references recalled, etc.

II. Notes on the Organization and Structure of the Book

The book consists of one-page essays, with each essay introducing at least one (new) idea. From a structural standpoint, the book leans towards light reading. However, that is not the case based on its content. The amount of new ideas the book throws at you makes it a challenge to read at times, but it is also part of its beauty. One problem I have with the organization of the book is that there is no bibliography or required readings at the end of each essay. The bibliography is included with the glossary, at the end of the book.

III. Notes on the Book

1. Notes on Chapter 1: Prologue

1.0. Prologue

Minsky sets out to explain how the mind works. He proposes a scheme he calls "Society of Mind": the mind is made of agents (small, thoughtless processes). By itself, an agent can only do simple things. But if these agents get organized (in special ways) into societies, they can give rise to intelligence.

Notes: There is something to be said about why the agents needs to be thoughtless/mindless. Else we fall into the homunculus problem: explaining a phenomenon in terms of the very phenomenon we want to explain, thus getting stuck in an infinite loop. An example from vision:

A: How does human vision work?
B: Think of a little guy inside your brain watching the images of the outside world projected by the retina as if it were a movie.
A: What happens with the little guy inside the brain?
B: Think of a little guy inside his brain watching the images…

1.1. The Agents of the Mind

Minsky provides a list of questions to consider when thinking about the smaller parts that compose the mind (agents, in his case).

"**Function:** How do agents work?
Embodiment: What are they made of?
Interaction: How do they communicate?
Origins: Where do the first agents come from?
Heredity: Are we all born with the same agents?
Learning: How do we make new agents and change old ones?
Character: What are the most important kinds of agents?
Authority: What happens when agents disagree?
Intention: How could such networks want or wish?
Competence: How can groups of agents do what separate agents cannot do?
Selfness: What gives them unity or personality?
Meaning: How could they understand anything?
Sensibility: How could they have feelings and emotions?
Awareness: How could they be conscious or self-aware?" (1)

Notes: I think it's important to come back to these questions after reading "The Society of Mind" to evaluate how the book answers each of them. I will return to these questions in the Conclusion section. At the same time, a theory of the mind that uses the notion of agents, or that is inspired by the society of mind approach, will have to confront these questions.

1.2. The Mind and The Brain

"How could solid-seeming brains support such ghostly things as thoughts?" (1)

The mind-body problem still haunts us. Can machines be conscious? Can they think? Can they feel? According to Minsky, there would be no problem, if we had the right kind of theories about how thinking works.

Notes: I have to agree with Minsky on this one. There is no reason to believe that science cannot do away with the mystery of consciousness the same way it did away with animism when it comes to living things. It puzzles me why people are so attached to the idea that the mind has to be incomprehensible and beyond our understanding.

It also puzzles me that people are opposed to the idea of seeing themselves as a machine. To be more precise, we are fine with viewing our bodies as machines, just not our minds. Take for example, a mechanical philosopher (considering the universe as a large scale-mechanism) like Descartes. He saw animals and human bodies as nothing more than complex machines. Except that according to him, humans possed an animated soul that controlled their bodies.

1.3. The Society of Mind

Minsky gets the reader to start thinking about the mind as a society by using an example of picking up a cup of tea. The act requires the interaction of many different agents: GRASPING, BALANCING, THIRST, MOVING, etc. None of the agents appears to be requiring much thought. Same goes about talking. However, "these processes actually involve more machinery than anyone can understand all at once (1)". Minsky seems to point out that we need to analyze ordinary activity, in order to get insight into how our minds work.

Question: How are agents activated in the first place?

1.4. The World of Blocks

To illustrate the blocks example, Minsky introduces the agent BUILDER, which takes a set of blocks and builds a tower using them. The job is too complicated for a single, simple agent like BUILDER to do alone. BUILDER calls for help the agents: BEGIN (choose a place where to start the tower), ADD (add a new block to the tower), and END (decide whether tower is high enough). Each of these agents relies on other agents, too. For example, ADD uses FIND (find a new block), and PUT (put it on the tower), and so on.

Notes: In this section we see the notion of decomposing complex processes into simpler processes. It is also suggested that the organization of these simpler processes is hierarchical. In later chapters, we will see that it is not a perfect hierarchy.

1.5. Common Sense

Even "simple" agents can raise complicated questions. Two important questions raised here, are how an agent understands what to do, and, how to define a goal for a machine. We rarely think about things we do commonly. However, "common sense is not a simple thing. Instead, it is an immense society of hard-earned practical ideas -- of multitudes of life-learned rules and exceptions, dispositions and tendencies, balances and checks (1)". Why is common sense so natural to us then? It is because we keep building new layers on top of old skills. As time goes on, the bottom layers become more remote. This could explain why as adults we find childhood games (like playing with blocks) so dull.

1.6. Agents and Agencies

When thinking about agents, it is not enough to explain only what each separate agent does, but also how agents interact with each other. For example, what a group of agents can accomplish. Minsky takes some time to explain why agents are also called agencies in the book, and vice versa. I don't think the distinction is that important, but the short answer is that it depends on the point of view. An agency seems intelligent, but an agent does not. For example, if you see BUILDER from the outside, you would think it knows how to build towers, thus an agency. If you see BUILDER from the inside, as a collection of other agents, you'd see it nothing intelligent about it, thus an agent.

2. Notes on Chapter 2: Wholes and Parts

2.1. Components and Connections

Section 1.6. told us we need to know what each agent does, and how it interacts with other agents. This section points out that agents need to be linked to one another using a suitable network of interconnections. Minsky draws a parallel between the way of understanding how complex things/systems work, and how the brain works:

1. Understand how each separate part works -- understand how different types of brain cells work
2. Understand how each part interacts with other parts to which it is interconnected -- understand how brain cells of each type interact with other cells to which they are connected
3. Understand how local interactions combine to accomplish what the system does as seen from the outside -- understand how brain cells are organized into societies.

Notes: Something to think about is what level of smaller parts will be suitable when trying to understand the mind: molecules, neurons, group of neurones, etc.

2.2. Novelists and Reductionists

According to Minsky, reductionists build on old ideas, and novelists create new ones. If we look at a science like physics, it is amazing how they depend on so few explanations. However, Minsky rejects the idea that psychology has to be like physics. Of course, we are not going to reject the laws of physics. But, “the `laws of thought' depend not only upon the properties of those brain cells, but also on how they are connected (1)”.

Notes: Perhaps way more on the connections, than the substrate.

2.3. Parts and Wholes

For Minsky, “gestalt” is a term we use when things combine to act in ways we can’t explain. These kind of terms do not help when it comes to our understanding of the mind. Minsky also talks about the difference between objective and subjective questions. According to him, subjective questions (about art, traits, styles of life, etc) are in fact quite technical. It’s just that we and our sciences haven’t learned very much about these subjects.

2.4. Holes and Parts

"Are life and mind so much more than the ‘sum of their parts’ that it is useless to search for them?" (1)

As we get better theories, we can do away with old notions of mystery. We can explain the property of “containment” of a box in the way each wall prevents movement of a mouse in that direction. We will be able to do away with the mystery of the mind, once we form better theories.

2.5. Easy Things are Hard

This is one of the key insights provided by AI. In trying to make robots work, we realized how complicated everyday problems are:

"In general, we're least aware of what our minds do best." (1)

2.6. Are People Machines?

People don't like to view themselves as machines. Minsky points out that one of the reasons is that our definition of a machine is outdated, and introducing a new definition will do away with the negative connotations of the word.

3. Notes on Chapter 3: Conflict and Compromise

3.1 Conflict

One of the ways agents interact with each other is conflict. For example, the agent WRECKER can conflict with BUILDER, when a child plays with blocks. Minsky assumes "that conflicts between agents tend to migrate upward to higher levels (1)". Why would that be the case? Because prolonged conflicts between agents tends to weaken their mutual superior. The conflict between BUILDER and WRECKER will weaken their superior PLAY-WITH-BLOCKS.

Questions: How does the conflict arise in the first place?

3.2. Noncompromise

When agents conflict with each other, how is the conflict settled? One way is the "Principle of Noncompromise":

"The longer an internal conflict persists among an agent's subordinates, the weaker becomes that agent's status among its own competitors. If such internal problems aren't settled soon, other agents will take control and the agents formerly involved will be `dismissed'." (1)

For example, the conflict inside PLAY weakens it, and makes it easier for SLEEP or EAT to take over.

Notes: However, I don't see an answer about how long it takes for other agents to take over. The longer SLEEP and EAT wait, the stronger they will get, and the weaker PLAY gets. If they wait too long, there's also a chance that PLAY will start to get stronger (stop the conflict), or some other agent will take over. The question of how long to wait remains unanswered.

3.3. Hierarchies

The notion of agents being organized in hierarchies is introduced again. Just like in human societies, when something becomes too complex and large, we need to introduce bureaucracy.

Noes: Although, when it comes to humans, bureaucracy is a nuisance and inefficient, it might not be so bad for the brain. Redundancy makes the brain robust against failure. It would be bad if an agent failed and there was no duplicate or no way to retrieve information about it.

3.4 Heterarchies
The agents cannot be organized in perfect hierarchies.

"When two agents need to use each other's skills, then neither one can be `on top'" (1).

For SEE to find out what's behind block A, it needs to MOVE block A from the view. But MOVE will have to SEE if there are any obstacles in the way. We might also need agents to do several different jobs at the same time. This introduces the need for cross-connected rings and loops.

3.5. Destructiveness
Let's go back to the conflict between PLAY and SLEEP, at the moment when SLEEP takes control. That doesn't mean that all the agents of PLAY suddenly stop. Strange things, like a child destroying the block of tower they built, may be explained through agents acting on their own when they're free from control (WRECKER destroys the tower, when free from PLAY's control).

"Destructive acts can serve constructive goals by leaving fewer problems to be solved. That kick may leave a mess outside, yet tidy up the child's mind." (1)

3.6. Pain and Pleasure Simplified
"Pain simplifies your point of view. When something gives you pleasure, then, too, it's hard to think of other things" (1).

In a way, both pain and pleasure distract us from other goals. Pain and pleasure appear as opposites because they "serve related goals -- or otherwise engage the selfsame agencies (1)".

4. Notes on Chapter 4: The Self

4.1. The Self
The self is another thing we do not have a good definition for. Minsky suggests that forcing definitions of these kind of concepts may cause more harm than good. Still, "we can learn a lot by trying to understand why we believe in them (1)". For example, what psychological function does the idea of a self serve?

4.2. One Self or Many?
Minsky seems to deny we even have a self at all. All we do is a result of agents interacting in our brain. However, the self (or selves) can still be a useful concept, when talking about a person's self-images and self-ideals.

4.3. The Soul
"People ask if machines can have souls. And I ask back whether souls can learn. [...] Why try to frame the value of a Self in such a singularly frozen form? [...] The agents, raw, that make our minds are by themselves as valueless as aimless, scattered daubs of paint. What counts is what we make of them." (1)

4.4. The Conservative Self

“To understand what we call the Self, we must see what Selves are for. One function of the Self is to keep us from changing too rapidly." (1)

Notes: I still find the take Minsky has on the self unsatisfactory. It doesn’t explain why people experience the self as unity.

4.5. Exploitation

Exploitation is another form of interaction agents can have. For instance, WORK can exploit ANGER to stop SLEEP. Why use indirectness? Because directness is dangerous. If WORK could simply turn off SLEEP, we'd quickly wear our bodies out. If WORK could simply switch ANGER on, we'd be fighting all the time. For this reason, evolution has replaced direct links between agents. and fantasies can be used to replace this links. Learning how to use fantasies is a skill we need to acquire.

Note: What I find interesting in this section, is the idea that emotion agents can be used to exploit other agents. A question I have here, is where would consciousness come into play?

4.6. Self-Control

Minsky provides a list of tricks we can use to force ourselves to work when tired: willpower, activity, expression, chemistry, emotion, attachment. In a sense, self-discipline is a skill that takes years to learn.

Questions: Will self-control or self-discipline mean building a new agent? Or are control structures between existing agents changing? Or both?

4.7. Long-Range Plans

“How can any long-range plan succeed? [...] We may need some way to make changes that won't let us change ourselves back again. I suspect that, in order to commit ourselves to our largest, most ambitious plans, we learn to exploit agencies that operate on larger spans of time." (1)

According to Minsky, slow-changing agencies help shape what we call character.

4.8. Ideals

Minsky uses the term “ideals” to denote “the standards we maintain -- consciously or otherwise -- for how we ought to think about ordinary matters” (1). Two ways to acquire ideals are conflict resolution within agencies and impriming.

"In childhood, our agencies acquire various types of goals. Then we grow in overlapping waves, in which our older agencies affect the making of the new. This way, the older agencies can influence how our later ones will behave” (1).

