

Imagination – a data science perspective

Dear reader, I do not know what you know or whether you are familiar with the concepts I am about to juggle before you. Consequently I feel the need to initially lay some basic foundations in your minds from which I can build a data science perspective of the topic, so please bear with me while I set you up to be shepherded, rather brusquely, through a pipe of my own bending. I hope you enjoy the ride.

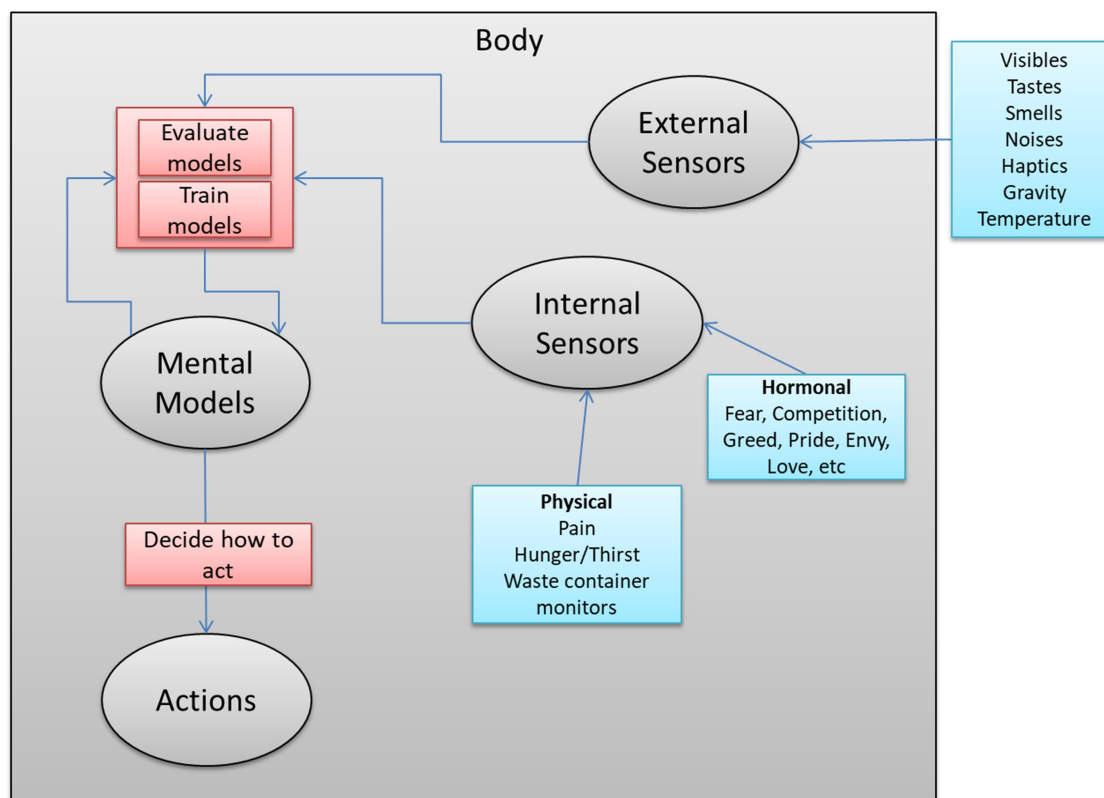
Context

The human machine has capabilities determined by its genes, one of which is the ability to create mental models of the way the world works, which help us navigate life.

At birth our models seem empty, trained simply to yell when unidentified inputs appear. Inputs arrive from our bodily sensors and we try to make sense of this incoming onslaught by building models that identify, categorise and organise seemingly disconnected sensory streams and events into connected relationships of cause and effect and situational contexts. These are things we can recognise and learn to interact with.

As we develop, we train basic models that allow us to differentiate hunger from a need to poop, crawl, catch a ball, identify peers and label them, smile and scowl appropriately. In our early years we must test our models constantly through trial and error and retrain them to be more accurate. We experiment with basic leg movements and vocal gymnastics and rely on feedback from our surroundings and peer group actors to determine whether our updated models are improvements or failures.

