

Augmenting Tacit Awareness

Accepting our responsibility for how we shape our tools

In 1816, Christian Jurgensen Thomsen, the first curator of the National Museum of Antiquities in Denmark, put the prehistoric artefacts in his collection into one of three technological eras, the Stone Age, the Bronze Age, and the Iron Age. In 1865 Sir John Lubbock made a further division between the old stone age (Palaeolithic) and the new stone age (Neolithic). After the Second World War, some suggested that we call our era the Atomic Age. These days we are more likely to encounter the designation Information Age. In his reflections upon information technology, the Canadian philosopher Marshall McLuhan points out that we not only shape our tools, our tools shape us. From an information technology perspective human history is sometimes divided into five eras. The first, which on genetic and cultural evidence some date to be no older than the emergence of modern humans about 150,000 years ago, was the creation of spoken language. The next (accompanying the first human civilizations) is the invention of writing. The third era is connected with the invention of the alphabet. The modern era is associated with the printing press. In our time, information technology is shaped by personal computing. Many different calculating devices have been invented. A metal object discovered in an ancient shipwreck in 1901 has been described as the first known analogue computer. Described as the Antikythera Mechanism, its thirty cogs tracked the movement of the planets and constellations, and predicted eclipses and equinoxes. In 1645 Pascal completed a device that he had begun building when he was 16. It was designed to spare us the drudgery of calculation. In 1672 Leibniz presented an improved version to the Institute of Science in Paris and to the Royal Society in London. In 1819 Charles Babbage outlined his plans for a “Difference Engine” which he intended to generate mathematical tables.

Either side of the entrance to the Information Age loom two giants, Alan Turing and Vannevar Bush. In 1936 Turing wrote a paper in which he demonstrated the theoretical possibility of creating digital “Universal Machines”. In 1945 Vannevar Bush wrote a paper in which he speculated about the possibility of creating a “Memex” machine that would help us to link between records. The first vision was called Artificial Intelligence by Marvin Minsky - AI, and it seeks to mechanize any reasoning process that can be replicated by an algorithm, replacing humans. The second vision, outlined by Engelbart, has been called Intellect Augmentation - IA, the quest to bootstrap and co-evolve a man-machine symbiosis in order to enhance our knowing. The information technology which brings about the Information Age is personal computing. In his “Mother of All Demos” in San Francisco in 1968 Douglas Engelbart demonstrated various ways in which a computer, designed for personal and collaborative use, can facilitate, not only the means but also our responsibility for, the way we shape our understanding of the world. In retrospect it was an event of civilizational level importance. He used digital machines (as envisaged by Turing) to create machines (as envisaged by Bush) which give us tools that augment our ability to solve problems. Engelbart emphasised the

importance which the framing and conceptualisation of a problem has upon our ability to solve it. He envisaged humans and computers participating in a co-evolutionary emergence which shapes tools for thoughts whose ultimate end is not information control but human flourishing. He was influenced in this task by the writings of the Polish philosopher Alfred Korzybski, who in his 1933 book *Science and Sanity* traced the insanity of the C20th back to our failure to understand the role which time binding plays in human flourishing..