Impriming is introduced in “The Emotion Machine” (2), where an Imprimer is one of those person to whom a child has become attached. This makes sense from a societal point of view as well: social and cultural principles learned across centuries are stored and passed down to the new generation in form of ideals.

5. Notes on Chapter 5: Individuality

5.1. Circular Causality

"There are countless different types of networks that contain loops. But all networks that contain no loops are basically the same: each has the form of a simple chain. Because of this, we can apply the very same types of reasoning to everything we can represent in terms of chains of causes and effects." (1)

However, in real life, causal relations are rarely simple, in many cases involving circular causality. We can simplify by removing recurrent loops. But by "straightening out" we are ignoring important interactions and causal links that move in the opposite direction.

5.2. Unanswerable Questions

There are questions that we can see no way to answer. What caused the universe and why? What is the purpose of life? Minsky calls these “basic questions”. Such questions are basic because they are circular. They are also the kind of questions that children ask. Every culture has come up with special ways to deal with these questions, so adults don’t have to think about them endlessly (for example, religion). However, Minsky sustains that circular thinking can also lead to growth when it results in deeper and more powerful ideas.

5.3. The Remote-Control Self

(Note: The notion of having a little person inside the head takes us back to the homunculus problem.)

“The idea of a single, central Self doesn't explain anything. This is because a thing with no parts provides nothing that we can use as pieces of explanation." (1)

Why do we readily accept this idea of a single-self then? "Because so much of what our minds do is hidden from the parts of us that are involved with verbal consciousness” (1).

Notes: Like I said in 4.4, there is still the problem of why we experience the self as unity.

5.4. Personal Identity

The view of a central Self is paradoxical. One one hand, it makes our lives simpler: in the physical world we must plan decisions as a single body, we can’t think two different thoughts at once, etc. On the other hand, it stops us from advancing the understanding of the mind.

5.5. Fashion and Style

Notes: I think in this section Minsky is trying to get to why we have common things we agree on. Like uniform styles of furnitures, making it so that we don't waste too much time on thinking how to furnish a room. Or societal rules, like driving on the right or left side of the road. I'm assuming he's using fashion and style as an example to illustrate this. I can't see any special reason for them in particular.

One thing I don't agree with is considering art as only something we commonly share. I'm not sure if the tastes for particular arts are common. But, I feel we need more consideration for art (particularly music) when it comes to understanding the mind. It doesn't seem to be the case that animals can compose music like humans do (bird songs are limited and the sounds are used mostly for communication). Storytelling can also be considering as a type of art, but Minsky touches on this later on.

5.6. Traits

What permits a writer to depict seemingly real personalities? The fact that we seem to agree upon traits. We also seem to agree on many generalities, what we call "commonsense knowledge" and "human nature". Why would this be the case? Minsky gives some reasons, like selectivity, predictability, self-reliance. Still, "a personality is merely the surface of a person. What we call traits are only the regularities we manage to perceive" (1). We never really know ourselves because many processes that affect our behavior from behind the scenes.

5.7. Permanent Identity

We experience a sense of changelessness, even though we keep changing. We are biased toward the notion of a single self, but, we might have to do away with that bias in order to understand the mind and ourselves.

Notes: Something to think about it, is why do we experience this sense of continuity.

6. Notes on Chapter 6: Insight and Introspection

6.1. Consciousness

According to Minsky, conscious thought is the signals we give to start complicated processing of which we are not aware:

"Our conscious thoughts use signal-signs to steer the engines in our minds, controlling countless processes of which we're never much aware. Not understanding how it's done, we learn to gain our ends by sending signals to those greater machines, much as the sorcerers of older times used rituals to cast their spells." (1)

Notes: Here's something to think about: Can consciousness be the sense we get from the internal workings of an agent (have the notion of the organizational structure, but don't know how individual agents work)?

6.2. Signals and Signs

Analogy seems to be key to understanding: we can understand new things, by thinking of it in terms of what we already know. Minsky also continues to expand on the use of signals.

"There are no doors inside our minds, only connections among our signs. To overstate the case a bit, what we call `consciousness' consists of little more than menu lists that flash, from time to time, on mental screen displays that other systems use. It is very much like the way the players of computer games use symbols to invoke the processes inside their complicated game machines without the slightest understanding of how they work." (1)

6.3. Thought-Experiments

Minsky seems to imply that we cannot use methods like introspection when it comes to understanding the mind: "Thinking affects our thoughts" (1).

Notes: This means we need new methods for studying the mind. I think it would make sense to use something inspired by David Marr:

- Start with the competence we are trying to understand
- Formulate it was a computational problem/theory
- Develop an implementation
- Test
- Review the theory

I'm not sure if we can completely reject introspection as a tool for studying the mind. For example, doesn't introspection play a role when thinking about different models of the self (Emotion Machine) : physical self, reflective self, self-reflective self, etc?

6.4. B-Brains

Minsky provides a scheme for the mind to watch itself:

"Divide the brain into two parts, A and B. Connect the A-brain's inputs and outputs to the real world -- so it can sense what happens there. But don't connect the B-brain to the outer world at all; instead, connect it so that the A-brain is the B-brain's world." (1)

Now, the A-brain can see and act upon the outside world, while the B-brain can monitor and influence the A-brain. In this sense, the B-brain is "reflective", and the system (A-B) is to an extent "self-aware".

Notes: There is some similarity to a multi-layer perceptrons here.

6.5. Frozen Reflection

In any situation, we make selective observations, and need a way to maintain records of those observations. How can the mind do this? Through some kind of memory that keeps these records

safe (prelude to K-lines). Similar to how scientists keep track of their experimental records (frozen phenomena), our minds keep track of our observations and reflections.

6.6. Momentary Mental Time

"It takes some time for changes in one part of a mind to affect the other parts. There's always some delay." (1)

“It is simply impossible, in general, for any agent P to know for certain what another agent Q is doing at precisely the same time. The best that P can do is send a query straight to Q and hope that Q can get a truthful message back before other agents change Q's state -- or change its message along the way." (1)

“Because of this, each agent lives in a slightly different world of time.” (1)

Questions:

- Why don’t we experience the delay; in what level of time does our consciousness operate?
- Are there particular agents that operate in the same world of time as “we” do?
- When can an agent operate in the same world of time as us?
- What is the nature of the messages/queries sent?

6.7. The Causal Now

"Our everyday ideas about the progression of mental time are wrong: they leave no room for the fact that every agent has a different causal history” (1). Different agents operate in different times, and every agent is influenced by what has happened in the common, remote history of the society. This is what Minsky considers the “now” time.

Notes: This explanation still does not answer my questions on 6.6. Even from a high-level description it’s not complete.

6.8. Thinking Without Thinking

We are not aware of how we think. Because we can not sense the workings of our minds, we have varying and conflicting theories about psychology. There are two key points from this section:

1. We need better theories (“Society of Mind” being an attempt at one)
2. We cannot say machines cannot be sentient, because we don’t much about sentience or our minds.

6.9. Heads in the Cloud

"The secret of what anything means to us depends on how we've connected it to all the other things we know. That's why it's almost always wrong to seek the `real meaning' of anything.” (1)

Question: Do individual agents have meaning?

6.10. Worlds Out of Mind

Each mind evolves its own set of meaning. The world may seem orderly because the mind has found a way to simplify itself.

“This is what we must suspect whenever some idea seems to explain too much. Perhaps no problem was actually solved at all; instead, the mind has merely found some secondary pathway in the brain, through which one can mechanically dislodge each doubt and difference from its rightful place!” (1)

Notes: I’m not sure I understand what Minsky is trying to say here. I’m assuming by “ideas that explain too much”, he refers to notions of revelation, or religious feelings. The question then becomes when does an idea explain too much? How to recognize when a problem is solved?

6.11. In-Sight

The first main point is, that we cannot understand the brain just by watching brain signals. We need a theory about how all fits together. "One cannot use data without having at least the beginnings of some theory or hypothesis” (1).

(Note: It seems to suggest neuroscience will not lead to understanding the mind. That’s why we need AI.)

The second point is that we get our ideas from communities; communities we were raised in and our own society of agents in our head.

Notes: The problem here is that most agents can't communicate with one another. The book raises this point, but doesn’t solve it.

6.12. Internal Communication

If agents cannot communicate with one another, then how can people? According to Minsky, we overestimate how much we actually communicate. The only way we communicate is through using common knowledge and experience.

Minsky continues with the point of most agents not being able to communicate with one another. He explains that two agents cannot communicate with one another if they lack common experience.

"The smaller two languages are, the harder it will be to translate between them. This is not because there are too many meanings, but because there are too few. The fewer things an agent does, the less likely that what another agent does will correspond to any of those things. And if two agents have nothing in common, no translation is conceivable." (1)

Questions: How do those agents that can communicate, do so? What languages do agents use? How do we get from agent-language to natural language?

6.13. Self-Knowledge is Dangerous
Why is what our mind does hidden from us? To protect us. If our minds weren't bound by many self-constraints, bad things would happen. We could change our self-ideals too often. Or deliberately take control of our pleasure system, and reproduce the pleasure of success without accomplishing anything at all.

6.14. Confusion
Confusion is a way to tell us that something went wrong, and we need to do something about it. For example, change our goal, change the way we are trying to reach the goal (change the agents we're using for it), etc.

7. Notes on Chapter 7: Problems and Goals

7.1. Intelligence
"Our minds contain processes that enable us to solve problems we consider difficult. `Intelligence' is our name for whichever of those processes we don't yet understand." (1)

Notes: Something to consider here, all the time spent trying to define terms like "intelligence", "consciousness", etc, and how that time could have been used to understand the mind.

7.2. Uncommon Sense
"To be considered an`expert', one needs a large amount of knowledge of only a relatively few varieties. In contrast, an ordinary person's `common sense' involves a much larger variety of different types of knowledge -- and this requires more complicated management systems." (1)

Commonsense is far more complicated that expert thinking (uncommon sense). This because commonsense requires a greater variety of representations.

Questions: Do we all form different representations? Then how can we make mental models of each other? Are the representations we make for "commonsense" similar between people?

7.3. The Puzzle Principle
Minsky introduces the "Puzzle Principle" (the random generate and test algorithm) to show it is easy for to write a program that can do what the programmer didn't imagine in advance. We cannot claim machines cannot be original or creative.

"We can program a computer to solve any problem by trial and error, without knowing how to solve it in advance, provided only that we have a way to recognize when the problem is solved." (1)

This algorithm is slow and inefficient, however, this might be fixed if we have algorithms that have better heuristics and are goal-oriented.

7.4. Problem Solving

In section 7.3, we concluded that generate and test is inefficient because of its trial and error nature. Minsky suggests an improvement: his “Progress Principle”:

"Any process of exhaustive search can be greatly reduced if we possess some ways to detect when `progress' has been made. Then we can trace a path toward a solution, just as a person can climb an unfamiliar hill in the dark -- by feeling around, at every step, to find the direction of steepest ascent." (1)

But, progress may be hard to recognize for difficult problems. One solution is to decompose the problem.

"The most powerful way we know for discovering how so solve a hard problem is to find a method that splits it into several simpler ones, each of which can be solved separately." (1)

Another solution is to embody knowledge in machines:

“The most efficient way to solve a problem is to already know how to solve it. Then one can avoid search entirely." (1)

Embodying knowledge has its own problems:

"We must discover how to acquire the knowledge we need, we must learn how to represent it, and, finally, we must develop processes that can exploit our knowledge effectively. To accomplish all that, our memories must represent, in preference to vast amounts of small details, only those relationships that may help us reach our goals." (1)

Notes: It is interesting how “Society of Mind” discusses general problem solving strategies that can be used outside the field of AI as well.

7.5. Learning and Memory

Minsky attacks behaviorism, saying it is not adequate to explain cognition. When it comes to complicated problems, we cannot only learn what we are rewarded for.

"Those twin ideas -- reward/success and punish/failure -- do not explain enough about how people learn to produce the new ideas that enable them to solve difficult problems that could not otherwise be solved without many lifetimes of ineffectual trial and error. The answer must lie in learning better ways to learn." (1)

7.6. Reinforcement and Reward

Minsky is not denying that reinforcement and reward is not important in learning. For example, we find it easy to do things if we’ve done them before.

"So let's take `reward' to mean that if agent A has been involved in arousing agent B, the effect of reward is, somehow, to make it easier for A to arouse B in the future and also, perhaps, to make it harder for A to arouse other agents." (1)

However, reinforcement is not enough when considering hard problems.

"We cannot learn to solve hard problems by indiscriminately reinforcing agents or their connections. Why is it that among all the animals, only the great-brained relatives of man can learn to solve problems that require many steps or involve using the same agencies for different purposes? We'll seek the answer in the policies our agencies use for accomplishing goals." (1)

Minsky not only says associating agents is not enough, but also proposes where to look for the answer: the policies in our agents for accomplishing goals.