In short we embody functionality that can be roughly represented as

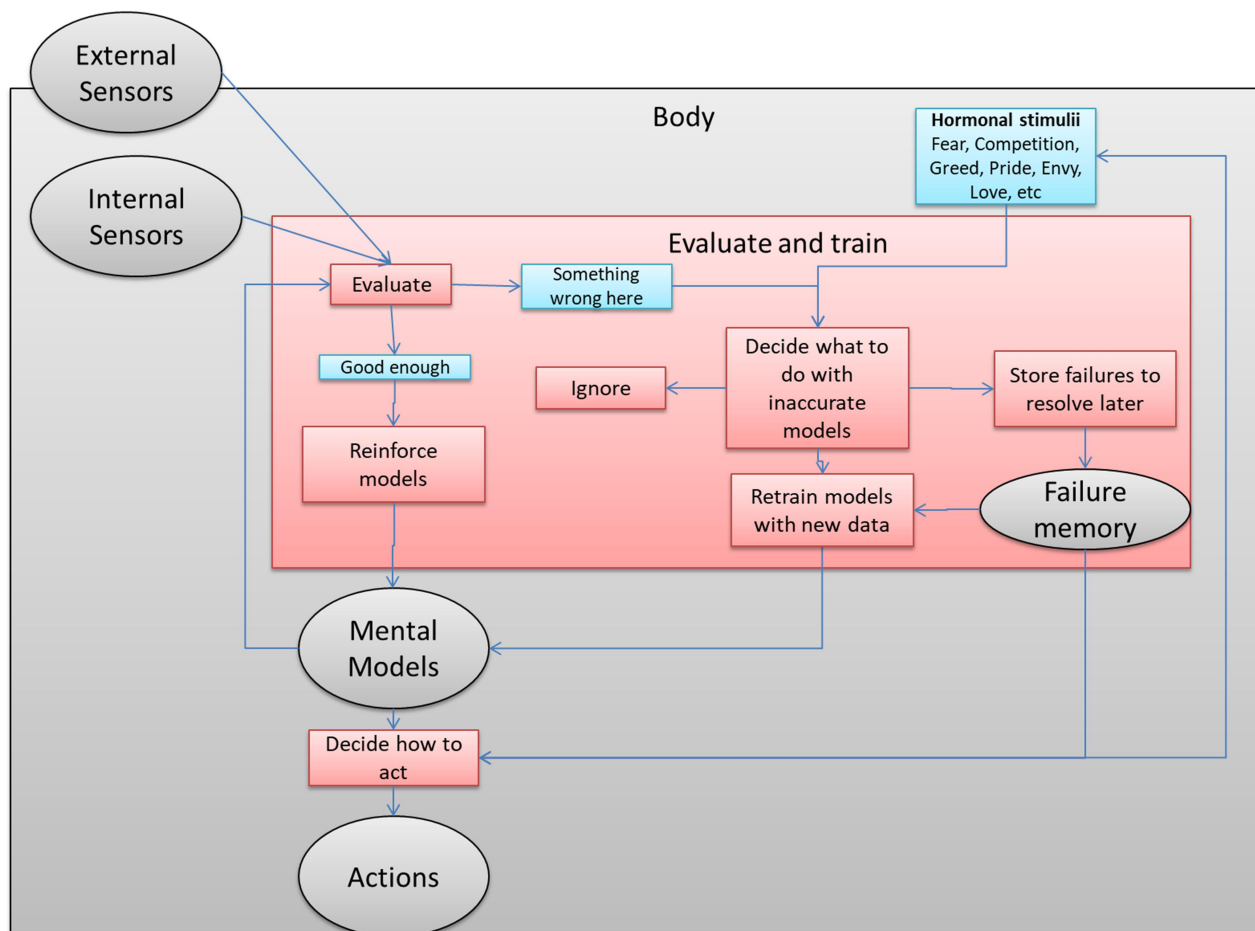


Our models are initially trained using sparse data, and they need to be tested and rapidly retrained as our experience of the world enlarges. Our training data comes from a mixture of direct sensory experience and indirect learning from what we are told, read and hear. Indirect data may initially be assigned a lower trust score that needs to be proven through direct experience to gain validity. If we have no experience of something we have been told, then it's 'weight' is determined by the quality and trust we assign to its source. Initially we might trust those who raise us and those who appear to have 'authority'. Our peers also provide simulated data, which might appear to be directly experienced through our sensors but is delivered in modern times by video, film and radio and historically by plays and pictures. Simulated data may have more weight than indirect data as it is received through different mechanisms.

We build observational models to identify and categorise actors, things, places and events, which we then use to connect together into contextual or situational awareness models. By recognising repetitive situations we can bring into play our automatic response mechanisms that can use previously trained models to respond successfully to events without having to think about it.

As we experience life, we assess our models by comparing predicted outcomes against real outcomes and then reinforce our models if they work 'just well enough', and if they appear faulty we can adjust and retrain them (and the weight given to our sources) or in a fit of pique ignore any conflicting result.

The above can be roughly represented functionally as



We have evolved to be efficient machines and as such we do not need the details of exactly how things work, just that they do. In short, most models need to be 'just good enough' to be of use in our competitive environment. We only seek out next level of detail when we have to, or through our natural curiosity.

Throughout our lives, but especially during our early stages, our contextual models are often disconnected, have dead ends, have vague shadows of unproven connections, are connected by a host of badly interpreted experiences or a weight of false indirect data and are utterly unique to each of us. Many of us have trained similar models due to common inputs from peers, education and the environment, and we may be able to use them to agree at a generic level, but often disagree when we drop down into detail or test wider contexts. We often disagree because our models are never equal, just similar.