In his concept of ‘time binding’ (the notion that the human emerges via a process of cultural accumulation) Korzybski advocates a Non-Aristotelian epistemology which he calls ‘General Semantics’. In his focus on problems that are caused by shortcomings in the conceptual frameworks we live by, Korzybski’s approach is commensurate with Michael Polanyi’s Post-Critical philosophy. Their approaches also resonate with the philosopher Robert Pirsig’s conception of Quality, the Integral Philosophy of Ken Wilbur, and Maria Popova’s ‘labour of love’ *Brain Pickings*, in which she declares that we need to re-learn the spiritual truths our parents have failed to transmit to us. Instead of Engelbart’s ‘conceptual problem’ the AI research programme issues dictums such as: "code wins argument", "move fast, and break things" and "try first, regret at leisure" - in short shaping AI tools irresponsibly. Trying things first and understanding them later is elevated into an "epistemology for deep learning" [The Epistemology of Deep Learning - Yann LeCun - YouTube](#) But if the dictum ‘code wins arguments’ is correct, we would now all be running Smalltalk. As Alan Kay (who at Xerox Parc declared that you can predict the future by helping to invent it) demonstrated in a [talk](#) in 2012, the Smalltalk system had all its capabilities available via hyperlinks in the browser. Since everything was a true objects, every object was interoperable and fully tinkerable. If code wins arguments this was the vision which should have been the future. Instead, on the basis of his glimpse of what was going on at Xerox Parc, Steve Jobs made it a commercial success with Apple. The new personal computing, despite a clear understanding of the PC’s potential to be a bicycle for the mind, ended up focusing more on shareholder value rather than human potential. We are now living with the consequences of that approach. Our systems are orders of magnitude more complicated and less usable. Despite the fact that we had a code which demonstrated that we can access tinkerable capabilities as easily as clicking on a hyperlink, that ‘argument’ has been lost, because vendors have locked us in with applications which sacrifice interoperability and tinkerability on the altar of being self-contained. HyperCard, which linked cards, along with Visicalc, one of the killer applications for Apple (which can be said to have made people buy into Apple technology in the first place), had tinkerability built into it, very much like the way it is done in MindCard. With Hypercard, users could hyperlink capabilities to serve their specific needs. Clearly, if users could create solutions for themselves, what’s in it for the system and application vendors?

Current AI revives the idea of ‘Reinventing Man’ from the 80s. When learning becomes synonymous with ‘machine learning’, and Google conceives itself as an AI first company, the level of centralisation in computing gets to the point where a wholly centrally directed society becomes a real and present danger, exhibiting all the characteristics that Polanyi identified in the centrally directed catastrophe which was the Soviet Union. Current computing embodies an epistemology which claims to know more than we can know, and it denies what we do know

but cannot explicitly justify. It is wedded to a hierarchy based [Cosmology for the Computer Universe](#), that is fixated on centralised control; with a Transhumanist promised land at the end of history replacing any belief in human potential. This approach not only damages the freedom of the individual to make new discoveries, it also generates a spiritual (or as Polanyi puts it moral) inversion in our understanding of what it is to be a human being. The current hype around AI promotes the idea that the human mind can be reduced to processes that can be replicated and run in ways that are akin to advanced deep learning algorithms in the cloud. In an AI vision of the brave new world all human expertise is replaced with machines, while the human struggle to cope with information overload is neglected. In a post-truth world our ability to make sense of the world is diminished, together with our capacity to assume responsibility for our own judgements, leading to enslavement.

It was Polanyi's scepticism in a 1949 symposium at Manchester University about the quest to wholly formalise human intelligence which prompted his friend and colleague Alan Turing to write his paper "Computing Machinery and Intelligence" (1950). Polanyi argued that Symbolic AI relies upon an approach which Kant designated as "Critical Philosophy" – that it is only the wholly explicit which can count as knowledge. He claims that although the emergence of the human becomes possible by indwelling within externalized symbolic forms, these symbols become meaningful as a consequence of us relying upon the ground supplied by our personal commitments and tacit awareness. We indwell within tools as extensions of our body. This approach was cited (for example by Mark Weiser) as an influence upon researchers at Xerox Parc - the place where the personal computer was invented. Personal computing exists within the context of a man-machine symbiosis, whose feedback loops augment our ability to index, make sense of, and share our experience. In the current hype, Machine Learning or Neural Net based AI is extrapolated to reach the Singularity, which is the inspiration for Transhumanism. We consider the aspiration that Transhumanism could bring about a new stage in the evolution of life as the product of a false, reductionist epistemology. What is required is not a devout reductionism, but an enhancement of the life of the spirit. Inspired by the insights of Engelbart and Polanyi, we see Transhumanism as an outgrowth of computing pop culture, a cargo cult characterised by lack of understanding and irresponsibility; whose failures will dehumanise and possibly even destroy humanity. We offer instead an augmented phenomenological vision in which computers function as an extelligence which helps us to map the gradients of our tacit awareness, by supplying us with tools which, via indwelling within our re-presentations ("vorstellung") facilitate our ability to manage and structure and communicate our experience. A personal computer has the capacity to augment our tacit awareness. By bootstrapping this awareness we can discover the connections, which Ted Nelson calls the "intertwingularity", that ground our quest for meaning.

