Introduction to Complexity and Applied Complexity, Spring 2021

Module 11 — Perception, Action and Mind

Notes by Sav Sidorov

Readings

- Gerd Gigerenzer — **Gut Feelings, Chapter 1**
- Maturana et al. — **What the Frog's Eye Tells the Frog's Brain**
- James J. Gibson — **The Ecological Approach to Visual Perception, Chapter 1**

Summary of the Readings

Let's take a look at the readings for this module. The first comes from Gerd Gigerenzer. His work is very focused on pushing back against the popular idea that humans do irrational things. He recontextualizes the kinds of things people do and asks whether there's some hidden logic in how people behave where there doesn't appear to be any — maybe it's your model of the world that's flawed, not the behavior that so-called "irrational people" partake in. A quote from Gigerenzer:

> *"Many skills lack descriptive language."*

In addition, even if you can describe a skill, your description might be in opposition to how you act. You might think you do one thing, but actually you do something completely different to achieve the outcome you're going for without realizing it. You could say that there exist two worlds: the world of embodied action and knowledge, and the linguistic world. We often try to connect them to each other, but we're not always successful.

Another quote from Gigerenzer:

> *"Good expert judgement is generally of an intuitive nature."*

That's a very interesting point, considering the scientism we're exposed to on a regular basis. In decision making, for instance, we often get told that we need some long formula, or a complicated and intricate model. But in fact, the best of us operate on intuitions developed over years of experience.

Another paper from Maturana, '*What the Frog's Eye Tells the Frog's Brain*':

> *"The operations [of the eye] thus have much more the flavor of perception than of sensation, if that distinction has any meaning now."*

Often, a distinction gets made between perception and sensation. The idea of sensation is that raw, unprocessed information (for example) hits your retina, and from there the information gets processed by your brain. It actually turns out that the two are much more entangled, and there's never a raw, unprocessed information layer. The moment the photons hit the retina is the moment some processing begins to take place, and patterns begin to be extracted.

This last quote is from Gibson, from *'The Ecological Approach to Visual Perception'*. For him, *ecological* means that the organism is situated in its natural environment. It doesn't refer to flows of matter and energy — the way in which we usually think about ecosystems. Gibson:

> *"The environment is not the same as the physical world."*

What does he mean by that? What exactly are we perceiving? We'll look into that as we go.

Motivations

- We looked at the organization of living systems, and we know they must reach requisite variety to remain viable in a fluctuating environment.

- Many organisms achieve this in part through acting in and on the environment.

- The ability to behave in this way demands the ability to perceive and act.

- We posit that perception and action are really two parts of an inseparable whole. For example, we can scope to the heart or to the liver when analysing an organism, but they both come from a unified process.

- Further, we posit that the mind emerges out of the *perception-action-environment* relationship, and that most attempts to treat the mind in an abstract, isolated fashion suffer from a scoping problem. For example, the mind is not just a product of the brain, even though the brain is certainly involved in that process.

- Similarly, we see that the *organism-environment* system comes as a whole. This gets at what Gibson said about the environment not being the same as the physical world. The environment of a given organism includes information about the structure of the organism. You can't only look at the environment, you can't only look at the organism.

- Connecting to our concepts of self-organization, behavior is not facilitated by a central controller. It might certainly seem that way, given our brains and bundles of neurons, but it would be wrong to think of the brain as a central controller that receives signals and sends out instructions. In reality, behavior is something that is a decentralized and distributed process that involves the brain, the body, and interaction with the world.

Closure and Information

Last module, we talked about **autopoiesis** — a kind of topological closure over some subsystem of the larger system. To revisit this quote from Ashby:

> *"Cybernetics might, in fact, be defined as the study of systems that are open to energy but closed to information and control — systems that are 'information-tight'."*

Again, we're seeing this interplay between closure and openness of systems. When talking about perception, it's odd, in a way, to talk about closure of the information of a system because the system has to get some kind of information from the environment in order to act. We'll seek to reconcile that.

Another quote, from Bernard Scott:

> *"...an organism does not receive "information" as something transmitted to it, rather, as a circularly organized system it interprets perturbations as being informative."*

Scott is getting at the notion of perturbing existing parts of the system, instead of consuming prepackaged information from the outside. There's a two way exchange that occurs between the organism and the environment through perturbation:

Brain and Mind

Often, when cognition and the mind is discussed, the brain is focussed on. That's a reasonable thing to do — the brain is, in and of itself, massively complex. The human brain, for example, has about 100 billion neurons, and about 100 trillion connections. And if you think about the total possibile state-space for this system, you'd have close to 2^{100} billion possible states, which is unfathomably huge.

But do you need a brain to have a mind? Does a mind demand a nervous system? Perhaps not. Cognition, or what we call *the mind*, might be an emergent property of an autopoietic system being coupled to an environment, and it's possible that nerves and neurons are incidental to that.

