HUMAN SUPERINTELLIGENCE:

How you can develop it using recursive self-improvement

Chapter 4: Stumbling onto the path of recursive self-improvement

Getting 12 out of 20 on my first physics test in Year 11 at school was the trigger that launched me onto the path of recursive self-improvement. Up until then, I had found maths and science easy. Generally, I got 100 percent or close to it in maths and science exams throughout my schooling. Usually, I did not get full marks only when I made a silly mistake.

I had not done as well in other subjects such as English, History and Geography. This was particularly the case when examinations in these subjects started to focus on essays rather than the mere regurgitation of facts. Looking back, I was exceptional in any subject with clear and identifiable rules for generating the right answer, including where straight memorization was required. In contrast, in subjects like English, there was often no formula that I could discover and apply that would produce what was valued by the teachers. Unlike in maths and science, 'correctness' in these subjects seemed to me to be subjective to some extent. It was not always black and white.

It was only years later that I saw that these characteristics of mine were symptoms of an underlying condition.

But this reliable pattern was seriously disrupted by the first physics test in Year 11 in 1968. It was immediately obvious to me that the cause was associated with the introduction of a new kind of physics course into the Queensland education system that year. This new approach to physics teaching, known as PSSC physics, focused on imparting an understanding of the underlying principles and theories of physics. The course then tested the extent of a student's understanding by getting them to apply their knowledge to solving relevant physics problems. The proponents of PSSC physics proudly proclaimed that the rote learning of facts about physics would no longer be sufficient to do well in physics in years 11 and 12.

This shift in emphasis from memorization to understanding was reflected in the fact that in a number of education systems around the world, the PSSC exams were 'open book'. You could take the PSSC text book into the exam with you. This gave all students easy access to the relevant physics facts, but it would not necessarily get them far in the exam because of its focus on understanding. However, as it turned out, 'open book' exams were not implemented in our course, though they were for chemistry, which had a similar philosophy.

12 out of 20 was almost failing! What was I to do about it?

This serious challenge to my self-esteem brought to the forefront of my mind something I had read in a book a few years earlier. 'Psycho-Cybernetics' was the title of this book that my father had brought home from a library. The blurb on the cover suggested that it was some kind of bestseller. Key themes of the book were that an understanding of cybernetics and related scientific principles could be used to achieve success and happiness in one's life. I was particularly intrigued by the idea that as we grow and develop, we can intentionally change the characteristics we find in ourselves.

This approach resonated with similar themes in several books about the ideas of George Gurdjieff that my father had also brought me from a library. Gurdjieff was known as a spiritual master of sorts in the early part of the last century and also demonstrated a highly developed capacity for systems thinking.

A central theme of Gurdjieff's philosophy was that individuals can consciously work on themselves throughout their lives to develop self-mastery. We do not have to take ourselves as fixed and given, but can instead see ourselves as a work-in-progress. As we discover the existence of capacities and skills that would make us more effective, we can take steps to install them, continually remaking ourselves intentionally as we grow and develop.

However, as far as I could see from the books that I had access to, Gurdjieff did not provide much detail on how to work on yourself to enhance your capacities. But his key idea took root in my mind and predisposed me to be on the lookout for methods that I could use to enhance my capabilities.

What particularly caught my attention in Psycho-Cybernetics was its discussion about how some teachers in Russia were using ideas about algorithms to turbocharge their teaching methods. Instead of teaching facts about science and maths, or providing worked examples of how standard problems could be solved, they were teaching directly the algorithms that could be used to generate solutions to the problems. As I interpreted it, they were identifying abstract methods and strategies that could be used to solve whole classes of science problems and teaching these to their students.

My immediate reaction was: "Why aren't our teachers doing this already?" It would be so much more effective. And the long hours that I spent bored in class every day would be more stimulating. It would not be as good for me as pursuing my obsession with fishing, hunting, prospecting, and other outdoor adventures. But it would be preferable to sitting at the back of every class I attended, whispering to my mates about what we had heard on the radio the previous night and trying to avoid the teacher's gaze.

This failure of our teachers to implement a teaching practice that seemed to me obviously superior confirmed the poor opinion I held of them all, without exception. The school I attended (the all-boys Church of England Grammar School in Brisbane) was then reputed to be the top school in the state of Queensland in both academic and sporting achievements. A number of the senior teachers had written the textbooks that were used widely throughout the state. But all of them had physical and personality oddities of various kinds. The boys mercilessly exploited these. And they spent their lives teaching stuff that was generally obvious and boring.

My shocking result in the physics test brought me to an important realization. Waiting for my teachers to teach me the relevant problem-solving algorithms would be futile. They had no idea. Furthermore, there were no books available where I could read more about these approaches (this was 1968, computer science was in its infancy, cybernetics was little known publicly, there was no internet, and I was a somewhat odd 15-year-old boy living in

Australia). I had to do it myself. I had to discover and develop the relevant algorithms on my own.

How was I to do this?