The poet T.S. Eliot, who was familiar to Polanyi because they participated in a discussion group called *The Moot* during the Second World War, declared "We shall not cease from exploration, and the end of all our exploring will be to arrive where we started, and know the place for the first time". The place where we start from is our experience, and our exploration takes place, in part, through our articulations; which becomes meaningful to us via a return to our experience. Articulation changes our experience, which is why, because of our

explorations, we know it for the first time. Hegel, in his objective phenomenology, claimed that the spiritual purpose of existence is to know itself. He saw Critical Philosophy as akin to the wise resolution of Scholasticus to learn to swim before entering the water. In his Post-Critical approach Polanyi reminds us that even our most explicit forms of knowing rely upon our tacit awareness. We augment this tacit awareness by pursuing abstract ideals, which although they are foreshadowed in more primitive types of life, are unknown outside the articulation that renders them possible. AI relies upon a bet. It is the bet that if you get your syntax (mechanism) right the semantics (meaning) will take care of itself. It is the hope that if computer engineers get the learning feedback process right, a new transhuman intellect will emerge. But a move fast and break things approach has potentially apocalyptic consequences. It is better if we seek to augment our intellect rather than trying to re-invent man in the hope that a mind will emerge. What rough beast, as the poet Yeats put it, is slouching towards Bethlehem to be born? We hope that humans would win the 'Butlerian Jihad' against thinking-machines prophesied by Frank Herbert in his 1965 science fiction novel *Dune*. But it would be better still to reject the path that leads to it. We should use the computer as a tool with the potential to enrich our lives, not create a spiritual wasteland by elevating machines into our masters. Instead of Transhumanism we should work to enhance human potential and pursue the goal of TransAI Humanism.

According to Polanyi the universe is not a uniform sea, it contains emergent comprehensive entities, with discernable levels of organisation, which via boundary conditions operate independently of the lower levels from which they emerge. These emergent levels realise various purposes. As Pirsig explains, each level of order has a characteristic defining Quality, material, biological, social, and intellectual. Indwelling within articulations enables us to pursue and reflect upon abstract ideals such as truth, goodness, and beauty. For Polanyi a computer, indeed every sort of machine, is contrived to achieve a purpose. A machine can be programmed to achieve this purpose if the desired behaviour can be described in a way that can be simulated by a universal machine. AI offers the prospect of machines which learn without us having to define every detail of its procedures. But they need to do the right thing in every circumstance. This returns us to our problem, because defining criteria for correct behaviour, for making the right choices, amounts to giving a machine "common sense". Humans excel in making choices in complex situations because we are guided by our tacit awareness. The lead AI scientist at Facebook when asked recently, "When will the true AI" revolution occur?" replied that it will take place when we create machines with common sense, and this may take up to 20 years. The goal of producing general purpose AI relies upon the assumption that the epistemology underlying this effort is sound. From a Post-Critical perspective it is misguided. Viable alternatives however are emerging that are consistent with the epistemology outlined by Polanyi. They not only acknowledge intertwinability and the metaphysics of adjacency, and the importance of tinkability and co-evolution, but also data sovereignty and the person and spontaneous order of decentralised, human first, free societies.

Tools for Thought: MindGraph and TrailMarks

[Brian Cantwell Smith](#) in his introduction to the [Age of Significance](#) wrote:

“I see!” said a voice at the back of the hall, breaking the silence. “You are trying to articulate, in language that only philosophers can understand, intuitions that only computer scientists can have.”

Like Brian Cantwell Smith, I (Gyuri Lajos) have spent over thirty years orbiting the "strange attractor" he aptly calls the "Age of Significance". Significance relates to the fact that computers deal with meanings. Douglas Adams defined computers as Universal Modelling Machines, but Smith intimates that they are best understood as Universal Meaning Machines. Armed with intuitions that only philosophers, such as Polanyi, can articulate, I approached the Age of Significance from the opposite direction. Instead of trying to articulate insights that only computer scientists can have, I built computer system prototypes while relying upon intuitions that only philosophers can have, embarking on long distance exploratory programming with the goal of creating machine capabilities which could augment philosophy itself. I chose to do this because I believed that we will need to be able to construct conceptualizations that without computer help we may not be able to accomplish. My motivation for making new "tools for thought" was to seek to apply the Kantian question “how is knowledge possible” to the question how is computer support for knowledge work possible, and on the basis of this reflection investigate how augmenting our conceptual capacity can be rendered possible. If computers could augment our epistemological intuitions, these could then inform our efforts to build tools which enable us to go where no computer programmer has gone before, enabling us to articulate our shared awareness better in what is not an “Information Age” but an “Age of Significance”. While I admire Smith’s work, I had a sinking feeling, and by now a growing conviction, that the existing language that philosophers can understand may not be up to the task of articulating these intuitions.