7.7. Local Responsibility

There are two ways to reward behavior: based on local success (agent did his job), or, based on global success (agent doing the job helped the system). Reinforcement on the basis of local success is easier to build on a system:

"It is harder to implement a global learning scheme because this requires machinery to find out which agents are connected all the way to the original goal by unbroken chains of subgoals. " (1)

Question: How to build global reinforcement schemes based on global success on a machine/system?

7.8. Difference - Engines

Minsky explains what it means for a machine to have a goal:

"A `goal-driven' system does not seem to react directly to the stimuli or situations it encounters. Instead, it treats the things it finds as objects to exploit, avoid, or ignore, as though it were concerned with something else that doesn't yet exist. When any disturbance or obstacle diverts a goal-directed system from its course, that system seems to try to remove the interference, go around it, or turn it to some advantage." (1)

An example of this are difference-engines (Newell, Simon, and Shaw).

7.9. Intentions

Intention is only in the watcher's mind. So what does it mean for an agent to have a goal? Minsky gives two characteristics: persistence (towards the goal) and the internal representation of the goal.

Question: Are difference-engines the right scheme for representing a goal? Minsky points out that as of 1986, that was the case.

"The difference-engine scheme remains the most useful conception of goal, purpose, or intention yet discovered." (1)

7.10. Genius

In Minsky's definition of genius, geniuses are different from ordinary people in having unusually better ways to learn.

8. Notes on Chapter 8: A Theory of Memory

8.1. K-Lines: A Theory of Memory

A theory of memory-stored-in-someplace raises many questions:

"How is knowledge represented?
How is it stored?
How is it retrieved?
Then, how is it used?" (1)

Minsky's answer to all these questions: we keep each thing we learn close to the agents that learn it in the first place. How? Using Knowledge-lines or K-lines.

A K-line is a wire-like structure that attaches itself to the mental agents that are active when you solved a problem or had a good idea. When this K-line is activated later on, it re-activates these agents, putting you in the "mental state" you were when you solved that problem. You can similar problems now.

Questions:

- How to decide what agents were relevant for the goal accomplished (we don't want k-lines to be attached to all agents that were active, if they didn't contribute to the goal). Can we delete K-lines? If we find better ways to solve that problem, do we just change the agents to which the k-line is attached to, or do we form a new k-line? Are k-lines special kind of agents? What are the algorithms K-lines use?

- To keep track of which agents were recently active, we need some representation of time in the system. How do we do this?

- (For the book overall) When Minsky goes into details about connectionist-like structures (like K-lines), he doesn't seem to bring anything new; most of the approaches fall into the category of perceptrons and neural networks. (Is that the case? Maybe I'm missing something.) However, those same structures can be viewed as an attempt to describe how connectionist structures/networks can give rise to high-level processes.

8.2. Re-Membering

From the definition of k-lines, remembering is re-activating members (agents) attached to the k-line.

"Your mind is in a new state, with agents Q aroused. Something in your mind suspects that Q is similar to P -- and activates kP." (1)

How do we balance between kP and KQ (k-line for a new problem Q)? Minsky proposes some ways to resolve this: maybe give priority to new k-line agents, or maybe let the old agents win, or maybe use the principle of noncompromise.

Problem: No good policy to handle k-line conflicts.

8.3. Mental States and Dispositions

K-lines work not only for problem-solving but for describing subjective states, like dispositions, attitudes, feelings. They can even explain some memory paradoxes: "Why do we find it easier to recollect our attitudes and feelings than to describe what actually took place?" (1). The answer: K-lines can reactivate the agents that put us in that mental state, which we can experience, but because of its complexity, find it hard to describe.

Problem: "Once we think in terms of K-line memories, it becomes easy to imagine, at least in principle, how a person could recall a general impression of a complex previous experience -- but it becomes hard to understand how a person can so easily comprehend a specific statement like `John has more candy than Mary'." (1)

8.4. Partial Mental States

Minsky introduces the notion that agents are binary (on/off). From this:

"A `total state' of mind is a list that specifies which agents are active and which are quiet at a certain moment. A `partial state' of mind merely specifies that certain agents are active but does not say which other agents are quiet." (1)

According to this definition, the mind can have exactly one total state at any moment, but many partial states, because partial states are incomplete descriptions.

8.5. Level-Bands

Minsky's introduces the level-band theory.

"The basic idea is simple: we learn by attaching agents to K-lines, but we don't attach them all with equal firmness. Instead, we make strong connections at a certain level of detail, but we make weaker connections at higher and lower levels." (1)

This solves part of the question of how K-lines select which agents to activate. The weakly activated agents serve as default memories that are activated (to fill a slot), but which can

be easily changed.

"Default assumptions embody some of our most valuable kinds of commonsense knowledge: knowing what is usual or typical." (1)

Question: If agents are binary, how do level-bands work?

8.6. Levels

Level-band theory can be used to explain how adopt old memories to current problems (we can apply BUILDER for towers to BUILDER for houses).

"Most of the skills embodied in BUILDER's middle level-bands will still apply. These seem to embody the sort of knowledge that is most broadly and generally useful, whereas uppermost and lowest level-bands are more likely to be based on aspects of the problem that are specific to an older goal or to the particular details of the original problem." (1)

The advantage of level-band theory is the way we can easily change less-relevant agents.

Notes: Minsky seems to imply the notion of level-bands is important. We can use them in many ways: to describe things, to do things, etc.

8.7. Fringes

The level-band theory introduces the notion of fringes:

Lower Band/Fringe:

- Beyond a certain level of detail, increasingly complete memories of previous situations are increasingly difficult to match to new situations.
- Concerned with the structure of things.
- Objective details of reality
- Default assumptions

Upper Band/Fringe:

- Memories that arouse agents at too high a level would tend to provide us with goals that are not appropriate to the present situation.
- Concerned with the function of things
- Subjective concerns with goals and intentions

8.8. Societies of Memories

There are two ways to make new memories.

1. Attach a new K-line to all agents active recently
2. Attach new K-line to older K-lines that were active recently.

The second way requires few connection, and builds a kind of K-line memory tress.

"Specifically, you will tend to remember only what you recognized at the time. So something is lost -- but there's a gain in exchange. These `K-line memory trees' lose certain kinds of details, but they retain more traces of the origins of our ideas." (1)

8.9. Knowledge-Tress

By connecting together K-lines can form societies. Minsky applies the level-band theory to keep these societies orderly.

"When making a new K-line memory, do not connect it to all the K-lines active at the time but only to those that are active within a certain level-band. [...] Something like this will happen automatically, simply because the new K-line societies will tend to inherit whatever hierarchy already existed among the original agents that become connected to those K-lines." (1)

Minsky states that “this policy of connecting new K-lines to old ones must be used in moderation. Otherwise, no new agents would ever be included in our memories. Furthermore, it should not always be required to produce simple, orderly hierarchy-trees." (1)

Question: How do we decide when to connect K-lines to old ones, or form new ones?

8.10. Levels and Classifications

No single hierarchy can be correct because it will depend on what we want to use it for. Hierarchies do not reflect some kind of order of the world, but merely the way we order things in our minds. As we come across new things, we always find exceptions. However, when attempting a new task, we don’t like to start anew; we want to keep using what we already have.

“So we search around inside our minds for old ideas to use. Then, when part of any hierarchy seems to work, we drag the rest along with it." (1)

8.11. Layers of Societies

Let’s start with some original agents S-agents and their S-society. A K-society can be constructed from the S-society, by linking each k-line directly to S-agents.

"But this will lead to a different problem of efficiency: the connections to the original S-agents will become increasingly remote and indirect. Then everything will begin to slow down -- unless the K-society continues to make at least some new connections to the original S-society. That would be easy to arrange, if the K-society grows in the form of a `layer' close to its S-society. [...] If arranged this way, the layer pairs could form a curious sort of computer. As S-agents excite K-agents and vice versa, a sort of spiraling activity would ensue." (1)

How to stop this from becoming chaotic?

"By specifying which level-band should remain active and suppressing all the rest. Indeed, that is precisely the sort of coarse control that a B-brain might exercise, since it could do all this without needing to understand the fine details of what is happening inside the A-brain." (1)

Minsky makes the claim that this is the way mental capacities develop: creation of new layers.

9. Notes on Chapter 9: Summaries

9.1. Wanting and Liking

"The relation between wanting and liking is not simple at all, because our preferences are the end products of so many negotiations among our agencies. To accomplish any substantial goal, we must renounce the other possibilities and engage machinery to keep ourselves from succumbing to nostalgia or remorse. Then we use words like `liking' to express the operation of the mechanisms that hold us to our choice." (1)

Minsky talks about a similar point in previous chapters: the machinery underlying our thoughts is hidden to us. In this section, he also seems to say that neglecting details make it easier to stick to goals we want to accomplish. (Note: I'm not so sure about this.)

9.2. Gerrymandering

Two key points in this section:

1. High-level agencies are involved in making summaries about how we feel:

"In a complex human brain, a great many layers of agencies are interposed between the ones that deal with body needs and those that represent or recognize our intellectual accomplishments. Then what is the significance, in these more complicated systems, of those pleasant feelings of accomplishment and disagreeable sensations of defeat? They must be involved with how our higher-level agencies make summaries." (1)

2. Summaries of how things are going are important in problem-solving:

"The only way to solve hard problems is by breaking them into smaller ones and then, when those are too difficult, dividing them in turn. So hard problems always lead to branching trees of subgoals and subproblems. To decide where resources should be applied, our problem-solving agents need simple summaries of how things are going." (1)

Questions:

- Is making summaries a major goal for high-level agents?
- How do these simple summaries work? How are they communicated?

9.3. Learning from Failure

Minsky claims that it might be more important to learn from failure than from success: " whenever you try to improve an already working procedure", you risk damaging whichever other skills depend on the same machinery (1)".

How do we handle learning from failure? It might be dangerous to alter a method M after it fails, because it might fail in other ways. A safer method would be to use censor and suppressors instead (more later on about them):

"A safer way to deal with this would be to modify M by adding special memory devices called `censors' and `suppressors'...which remember particular circumstances in which M fails and later proceed to suppress M when similar conditions recur. Such censors would not tell you what to do, only what you shouldn't do; still, they prevent your wasting time by repeating old mistakes." (1)

Questions: How much does learning from failure hold when it comes to language learning? It seems that we don't get negative examples when learning a language. So how do we know not to make mistakes later on? An example of "mistake" would be in the case of forming a question for the sentence :"The dog that is in the corner is hungry". We do not end up asking: "Is the dog that in the corner is hungry?" Instead, we are able to ask: "Is the dog that is in the corner hungry?"

9.4. Enjoying Discomfort

"There is more to motivation than immediate reward. When we succeed, a lot goes on in the mind. [..] These kinds of complications make it impossible to invent good definitions for ordinary words like `pleasure' and `happiness'. No small set of terms could suffice to express the many sorts of goals and wants that, in our minds, compete in different agencies and on different scales of time." (1)

Notes: So basically, we don't know much of what is going on in our minds in this case. What changes happen when people develop obsessive disorders? I also wonder if we can train ourselves to enjoy discomfort. It would help when learning a new skill, etc.

10. Notes on Chapter 10: Papert's Principle

10.1. Piaget's Experiments

Watching children might be a way to see how mind-societies grow. Example given: Piaget's "conservation of quantity" experiments.

"In the next few sections we'll examine the idea of `more' and show that it conceals the workings of a large, complex Society-of-More -- which takes many years to learn." (1)

10.2. Reasoning about Amounts

Minsky's account of why young children fail to conserve amounts:

"The younger children possess the ideas they need but they don't know when to apply them! One might say that they lack adequate knowledge about their knowledge, or that they have not acquired the checks and balances required to select or override their hordes of agents with different perceptions and priorities." (1)

10.3. Priorities

Minsky proposes example agents: TALL, THIN, CONFINED, in a way that they always conflict. Thus the failure of the child to conserve is an attempt at resolving such conflicts.

Notes: I assume it implies that conflict resolution schemes will be learned by the child later on in life? Or do they emerge later on?

10.4. Papert's Principle

"Papert's Principle: Some of the most crucial steps in mental growth are based not simply on acquiring new skills, but on acquiring new administrative ways to use what one already knows." (1)

In our case:
"The new APPEARANCE administrator is designed to say `more' when the agent TALL is active, to say `less' when the agent THIN is active, and to say nothing at all when something appears both taller and thinner. Then the other new administrator, HISTORY, makes the decision on the basis of what CONFINED says." (1).

Question: Can hidden layers in connectionist networks serve a control structure function similar to middle-level-agents here?