I decided to interview the top kids in the physics class in order to identify the problem-solving algorithms they used, and then install them in myself. A simple and powerful plan! What could go wrong?

I started to question the top students about the strategies they used to solve the problems that we had been set in our first physics test. However, I was surprised to find that they had very little conscious awareness of the methods they used. They could show me the formulas, calculations, and other workings they used to figure out the answer to a particular problem. However, they could not describe the abstract mental strategies that they had used to come up with the specific and concrete steps they took to solve a problem.

They could not identify any general problem-solving strategies they used, although they could set out the concrete steps they took, just like the teachers did in the worked-examples they used in their lessons. When questioned about why they took the particular approach they did in specific instances, they were unable to provide much insight.

It was not that they wanted to keep secret their problem-solving techniques. They were actually unaware of the strategies they used.

However, the interviews and interrogations were not completely useless. I picked up a few hints here and there.

My second strategy was to attempt to become aware of my own thought processes when I was confronted with a difficult physics problem. I intended to identify cases in which my thinking strategies worked and then reflect on why the strategies worked in those particular instances. My next step was to attempt to generalise the successful thought strategies as much as possible, in order to develop approaches that could be applied successfully to a wider range of problems.

When a strategy did not work in a particular instance, I would try to understand why it failed and then attempt to develop an improved strategy that did work. Whenever an amended strategy worked, I tried to identify why it did and how to improve it. Whenever I developed a more effective strategy, I would then try to generalise it further, and so on.

Eventually, I became increasingly able to reflect on the effectiveness of the meta-strategies that I was using to develop better problem-solving algorithms. As this occurred, I also began to work on systematically evaluating and improving the meta-algorithms. And so on, and so on, at all levels, recursively self-improving my problem-solving capabilities.

At first, I spent a lot of time looking inward in my attempts to identify the problem-solving strategies and thought processes that I was using. However, I did not see much initially. I could see why my interviews with the top students did not get me very far. But gradually, my problem-solving thought processes emerged out of the mist.

Increasingly, I became consciously aware of them, and could consciously amend them and try out different approaches mentally, in my head. The more I succeeded in making some of these thought processes conscious, and the more I started playing with ways to amend and improve them, the better I became at being aware of my problem-solving algorithms.

It is difficult to put into words how I experienced this developmental process. But at the time, I remember a metaphor rattling around in my head that somewhat captured the actual experience of what I was doing. It was as if I was reducing the size of the gap that unconscious intuition would have to spark across in order to solve a problem. My conscious development of problem-solving strategies was reducing the need for unconscious insights that were beyond my control.

Ultimately, my goal was to maximize the effectiveness of my conscious problem-solving algorithms and minimize or even eliminate the role of intuition and unconscious insight.

This approach produced positive results fairly quickly. Soon, I started to get marks that put me near the top of the physics class. And I achieved this with less work. Before long, I stopped doing homework insofar as it required me to write down the solutions to problems and show how I worked them out. Instead, I would restrict myself to looking at each problem set for homework, working out in my head an appropriate problem-solving strategy for it, checking mentally that it would work, and moving on if it did.

In the increasingly rare instances in which my accumulated problem-solving strategies failed, I would try to develop a new strategy, test it, and then amend my accumulated algorithms as necessary so that they could now solve the new kind of problem.

Homework was a breeze, and I did not need to pay much attention during lessons. I was a happy little chappy, sitting up the back of the physics class, talking covertly to my mates about things that really mattered to us in our lives then, not physics.

Soon, I extended the meta-thinking approach that I had developed for physics to all other subjects at school. Maths and chemistry did not put up much resistance. I also started doing better at English, but there were things required in English exams that were still largely impervious to my strategies. For English, I could not sufficiently reduce the size of the gap that the spark needed to cross. Intuition and insights were still needed to bridge the gap.

In particular, I was unable to devise algorithms that could fully produce the creative essays required for English exams. It was not until nearly forty years later that I realized that this could be done more or less comprehensively and that James Joyce and some other literary giants had, in fact, accomplished this to varying degrees. I discovered that James Joyce, in particular, seemed to have been propelled in this direction by the same underlying condition that probably played a major role in this aspect of my own cognitive development.

This left zoology. Unlike physics and chemistry, the teaching of zoology in Queensland had not yet been radically updated. The year 11 and 12 zoology courses did not even attempt to teach the problem-solving approaches and theoretical understandings that are at the center of most scientific advances in the real world. Instead, zoology still focused on the memorization of facts about zoology.

However, I had an exceptional memory. I also had plenty of spare time because I did not need to do much work on the other subjects and because my parents refused to buy a TV set because it would distract their children from homework and other things that were more important to my parents. So I swallowed the zoology textbook, and even now, I would still be able to pass a thorough exam about its contents easily. In addition, this memorisation process was greatly aided by my development of mental strategies that made memorization easier. Central to this was developing schema and connections that linked up the disparate facts, thereby significantly reducing the information content that needed to be memorized, and making it easier to retrieve relevant facts from memory.