I am not sure these briefest of glimpses into the world of machine learning are enough for my purposes, but let's dive into the topic and find out.

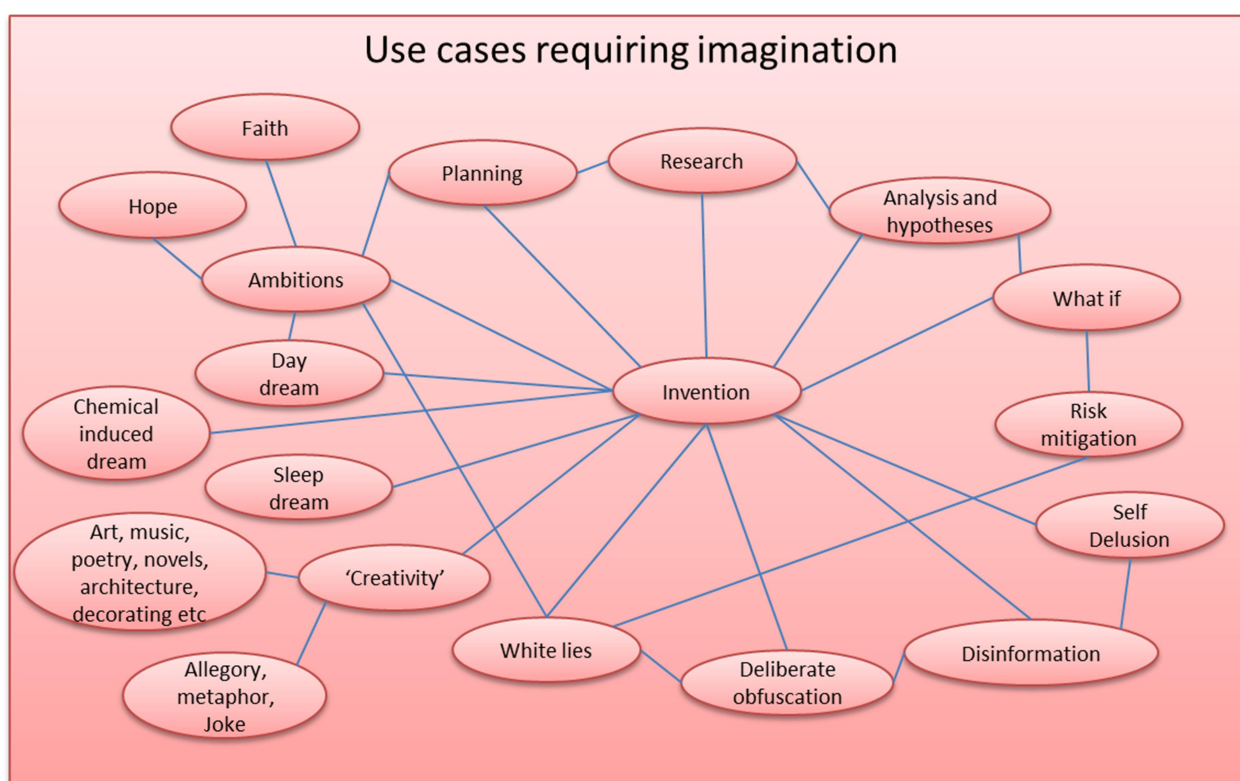
Main

What part does imagination play in all this?

We have several tools that help us to develop our models.

The first is 'Why?' which we deploy at an early age to assist the transfer of trained models from parents and peers to ourselves, and to discover the limits of their situational awareness and detail. As the machinery available to parents and other supposedly knowledgeable folk has evolved through competition, hormonal chemistry may cause us, when confronted with a 'why?' that our models cannot answer with any degree of confidence, to make things up. After all, knowledge is power, in the land of the blind the one eyed man is king, and maintaining an authoritative and trusted position gives an edge... but I digress.

'Make things up' surely requires functionality that can be placed firmly into the Imagination category. We have developed many words and terms for use cases that require imagination and I am sure the items in the following diagram will be present in many discussions. It is worth initially lumping some into the box below with the aim of navigating through them to tease apart differences and commonalities in their implementation.



