

Levels of Distributed Cognition

The US military budget for 2017 allocates three billion dollars for the study and development of what they call “Centaur Warfighting” — the collaboration of humans and machines.

This reflects something that Paul Tudor Jones recently said: “No man is better than a machine. And no machine is better than a man with a machine.”

The real frontier here is not a computer program that runs by itself, or a software program that executes by hitting an enter key. That would be too easy.

No, the frontier is a true synthesis of wetware (the mind of man) and hardware (the capabilities of the machine). Integrating the cognitive capabilities of the creative trained mind with the speed, precision, organization and calculation capabilities of digital circuits.

This is not an easy thing to do. A lot of mental work is required to truly integrate man and machine, with no substitute for training.

Think of practicing with, say, a samurai sword.

If you train with a sword long enough to develop expert level proficiency, the sword is no longer a separate entity. Your mind learns to adopt the sword as an extension of your arm.

In both psychological and physical terms, with enough training and familiarity, the sword becomes an aspect of your being.

Anyone with years of experience driving a car has known at least a basic version of this feeling.

If you drive a car for long enough, your sense of physical space expands to include the vehicle. The three-dimensional space of the car becomes “you.”

Your brain then interprets the surrounding area of the car, when driving it, the way you assess the physical space of your own body. If someone drifts into your lane, it is as physically alarming as a person walking directly in your personal space.

But this feeling only comes through familiarity and training. Teenage drivers will routinely be at high risk of backing or bumping into things, or being oblivious to risks from other vehicles on the road — until they have also driven long enough to incorporate an expanded sense of the car as an extension of their own physical form.

Another way to register this is to go from driving a small car to a large sport utility or vice versa. The feeling can be alien — until the brain adjusts to it, which again happens naturally via experience.

In various sci-fi movies you see people fighting in mech suits.

The “Iron Man” concept from the Marvel Comics series is the same thing — a hyper-advanced exoskeleton with auxiliary support from an AI supercomputer.

The thing that is not discussed in movies or comics is that full use of a mech suit, in high pressure combat situations, would probably require thousands of hours of training.

This is because the mind would have to learn, through extensive practice, to subconsciously adjust to new dimensions of physical space, new levels of speed and strength, modified centers of gravity, the difference between gripping with a human hand and a bulkier robotic one, and so on.

Learning to use a mech suit would thus require massive changes inside the user's brain – a specialized rewiring born of practice and training.

The neural network between the user's ears would have to build a new series of pathways that allowed them to sense the physical parameters and intuitive if/then propositions of the mech suit as if it were a second skin.

And as we've discussed, there are no shortcuts to architecting a neural network inside someone's head. We don't have the science, and even if we did have the science, the software part is far harder.

So when the military puts billions in a centaur program, this is what the focus is: Figuring out how to integrate mind and machines to maximum effect, which requires building the machines a certain way and building the process a certain way, then creating repeatable training programs so that results can be replicated.

But this concept is not just for super powerful technology or exotic machines. It applies to the simplest technology available. You could apply it to a stick or a rock.

There is a concept known as "Distributed Cognition."

Distributed Cognition can be roughly defined as "emphasizing the ways that cognition is offloaded into the environment through social and technological means."

Distributed cognition is about taking knowledge and process and information outside your brain, then accessing it through habits and training and tools.

How does this apply to rocks?

Well, imagine that a hunter gatherer band 50,000 years ago developed a system where the placement of four medium-sized rocks in a straight line meant "danger."

This four-rock symbol could then be placed, say, in front of the hidden entrance to a rattlesnake den near an area of fruit trees. If another member of the group saw the rocks, he or she would know to instantly be on guard. Communication would occur through preexisting shared knowledge of a symbol and its significance. That is distributed cognition in one of its most basic forms.

So distributed cognition starts with something as simple as a symbol or a learned gesture. This imparts knowledge to the external environment, taking it outside the brain.

Distributed cognition technology tools enhance the power of the mind by off-loading mental tasks that would have consumed energy otherwise. The more that these tasks are off-loaded, the more here-and-now energy is available.

We can perceive different levels of distributed cognition capability.

Think of these differing distributed cognition levels like levels of power-up in a video game or Dungeons and Dragons style role-playing game. The higher up you go in the levels, the more powerful your mind becomes, through add-on capabilities of external integration.

In terms of distributed cognition levels, Level “Zero” would be relying on memory alone – no outside tool use, just using the brute force of the conscious mind to keep information handy. This would be the person who relies on short term memory to juggle all the things they have to remember or focus on or get done.

Many people are at DC Level Zero not all the time, but a significant percentage of the time, for a lot of things. “Oh I’ll just remember that,” or when someone tells you something, telling yourself to remember it but not writing it down.

As we know, level zero is not very efficient, because the mind was not built for short term memory storage in a complex day-to-day world. So this is akin to juggling a bunch of tennis balls and dropping them most of the time. Juggling is a neat trick, and it can be done with practice and concentration, but we weren’t built to do it easily or under pressure while moving fast.