10.5. The Society-Of-More

Minsky presents the notion of different meanings for "more" using different agencies. Also, an answer to why bureaucracy might be beneficial in the brain:

"You might complain that even if we needed these hordes of lower-level agencies to make comparisons, this system has too many middle-level managers. But those mountains of bureaucracy are more than worth their cost. Each higher-level agent embodies a form of `higher-order' knowledge that helps us organize ourselves by telling us when and how to use the things we know. Without a many layered management, we couldn't use the knowledge in our low-level agencies; they'd all keep getting in one another's way." (1)

Question: Are there different agencies for every different meaning a word can have?

10.6. About Piaget's Experiments

"What is the significance, then, of evidence that the young children do possess methods that could give correct answers -- and yet they do not use those abilities? As far as I can see, such evidence would only further support the need for explanations like those of Papert and Piaget." (1)

10.7. The Concept of Concept

"Instead of assuming that our children come to crystallize a single `concept of quantity', we must try to discover how our children accumulate and classify their many methods for comparing things." (1)

According to Minsky the secret is to always try to combine related agents first. A way to do this is to be close in brain locations, and sharing many other agents in common.

Problems:

- There are cases when we learn, not just combine (ex: societies-of-more).
- In what other ways can we relate agents?

10.8. Education and Development

Can we speed up the learning process (like acquisition of conservation)? No, because mature agencies require many special rules, and a child's mind cannot manage that kind of complexity.

Question:

- Can a child come across some special ways to learn that can make it possible to achieve mature-like agencies bypassing the complexity issue?
- Are there methods of conflict resolutions that can scale up? (the examples given in the book are relatively simple)

10.9. Learning a Hierarchy

How does the brain continue to function while changing? Minsky's approach is to build a detour around old systems, and not change the old system until you're sure the new route works. This detour approach results in building new intermediate layers.

Questions:

- When new agents and connections form (new route), how do the old links fade away?
- How do we know the new route works?
- How often do we change to new systems?
- Doesn't "Society of Mind" need a chapter on "Forgetting"?
- What happens if the new route added doesn't fix the problems the old system had?

11. Notes on Chapter 11: The Shape of Space

11.1. Seeing Red

Many different brain processes must be correlated to the meaning of an ordinary word. This doesn't mean that no machine can ever have the range of sensibilities people have. It merely means that we are not simple machines.

Questions:

- Will consciousness emerge from sufficient complexity?
- What does meaning mean: for agents, for a system, for us?

11.2. The Shape of Space
Sense is a complicated illusion:
"We never actually make any direct contact with the outside world. Instead, we work with models of the world that we build inside our brains." (1)

What does the sense of touch mean?
“"There is little that one could say about any `single touch' -- or about what any single sense-detecting agent does. However, there is much more to be said about the relations between two or more skin touches, because the closer together two skin spots are, the more frequently they'll both be touched at the same time." (1)

Question: Even if an illusion, how does the sense of touch arise at all (from neurons to mental states)?

11.3. Nearness
If two stimuli arise the same sensors to a great extent, the more alike the partial mental states they produce will be, and the more similar these two stimuli will seem. Nerves from skin to brain tend to run in parallel bundles, so stimulating nearby spots of skin usually will lead to similar activities inside the brain. If there is enough spatial regularity in the sensors, agents will be able to learn topographic maps (ex: what points are near each other).

Problem: This explains how a sense for a child’s own skin can arise. But how does the child learn the spatial world beyond his skin?

Question: What is the balance the “mechanism” of learning your own skin requires in terms of outside input and inside input? If nerves were damaged how would that affect the learning of skin-world boundary for a child?

11.4. Innate Geography
Minsky addresses how to fix the problem in section 10.3:
“The nerve pathways that preserve the physical nearness relations of our skin-sensors can make it easy for inner agencies to discover corresponding nearnesses about the outer world of space." (1)

According to Minsky, we learn “space” as a society of nearness relations among places. The question remains, how does the brain construct these maps.

Question: I wonder if Minsky thinks if there are innate principles for language.

11.5. Sensing Similarities
The qualities of signals sent to the brain depend only on relationships (between agents):
"For just as there is nothing to say about a single point, there's nothing to be said about an isolated sensory signal. When our REDNESS, TOUCH, or TOOTHACHE agents send their signals to our

brains, each by itself can only say `I'm here." The rest of what such signals `mean' to us depends on how they're linked to all our other agencies." (1)

Minsky also seems to suggest that what we learn early in life depends mostly on wiring.

Question: How similar is wiring between people?

11.6. The Centered Self

How do we learn about the world? It seems that our brains have evolved with some special mechanism to help us compensate for motions of the body, head, and eye. But we still don't understand how it works.

"Perhaps we start by doing many small experiments that lead to our first, crude maps of the skin. Next we might start to correlate these with the motions of our eyes and limbs; two different actions that lead to similar sensations are likely to have passed through the same locations in space. A critical step would be developing some agents that `represent' a few `places' outside the skin. Once those places are established, one could proceed to another stage: the assembly of an agency that represents a network of relationships, trajectories, and directions between those places". (1)

Question: Is the notion of 3D bootstrapped?

11.7. Predestined Learning

There is not clear-cut boundary between heredity and environment. However, if the process of learning is so tedious, why not make it (part of it) predestined?

"We acquire our conceptions of space by using agencies that learn in accord with processes determined by inheritance. These agencies proceed to learn from experience -- but the outcomes of their learning processes are virtually predestined by the spatial geometry of our body parts." (1)

Questions:

- How much is built-in, how much predestined, and how much learned for our minds?
- How does predestined learning apply to language?

11.8. Half-Brains

We are familiar with the notion of a left-brain and a right-brain. But, for Minsky the brain has many parts, not two. The brain is built of many cross-connected pairs of agencies.

"My own theory of what happens when the cross-connections between those brain halves are destroyed is that, in early life, we start with mostly similar agencies on either side. Later, as we grow more complex, a combination of genetic and circumstantial effects leads one of each pair to take control of both. Otherwise, we might become paralyzed by conflicts, because many agents would have to serve two masters." (1)

11.9. Dumbbell Theories

Although, we dichotomize a lot, there are flaws to such divisions. However, they might provide us with useful ways to think.

“Dividing things in wo is a good way to start, but one should always try to find at least a third alternative. If one cannot, one should suspect that there may not be two ideas at all, but only one, together with some form of opposite." (1)

12. Notes on Chapter 12: Learning Meaning

12.1. A Block-Arch Scenario

The process of learning includes changing mental descriptions. For instance, a child playing with blocks is trying to figure out the meaning of a Block-Arch. So far, he has: two standing blocks and a lying block. But this description applies to both an arch, and to when the bottom blocks are touching each other. The new representation of an arch requires that bottom blocks don’t touch.

What ways do we use to change mental descriptions? We find a way to attune the structure and function. In the child’s case, we had the structural description of a block-arch, and the functional description of hand--change (to make a toy car pass through the arc, we need to change the hand holding the car).

12.2. Learning Meaning

It’s hard to find a single definition for “learning”, because it spans so many kinds of processes: "The problem is that we use the single word `learning' to cover too diverse a society of ideas." (1)

In this section, Minsky introduces 4 different ways to learn:

- Uniframing: combining several descriptions into one
- Accumulating: collecting incompatible descriptions
- Reformulating: modifying a description’s character
- Trans-framing: bridging between structures and functions

By introducing these ways of learning, Minsky also introduces the need a new vocabulary in AI (with respect to psychology), because ideas in AI are new and deserve new names. Another thing Minsky considers important in learning is motivation to learn.

Questions

- What’s the difference between societies of ideas and societies of agents?

12.3. Uniframes

Uniframing: combining several descriptions into one (observing that all arches have common parts). Uniframing is particularly important, because it helps us generalize about the world. The problem with uniframing is how to decide with aspects/descriptions are important. We don’t

have an answer for that.

Question: Is an uniframe an agent?

12.4. Structure and Function
A way to do learning is to connect to preexisting structures:
"To learn new words or ideas, one must make connections to other structures in the mind." (1)

But, we can't learn just by tying things to names! The way to do these connections, is to build analogies between structural descriptions (physical properties) and functional descriptions.

12.5. The Functions of Structures
Many things we consider as physical (a chair) are actually psychological (function: a place where we sit). The association between structural and function descriptions is not simple (it includes many highly specific links).

"With that knowledge we can do amazing things, like applying the concept of a chair to see how we could sit on a box, even though it has no legs or back!" (1)

We incorporate this kind of knowledge using uniframes.

Question: How are the mappings between structure and function initiated?

12.6. Accumulation
Accumulation: simplest form of learning, where you simply remember each example or experience as a separate case. We need accumulation learning because uniframing doesn't always work.

But, then, why not simply accumulate everything? Minsky's theory is that there are architectural constraints on how many K-lines are directly accessible to various types of agents. This will constraint how much an agent can accumulate. In that case, uniframing can be used (agencies are forced to merge some examples into uniframes). Of course, we need to remember that the brain will use more than one strategy, thus, many types of learning are involved.

Question: What is the connection between uniframes and K-lines?

12.7. Accumulation Strategies
"We all use mixtures of different learning strategies -- accumulations of descriptions, K-lines, uniframes, or whatever." (1)

12.8. Problems of Disunity
Key point: "We know only a very few -- and, therefore, very precious -- schemes whose unifying powers cross many realms." (1)

12.9. The Exception Principle
“Exception Principle: It rarely pays to tamper with a rule that nearly always works. It's better just to complement it with an accumulation of specific exceptions." (1)

Questions:

- What does “rule” mean in the context of Society of Mind?
- How to balance exceptions? If we keep accumulating, we cannot come up with new concepts.

12.10. How Towers Work
“An idea will seem self-evident, once you’ve forgotten learning it!” (1)

Obvious things are not simple. "Many such things are done for us by huge, silent systems in our mind, built over long forgotten years of childhood" (1).

12.11. How Causes Work
Causes are partly a mental invention, and partly a consequence of the constraints of the world we live in.

"Causes are indeed made up by minds -- but only work in certain parts of certain worlds." (1)

"Why does a block retain its size and shape when it is moved? It is because we're fortunate enough to live within a universe in which effects are localized. ... This can happen only in a universe whose force laws work in close accord with the `nearness' of time and space -- in other words, a universe in which entities that are far apart have much less effect on each other than ones that are close together. In worlds without constraints like that, there could be no things or causes for us to know." (1)

It is because the the way the universe we live in is, that we can speak of causes.

12.12. Meaning and Definition
(Note: In this section, Minsky borrows Wittgenstein’s most famous example: what is a game?) Minsky seems to imply that the answer lies in relating physical structures to a psychological function.

Question: What’s new here? The common traits for family resemblance could have any nature (physical or psychological).

12.13. Bridge-Definitions
“Our best ideas are often those that bridge between two different worlds!" (1)

Minsky is quite strong on the idea of bridging structure and function. An example of this is learning.

13. Notes on Chapter 13: Seeing and Believing

13.1. Reformulation
Reformulation is a powerful tool to both solve problems and create uniframes. Example given: we could reformulate the notion of an arch by dividing into portions: BODY and SUPPORT.

13.2. Boundaries
Reformulation plays a strong role in creativity: finding new ways to look at things. We're always changing (imaginary) boundaries!

Question: How does the mind create these boundaries in the first place?

13.3. Seeing and Believing
"We normally assume that children see the same as we do and only lack our tricky muscle skills. But that doesn't explain why so many children produce this particular kind of drawing (heads with sticks for arms and legs), nor why they seem so satisfied with them." (1)

Minsky suggests that the child does not have a picture in mind, but only some network of relationships that various "features" must satisfy.

Question: How do these "scripts" evolve as the child develops into an adult?

13.4. Children's Drawing-Frames
Minsky provides some answer to my question, in terms of the "script" for drawing people.

"I suspect that after children learn to make recognizable figures, they usually move on to face the problems of representing much more complicated scenes. As they do this, we should continue to appreciate how well children deal with the problems they set for themselves. They may not meet our own grown-up expectations, but they often solve their own versions of the problems we pose." (1)

Question: Why does this happen?

13.5. Learning a Script
How does practice speed things up? According to Minsky, what happens is learning by bridging, which links an existing program to a set of new and simpler processes (script). The script provides the essence of the program and is more efficient.

"The people we call `experts' seem to exercise their special skills with scarcely any thought at all -- as though they were simply reading preassembled scripts. Perhaps when we `practice' to improve our skills, we're mainly building simpler scripts that don't engage so many agencies." (1)

Questions: Is this an instance of reformulation?

13.6. The Frontier Effect

"Frontier effect -- the tendency to place new features at locations that have easily described relationships to other, already represented features." (1)

We notice this effect in children, because copying (physical properties) is too hard for them. Instead, children "copy" abstract descriptions of the scene (relationship between blocks). It is easier to produce an abstract representation, than to produce a simple copy or imitation.