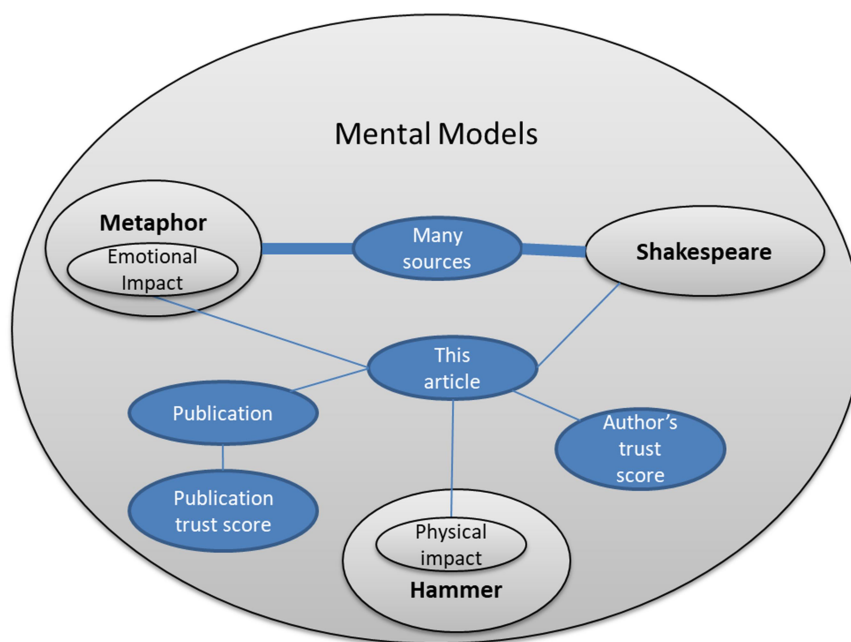
It might also be useful to lay out a few primary triggers that make use of these use cases. Triggers such as curiosity, education, fear, competition, necessity, self protection and self promotion.

More entries could be added to the above, and the clever could categorise and better connect them, but this will 'do well enough' for a discussion on the central item required by all these use cases, which is to invent something that does not yet exist in any of our mental models.

Let us investigate a few and try to identify some patterns. Let us start with **Metaphor** (bottom left) for its simplicity. Rather than reiterate a well-known example I am going to generate one, which on reflection seems more than a little strained, but let's run with it. We are going to examine the great Shakespearian hammer 'Create a **Metaphor**'. Most readers will have already trained a model to identify a hammer and understand the context in which it is used. Most readers will have trained a

model to identify the label ‘Shakespeare’ and associate it to texts that skilfully use metaphors to generate emotional impact. All we are doing is joining two incongruous models that have at some level of abstraction a common essence, and a reliance that we have all been trained on similar data to be able to make a similar connection in another’s machine through the process of reading words. Simple, but effective. Well, simple in that an originator has to distil models to identify common attributes, of which a hammer’s physical impact related to the emotional impact of a good metaphor is a poor example, but it is just-good-enough to illuminate better examples such as ‘winter of discontent’.

In fact you might have now created in your own mind something that looks roughly like



At this stage it is worth considering what happens to this incongruous connection. Does it become a permanent pathway? Does the weight we assign to it reflect the flippancy of the cause? Can we ever truly forget it? It seems to me that repetition reinforces connections however they were formed, and unused pathways become overgrown, hidden and rarely navigated, but never truly disappear. ‘Motorways’ and ‘byways’ may have more resilience than the example above. Enhancing the mental image by connecting additional models along with a dash of emotion can also increase the likelihood of retention. The hammer could be Thor’s, wielded by a wide eyed, tousle haired Shakespeare, smashing a metaphor into his audience. Shock and awe, and all in black and white copper plate suitable for the age. Is this pathway now better lit? More memorable?

Pointing out abstract connections between things can explain several use cases on our previous map. Metaphors used as embellishments can invoke stronger emotions in readers and listeners of our tall tales, an incongruity can generate energy in us that we express in many ways. Comedians, artists, writers, playwrights all use this artifice.

Fiction however requires something extra in the invention department. Creating similarity relationships between more complex situational models can result in allegorical stories such as those used in aspirational advertising and marketing, religion and education, but what does it take to invent a story with characters and plot? Let me introduce Frederick. I imagine that you have just

created a mental model to represent this actor, and assigned it a label, a sex and a loose connection to other models labelled Frederick that may be in your store of models. You may even have assigned a culture, and generated a candidate image. Now watch how your model morphs as I embellish with 'The great'. Have you assigned your model to a historical timeframe or adjusted your image of the actor whose peers consider him Great? There is much we do not know about Frederick, but every extra detail and interaction enriches our model of him. This is exactly the same process whether the actor is real, who we have come across on the news or from an email, or is fabricated by fiction. Actually seeing a representation of Frederick on a billboard or video, or being introduced to the real Frederick seeds our model with different and richer criteria, but it is the same process, and all of us can convey a model of an actor in words or graphics that recipients build in their own repositories. Each of our models of Frederick is unique and depends on features drawn from the existing models in our own repository. As a reader of the same fictional book we would be subject to the same layers of training data concerning Frederick, and our models should gradually become more similar as further detail is conveyed. Now I want you to destroy the model I just caused you to create by adding 'Frederick the Great is a cat'. Is there some annoyance with this deceit? Did you predict it? Was some emotion generated by leading you down this path, or did you seamlessly adapt by simply swapping out your image of him for a furrer one?