Nowadays, for example, it's pretty easy to find arguments that trees and other biological systems have something resembling a mind. Trees don't have nervous systems, but they have root systems that engage in a lot of interesting chemical signaling within their own root networks and among other trees.

The Scope of Mind

How should we scope the mind? We might start by looking at the brain in isolation.

But then, we might realize that the brain relates to the body, and start looking at the mind as the brain plus the body.

Then we might say: “well, the *brain-body* system exists in an environment, and actually couldn’t come about and thrive if the environment around it didn’t exist”. So then, we might think about *the mind* as the brain, the body, and the world at large outside of the bounds of the autopoietic system.

You’ll reach different insights depending on how you set the scope of your analysis. It’s possible to miss-set the scope, and thus not be able to answer the question you’ve set out to answer. For example, if you ask “how does a person catch a ball?”, you’ll get no answer if you study the brain in isolation.

The Homunculus

Here’s a thought experiment to show why this notion of a “central controller” is wrong. Central to it is the idea of the **homunculus**.

Imagine a little man that sits in the machine and acts on the world. That can’t possibly be the case. Because then, you’d have to have another little man who sits in the head of the first little man and dictates all of his actions. Hence, an infinite regress of central controllers. Turtles all the way down. You never get closer to answering the question of how actions and behavior is controlled.

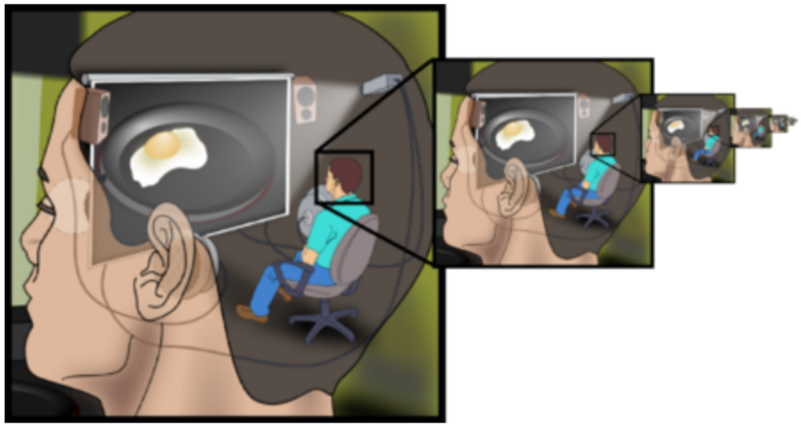

What Does the Mind Do?

So what *does* the mind do, then? The classical view is that the mind houses mental representations. The best thing that those neurons can do is try to figure out the material properties of the thing being looked at and represent that thing with as high of a resolution as possible as a model in the brain, *as that thing exists in nature*.

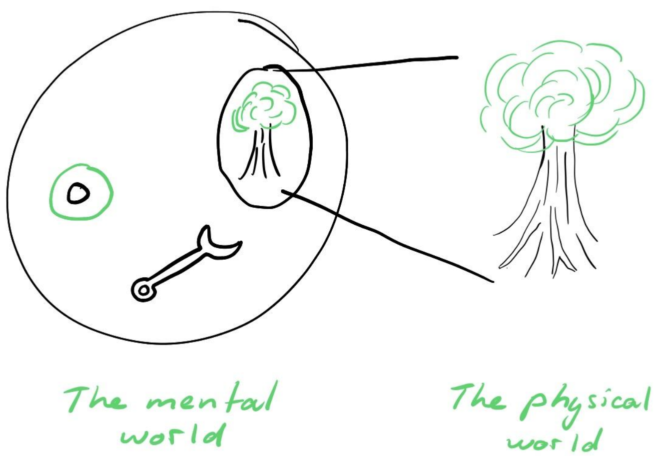

That classical view is almost certainly wrong. There's way too much going on in the world for our brains to represent it in detail. Every piece of information gets filtered as it goes to your brain. We can say that there's something objectively out there, and that our brain creates some kind of models to help you act in the world, but how exactly the physical and the mental world relate between each other is a topic of hot debate for those in the field.

Symbolism vs Connectionism

The Artificial Intelligence community is a group of people who think about the mind and how it works in a rigorous way. Let's briefly explore their ideas. There's a major dichotomy that exists in this space — that of the **symbolists** and the **connectionists**. There is definitely a lot of crossover that happens between these two schools of thought, but for our purposes, let's focus on the distinction between them.

Symbolists, basically, want to build AI systems that resemble our representation model. They think of systems as having discrete symbols that have a clear representation, and that live in some network of relations.