13.7. Duplications

Sometimes, it makes sense to count features more than once: "we seem to see two complete arches, despite the fact that there aren't enough legs to make two separate arches (1)". We can do this if we consider the functional description of archs, instead of the structural one.

14. Notes on Chapter 14: Reformulation

14.1 Using Reformulations

Reformulation is a powerful tool for solving problems.

"In the long run, the most productive kinds of thought are not the methods with which we solve particular problems, but those that lead us to formulating useful new kinds of descriptions." (1)

14.2 The Body-Support Concept

As we saw in previous sections, reformulation introduced the body-support concept in block-archs. What we see in this section, is that this concept can be extended to many domains.

Question: Is body-support a fundamental idea for Minsky?

14.3. Means and Ends

We have many ways to connect our resources to goals. For each of these ways, we have to identify essential parts. We can also use more many approaches at once.

"The quality of our understanding depends upon how well we move between those different realms. In order to translate easily from one of them to another, we must discover systematic cross-realm correspondences. However, finding these is rare. ... What is remarkable about the body-support concept is how often it leads to systematic cross-realm correspondences." (1)

14.4. Seeing Squares

"We often self-impose assumptions that make our problems more difficult, and we can escape from this only by reformulating those problems in ways that give us more room." (1)

As an example of this, Minsky offers a dot puzzle, that is hard to solve if you perceive the dots as forming a square.

“Which comes first in recognition, specific features or global shapes? It must depend upon one's state of mind." (1)

Question: Does it?

14.5. Brainstorming
Brainstorming is a way to break out of old ways of thinking, and search for new reformulations.

14.6. The Investment Principle
"The Investment Principle: Our oldest ideas have unfair advantages over those that come later. The earlier we learn a skill, the more methods we can acquire for using it. Each new idea must then compete against the large mass of skills the old ideas have accumulated." (1)

14.7. Parts and Holes
Minsky illustrates reformulation, by reformulating the arch as a container.

Question: At the end of the section, Minsky hints that the container concept is an important concept (just like the body-support concept). Why is that the case? Does it have to do with our notion of space?

14.8. The Power of Negative Thinking
To know how boxes keep things in, aside from geometry, we need to know how moving works. The power of negative thinking comes in when we think about how escaping is impossible.

"The diagram below depicts an agency that represents the several ways an arm can move inside a rectangle ... If we connect each of these sub-agents to the corresponding side of our four-sided box frame, each agent will be able to test whether the arm can move in the corresponding direction." (1)

Notes: Minsky introduces inhibitory connections between agents.

14.9. The Interaction-Square
"What's so special about moving left or right or up or down? At first one might suppose two dimensional space. But we can also use this square-like frame for many other realms of thought, to represent how pairs of causes interact." (1)

Minsky also suggests that we use interaction-square arrays when learning to control our movements:

"An interaction-square array provides a convenient way to represent all the possible combinations. If square- arrays can represent how pairs of causes interact, could similar schemes be used with three or more causes?" (1)

The answer is: no. Because it would be too complicated.

15. Notes on Chapter 15: Consciousness and Memory

15.1. Momentary Mental State

"Consciousness does not concern the present, but the past: it has to do with how we think about the records of our recent thoughts." (1)

How can thinking about thoughts be possible at all? There must be agents that learn to recognize events inside the brain. Of course, we need a notion of memory. But our capacity to keep records is limited. When we run out of room, the records of recent thoughts must displace those of our old ones. This gives us the sense of a serial stream of consciousness.

15.2. Self-Examination

Why is it so hard to talk about our present state of mind? Minsky gives several reasons for this:

- There are time-delays between different parts of the mind, so we cannot talk about a "present state".
- Our attempts to think about our mental states change our mental states.
- It's unlikely any part of the mind can obtain complete information of what happens in other parts, because our memory-control systems have too little temporary memory. They can't even represent their own activities in much detail.

Note: Minsky is criticizing introspection again, I think.

15.3. Memory

We still don't have an established theory about how memories are formed.

"Our various agencies selectively decide, unconsciously, to transfer only certain states into their long-term memories -- perhaps because they have been classified as useful, dangerous, unusual, or significant in other respects." (1)

Question: How is this selection (of mental states to transfer to long-term memories) made?

15.4. Memories of Memories

According to Minsky, memories are reconstruction of previous states of mind (so we don't retrieve memories, but mental states?). Memories are processes that make some of our agents act in much the same ways they did at various times in the past.

"I suspect that this `amnesia of infancy' is no mere effect of decay over time but an inevitable result of growing out of infancy." (1)

What this means is that the reconstruction of our childhood memories fail, because they are written in scripts we can no longer read.

Questions:

- Do we delete any of these scripts? Didn't we already change some of them into new ones? Does the brain store all the scripts we come up with? Wouldn't reconstruction also fail because our agents change as we grow up?
- Is what we remember limited by the number of agents we can have active at a time?

15.5. The Immanence Illusion

"Immanence Illusion: Whenever you can answer a question without a noticeable delay, it seems as though that answer were already active in your mind." (1)

15.6. Many Kinds of Memory

"A brain has no single. common memory system. Instead, each part of the brain has several types of memory-agencies that work in somewhat different ways, to suit particular purposes." (1)

15.7. Memory Rearrangements

Minsky starts with the example of moving mental furniture. First, we need some way to represent the room and how objects are arranged in space (symbolic representation?). Next, we need some techniques for manipulating these representations (symbol manipulation?). The manipulation is described in terms of swapping the states of two agents. For this to succeed, we need memory buffers. The buffers should be able to store the information to be swapped, and perform the swap at the correct time.

Questions:

- Can we talk of meta-information here? How do we recognize/store the information we need to retrieve?
- What is the nature of memory buffers?

15.8. Anatomy of Memory

Minsky proposes a memory-machinery for a large agency. The agency has micromemory units: temporary K-lines that can quickly store or restore the states of many agents in an agency. Short-term memory units store the memories of the micromemory units themselves. Information from either of these types of memories can also be stored in more permanent memory systems.

Problem: Memory control structures need memory to work. Aren't we falling into the homunculus problem?

15.9. Interruption and Recovery

For a system to return to its previous state after it is interrupted, it requires memory.

"Why do we so often get confused when we're interrupted? Because then we have to keep our place in several processes at once. To keep things straight, our memory-control machinery needs intricate skills." (1)

15.10. Losing Track

Why do we lose track in language, but not in vision? One reason is that our visual systems support more simultaneously operating processes than our language-systems can, and this reduces the need for any process to interrupt the other one.

15.11. The Recursion Principle

"The best way to solve a hard problem is to break it into several simpler ones, and break those into even simpler ones. Then we face the same issue of mental break it into several simpler ones, and break those into even simpler ones. Then we face the same issue of mental fragmentation. Happily, there is another way. We can work on the various parts of a problem in serial order, one after another, using the same agency over and over again. Of course, that takes more time. But is has one absolutely fundamental advantage: each agency can apply its full power to every subproblem." (1)

"The Recursion Principle: When a problem splits into smaller parts, then unless one can apply the mind's full power to each subjob, one's intellect will get dispersed and leave less cleverness for each new task." (1)

16. Notes on Chapter 16: Emotion

16.1. Emotion

Minsky views emotions as types of thoughts. In fact in "The Emotion Machine", he gives this definition: "Each of our major "emotional states"/"ways to think" results from turning certain resources on while turning certain resources off; thus changing the ways that our brains behave."

An interesting point: "The question is not whether intelligent machines can have any emotions, but whether machines can be intelligent without any emotions." (1)

16.2. Mental Growth

"How do our minds form? [...] We'll start by envisioning a simple brain composed of separate `proto-specialists', each concerned with some important requirement, goal, or instinct like food, drink, shelter, comfort, or defense. But there are reasons why those systems must be merged." (1)

Reasons for merging include: dealing with conflicts, components exploiting the knowledge of others, learning new goals, etc.

"Consequently, our genes must build some sort of `general- purpose' machinery through which individuals can acquire and transmit goals from one generation to another." (1)

16.3. Mental Proto-Specialists

"Most animals economize by having all their proto-specialists share common sets of organs for their interactions with the outer world. ... Another kind of economy comes from allowing the proto-specialists to share what they learn. Whether you seek warmth, safety, nutrition, or companionship -- eventually you'll have to be able to recognize and act in order to acquire the objects you need." (1)

Notes: Later we see these proto-specialist can serve as basis for building more complex machinery?

16.4. Cross-Exclusion

How do we select a goal from many competing goals? Minsky suggests a mechanism called cross-exclusion. In such a system, each member of a group of agents is wired to send inhibitory signals to all the other agents of that group (therefore competing with them). Whenever any agent of the group is aroused, its signals will inhibit the others. This leads to an avalanche effect: as each competitor grows weaker, its ability to inhibit competitors also weakens. Thus, even a small initial difference between competitors can quickly lock out others.

Cross-exclusion can provide the basis for the principle of noncomprise in regions of the brain where agents lie close together.

"Cross-exclusion groups can also be used to construct short-term memory units. Whenever we force one agent of such a group into activity, even for a moment, it will remain active (and the others will remain suppressed) until the situation is changed by some other strong external influence." (1)

Question: What's the external influence?

16.5. Avalanche Effects

A typical agent can arouse several others. Each new active agent will arouse others, and so on, more and more things turn on. This can lead to interferences in the network. To solve this problem, Minsky suggests more methods (aside from cross-exclusion): conservation, negative feedback, censors, supressors.

We can think of these methods as applicable only to small societies. Then what about more complex societies? They will formulate their own methods of dealing with self-regulation, as we see in the next chapters.

16.6. Motivation

How do we keep what we learn for satisfying different goals separate? One solution is to maintain a separate memory bank for every distinct goal. Wouldn't it be better if all these specialists shared a common, general-purpose memory (cross-exclusion could be used to provide the needed separation)? It will cause some problems.

"Because each separate specialist is much too small and specialized to understand how the others work, the best each can do is learn to exploit what the others can do, without understanding how they do it." (1)

Questions: How would specialists cooperate in this case? How many proto-specialists do we need to start with?

16.7. Exploitation
How could any specialist cooperate when it doesn't understand how the others work? Minsky's answer is that you don't need such understanding:

"Each part of the mind exploits the rest, not knowing how the other parts work but only what they seem to do." (1)

Questions: What stops a specialists for paralysing the system? Are the anti-avalanche methods enough to prevent this from happening?

16.8. Stimulus vs. Simulus
One agency can activate another by merely imagining a stimulus. A simulus is a reproduction of only the higher-level effects of a stimulus.

16.9. Infant Emotions
Baby's minds are made up of nearly separate agencies.

"One explanation of those striking shifts in attitude is that one agency attains control and forcibly suppresses the rest. Another view is that many processes continue at once -- but only one at a time can be expressed." (1)

What is the advantage of this design?

"Perhaps that artificial sharpening promotes the child's welfare by making it easier for the parent to respond to whichever problem has the greatest urgency." (1)

16.10. Adult Emotions
"Our earliest emotions are built-in processes in which inborn proto-specialists control what happens in our brains. Soon we learn to overrule those schemes, as our surroundings teach us what we ought to feel." (1)

Questions:

- Should we consider what is taught to us by our surroundings as “common” emotions?
- Does that mean that in the absence of a very complex environment (modern civilization), emotions like “depression” wouldn’t exist?

17. Notes on Chapter 17: Development

17.1. Sequences of Teaching-Selves

Human development proceeds through many stages, where each new stage first works under the guidance of previous stages, to acquire knowledge, values, goals. Then it proceeds to change its role and becomes a teacher for subsequent stages.

How could so many steps and stages lead to any sense of unity? Minsky speculates that the old stages still remain.

"One's present personality cannot share many of the thoughts of all one’s older personalities -- and yet it has some sense that they exist. This is one reason why we feel that we possess an inner Self -- a sort of ever-present person-friend, inside the mind, whom we can always ask for help." (1)

Questions:

- Why do new stages emerge?
- How are they tested?
- We need space to store all the old stages? How do we deal with this?

17.2. Attachment-Learning

Using a difference-engine formulation, Minsky describes how different types of learning are activated by different contexts:

- ordinary failure context: learner changes the methods used to reach the goal
- being scolded by a stranger (fear context): learner changes the description of the situation
- being scolded by a parent (attachment context): learner changes which goals are worth pursuing

Depending on the context, the effects of learning are switched to different agents.

17.3. Attachment Simplifies

What is the function of childhood attachment? Considering how complex human society and culture is, it would be impossible for a child to learn all by experience. So attachment is a way to pass down experience and cultural norms to children (see section 4.8). It is also a way to keep the child safe.