Of course I cannot see you, I do not even know who will read this, and in order to come up with my textual fiction I have had to do several things.

- I had to invent an actor. It did not really matter who the actor was so I just dipped into my own models and randomly picked some data to seed a model of an actor.
- I wanted you to observe your own functionality at work so I **planned** to make you embellish your model, then destroy it and replace it with a completely new version. In order to do this I had to invent a course of sequenced events that I predicted would lead you to a desired outcome.
- Without feedback, I have no idea if this will be successful, so for me it is a predictive model to which I have assigned a strong success score, but also attached some doubt as to its effect in all readers.

Can all 'Creativity' use cases (**Art, music, poetry, novels, architecture, decorating etc.**) be described as the ability to shape the models of readers, viewers and listeners by first building a model in an author's own mind, then conveying that model by controlling the inputs seen, heard and felt in recipients?

To be truly creative, there must be some understanding of the emotions caused by the media used to convey a model. As creatives develop their skillset models, they may learn to vary the outputs to achieve different strengths of emotion in recipients, by acting in slightly different ways to shape the artefacts conveyed. Some evocative artefacts may be the result of random combinations of existing model components but these accidental creations fall outside our topic. Others are variations of things known to work, that just require some model re-organisation or enrichment to result in 'new' plagiarisms. Creating a desired outcome in a recipient can be difficult using a single image, music or shape, the outcome relying on a recipients 'interpretation' to achieve the closeness to the model imagined. Creatives often use their chosen media to evoke a model essence and rely on the recipient to embellish it themselves by navigating their own motorways and byways. The difference between

abstract and detailed art is in the level of participation required from a recipient. Greater information given to a recipient can result in a more accurate recipient model.

Some effects seem to be formed at the very edge of expression like a Jazz artist, orator or racing driver whose next action swerves from a vague plan to react to current inputs. These 'swerves' seem not to require prior imagination. Actions here seem to be governed by a different mechanism, something more instinctive, something that can act fast without having to build a mental model, but can still forge new connections that present themselves in the moment. In these cases there seems to be collaboration between an imagined plan that gives structure to events, then hands over to an implementation process to add in the detail, the embellishment, the ornamentation.

Situational or contextual models generally have predicted or desired outcomes. If a predictive model is to be communicated, then we can guide a predicted outcome to the recipient. This capability governs much of our life. I can tell a child what might happen if he hits his hand with a hammer. He does not have to smash his finger, he can imagine the conveyed predicted outcome, although if he does hit his finger, the model should be reinforced. When we do come across a situation for which we have a predictive model, whose risk trigger we did not recognise in time, we tend to express the resulting energy in regret, self criticism, embarrassment and/or a strengthened likelihood to use the model in future, but I digress.

All education and all religions rely on conveying predictive models to recipients. Ethics, morals, knowledge are all predictive models strengthened and proven, or weakened and disproved, by experience. All societies use them to build foundational models in emerging generations in an attempt to shape behaviour and build the tools that govern social collaboration and manage raw competition. That last statement may appear to be too large a step to build from the meagre components so far supplied. So let us continue our tour of use cases requiring imagination and see if we can add the detail to your models that bring them closer to my own versions.