As concepts are indices of our experience, when faced with the radically new domain of experience opened up by computers, it is not surprising that existing conceptual structures do not have words for them. Worse still, much of the epistemology that is being relied upon when articulating intuitions that programmers have, is wedded to the paradigm of logic. Logic provides us with a universal mechanism that guarantees truth preserving methods of inference. But the price we pay is that we abstract away from the meaning of the terms, and lose comprehension of the meaning of the content. The rules of inference are encoded in syntax, and they enable us to render the surface structure of declarative knowledge explicit, even if the meaning of the terms in that structure gets lost. At lower levels however the logic embodied within the circuits of a computer create the primitive capabilities which ground the emergence of a computer as a comprehensive entity that can go way beyond logic. Computers can have a procedural semantics, they can not only preserve but also express the declarative meaning of the content, and manipulate the situated meanings as rich content with depth. Through

programming computers we can also capture procedural, or non declarative, “how to” knowledge. These programs can exhibit desired behaviours capturing intents. In other words, through software, we can create new meaning. The Computer in fact opens up new Multiverses. Logic can be viewed as the poor man’s computer, in that it allows mechanical manipulation of declarative meaning without actually having one available. My preoccupation with the Age of Significance grows out of an early insight which lead me to put forward an improbable, and for many perhaps implausible notion, that we may be able to produce a new kind of theory of descriptions that will eventually supplant current theories of effective computation. I felt deeply that “[we really do not know how to compute](#)”, and the existing theories of computing are of not much help. A view that Brian Cantwell Smith eloquently articulated on many occasions, and which I tried to explore by developing “Language Oriented Programming” as a new paradigm.

For half a century Ted Nelson has been trying to open our eyes about the perniciousness of the current [computer tradition](#). We need to understand that the universality of computing machines implies that we have a God level capability to exercise the freedom and responsibility of turning anything we can coherently envision into a new Computer Universe, accompanied by its distinct Cosmology. The significance of this goes way beyond what Charles Simonyi says about software: that it is the closest we can get to practical magic. As [Ted Nelson explained](#), it all comes from the concept of the “lump file”. Files are put in directories, or folders, which can contain other folders, all arranged in a hierarchy. Applications are themselves executable files that operate on documents in files. Everything on your computer is arranged in a hierarchy, therefore, and thus as Ted Nelson puts it: “hierarchy has been imposed on the content and in the world, and on the Universe” and the alternative is: “Throw away the lumps, or rather create a new Universe of small portions which are addressable and which can be connected simultaneously in many different ways. We need addressable tiny pieces, we need changes to be addressable, we need to file the changes to a document, so they can be run forward and backward. It is time to re-examine the entire computer world.” and come up with a better Computer Cosmology. It should be based on a Graph or Network. A hierarchy is actually a special case of a graph or network. With our project, which we call MindGraph, we attempted to create a graph based computing cosmology which realises all the desired characteristics Ted Nelson talked about. I would venture to suggest that there may only be one true Computer Cosmology and it is Graph-shaped, and based on the “[Metaphysics of Adjacency](#)”. Based on this, hunch, I have spent the last 15 years seeking to create systems based on a graph based cosmology for the computer universe.