17.4. Functional Autonomy

“Functional Autonomy: In the course of pursuing any sufficiently complicated problem, the subgoals that engage our attention can become both increasingly more ambitious and increasingly detached from the original problem.” (1)

So one way to invent new goals is decomposition.

17.5. Developmental Stages
Why does development appear to be stage like?

"I'll argue that nothing so complex as a human mind can grow, except in separate steps. One reason is that it is always dangerous to change a system that already works. [...] Another conservative strategy is never to let a new stage take control of actual behavior until there is evidence that it can outperform its predecessor.” (1)

17.6. Prerequisites for Growth
The pace of mental growth is constrained by the requirement for prerequisites: some processes cannot be learned until other processes become available.

17.7. Genetic Timetables
"Once it becomes too hard to change an old agency, it is time to build another one; further progress may require revolution rather than evolution. This is another reason why a complex system must be grown in a sequence of separate steps." (1)

The timing for adding new agencies can be controlled genetically.

17.8. Attachment-Images
"Our attachment mechanisms force us to focus on our parent's ways, and this leads us to build crude images of what those parents themselves are like. That way, the values and goals of a culture pass from one generation to the next. They are not learned the way skills are learned." (1)

17.9. Different Spans of Memories
"I suspect that attachment bonds involve memory-records of a type that can be rapidly formed but them become peculiarly slow to change." (1)

17.10. Intellectual Trauma
Social failures can inform us about taboo violations, while intellectual failures inform us about our own deficiencies. Both can produce fear (we can fear our states of mind).

17.11. Intellectual Ideals
"Human though is not based on any single and uniform kind of `logic', but upon myriad processes, scripts, stereotypes, critics and censors, analogies and metaphors. ... We can make intellectual attachments, too, and want to think the way certain other persons do. ... I suspect we

depend as much on images of how we ought to think as we do on images of how we ought to feel." (1)

18. Notes on Chapter 18: Reasoning

18.1. Must Machines Be Logical?

Machines don't have to think with perfect logic. The design of the machine can be based on neat, logical principles. But we cannot expect the machine to act in a similar neat and logical fashion. According to Minsky, we use logic to simplify and summarize our thoughts after the problem is solved. However, logic does not explain much about how we think.

18.2. Chains of Reasoning

"For generations, scientists and philosophers have tried to explain ordinary reasoning in terms of logical principles -- with virtually no success. I suspect this enterprise failed because it was looking in the wrong direction: common sense works so well not because it is an approximation of logic; logic is only a small part of our great accumulation of different, useful ways to chain things together." (1)

18.3. Chaining

Chaining is important because it can be used in many different domains. The analogy here is: breaking the chain -- failure of reasoning.

18.4. Logical Chains

How are chains and logical arguments different?

"The difference is that in logic there's no middle ground; a logic link is either there or not. Because of this, a logical argument cannot have any `weakest link'." (1)

Questions and Notes:

- Why do we assume thinking is logical, when there's scientific evidence that it's not the case?
- A complaint I have here is the argument "against" logic is narrow ("weakest link"). It seems that in the original version of the book there was a section that talked about it in more details: A Framework for Representing Knowledge (3). In a nutshell, Minsky sees logic as a collection of heuristic methods, effective only when applied to simplified schematic plans; it cannot discuss at all what ought to be deduced under ordinary circumstances.

18.5. Strong Arguments

What does "strong argument" mean? In logic, we only have right or wrong. In real life, few arguments are absolutely sure. We can use different methods to make our chains of reasoning harder to break. One method is to use different arguments that prove the same point in parallel.

18.6. Magnitude from Multitude

Usually we don't have to choose between what's absolutely right or wrong, but choose the best alternative. Minsky suggests two strategies for doing this:

- Strength from magnitude: look at cooperative or competitive sum of forces.
- Strength from multitude: look at number of reasons in favor of a decision

Questions: What do we know about decision-making heuristics? And how do these strategies apply?

18.7. What Is A Number?

"In order for two minds to agree perfectly, at every level of detail, they'd have to be identical." (1)

Math may be the domain where we mostly agree about meaning, but the agreement is still not perfect.

"Even something as impersonal as `Five' never stands isolated in a person's mind but becomes part of a huge network." (1)

"The really useful `meanings' are not the flimsy logic chains of definitions, but the much harder-to-express networks of ways to remember, compare, and change things. A logic chain can break easily, but you get stuck less often when you use a cross-connected meaning-network; then, when any sense of meaning fails, you simply switch to another sense." (1)

Question: How do we build this cross-connected meaning-network in our minds?

18.8. Mathematics Made Hard

"Why do so many schoolchildren learn to fear mathematics? Perhaps in part because we try to teach the children those formal definitions, which were designed to lead to meaning-networks as sparse and thin as possible." (1)

18.9. Robustness and Recovery

How do minds continue to function, even when they change? Several reasons: duplication, self-repair, distributed processes, accumulation.

19. Notes on Chapter 19: Words and Ideas

19.1. The Roots of Intention

"We must discard the usual view that words denote or represent, or designate; instead, their function is control; each word makes various agents change what various other agents do." (1)

Question: So what is the root of intention?

19.2. The Language-Agency

Minsky's approach to language is to divide the system into three different regions:

- agents concerned with words
- agents concerned with how words engage mental processes
- agencies affected by words

Question: How does this relate to syntax and semantics?

19.3. Words and Ideas

Minsky introduces two new agents that help integrating language and thinking:

- polynemes: a type of K-line that sends the same simple signal to many different agencies, each of which must learn what to do when that signal is received
- isonomes: controls a short-term memory in each of many agencies

How do words (apple) evoke states of mind (thinking of the real thing: red color, round shape, etc)? When you hear the world “apple” a certain polymene is aroused, and the signal sent will put the COLOR agency in the state that represent redness, SHAPE agency in the state or roundness, etc.

Question: Does semantics boil down to experiences associated with the word?

19.4. Objects And Properties

"We derive a wonderful power from representing things in terms of properties that do not interact: this makes imagination practical. It lets us anticipate what will happen when we invent new combinations and variations we've never seen before." (1)

Because different agencies can represent different properties. That way, a single word can activate many different kinds of thoughts at once: word “apple” can set COLOR agency into “redness” state, and so on.

19.5. Polynemes

- permanent K-lines. They are long-term memories.
- invoke partial states within multiple agencies, where each agency is concerned with represented different aspects of a thing.
- Ex: ‘apple-polyneme’, invokes properties of color, shape, taste, etc.

"To understand a polyneme, each agency must learn its own specific and appropriate response. Each agency must have its private dictionary or memory bank to tell it how to respond to every polyneme." (1)

Questions: What is a signal in the context of the polynemes? Do different polynemes send different signals? If not, how do agencies differentiate the signals received?

19.6. Recognizers

How do we recognize things?

- We could verify that it has certain properties. This can be done with AND gates. Problem: no recognition-scheme in real life works on absolutely perfect evidence.
- Require enough properties to be detected. Problem: it will miss, if many features are out of sight, or present but in the wrong arrangement, etc.

Problem: The schemes presented aren't able to recognize relationships among features/properties.

19.7. Weighing Evidence

"All feature-weighing machines have serious limitations because, although they can measure the presence or absence of various features, they cannot take into account enough of the relations among those features." (1)

Questions:

- Isn't the purpose of hidden units to find these relations?
- I think the problem I have with connectionist networks is: Who is reading the weights?

19.8. Generalizing

So far we have seen several ways to generalize: uniframes, level bands, and polynemes (trying to guess the character of a thing by combining expectations based upon independent properties).

Problem: We still haven't solved the problem of generalization.

19.9. Recognizing Thoughts

How do we recognize our own ideas? This section answers this, in a limited context: when an idea has some kind of physical version (like the idea of an "apple"). In this case, input from memory can engage similar representations and processes as the sensory input.

Question: How do we deal with ideas that don't have a corresponding physical equivalence (like freedom)?

19.10. Closing the Ring

Minsky includes recurrent connections to the language agency. Why?

"If you start with enough clues to arouse one of your apple-nemes, it will automatically arouse memories of the other properties and qualities of apples and create a more complete impression, `simulus', or hallucination of the experience of seeing, feeling, and even of eating an apple. This way, a simple loop machine can reconstruct a larger whole from clues about only certain of its parts!" (1)

20. Notes on Chapter 20: Context and Ambiguity

20.1. Ambiguity

Not only words, but thoughts themselves, are ambiguous. This because we have limited access to the agents of our thoughts, and while we try to express a state of mind, we're already in the next state.

How come we have a sense of what sentences mean? We use context. We also are adapted to coping with ambiguity.

20.2. Negotiating Ambiguity

How does context clarify ambiguities (John shot two bucks)? It can activate specific polynemes, that in turn activate specific agents ("outdoors" polyneme produces bias towards "deer" than "dollar", and so on). In fact, many polynemes will be activated, mutually supporting a consistent interpretation.

Won't this cause an avalanche that arouses all the agents of the mind? No, because cross-exclusion can prevent this.

20.3. Visual Ambiguity

We also have ambiguities in vision, but, they are usually resolved so quickly, that we are not aware of there being a conflict at all. To be fair, sometimes we have the sense of perceiving the same structure in several ways at once, but we usually lock in one one particular interpretation. Minsky seems to favor the resolution of ambiguities by high-order processes.

20.4. Locking-In and Weeding-Out

We can view context as a way to weed out ambiguity. Ambiguity activates many polynemes. Context also activates agents, that will support only some of these polynemes (and inhibit the rest).

What happens when the interpretation locked in is wrong? How to correct the mistake? Record the meaning-senses adopted in the previous cycle, suppress them temporarily, and start the new cycle. A high-level agent could do this.

20.5. Micronemes

- inner mental context clues that shade our mind's activities in ways we can rarely express (most flavors, aromas, gestures, etc)
- implemented K-lines
- agents that produce only small or subtle signals

"There is a somewhat different microstructure to each person's thoughts; indeed, their inexpressibility reflects our individuality." (1)

20.6. The Nemeic Spiral

"Our polynemes and micronemes grow into great branching networks that reach every level of every agency. They approximate the general form of a hierarchy, but one that is riddled with shortcuts, cross-connections, and exceptions." (1)

Something similar to B-brain can be used to control levels of activities in agents (lower-level agents?).

20.7. Connections

"To learn the proper use of a single word must involve great numbers of connections between the agents for that word and other agents." (1)

Questions: What causes these connections? How are they physically instantiated?

20.8. Connection Lines

Minsky describes a connection-scheme that permits many agents to communicate with one another, using few connection wires (Mooer's scheme, 1946).

"The trick is to make each transmitting agent excite not one, but five of those wires, chosen at random from the available ten. Then each receiving-agent is provided with an AND-agent connected to recognize the same five-wire combination." (1)

Minsky's intuition of what these networks do, is that they are similar to Perceptron-like learning machines.

20.9. Distributed Memory

Interconnected agents, will have multiple inputs and multiple outputs. Represented in this way, we can see them as simple evidence-weighing agents, with different threshold values. Agents can learn to recognize signals by altering the connection weights.

Notes: Starting with K-lines, many elements of SoM are connectionist in nature (connection lines, distributed memory, inhibition networks, etc). The question remains how to go from connectionist networks to high-level symbolic systems/processes. We don't have a good theory for that.

21. Notes on Chapter 21: Trans-Frames

21.1. The Pronouns of the Mind

"Pronouns do not signify objects or words; instead, they represent conceptions, ideas, or activities that the speaker assumes are going on inside the listener's mind. But how can the listener tell which one of the activities is signified when there are several possibilities?" (1)

To explain this, Minsky introduces the term "pronome" -- a temporary handle for taking hold of, or moving around active fragments of mental states.

21.2. Pronomes

- temporary K-lines. They are short-term memories.
- associated with a specific role (more in 21.6)

21.3. Trans-Frames

- a compact way of representing conceptual dependencies

21.4. Communication Among Agents

Transframes provide context for control structure. Agents communicate without explicit messages by using pronomes. For example, if we want to get an apple, we create a pronome for apple in a transframe. In fact, “nomes” control how representations are manipulated!

21.5. Automatism

Simple agents need simple messages to control them. The intervention from high-level agents is minimal. For example, once you notice an apple (agents that are in touch with the real-world notice an apple), the polyneme for apple will arouse certain agents automatically (COLOR, SHAPE, etc). In a sense this is all automatic.

21.6. Trans-Frame Pronomes

In this section, we get a better elaboration of a trans-frame. A transframe can include many pronomes (Actor, Cause, Destination, etc).

"I suspect that Trans-like structures have a special prominence in how we think. One reason is that some sort of bridge like scheme seems indispensable for making those all-important connections between structures and functions." (1)

Problem: I still don’t have a good idea of what a transframe is. In fact, many of the SoM components are very little specified.