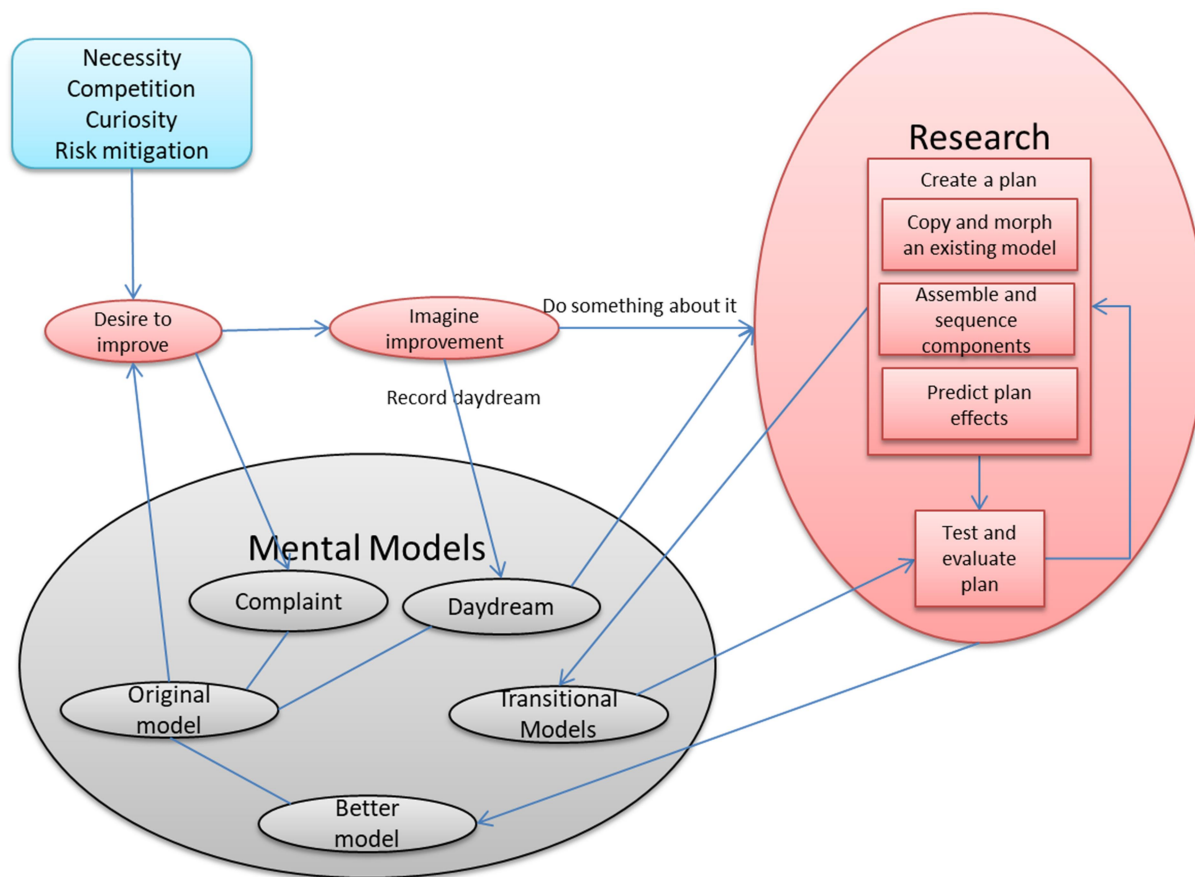
Deliberate **lying**, **disinformation** and **obfuscation** all use the same components. I have imagined the inputs that you require to build a model to generate my predicted result, which I must now convey to you in enough detail and with enough authority to have you trust it. This entire document with its concoction of words and diagrams could be one big whopper of a lie, a figment of my own imagination with no relation to reality that I have generated to trouser a prize, and as the author is an unknown entity to all of you, I have predicted that a small proportion of you may assess it as such.

More of you may conclude that this is an example of **Self delusion**, a use case also constructed from the same functional components. Our models only represent what we have constructed for ourselves or been guided to construct. Tenuous relationships created from naïve analysis of cause and effect, trusting the models of others, retaining just-good-enough models without the effort of research and improvement, lack of experience to prove or disprove a model are all causes of our collection of half-baked models that we regard as today's 'I' (or at least one of our I's – but this is also out of scope). From birth, our models race from zero till we can act in the world in a just-good-enough predictable way and this seems to be driven by an instinctive reward for building and confirming our models, being right in our predictions contributes to the sheer joy of learning. We can never learn everything about everything and at all stages of our lives we have much which is dodgy. All of us have dead ends, erroneous, vague and unproven connections, and clusters of

specialisms where detailed proven models sit. A few constantly test their models and seek improvement, but humans have always truly believed in untruths. Religions have come and gone. Stories of what is over the horizon have been imagined and dispelled. Theories of geology, astronomy, philosophy and biology are in a constant state of proposal, testing, improvement and replacement. Humans have always truly believed in a mix of truth and utter gibberish. The truth of our beliefs is of no importance except they help us navigate life and provide springboards for the next steps in our collective evolution.

Research, analyses, hypotheses, curiosity are also on our map of use-cases-requiring-imagination, and these few labels represent an entire category that add up to the ability to find out how things work so we can exploit them in our role as a competitive species. Faced with a need to cook some meat, a living person may just strike a match to light a fire, having learned a model from their peers that predicts reliable fire from this activity. They do not have to understand the chemistry involved. It is just-good-enough to have modelled the activity of striking and lighting. Somewhere in the past an individual became bored with the unreliability of rubbing sticks together and had to imagine a better way, an aim, a vision of a brighter and easier future. How to get there? There is already an existing model that can be improved, so gaining an understanding of how the current technology works then researching different friction causing materials, and different processes used to create the work that causes friction seems like a useful set of activities. Additionally, a background understanding of the properties of phosphorous and equating its essence to that of fire creation, probably aided the new invention. Hypotheses, research, testing, adjustment and iteration all require imagination at each stage to form plans of activities that achieve specific results. The result of multiple iterations of trial and partial success is perhaps a working prototype and an altered but achievable aim. Then there could be further rounds for productionisation and commercial exploitation. In this way, imagination has helped kick-start our entire industrial, agricultural and socially organised existence. Well, imagination and the desire to do something better or different. That initial desire is worth considering. While not strictly 'imagination' it is a trigger whose next step requires it, otherwise it just becomes complaint and revolution against that which is, rather than a search for improvement. This desire seems to be instinctive, built into our very core, which comes to the fore at specific stages in our development as we struggle to find our ways in life, then moderates as we gain control having built models that work just-well-enough. It seems to be stimulated by competition whether commercial, sporting or cultural, and by necessity (being the mother of invention). Curiosity and the rewarding light of discovery also seem to drive some to continually test the outer edges of what we consider possible. But as these are triggers which start processes that require imagination, I digress.