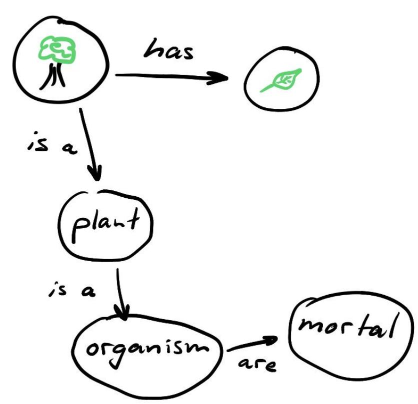

Symbolists envision the mind as a kind of logical process over these networks. So what you have is a graph and a set of reasoning rules to make inferences internally in 'the mind'. Let's look at one example of such a graph, shown above. Starting at the node that represents a tree, one can determine that a tree is mortal, even though there is no direct arrow that goes from *tree* to *mortal*.

Historically, graphs like this have been called **expert systems**. The more contemporary term is **knowledge graph**. There's usually an **inference engine** involved as well — it tells you how to make new propositions over the graph.

Expert systems are really useful in a lot of settings. They've had success in medical diagnostics, for example. They're also used quite extensively by Google.

Expert systems tend to go in and out of fashion, and are not without their shortcomings. These graphs are incredibly fragile to the assumptions they make. If you have a wrong link, then your conclusions are wrong. If you present something that's not on the graph, the system doesn't know what to do with it. The strength, of course, is that these systems are human-readable.

Connectionists, on the other hand, are the people more interested in what we call **artificial neural networks (ANNs)**.

The brain is composed of neurons — cells that fire electrical signals down a tube — and once you connect all of those neurons together, you have a brain. Neurons either fire at any given moment (associated with *on*) or not (associated with *off*). A neuron "decides" whether to fire or not based on whether the neurons connected to it are firing into it or not. That's some quick and dirty neuroscience for you.

Connectionsits decide to build systems that use this idea of neural networks within the brain as an analogy.

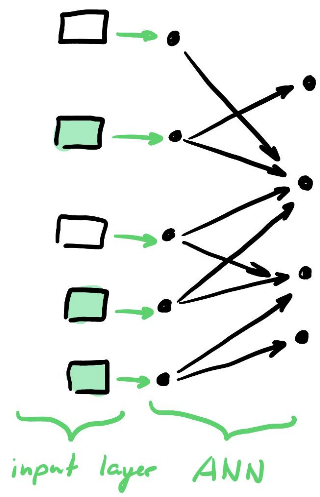

In an ANN, you have “neurons” that are very internally simple — they can either be on or off — and are connected to other neurons. You also have some kind of input layer — maybe a picture with pixels, some pixels being on and others being off — that gets fed into the neural network. In an ANN, your focus is on what the connections between the neurons should be — which neurons should be connected to which neurons, and how strongly should they be connected? In other words, you’re concerned with the organization of the system.

Because none of these nodes — the “neurons” in the ANN — seem to have any obvious correspondence to anything we’re interested in — trees, leaves, plants, etc — an obvious question to ask might be: do they represent anything at all? Often, these systems are referred to as **subsymbolic**. If there’s something in these systems that’s akin to mental representation, it can only be on the collective scale of many neurons acting together in coordination — maybe those patterns of activity correspond to things in the world. That’s a very different level of analysis, compared to the symbolists. With ANNs, what we focus on is the collective behavior of many simple units, and how to organize these simple units to do useful work for us.

These systems can be set up to learn about certain kinds of patterns. Based on the output, you either tell the network to change, or reinforce its connectivity. This is what Machine Learning is all about.

The upside of connectionist systems is that they handle changes and differences in input very well. The downside is that these systems are basically a black box. There are many neurons and many connections, and it can be very hard to tell — if not impossible — why you're getting the output that you're getting.

Types of Artificial Neural Networks

Within the realm of ANNs, we have two major classes.

Feedforward networks are like the one we drew above: they have an input of some kind, and multiple layers of "neurons" which eventually converge on an output. The point of a feedforward network is that the flow of activity is going one way. A neuron to the right cannot affect a neuron to the left of it.

In **recurrent networks**, circularity is possible. You can have loops in the system.

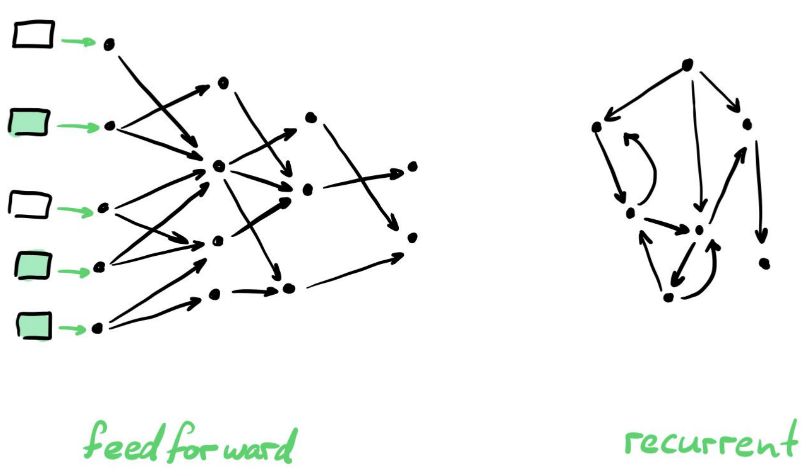

Recurrent systems can generate and maintain their own activity. They have internal dynamics and lawful behavior of their own. Feedforward systems have a clear input, a set of transformations that take place, and a clear output. In a recurrent system, you might have input (i.e. perturbations) coming from the outside, and you might have some effectors that it's influencing, but in general these systems resemble the Bernard Scott picture of internal dynamics being perturbed much more.