The main lessons to be learned from AI is that graphs, which are networks of things which are connected in some way, are ubiquitous. Depending on the task at hand, there are vast numbers of different ways of conceptualising a graph as a network of things. Graphs are formed from two sets of things. From a set of nodes and a set of links that connect these nodes. Graphs have a huge literature in Mathematics and is now a subject of a new field called Network Science. This simply defined mathematical structure can be embellished with many properties relating to the nature of the nodes and the properties of the links e.g. the links can have directions and labels. Any restrictions can be applied regarding the kinds of links which can be

introduced. Both nodes and links can have properties, or whatever else you wish to put into them or specify. What is a graph, depends on what you put in the nodes and the links, how you establish their identities or assign names to them. They are up for grabs, you can create them any way you deem fit. The design space for universally applicable graph based computer cosmologies is therefore huge, and as ever too much combinatorial freedom leads to chaos. Even if we limit ourselves to the issue of knowledge representation, we still have numerous different conceptualizations for it as a graph. None of them seem to cut it. That's why we had the first AI winter. The next one is due soon. The root of the problem is a commitment to epistemologies which are not appropriate, or worse still, fail to acknowledge that there is a Ying-Yang "mutual arising" here: because not only is it the case that how we conceptualize our graph based representations has epistemological import, it is also the case that epistemological commitments influence the cosmologies which underlie our representations. The way we constitute our graph has implications for the kind of knowledge which can be captured by it.

Most of the better known knowledge representations I am aware of, share fundamental epistemological commitments which can be characterised by labels such as: positivist, objectivist, and impersonal. The best practitioners in the computing field are aware of the relevance of epistemology to their endeavours, but they (it must be said, rightly) treat mainstream epistemologies with suspicion. They see a great deal of wooliness, so they prefer instead on focus on issues at hand. [What's in a Link - Revisited on Vimeo](#). It is painful to see how slowly good ideas gets adopted in current linguistic theory research, works such as [Rethinking Conceptual Metaphor Within an Integral Semantics Framework](#) are slowly inching towards an [integral theory](#) and a recognition of the personal, and by implication the tacit lifeworld "the world of lived experiences shared with other human beings". It is the epistemology of Polanyi's "Personal Knowledge" which informs our efforts to build a Personal Knowledge Augmentation Engine. Based on a graph-based framework for representing person first, emergent knowledge that we call MindGraph. We start with an epistemology when we build our "tools for thought". Our experience with using the tools we build, supplies us with the feedback we need in order to bring about improvements in the way we conceptualize our tasks, these in turn help us improve the epistemological understanding from which we started. This can be conceived as an experimental epistemology. Let us give a brief account of the epistemology which informs our re-conceptualization of the concept of the graph as a MindGraph.

MindGraph is a framework for representing person first emergent knowledge. The MindGraph based Personal Knowledge Augmentation Engine which supports the co-evolution of the hyperlinked capabilities furnishes us with tools that can augment our knowledge. It has integrated capabilities which help us to contextually recall existing knowledge, at resolutions that are under the user's control, enabling us to improve our comprehension and organization of the existing material. In developing an augmented authoring component for MindGraph, we are following the path which Engelbart opened up in his account of the [Authorship Provisions in AUGMENT - 1984](#). It accords with his overriding goal of "taking a new and systematic approach to improving the intellectual effectiveness of the individual human being" (see the

abstract in [AUGMENTING HUMAN INTELLECT](#)). MindGraph provides ways of self-sovereign sharing and collaboration in a decentralised network of people networks where each person can either operate as an island, or as a hub for their own decentralized network of peers. This generates a global emergent, self organising spontaneous network of networks of knowledge. Although it is rooted in the personal, it is possible to anchor it in notable things harnessed from the Linked Data Cloud, forming a giant global distributed Knowledge Graph. MindGraph gives us sources for discoverable, notable, codified, federatable machine-readable knowledge on the Semantic Web. We are committed to the goal of Bootstrapping and co-evolving a Knowledge Augmentation Engine that is able to give us a machine supported framework which facilitates personal knowing and the augmentation of our tacit awareness, through augmenting our articulations, and providing a means for collaboration in teams, communities and institutions. We are preparing for release a minimal workable subset of these capabilities called TrailMarks.