A diagram of the above functionality looks roughly like



Other use cases requiring imagination include our legal and legislative tools of social agreements that are the foundation of all societies. Law and contracts with their interminable clauses that aim to clarify vagueness, mitigate risk, gain trust, or gain an advantage through obfuscation, all require imagination to predict possible divergence from a desired outcome. To imagine divergence our machine runs predictive models varying inputs and assessing probable results against our acceptance criteria, often inventing fictional actors to play adversarial roles. The machine then invents clauses, paragraphs, sub clauses and codicils to ensure compliance or manage non-compliant behaviour in contract parties. This requires the same functionality as in the process above.

‘Agreement’ is itself an imaginary concept. An agreement is something we willingly enter into, are born into, are forced into, or accidentally fall into. We imagine we are bound by an agreement, and we should look at what binds us to the roles we agree to play and what can cause breakage. Are we obligated to play our roles through simple forces such as risk aversion and the joy of reward, or is there something else? We have been trained to play roles on many types of agreement such as employment contracts, familial roles such as parent, marriage, citizenship, financial loans, and even the concept of ownership. These are all imaginary concepts that we try to realise and sustain through contracts and their obligations.

Looking wider, we have been stuffed with imaginary concepts. Nation, culture, money, debt, border, government are all entities that only exist because we agree that they do. Even basic morality and ethics are agreement models that require us to conform to rules of behaviour in order to be part of a ‘civilised’ collective. Common behaviours reinforce social agreements. They enable trust, cultivate a sense of ‘We’ in a competitive world and promise the reward of collaborative strength while

retaining a tension with our individualism that allows an acceptable degree of freedom. Civilisation is a thin veil covering our individual competitiveness that is fabricated from our imaginary concepts and an instinct that has proved collaboration and brotherhood are useful tools. Our judicial services spend much effort monitoring, investigating, prosecuting and separating those who break our social agreements. The evaluation of models that predict reward and risk of an action can be the cause of much corruption, fraud, theft and worse. If I predict the likelihood of being caught or punished is low, or I predict that if I do not do something my peer group will punish me, and if I predict that my reward is of high enough value, then I might be tempted to break obligations and form new ones. Of course our machines also have endocrine systems which can interfere with model evaluation or at least the actions we take, but I digress.

We have developed so many uses for imagination.

All our organisations are imaginary. Companies, institutions, nations, communities exist as imagined entities only because we agree they do. Our governments negotiate treaties and trade agreements, and develop defensive and offensive capabilities to mitigate collectively imagined risks. Economists, actuaries and legislative representatives predict our behaviours in this ever changing soup of life and death then adjust agreements that we have been trained to truly believe we are party to.

So many conflicts arise from differences in our models of nation and culture, their ties to geography and the obligations that surround citizenship, border, treason and heresy. Many conflicts arise from applying inherited models of group victimisation and the conveyed memory of historical activities to present day behaviour. Our economic concepts of capitalism, communism and socialism, which each of us has modelled slightly differently, along with GDP, interest, currency exchange and many other terms cause trade wars, tariffs, protectionism and prevent weaker nations from developing by miring them in our agreed concept of national debt.

The historical human collective has dreamt up all these things, conveyed, refined and altered them over generations into the models we hold in our bio-machines today. They are all examples of higher level use cases that require imagination, but back to our smaller map.