Perceptrons

The most basic version of a feedforward network consists of a type of artificial neuron called the perceptron. A perceptron can only have a binary output. It works like this: there are weights associated with the neuron — x_1, x_2, x_3 for example. Each weight can have a certain value. Depending on the values of the weights, the neuron determines whether to fire or not (i.e. whether the output is 0 or 1).

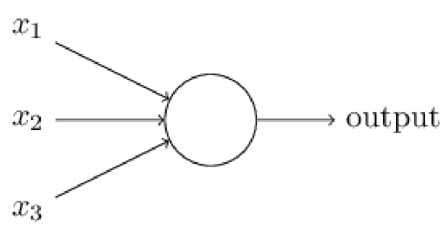

If the sum of the weights is above some threshold value, the neuron fires. If below, it doesn't. We can write this out in an equation:

$$f(x) = \begin{cases} 1 & \text{if } x \cdot w \geq b \\ 0 & \text{otherwise} \end{cases}$$

The variable *x* is our input, *w* is the weight assigned to that input, and *b* is the threshold value.

Hopfield Networks

A hopfield network is the simplest case of a recurrent network. In the network below, for example, everything is connected to everything, making a complete graph. If we have 5 nodes, for example, we can refer to the complete graph as C_5.

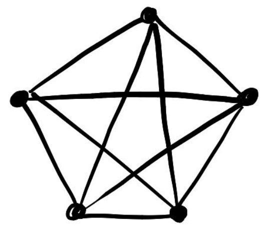

As opposed to feeding it an input and getting an answer, what you're trying to construct is a system that'll have a set of attractors corresponding to things that you want the system to remember and recognize. Like in a feedforward system, we're looking for a set of weights, except now we're looking for a set of weights such that there's a prototypical representation of, lets say, the notions of *1, 2, 3,* etc (if we're looking to build an ANN that recognizes handwritten numbers). The individual example of handwritten digits may be different, but the idea here is that they land somewhere within these basins of attraction, and will evolve under their own dynamics to a stable fixed point. The fixed point acts as a sort of "memory" of a 1, a 2, a 3, etc.

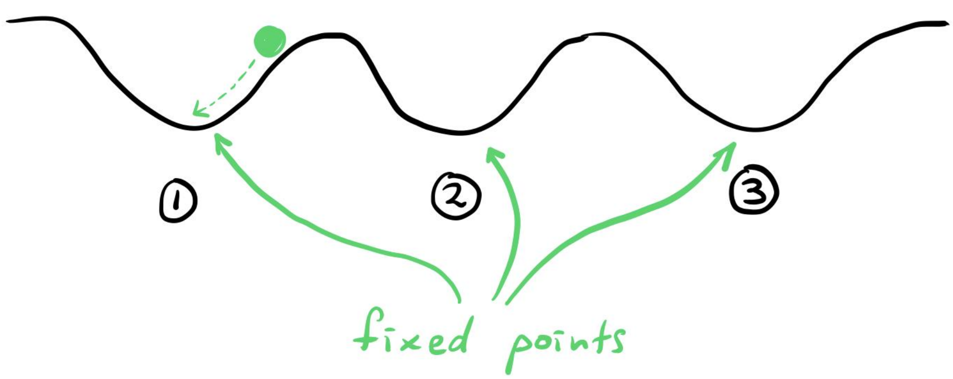

This is sometimes called **content-addressable memory**, or **associative memory**. Instead of searching solely by address or location, you look for places where the system resolves to a common state. You can evolve the state with a set of difference equations, like we've seen before, resulting in some binary output:

$$s_i\,[t+1] = \begin{cases} 1 & \text{if } \sum_{j=1}^{N} w_{ij} s_j \leq \theta_i \\ -1 & \text{otherwise} \end{cases}$$

The state of neuron i — s_i — at the next point in time is determined by summing over all of the neurons (from 1 to *N*). With each neuron you take the weight of the connection between our neuron i and a neuron j, and multiply it by the state of j. If the value of the sum is greater than some threshold value for that neuron, Θ_i, then the state is 1. Otherwise, the state is -1. Why not 0? It has to do with how neurons reinforce each other. If the connection between two neurons is positive, the neurons will synchronize (if neuron i is positive, it will push neuron j to also be positive). If the connection is negative, the neurons will push away from each other (the positive neuron will push the other into a negative state, just as a negative neuron will push the other into a positive state).