But there was a price to pay for refraining from actually working out the answers to math, physics and chemistry problems. It hit home three years later when I was taking a university end-of-year physics exam. In examinations, I had to produce actual written workings for the problems I solved. I could not just check that I had the mental strategies that would solve the problem and then move on to the next one.

I froze in this exam when I found that I was having difficulty remembering some of the basic multiplication tables that I needed to use to work out the actual answers to the problems. This memory deficiency also extended to complex calculations such as long division. I was out of practice. I had hardly done any calculations for years.

I wondered why the examiners were asking physics questions that required complex computations when their job was to test whether students understood the relevant physics, not whether they could do basic maths. Surely?

But I got stuck on the details of a number of problems. I could not complete the calculations. My solution was to write down the procedures that I would use to solve each problem, step by step. But I did not (could not) actually apply those steps in practice. I explained why I did this in a note to the examiners in my exam paper. I wrote that I had demonstrated mastery of the relevant physics, if not of basic computation, and that surely that qualified me for full marks.

Apparently, they agreed and gave me a High Distinction for the exam. But after that, I practised my multiplication tables and other computational basics before each exam.

By the second half of year 11 and throughout year 12, I was ranked first in the school academically out of around 200 students. In the previous years, I was ranked around 15th.

In the final physics exam covering years 11 and 12, I got 92 percent, while the second student got 81 percent. Interestingly, both physics teachers also did the exam. Apparently, they wanted to provide some kind of independent test of the difficulty of the new PSSC course and of the exam they had set. The head physics teacher got 68 percent, and the other got 39 percent.

I mentioned earlier that initially, I was extremely disappointed with my teachers for failing to teach us problem-solving algorithms directly rather than taking us through worked examples and getting us to read the textbook. But their exam results underlined that their failure to do so was not due to laziness or ignorance. It was not their fault. They simply lacked the capability to do so.

It was not just that they lacked the ability to be conscious of their own problem-solving mental processes so that they could impart them directly to their students. As demonstrated by their results in the physics test, they also lacked effective problem-solving techniques. They were at least two vertical cognitive levels below where they needed to be if they were to teach algorithms directly and effectively.

Even if they had possessed highly developed problem-solving abilities, that would not have been enough. In order to teach students the algorithms that underlay their own abilities, they would need to be conscious of them. They would need to be able to ‘see’ their own problem-solving strategies as objects in their awareness.

Without this, they would have little capacity to describe to their students how they went about solving particular problems. Until they reached this level, their problem-solving methods and algorithms would not exist for them consciously and could not be used consciously to teach. Nor were the teachers capable of developing with their students a commonly-understood set of words and concepts that could be used to describe to students their problem-solving strategies.

The teachers were no more able to look inside their own minds and ‘see’ their own problem-solving algorithms than dogs are able to ‘see’ things outside their visual field, or Maggie Thatcher was able to ‘see’ a society and other large-scale systems, or First Enlightenment scientists were able to ‘see’ complex phenomena that cannot be captured by reductionist, mechanistic thinking. However, none of these individuals are responsible for their particular blindness.

At the time, I was not very charitable to my teachers. But they were acting in the only way they could, given their level of development. I now see that berating them for their failings would have been as senseless as abusing a dog for failing to be able to walk around upright on its hind legs.

Furthermore, if my teachers had been able to teach algorithms directly, I would have been deprived of a powerful motivation to look inside my own mind, develop the ability to see my own thought processes, and see how I could improve them recursively. I might never have been propelled along the path of recursive self-improvement.

Schooling in Queensland ended at the completion of year 12 with public examinations covering what had been studied in years 11 and 12. These public examinations were the gateway to university. Over 10,000 students sat for the exams in 1979, and I was ranked fourth overall. This was even though the ranking system compulsorily included the student's marks in English. As I indicated earlier, my meta-thinking techniques had not worked as well for English. On the scale of 1 to 7 that was used to indicate how well a student performed in individual subjects, I got a 6 in English and a 7 in each of the other subjects.

The events that surrounded my poor performance in a physics test at the beginning of year 11 launched me on a path of conscious, recursive self-improvement that has continued throughout my life.

Central to this path is the attitude that the individual is a work-in-progress that can be continually remade consciously and intentionally. Each enhancement of intelligence can be used to enhance further one’s capacities, including by enhancing one’s capacity to improve one’s abilities, and so on.

My final years at school were just the beginning.

Try out this approach on some challenges that are of vital interest to you. Do it recursively. Extend it to other classes of challenges. And each time you encounter a problem that cannot be dealt with effectively, celebrate because it is an opportunity to discover how to adapt your strategies so that they can now solve wider classes of challenges.

You may not find it easy to make conscious what it is that you currently do intuitively and unconsciously, and then to enhance these processes consciously. However, as the book proceeds, I will set out specific techniques and practices that will enable you to scaffold these capacities.