Most of the use cases requiring imagination covered so far require a certain purpose. I was going to write 'require consciousness' but many of them can still happen in an automatic state, so 'purpose' will have to do. There are however more random seeming activities such as **day dream, sleep dream** and **substance induced ramblings** that we should also examine. If we embody functionality that allows us to assemble actors and assign features to our fictional characters, a more random process could use the same mechanism, but whereas the features of fictional characters are assembled according to some plan we have of what we want to create, a random process has to select features using different criteria. The first thing to look at is whether these features already exist in our store of known things and we just assemble them into dream actors or do we actually invent something we have no prior knowledge of? Let us go back to fiction for a minute and consider the genres of science fiction, horror, religions and cartoons, all of which at some stage of their description require actors, places, capabilities and events that are highly unlikely to be experienced. Now consider how many are assembled from features we have experienced, just connected in a different way. Martial artists and angels fly like birds, scaled up scorpions with pterodactyl wings pursue us through caves, heaven is all cloudy and white and warm and welcoming and we had to wait till our understanding of astronomy had developed before creating interplanetary fiction. Limited initially to sun kings and

moon goddesses, but now running to intergalactic instant travel. I would say that most of these are assemblies of the known or extrapolations of the known. What about time travel? Can we experience it? No. Yet we model sequences of events to understand past and future, model our history and plan our activities. We have the concept of time, and travel, so it is not so hard for a fiction writer to stir events into a zig-zag timeline to achieve time travel and predict it can be understood by a recipient. So do we actually invent anything new at all? I suggest not. I suggest that we can only assemble features and concepts that we already have access to, and we get those through our senses from the activities and peers around us, as well as those conveyed to us.

Back to the question of what guides the assembly of components into a dream, or hallucination.

Are dreams a random mash up of temporary fictional actors, mixed with models of actors we have real experience of, interacting in a sequence of events generated from our situational awareness models? We can assign all manner of behaviour to our fictional actors, and these can generate surprising dream sequences, but is there more to it? Here we run into a number of concepts. The first is 'sleep on it' problem solving, which suggests it is possible to set up a topic for our machine to work on during sleep as an automatic improvement process. Then there is lucid dreaming, which suggests we can inject a degree of consciousness into the process of sleep to guide model generation. Thirdly, we harbour a pocket industry of dream interpretation which suggests that in our sleep, much like abstract artists, our machine distils essence from our models and creates its own metaphorical actors and allegorical stories. These all point to the ability to queue up topics for model training and evaluation while the body is otherwise engaged. Stress may be something that prioritises that queue. Our first diagram showed that we test models and if they fail we have to retrain them, but we do not have to do that immediately. Test results seem to be stored either as successes that strengthen a model or as a failure. In this world of models, stress appears to be a badly performing model that we have not been able to improve due to various factors, and if the model has been assigned high importance, repetitive failures can result in hormonal activity that heightens the need to improve it, consciously or unconsciously. In short, model clean up and improvement processes run during some stages of sleep, as do model similarity comparison and deduplication. We maintain a prioritised queue of models for refinement, and some of us can inject a consciousness wakeup event into the process.

What about psychotropic ramblings? I am not experienced in their use so can offer no personal experience. You will have to judge for yourselves, but I bet you, not with anything I value, that they supercharge the processes we have described above.

Recap

We have covered most of the use cases requiring imagination on our original map, and have discovered more on the way. We have identified a core set of system components and processes that can support them all. Everything we have discussed has been described as a collection of models, which are trained, evaluated, tested and retrained, connected together with scored relationships then assembled into situational awareness models that can be navigated to deliver predictions and control our actions. Action control is outside our topic.

This playbook is straight from the world of machine learning and Artificial intelligence, not the probabilistic word salads nor the generative models trained on sounds or graphics, but AGI, or Artificial General Intelligence, the hard one, the one we have yet to build, the one where machines can model the world like we do but lack both the endocrine system that layers hormonal activity on top, and the slightly error prone mechanisms of gene expression and replication that give rise to a truly adaptable and evolutionary system. They would also lack the consciousness thing, but that is another story.

I have declared that imagination and imaginary concepts such as social organisations, agreements and obligations underpin our civilisations and help us compete as both collaborative groups and individuals, and I have declared that inherited imaginary artefacts are the cause of much human conflict and social strife.

I could declare that Global societies could be totally reconfigured if we did not teach existing concepts to our children and spent some effort redefining those which are useful and binning the rest, but that would require complete control of the education system. Some societies already attempt this to varying degrees in order to maintain their existing governance systems.

However, this is evolution at work, and the capability of imagination has evolved with us. It helps us learn and adapt to change, it helps us conceive abstract entities and be bound into collaborative groups, it helps us create visions of the future and create plans to try and achieve them. As our evolutionary journey continues, I feel we might have to learn to manage our set of inheritable imaginary artefacts better than we do now and globally agree some that are common, but that has to start with educating ourselves better in the way it all works.

Subtext

The above is a snapshot of my own models trained from a lifetime observing self and others, while attempting to maintain a persistent process of model assessment and retraining. The current versions seem to me to be relatively robust. For most readers this article is likely to represent indirect learning that needs to be proven, a few may use it to reinforce data they already have, and I expect arguments (whether internal or expressed) from those who know better!

I should also add that I have no idea how our biochemistry works to provide our capability of imagination, I just know that it does, and its functionality looks like the above.