The name of the game here is determining the weights such that they generate an attractor dynamical landscape, where the memories are of meaningful things.

The Retinal "Image" and the Frog's Eye

Let's continue with the example of visual perception, since it's a very natural and intuitive thing for us to think about.

There's this idea of the **retinal "image"**. The thinking is, the eye acts like a camera. The light enters the eye, activates certain nerves at the back of your eye (which can be thought of as pixels), and this created bitmap gets sent to the brain for processing.

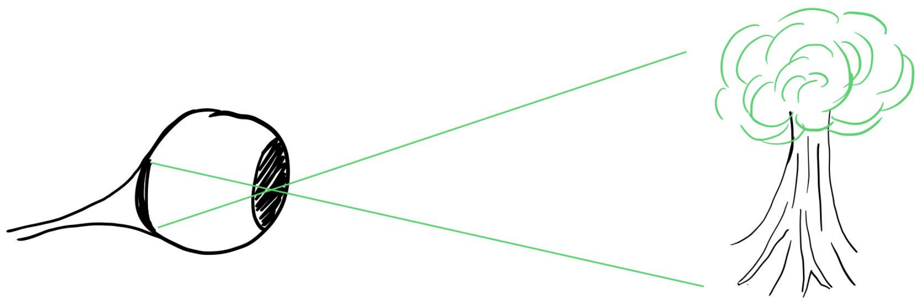

It turns out that this notion of a retinal image is wrong. What the Maturana paper shows is that there never exists a retinal image — this raw unprocessed bitmap that then gets interpreted by the brain. The moment that light hits the retina, there are lateral and feedforward connections that determine what gets sent down the nerve and to the brain.

There's an organization that our retinas follow, called **center-surround organization**. Center-surround organization responds to the spatial pattern of light. It looks something like this:

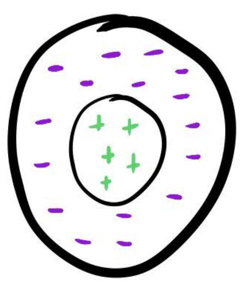

This should look familiar, as it's similar to the **local activation, long range inhibition** we've seen before. Center-surround is organized in basically the same way. There's an inner ring of excitatory activity (+) and an outer ring of inhibitory activity (-).

How does it work? Imagine this center-surround ring is scanning an area of black and white pixels — the black representing positive numbers and the white representing negatives. Let's see how our ring responds to a black square. If the square is big enough to cover both rings, the +'s combine with the positive black square to get a positive result, and the -'s combine to yield negative. Taken together, you get zero:

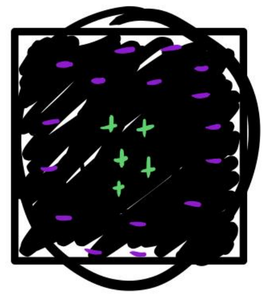

However, if only the inner ring was sitting on the square (+'s combine with positive pixels), and the outer ring sitting on the white (-'s combine with negative pixels), it would yield positive in both rings:

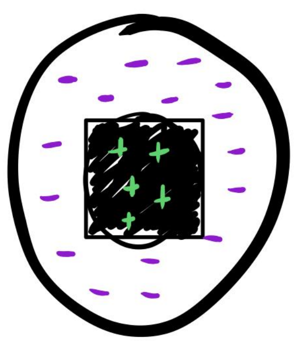

So our detector is not responding to the pointwise light of the retinal "image", but rather the spatial pattern of light.

Even this model is massively oversimplified. It turns out that there's also a bunch of information coming down from the central brain to the retina that modulates activity as well, along with other mechanisms. One interesting consequence of this: color is always thought of as wavelengths, but you can induce the perception of certain colors without the wavelengths being there. This is because perception really comes out of these relations and patterns — information getting communicated back and forth between the brain and the eye.

Gestalt

There's a whole field of what's called **Gestalt psychology** that opened up — looking at how the wholes of what we perceive relate to the component parts, and the implications that has for perception. Here are some of the concepts they looked at:

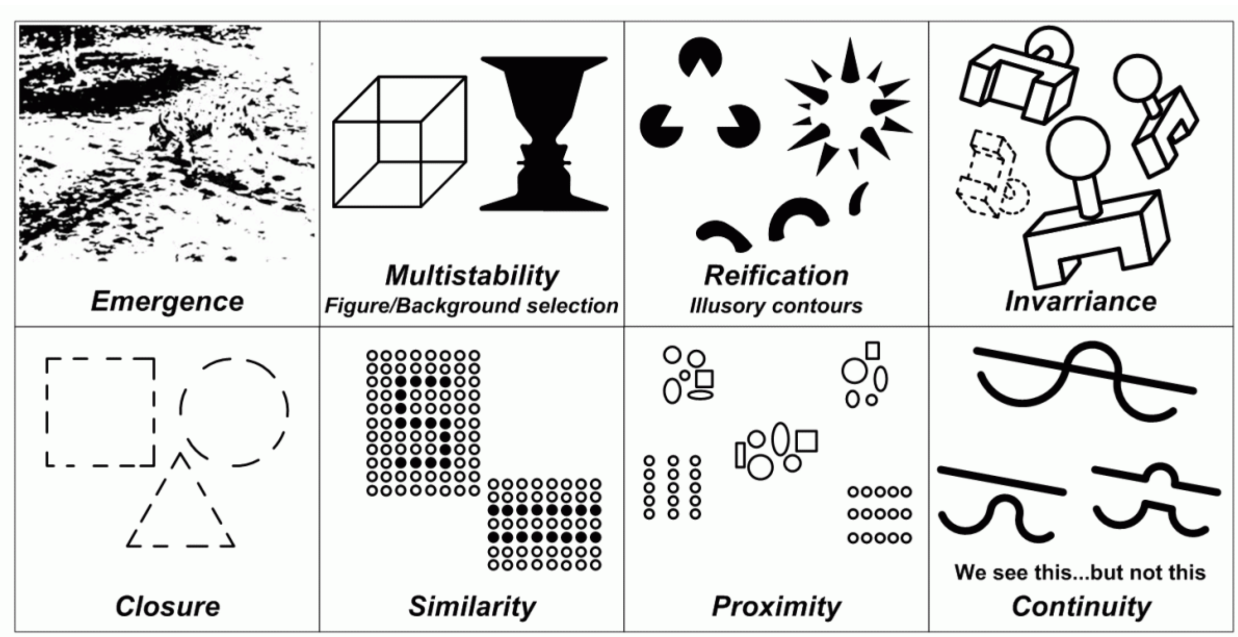

For example, with **reification**, you can see the edges of a triangle where none exist, simply based on the strong implication of continuous edges you get from other parts of the image.

One-Shot Learning

There's another interesting property of our perception called **one-shot learning**. When you experience something, you get it forever — you don't need a training set to learn it over time.

See if you can make sense of what's going on in this picture:

Once you see it, you can't unsee it (see the answer here).

Detection of Constraints

Your brain can also pick up on the constraints between different parts in a dynamic image — how do the parts move relative to each other.

This is highlighted in the Johansson experiment. Click on the video and see how much easier it is to tell what is being communicated when the dots are moving, as opposed to them being paused:

One might think that movement complicates things and makes it harder for us to perceive and interpret objects in the world. However, what we often find is that movement often gives you information. It's not an additional problem to be solved, it's actually telling you how to solve the problem.

Closure of the Nervous System

Imagine a child sitting and doing math on paper: does that violate informational closure? Actually it doesn't — the nervous system doesn't care if two neurons talk to each other through synaptic connections, or if that information passes through the world (say, neuron → hand movement on page → retina → neuron). We often do our thinking through the world, without breaking the assumption of informational closure of the nervous system.

Sensorimotor Contingencies

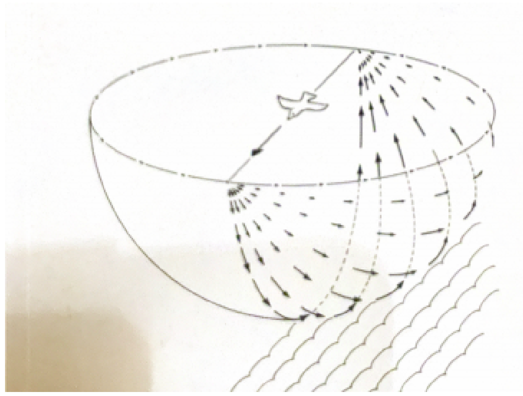

In the above picture, we see a flying bird. How it's flying is determined by how it's using its muscles. But when it uses its muscles, that has consequences on the sensory stimuli it receives, creating a kind of feedback loop.

When the bird moves *toward* some point, from its perspective it will see the **optic flow** come *out* of the point. It'll see a spreading out of the optic environment.

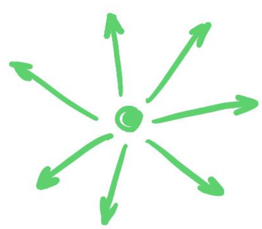

On the other hand, if it's moving *away* from something, the direction of the optic flow will be reversed, now flowing *into* the point. The environment will seem to clump together.

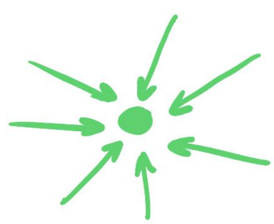

That's how birds — and all other animals — figure out what it's moving towards and away from — by looking at the optic flow. If you start walking towards a wall, for example, you'll see an outward optic flow. If you stop walking, the optic flow will stop. The sensation is completely coupled to the motor behavior.

And that's the idea of the **sensorimotor contingency** — how you move affects what you sense. This is how we learn to perceive the world.