21.7. Generalizing With Pronomes

What’s the advantage of using trans-frames? If we defined objects/actions using pronomes in trans-frames, one control structure can be applied to many situations.

21.8. Attention

Attention is limited by short-term memory. At some point, we’ll also need to develop a machinery for keeping track of long-term memories. We can start by being able to keep track of simple polynemes. Then the machinery can evolve to keeping track of an entire trans-frame with various pronomes.

Question: What’s the difference between a pronome and a temporary polyneme?

22. Notes on Chapter 22: Expression

22.1 Pronomes and Polynemes

Pronomes: temporary K-lines. They are short-term memories
Polynemes: permanent K-lines. Long-term memories.

Question: What's the physical difference?

22.2. Isonomes

- similar built-in effect on each of its recipients. It applies the same idea to many different things at once.
- can account for common genetic origins for agents
- pronomes are a type of isonomes

Compare to polyneme: different, learned effects on each of its recipients.

Question: What are the physical and structural differences between nemes and nomes?

22.3. De-Specializing

One way to despecialize is to replace polynemes with less specific isonomes (replace polyneme "put apple into pail" with isonomes, so now we can put "onions into pails", or "umbrellas into suitcases")

The issue here is how to contain these replacement so it doesn't lead to absurdities.

"What we call 'generalizing' is not any single process or concept, but a functional term for the huge societies of different methods we use to extend the powers of our skills. No single policy will work for all domains of thought, and each refinement of technique will affect the quality of the generalizations we make." (1)

22.4. Learning and Teaching

"The power of what we learn depends on how we represent it in our minds." (1)

How do we acquire these representational skills?

"Each child learns, from time to time, various better ways to learn -- but no one understands how this is done." (1)

22.5. Inference

A simple theory for doing inference is to fit together transframes into chains. The problem here is that we have to recognize that transframes do not give us perfect matches. So, we need to learn to manipulate our isonomes and polynemes.

"By learning to manipulate our isonomes, we become able to combine mental representations into structures that resemble bridges, chains and towers. Our language-agencies learn to express these in the form of compound sentences." (1)

Question: How do we train ourselves to do this?

22.6. Expression

"Why do we 'thingify' our thoughts? One reason is that this enables us to reapply the wonderful machines our brains contain for understanding worldly things." (1)

"I suspect that, as they're represented in the mind, there's little difference between a physical object and an idea." (1)

Question: Could we create complex and abstract ideas using physical ideas as a basis? How do we evolve “the abstraction machinery”, in that case? How does it work (how do we extend the meaning of ideas)?

22.7. Causes and Clauses

"I wouldn't be surprised to find that brains have built-in tendencies to try to represent all situations in certain special ways: THINGS, DIFFERENCES, CAUSES, CLAUSES." (1)

22.8. Interruptions

Interruption engages agents that have control over short-term memories. Sentence clauses are a type of interruption (relative clauses interrupt main clauses to provide new information). Of course, to represent the sentence we need to use a transframe. Minsky also proposes that in English we use certain wh-words to interrupt a listener’s language agency and cause its short-term memories to temporarily store away some of their pronome assignments.

22.9. Pronouns and References

Some ways to deal with ambiguous pronouns is to use grammar, context, and expectations.

"What does 'expectation' mean? At each point in a dialogue, both parties are already involved with various concerns and desires. These establish contexts in which each new word, description, or representation, however ambiguous, gets merged with whichever short-term memory best matches it.” (1)

22.10. Verbal Expression

How do we know what to say to affect the other person’s agencies?

- Build a version of the structure in your own mind.
- Attempt to construct a similar structure in the other person’s mind, using verbal expressions.

"To be able to do that, Mary must have learned at least one expressive technique that corresponds to each frequently used mental operation. And Jack must have learned to recognize those expressive techniques -- we'll call them grammar-tactics -- and to use them to activate some corresponding isonomes and polynemes." (1)

22.11. Creative Expression

When communicating, we cannot hope to reconstruct the structure of sentences/thoughts perfectly in the listener's mind. But Minsky states that's not a bad thing, because it can lead to new insights.

"When we try to explain what we think we know, we're likely to end up with something new. All teachers know how often we understand something for the first time only after trying to explain it to someone else." (1)

23. Notes on Chapter 23: Comparisons

23.1. A World of Differences

Much of ordinary thought is based on recognizing differences. Familiar mental activities can be represented in terms of differences between situations: PREDICTING, EXPECTING, EXPLAINING, WANTING, ESCAPING, ATTACKING, DEFENDING, ABSTRACTING. We can also think about differences between differences (the height of an object is the difference of the locations for top and bottom). This is the basis for reasoning by analogy.

"The ability to consider differences between differences is important because it lies at the heart of our abilities to solve new problems. This is because these 'second- order-differences' are what we use to remind ourselves of other problems we already know how to solve." (1)

23.2. Differences and Duplicates

"The Duplication Problem: The states of two different agencies cannot be compared unless those agencies themselves are virtually identical." (1)

In turn, these agencies must receive inputs of near identical character. For that to happen, so must their subagencies. This will lead to the problem of requiring duplicate brains.

Questions: Is this a problem? Wouldn't it be good to have similar structures in the brain (for example for robustness)?

23.3. Time Blinking

Minsky provides a way to get around the duplication problem. We can compare two descriptions by presenting them to the same agency at different times. These agency can be an agency sensitive to changes in time, and we can also use it to detect differences.

23.4. The Meanings of More

Words with many meanings (“more”) also engage isonomes, because all the different meanings share a certain common character (disposition to make comparisons).

23.5. Foreign Accents

"Why do adults find it so hard to learn how to pronounce new word sounds? ... I suspect this particular disability is caused, more or less directly, by a genetically programmed mechanism that disables our ability to learn to make new connections in or between the agents we use to represent speech sounds." (1)

Question: What’s the relation between biology/genetics and agents?

24. Notes on Chapter 24: Frames

24.1. The Speed of Thought

What could explain the blinding speed of sight? The secret is that sight is intertwined with memory.

"The moment you sense the presence of a person, a whole world of assumptions are aroused that are usually true about people in general. At the same time, certain superficial cues remind you of particular people you've already met. Unconsciously, then, you will assume that this stranger must also resemble them, not only in appearance but in other traits as well." (1)

These structures acquired from previous experiences that get activated by perceptual experience are frames.

24.2. Frames Of Mind

A frame is a sort of skeleton, with many blanks or slots to be filled (terminals). Terminals are used as connection points to which we can attach other kinds of informations.

"In principle, we could use frames without attaching their terminals to anything. Normally, though, the terminals come with other agents already attached -- and these are what we called 'default assignments' when we first talked about level-bands." (1)

Questions: What methods could we use to instantiate frames (PDP, etc)?

24.3. How Trans-Frames Work

How might a frame actually work? A frame could consist of little more than a collection of AND agents, one for each of the frame's pronome terminals."

24.4. Default Assumptions

"But why use default assumptions at all, instead of simply seeing what's really there? Because unless we make assumptions, the world would simply make no sense. It would be as useless to

perceive how things 'actually look' as it would be to watch the random dots on untuned television screens." (1)

24.5. Nonverbal Reasoning

This kind of reasoning involves replacing particular things with typical things. We do this by manipulating our memories.

"Children must develop complex skills, not merely to replace one representation with another, but to compare the two representations and then move around inside them, making different changes at different levels. These intricate skills involve the use of isonomes that control the level-band of the activities inside our agencies." (1)

24.6. Direction-Nemes

We don't know how shapes and places are represented in the brain. Minsky's hypothesis:

" [...] many agencies inside our brains use frames whose terminals are controlled by interaction-square arrays. Only now we'll use those square arrays not to represent the interactions of different causes, but to describe the relations between closely related locations." (1)

We represent directions and places by attaching them to a special set of pronome-like agents, called "direction-nemes".

Questions: Many processes are computed automatically in the brain. How do we deal with the computational load?

24.7. Picture-Frames

Once we accumulated enough room-frames to represent most rooms we're likely to see, we could create a generic frame that fits almost any room, with terminals corresponding to ceiling, floors, and walls. Each terminal will have subframes that include direction-nemes to identify different regions. We also can have different rooms, sharing the same terminal. This to prevent mistakes (you can assume you walked into the living room, but you recognize the kitchen table), and switch easily from one frame to another.

24.8. How Picture-Frames Work

To build picture-frames we use the same approach as with trans-frames but we replace the pronomes by a set of nine direction-nemes. We can also attach an agent to the picture-frame responsible for turning on the frame itself.

Now that we build the frame, how does it work?

"Imagine that you're looking at some real-world scene. Your eyes move in various directions, controlled in some way by direction-nemes. Now suppose that every time you move your eyes, the same direction-nemes also activate the K-lines attached to the corresponding terminals of a

certain vision-frame. Suppose, also, that those K-lines are ready to form new memories. Then each time you look in a different direction, your vision system will describe what you see -- and the corresponding K-lines will record what you see when you look in that direction!" (1)

24.9. Recognizers and Memorizers

How do frames become activated?

"We'll simply assume that every frame is activated by some set of recognizers. We can regard a recognizer as a type of agent that, in a sense, is the opposite of a K-line -- since instead of arousing a certain state of mind, it has to recognize when a certain state of mind occurs." (1)

Very little is said about memorizes in the section. From the picture, we see that recognizers depend on information from memorizers to do their work.

Problem: The mechanism for frame recognition is not fleshed out.

25. Notes on Chapter 25: Frame- Arrays

25.1. One Frame at a Time?

Visually ambiguous figures include more than one interpretations, but we can only see one interpretation at a time. This is because our agencies can only tolerate one interpretation at a time. Why is that the case?

"Our vision-systems are born equipped, on each of several different levels, with some sort of 'locking-in' machinery that at every moment permits each 'part, at each level, to be assigned to one and only one 'whole' at the next level."

25.2. Frame-Arrays

“Frame -Arrays: When we move, our vision-systems switch among a family of different frames that all use the same terminals.” (1)

An example of this, is seeing a cube from different sides.

25.3. The Stationary World

Why does the world appear stable?

“This is because our higher-level agents don't 'see' the outputs of the sensors in our eyes at all. Instead, they 'watch' the states of middle-level agencies that don't change state so frequently.” (1)

What keeps the “inner models” of the world from changing all the time? Being implemented by frame arrays.

How, then, do we automatically compensate for changes in view? The system uses direction-nemes both the control motion and to select frames from our frame-arrays.

25.4. The Sense of Continuity

Frame-arrays let us 'visualize' imaginary scenes, such as what might happen when we move, because the frames for what we can expect to see are filled in automatically. This gives us a sense of continuity.

25.5. Expectations

According to Minsky, expectations are the content of the filled-in nodes of the frame-arrays.

25.6. The Frame Idea

One of the criticism the frame idea gets is being too vague. Minsky addresses this:

"In retrospect, it seems those explanations were at just the right level-bands of detail to meet the needs of that time, which is why the essay had the effect it did. If the theory had been any vaguer, it would have been ignored, but if it had been described in more detail, other scientists might have 'tested' it, instead of contributing their own ideas."

Questions:

- Why wouldn't you want a theory to be tested? I don't buy the "contributing their own ideas" answer.
- I wonder if this book receives similar criticism.

26. Notes on Chapter 26: Language- Frames

26.1. Understanding Words

Consider all the assumptions and conclusions we make when reading a story. These understandings are part of "common sense".

They are made so swiftly that they are often ready in our minds before a sentence is complete! But how is this done?" (1)

Minsky sees this as a consequence of activating frames. The default assignments for the frames come from previous experience.

26.2. Understanding Stories

Frames are used for story understanding as well. Terminals are filled in with information from the story. The story becomes comprehensible, because each phrase and sentence stirs in frames into activity, or helps already active ones to fill in their terminals.

26.3 Sentence-Frames

We also have sentence-frames, where the words in the sentence fills various roles in the frame.

26.4. A Party-Frame
We also form frames for events or customs. For example, a birthday-party-frame can include ARRIVAL, GIFT, GAMES, DECOR, etc.

26.5. Story-Frames
The story-teller activates particular frames for a story, as well as expectations for each frame (for example phrases like “once upon a time”).

26.6. Sentence and Nonsense
"A word-string seems 'grammatical' if all its words have hit quickly and easily into frames that connect suitably to one another." (1)

Questions: How many different sentence-frames do we have? Does it depend on the language?

26.7. Frames for Nouns
When we hear a sentence, we expect the words to arrive in a more or less definite order. This can suggest that we use frames to describe nouns, as well as verbs.