DC Level One would be using simple pen and paper. That alone is a form of distributed cognition.

A person who writes stuff on post-its and sticks them to their desk or car dash is already better off than someone who tries to just remember stuff with brute force. This is using symbolic information to reduce cognitive load. Writing “BUY MILK” on a post-it note that is visible on the dash to jog memory on the way home, driving by the grocery store, is a step forward from just plain old memory.

Level Two, moving past just pen and paper, would be going from just “writing things down” to developing some kind of rudimentary system. Breaking things into groups, or having a structure on a page or some other kind of super-basic routine, would qualify.

The levels then just keep going higher.

Next on the distributed cognition ladder might be incorporating a priority system. Then a basic calendar system. Then the ability to have alerts or automatic reminders, and so on.

By the time someone gets to a calendar based alert system, perhaps run by a software program that sends them reminders for tasks, they are probably around level ten.

This person has significantly less head clutter than the average individual. If they use the system and rely on it, they have off-loaded a lot of mental energy expenditure to a machine.

Of course, to really get the benefit of being at level ten, the system has to be relied on fully – and most people fail on this count as well. If the mind does not trust a system, then subconsciously the mind will not rely on that system. It will use the system somewhat, but also rely on lower level, more primitive distributed cognition systems as a back-up... or even revert to level zero, just trying to remember.

This in turn defeats the whole purpose, which is why most half-hearted attempts to upgrade distributed cognition systems, to the latest productivity tool du jour, wind up being abandoned.

To really get use of the distributed cognition upgrade, as discussed with the sword and the car and the mech suit, there has to be “full integration” between brain and system. A neural network between the ears has to be formed, from which automatic trust and confidence and routine execution is developed. This is the hard practice and training part. And as far as level ten on the distributed cognition scale goes – reliably having a system that spits out reminders for activities to keep them out of physical memory – that is a big upgrade for most people, to be solidly at level ten.

And yet, in the big scheme of things level ten is still almost nothing. It is barely getting started. There is a level fifty out there. There is a level one hundred. There is even a level five hundred.

Super advanced levels of distributed cognition go beyond task planning and calendar enabling and move into storing creative nodes of information outside the brain, in ways that radically enhance the mind’s ability to store, retrieve, combine and manipulate patterns of information.

The higher one goes in distributed cognition levels, the more one effectively moves from “human” realm toward “centaur” realm – becoming a kind of cyborg, though without the messy need for body modification or wires sticking out of one’s head.

Imagine the ability to take nodes of information – relating to creative planning, or trading research, or sports betting statistics, or anything at all really – and recreating them digitally outside the brain, within a system that the mind has full and complete access to, as such that the mind can then interact with such nodes effortlessly and at will, drawing on all the power of the machine (digital circuits) that is used while doing so.

By the time a person gets to, say, distributed cognition level 50, they are using technology to create a sort of solid state hard drive that exists outside the brain, but is wired directly into the brain. Not wired physically, mind you, but wired in a state of invisible connectivity powered by the plasticity of the neural network created between one’s ears.

The connections that allow for distributed cognition level fifty or one hundred and so on are comparable to those that let a samurai wield a sword like an extension of their arm, or a racecar driver intuitively handle a multi-thousand pound machine going 130 miles an hour into a turn.

The structure in the brain is what maps to the outside technology and connects it.

This relates back to Mushin – and “empty mind, trained mind” – because the more distributed one’s cognition is, the greater the ability one has to handle higher levels of informational complexity on the fly, fluidly and in real time, while keeping a clear and uncluttered mind.

It is also possible to build mental models and best practices of reasearch or analysis *directly into* the structure of a well designed distributed cognition system.

Doing high quality analysis on a trade idea, for example, means diligently following a series of steps. Those steps can be built into an external system.

We have said that there is no substitute for building the neural network -- the suite of models inside the brain, represented by the shaping of neuron clusters and whatnot. But that is not quite true. There is actually a form of shortcut, which is a combination of rapid learning and external guidance rails accessible in real time through interaction with the system embedded in the machine.

If one could have access to a distributed cognition system that had the right mental models built into it... and if a habit was developed of using this system to do analysis... then that would be a form of routinely accessed wisdom embedded in the machine – built into the hardwar itself.

The hardware of the distributed cognition system is a literal thing. There is a structure to it, based on software that runs in the cloud. It can be pen and paper, but it becomes potentially more powerful as more power is added to the system that hosts and delivers the externally stored information.

This describes our own “centaur program,” in a way, in which the nature of the distributed cognition systems as described, for the purpose of trading, are being evolved in real time.

As mentioned before, one of the motivations for daily commentary in the MMTP was to give time and room for next stage evolution steps in the distributed cognition systems used today... that already exist... that are being refined for public sharing and distribution across the community.

This is the technological shape of mastery.