Heuristics

Scoping out a bit, let's look at how people solve perceptual and action problems using the environment. One example that Gigerenzer brings up in *Gut Feelings* is how an outfielder catches a fly ball.

Richard Dawkins proposes an idea in *The Selfish Gene* about how this might work: what the person must be doing is running a model in their head that calculates the trajectory of the ball and where it'll land, and using that model to catch the ball. It turns out that *this is wrong*. Very, very wrong.

Instead, researchers identified a network of heuristics that people use to catch the ball. We won't get into nuances here, because it turns out to be quite an intricate network of heuristics, but let's look at one basic heuristic, when the ball is already in the air: look at the ball directly, then move your body such that the angle of your gaze becomes stabilized. If you need to drop your gaze to keep looking at the ball, that would indicate you need to speed up. If you start raising your gaze, that means you need to slow down. The constant angle puts you on an intercept path with the ball.

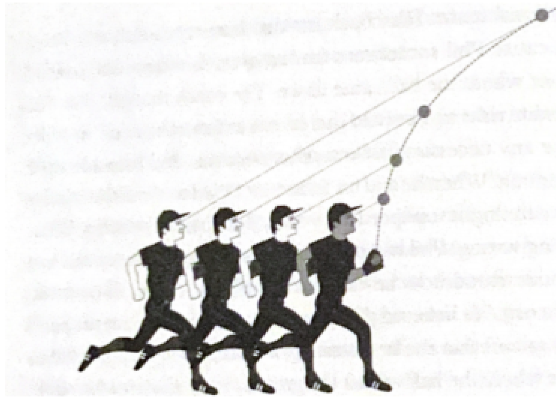

In general, heuristics trump abstract calculations.

Landing a plane is similar. J. J. Gibson's early work focused on WWII pilots and the kinds of heuristics they used for landing planes.

Pilots look out at the runway through their windshield. It turns out that they use smudges on their windshield to align themselves perfectly with the runway — if the smudges stopped moving relative to the runway, that put them on a landing trajectory.

Actually, this same phenomenon is responsible for air-to-air plane collisions. If two planes are on a collision course, the angle between them stays constant, and therefore the relative angle on the retina of the pilots stays constant, which means there's no movement on the retina. We perceive movement much better than we perceive static things. In fact, the retina actually filters out unmoving objects if they stay on the retina for long enough, because it assumes that it's some kind of artifact. Because of this, pilots often fail to see the incoming plane.

Asimo vs Boston Dynamics

Take a look at these two videos:

[Honda's Asimo Robot]

[Boston Dynamics' Atlas]

These two robots are great examples of different paradigms when it comes to perception. Asimo is based on precalculation — calculate everything in advance in your brain model and decide what to do afterwards. Atlas, on the other hand, uses a combination of neural networks and heuristics. He can deal with uneven terrain and people pushing him over — notice how he puts out his leg as he almost falls (0:45). As Atlas detects that he's losing his balance, he relaxes his leg and thereby lets his leg find gravity. He doesn't necessarily know where gravity is, but his leg does. It's easier for his leg to figure that out than his brain.

Asimo, on the other hand…

Affordances

J.J. Gibson's lasting contribution is in the concept of **affordances**. Affordances are defined as *opportunities for action*. What the opportunities for action *are* depends on the structure of the organism. How is it built? What does it know how to do? What are its previous experiences?

Imagine there's a tree and that there's a branch on that tree. A question might be: "Is this branch graspable?" Well, that would depend on who's asking. For an organism that has a hand, for example, that branch might be graspable. An affordance always evokes the agent and its structure.

Remember: *"The environment is not the same as the physical world."* Whether or not a branch is graspable is not a statement about the physical world. It is instead a statement about the interface between the agent and the physical world. Things in the *environment* are graspable, but *graspable* is not a concept that exists in the *physical world*.

If the branch is graspable, then the tree might be climbable. Other things in the environment might be sittable, walkable, or liftable. A kind of verb-centric approach to perception.

The idea of affordances has also found its way into design. Don Norman wrote a book called *The Design of Everyday Things* in which he explores this. Affordances can work well with design or work counter to the design. Here are examples of bad affordances:

Notice that the handle for the teapot is on the wrong side — it would be difficult to pour from such a teapot. It's not *pourable*. Although, if this is a chug pot, this might actually be a good affordance. Same with the doors — the handles on the door make you want to pull, when in fact you have to push to open the door.

Coordination

All of these heuristic-based control processes have a few things in common.

First, the control is decentralized. There's some concentration of control in certain places, like in the nervous system, but for the most part, the processing takes place in many locations instead of one central location.

We've seen pattern-forming processes, and we can think of coordinated behavior in pattern-forming terms.

We also have dynamical behaviors like the ones we've seen before, such as phase transitions.