"Many scientists have asked, indeed, why so many human languages used similar structures such as nouns, adjectives, verbs, clauses, and sentences. It is likely that some of these reflect what is genetically built in to our language-agencies. But it seems to me even more likely that most of these nearly universal language-forms scarcely depend on language at all but reflect how descriptions are formed in other agencies." (1)

26.8. Frames for Verbs
In the earliest stages of learning to speak, we simply fill the terminals of word-string frames with nemes for words. Later, we learn to fill those terminals with other filled-in language-frames.

Questions: How do we know which terminals to fill with which words? How do we learn these policies?

26.9. Language and Vision
"In vision, too, there must be similar processes (similar to language processes) involved in breaking seems apart and representing them as composed of objects and relationships." (1)

26.10. Learning Language
“We cannot learn meetings only by memorizing definitions: we must also 'understand' them." (1)

"Some language theorists have suggested that children learn to use grammar so readily that our brains must be born with built-in grammar-machinery. However, we've seen that our visual systems solve many similar problems in even earlier years." (1)

The question here is why language learning takes so long?

26.11. Grammar
It's hard to even speculate about early stages of language learning because we know so little.

26.12. Coherent Discourse
"Every discourse works on several scales. Each word you hear can change your state in a way that depends upon all the structures you have built while listening to the words that came before." (1)

"To understand what people say, we also exploit our vast stores of common knowledge, not only about how specific words are related to the subjects of concern, but also a how to express and discuss those subjects." (1)

27. Notes on Chapter 27: Censors and Jokes

27.1. Demons
Charniak, a student of Minsky, suggested that whenever we hear about a particular event, specific recognition-agents are thereby aroused. These then proceed actively to watch and wait for other related types of events. Because this recognition-agents lurk silently, they are sometimes called "demons". Charniak's approach raises many questions. We don't have good answers to those questions, and Minsky concludes that understanding must be a huge accumulation of skills.

27.2. Suppressors
Suppressor-agents: Wait until you get a certain "bad idea." Then they prevent your taking the corresponding action, and make you wait until you think of some alternative. lf a suppressor could speak, it would say, "Stop thinking that!"

Censor-agents: Need not wait until a certain bad idea occurs; instead, they intercept the states of mind that usually precede that thought. lf a censor could speak, it would say, "Don't even begin to think that!"

27.3. Censors
Censors are efficient because they intercede before a useless thought occurs. But, in order to work, a censor will require a substantial amount of memory.

27.4. Exceptions to Logic
"We search for 'islands of consistency' within which ordinary reasoning seems safe. We work also to find and mark the unsafe boundaries of those domains." (1)

27.5. Jokes
According to Freud, most jokes are stories designed to fool the censors. Freud's theory couldn't account for nonsense jokes. Minsky suggests that the purpose of the nonsense jokes is the same. Absurd results of reasoning must be tabooed just like social mistakes.

27.6. Humor and Censorship
Humor matters in learning, because it is involved with how our censors learn.

Questions: Is humor crucial to intelligence? Could we use the ability to make jokes or appreciate jokes as tests for AI?

27.7. Laughter
The function of laughing is to disrupt another person's reasoning. It focuses attention on the present state of mind. By preventing you from “taking seriously” your present thought (and proceeding to develop it), laughter gives you time to build a censor against that state of mind.

27.8. Good Humor
"To ask how humor works in a grown-up person is to ask how everything works in a grown-up person, since humor gets involved with so many other things." (1)

28. Notes on Chapter 28: The Mind and the World

28.1. The Myth of Mental Energy
“Machines and brains require ordinary energy to do their jobs -- and need no other, mental forms of energy. Causality is quite enough to keep them working toward their goals.” (1)

28.2. Magnitude and Marketplace
When comparing alternatives:
"We turn to using quantities when we can't compare the qualities of things." (1)

Agents could also use some sort of currency or magnitude to keep account of their transactions. But what could this currency be? Minsky suspects that what we call “pleasure of success” may be the currency of such scheme. When talking about agents this can mean having access to a limited chemical, etc.

28.3. Quantity and Quality
Quantitative descriptions conceal the structures that give rise to them. Quantitative descriptions ignore qualitative differences.

“We should never assume that the quality or character of a thought process depends directly on the nature of the circumstances that evoked. There is no quality of 'sweetness' inherent in sugar itself." (1)

28.4. Mind over Matter
Feelings are not inherent, they are engineered. Feelings like pain or hunger, are engineered to serve like warning sign to indicate dangerous conditions, and warn us before too much damage is done.

Note: An interesting point here is that we could overcome some feelings (pain, depression, fatigue, etc) by finding ways to rearrange our priorities.

Problem: I don't find the answer about feelings of depression satisfactory: "They are designed to prevent unproductive use of time". I think feelings like depression might have to do with the fact of having complex mental states and a complex environment/culture. How complex/intelligent does an entity (creature or machine) have to be to want to kill itself?

28.5. The Mind and the World

"Minds are simply what brains do." (1)

Whenever we speak about a mind, we're speaking of the processes that carry the brain from state to state.

28.6. Minds and Machines

The problem with minds is that the mind is a set of self-modifying processes.

"The principal activities of brains are making changes in themselves." (1)

The mind is a complex machine. Our difficulties with understanding the mind come from the fact that we have so little experience with machines of such complexity.

28.7. Individual Identities

"Modifying or replacing the physical parts of a brain will not affect the mind it embodies, unless this alters the succession of states in that brain." (1)

Will the new machine be the same as you? The problem here is the definition of "same". I don't think we have a clear answer to this question.

28.8. Overlapping Minds

There are many ways to draw in imaginary boundaries through brains. We say a person has a left-brain and right-brain, or as having a frontal-brain and back-brain. Do these subregions have a mind? It will depend on what we mean by "mind". The less an entity resembles us, the less it means for us to say that it has a mind. However, many agents have human-like capacities (like problem-solving). In a sense, some agencies are more "conscious" than us of what goes on, because they keep records of our own internal activities. Minsky concludes that it makes sense to think that inside our brains there is a society of different minds.

Questions:

- Can "minds" in the context of society of minds be considered high-level functions?

29. Notes on Chapter 29: The Realms of Thought

29.1. The Realms of Thought
We view the world as being divided into many different realms: Physical, Personal, Mental, etc.

29.2. Several Thoughts At Once
We can think in several mental realms at once**.** Take the sentence “Mary gives Jack the kite”. One way is to think in the realm of space (physical objects and motions). Another is to think in the realm of ownership. And so on.

How can we think in several realms at the same time, without interference from one another. Minsky’s answer is that the agents in separate realms do not compete with each other.

29.3. Paranomes
- pronomes that operate in several different realms at once
- forge links between parallel frames in different realms

29.4. Cross-Realm Correspondences
What enables us to turn skills to different other purposes? We can do so because of the systematic “cross-realm correspondences” tendencies embodied in our families of polynemes and paranomes. Every mental realm accumulates its own abilities but also discovers, from time to time, how to exploit the skills of other realms.

Questions:
- What methods can use to study this? Like simulation, etc.
- How do mental realms and frame-arrays work for bilingual or multilingual people?

29.5. The Problem of Unity
Why do we form many separate mental realms, instead of a unified view of the world? Because it wouldn’t be practical to do otherwise. Different realms require different principles, thus, different systems for representing these principles.

29.6. Autistic Children
Children start to distinguish between psychological and physical relationships by discovering different principles for each realm. Autistic children fail to do so.

29.7. Likenesses and Analogies
"We always try to use old memories to recollect how we solve problems in the past. But nothing is ever twice the same, so recollections rarely match. Then we must force our memories to fit -- so we can see those different things as similar. To do this, we can either modify a memory or change how we represent the present scene.” (1)

29.8. Metaphors
What is a metaphor? It is our ways to transport thoughts between various mental realms.

Question: What are these ways?

30. Notes on Chapter 30: Mental Models

30.1. Knowing
What does "knowing" really mean? In a sense, knowing about the world means assuming what things are typical (default assumptions).

30.2. Knowing and Believing
Why do we think that our beliefs are certain? Because, at times particular agencies dominate, even if for a short time.

30.3. Mental Models
According to Minsky, knowing something means we have a model of it inside our head. A model is good to the extent that it can be used to answer questions.

30.4. World Models
To make a model of the world, we have to add an additional component to represent the world itself. As a result, when we learn about the world, we are learning about our models of our models of the world.

30.5. Knowing Ourselves
We cannot examine directly our self; we can only examine our model of our self. Can we go on to make models of our models of our models of self? If we keep doing this, we'll be trapped in an infinite regress.

Question: How far can we take the "model-making" till we lose track?

30.6. Freedom of Will
There's no such thing as freedom of will for Minsky.

"Every action we perform stems from a host of processes inside our minds. We sometimes understand a few of them, but most life far beyond our ken." (1)

30.7. The Myth of the Third Alternative
Even though the physical world provides no room for freedom of will, the concept is essential to our models of the mental realm.

Notes: How is free will essential? We don't need free will to describe the brain. I see it more freedom of will more as a concept necessary from a religious or cultural perspective. But, how much should we bother with it?

30.8. Intelligence and Resourcefulness

Intelligence stems from our vast diversity.

Question: We have a diversity of choices, but how do we intelligently choose from them?

IV. Conclusion

Let's return to the questions of section 1.1, and see how the society-of-mind theory answers them.

Function: *How do agents work?*
From the theory it is not clear how agents work, or what exactly agents are. Even by going through the types of agents described in the book, I still cannot get an answer for this question.

Embodiment: *What are they made of?*
The theory doesn't answer this.

Interaction: *How do they communicate?*
The theory gives some examples for this, like K-lines, indirect links, frames, etc. However, we don't have a full answer, because there can be many ways of communications that are not covered in the theory and that we don't know about.

Origins: *Where do the first agents come from?*
No clear answer. The reason is that the main focus of the book is how agents interact and create more complex structures. But so little is given about the nature of the agents themselves. Also the term agent is quite vague: an agent can be composed of other agents. So with "agent" we can mean a single agent, or a collection of agents. To answer where the first agents come from, we'd have to know if there are basic agents, and if we're born with all agents. What we know from the theory is that we start with proto-specialists (mostly for instincts), and that there we have a genetic predisposition to build agents of some sort.

Heredity: *Are we all born with the same agents?*
Not clear. From what I got from the theory, I can speculate that we must be born with similar agents, because of similar genetics. However, here are two questions that the theory doesn't answer: Are we born with a set of universal agents (they don't have to be the same, just similar)? Are we different only the in way we arrange them?

Learning: *How do we make new agents and change old ones?*
There are ways to build new agents from existing agents: K-lines (if we count them as agents), A-brain/B-brain approach, layers of societies, etc. What I'm not clear about is if we build new agents from scratch (as in not from using other agents as a basis).

I'm unclear about changing agents. Minsky's theory is that we keep building on top of old agents or layers. At some point the new layer takes control, and the old one becomes remote. That could be seen as a way to change old layers: by building new layers on top. However, I'm not sure how we change old layers in the literal sense. It's not clear if we delete old agents or layers, or how agents change their connections to other agents.

Character: *What are the most important kinds of agents?*
Not clear. I cannot conclude that the agents mentioned in the book are the most important ones, or not.

Authority: *What happens when agents disagree?*
Some ways to handle this: higher-level agents interfere (B-brain can affect agents of the A-brain), non-compromise principle, economic priority (some alternatives have bigger "magnitudes" assigned than others), etc.

Intention: *How could such networks want or wish?*
We can't talk of wanting and wishing in "human" terms, but in terms of "goals" (reaching a goal). An example the theory provides is difference-engines.

Competence: *How can groups of agents do what separate agents cannot do?*
Many complicated things, from building a block tower to a mind. In fact, intelligence is a consequence of the way agents interact with each other (forming societies, etc).

Selfness: *What gives them unity or personality?*
The reason we have agree on personalities is because they are part of our commonsense knowledge.

The theory argues that we have different selves for different purposes. However, it doesn't explain why we have this notion of unity.

Meaning: *How could they understand anything?*
We don't have a complete answer for understanding yet. A way of understanding for agents is through recognizing differences, and for higher-level agencies -- analogy.

Sensibility: *How could they have feelings and emotions?*
Being mindless, they don't. But our feelings and emotions come as a consequence of their interactions.

Awareness: *How could they be conscious or self-aware?*
Are we? The theory seems to imply that what we see as consciousness or self-awareness is just an illusion. Also, see section 28.8 of this paper.

In my view, in “The Society of Mind”, Minsky attempts to propose a theory about the mind (what it is and how it works). The theory is not complete, and not well-specified (in fact it might be under-specified). I see the book as introducing research topics worthy of pursuit, and as a motivation for taking on the mind problem.

V. References

(1) The Society of Mind. Marvin Minsky. 1986
(2) The Emotion Machine. Marvin Minsky. 2006
(3) Society of Mind: a response to four reviews. Marvin Minsky. 1990