Horse Gaits

Look at this **video of horse gaits**. Notice that as the horse changes speed, you get distinct phase transitions between different modes of walking — qualitatively different coordination regimes, just like the phase transitions we've seen before. There's a continuous control parameter — velocity — but as you sweep it, you get qualitatively different coordination patterns. All motion contributes, in a decentralized coordination process, in order to change phases.

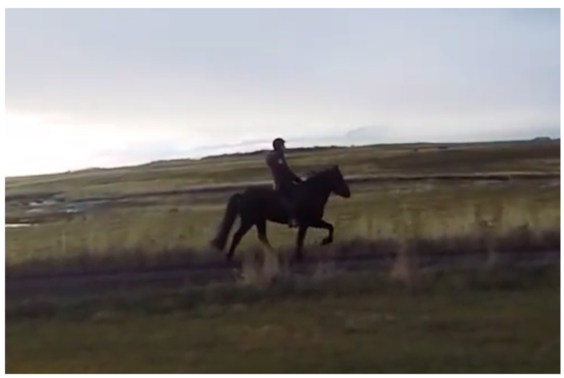

Coordination Dynamics

Neuroscientist Scott Kelso developed a subfield of dynamical systems that looks at coordination behavior. One of his experiments involved moving your fingers in-phase and anti-phase. What you start to notice, as you increase the speed of the movement, is that you'll typically spontaneously change into an in-phase pattern, because it's more stable.

The stable, in-phase pattern is not just a pattern in your nervous system — it's also a pattern in your ligaments. At different speeds, your ligaments optimize for different motions, like we saw with the horse gait.

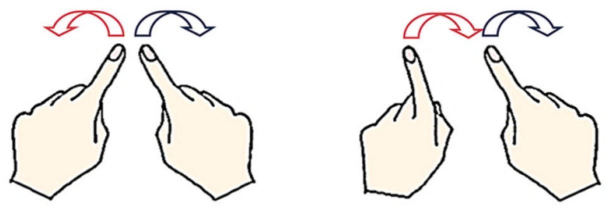

Countersteering and Explicit Knowledge

This video shows that when you steer a bike, to turn left, you actually have to first **countersteer** and turn the handlebars to the right. Yet if you were explaining it, you might say that to turn left, you turn the handlebars to the left. Here once again, we see a disconnect between explicit and implicit knowledge. This, like the other examples we've seen, serves as a strong indictment of explicit knowledge — you might think you know something and be totally wrong.

There's also an implication here for learning — in some cases it might be better to learn by doing something in the physical environment as opposed to listening to a lecture or reading a book, because you'd pick up on intuitions and heuristics that come with implicit knowledge. Of course, this breaks down with more abstract subjects, but wherever possible, learn by doing.

Target Fixation

Another brain quirk we come across when looking at the way people drive bikes is **target fixation**. This is a common way that people crash bikes into objects. Watch.

The gist is, ***you go where you look***. By setting a goal, the system organizes itself around that goal. Out of this phenomena comes the aphorism "*focus on the road, not the wall*".

Environment Coupling

There are two ways in which you can move your eyes: in saccades, where your eyes jump from one point to another, and smooth pursuit, where your eyes smoothly follow a moving object. It turns out that in the absence of a moving object, you can only move your eyes in saccades, even if you try to move them smoothly. Here we see a way in which you are coupled to your environment — the option of smooth pursuit doesn't become available to you unless there's something in the environment that prompts it.

Attention & Bandwidth Limitations

The classic **selective attention test video** demonstrates an important point about the way our attention works.

We are severely bandwidth constrained. Even things that are right in your face can be missed when you're focussed on something else. In experimenting on pilots flying in a simulator, it's been shown that even a small distraction can lead to failure.

There's a big implication here for the organization of systems. The reason why we're not centralized is because there's too much information to deal with, and a centralized system can only do so much.

Metacognition

Metacognition is the idea of evaluating your own states of knowledge. It's most apparent not when you know you know something, but when you know you don't know.

Going forward, as we start talking about risk and decision making, it turns out that knowing when you don't know can be quite helpful.

Strange Loops

This idea comes from Douglas Hofstadter. A **strange loop** comes about when a system begins to represent itself, forming a kind of strange feedback loop. In *Gödel, Escher, Bach* and later in *I am a Strange Loop*, Hofstadter uses this idea to develop a thesis that attempts to explain consciousness.

From *Gödel, Escher, Bach, Preface to the 20th Anniversary Edition*:

> *“In short, an "I" comes about - in my view, at least - via a kind of vortex whereby patterns in a brain mirror the brain's mirroring of the world, and eventually mirror themselves, whereupon the vortex of "I" becomes a real, causal entity. For an imperfect but vivid concrete analogue to this curious abstract phenomenon, think of what happens when a TV camera is pointed at a TV screen so as to display the screen on itself (and that screen on itself, etc.) — what in GEB I called a "self-engulfing television", and in my later writings I sometimes call a "level-crossing feedback loop”*