We conceptualize symbols and concepts as “indices” to experience [Voegelin *Anamnesis* p 182] that articulate our tacit awareness. It is through symbolic articulation that we convey meaning. In doing so we may extend the human record as part of the process of time binding. Bush warned that our ability to add to the ‘common record’ far exceeds our ability to consult it. The semantic Web can be seen as an attempt to mechanise the process of retrieving not only a record of a thing but all immediately linked things, enabling first machines, then humans, to be able to tell how and why they are connected. It is the potential for this kind of associative indexing and contextual recall that makes “graphs” so alluring. To capture the intertwingularity and mutual interdependence and mutual arising of the world and also our tacit awareness of it, we need to have the Graph as the basis for our computational cosmology. In our epistemology we reject the Objectivist concept of the impersonal “single version of truth”. The objectivist tendency and machine bias both promote a concept of knowledge which can exist without a knowing subject. The desire for control results in an overuse of hierarchies. This is not to deny the utility of hierarchies, the problem arises when they are extended beyond their usefulness. Let the warning be clear, “He who controls metadata controls much of the world”. Thoughts, like reality, are non-linear and intertwined, and so it is desirable for us to have systems that facilitate our ability to capture the richness and multiperspectivity of all things; capturing the butterfly in flight instead of turning everything into linear caterpillars. In reading papers, which is the main form of disseminating new knowledge in science, we scan the linear text in order to upload its meaning into our minds and then we reconstruct the intertwingularity of the ideas they contain in our tacit awareness. What we would like to have is a HyperMap of the paper’s content with the ability to supply us a glance with comprehension, searches, and navigation. These are the kinds of capabilities we require if we are to augment our intellect.

Our epistemological stance focuses on the needs of the human knower, not at the expense of the machine but in a mutual enriching symbiosis which takes account of our choices. From an Objectivist epistemological stance it makes sense to only talk about how we can create data structures which can be processed by machine, but we also need to produce things which people can know about and understand. The requirements of the machine makes us lose sight of the requirement to consider things from a human standpoint. If we are interested in augmenting

our intellect, we need to focus on how we can construct and curate and convey and create our knowledge. Our primary goal is not to find better ways of representing but to find better ways of presenting knowledge artifacts to the user, in interactions conducive to augmenting articulation of our tacit awareness. This was also the primary focus in Engelbart and his team's bootstrapping work. As explained earlier, the reason why we have applications at all can be traced back to the pernicious tradition of Lump files and hierarchy. Engelbart and Ted Nelson both had high resolution addressability and keeping links outside the file as ground rules in their vision. Both retained the notion of a document. In MindGraph we have high resolution addressability with no files at all. Both documents and "files" in MindGraph become virtual creations, resolved at the time of retrieving, a part of the graph containing the fraction which our attention focuses on, either because we followed a link or as a result of a search. In a sense these virtual documents are like photons in quantum physics, they become "particles" (i.e. standard, bounded texts "imitating paper") only when they are observed. As our focus changes in the graph new virtual document materializes. Everything is deeply rearrangeable and repurposable, in accordance with Ted Nelson's vision. Let us say goodbye to document management as such. Documents without deep bidirectional addressability are graveyards of ideas, and force us to say the same thing umpteen times. It should be possible for every person to have a single graph in a single place, that is capable of encompassing all of their digital life, with the ability to allow each other to have selective access to any part of it.

Alan Kay, Engelbart, and Ted Nelson all agree that information and capabilities for meaningful manipulation all have to be a first class object. Moreover, as knowledge grows and our ability for managing it grows we need to be able to manage all knowledge by relating it to the process itself. We need to be able to give all knowledge artifacts we defined in the graph their own direct manipulable morphic interfaces that are personalizable and even extensible, tinkerable by the user, so that we can organize and present knowledge and means to manage them within one system that rely on a common mental model on which users can depend under all circumstances. Just like it was in what was produced at Xerox Parc and the SmallTalk system which Allan Kay and his team created. This consideration should lead us to follow a path that avoids the main pitfalls of most computer support for knowledge work: which is to say high cognitive load, costs and low isolated benefits. MindGraph's design has grown out of a focus on our human need to answer questions such as: What are the most promising manipulable forms of presenting articulation that is conducive for both production and consumption of knowledge, which minimises cognitive load of both in the creation, recall, organisation and consumption of knowledge and the very means we deploy for its organization. It is here that we see that the Graph is essential but not sufficient. SmallTalk provided morphic direct manipulation, and capability linking, but it did not have the graph. MindGraph has the graph, and capability linking, but it restricts its attention to HyperText, and its growth through HTML is treated as general hypermedia. The test use case scenario is using Google Docs to write a paper. It gives us the ability to convert a draft into a MindGraph, research and link everything of interest, and create deeply addressable, and hence reusable structures. Once it is written down it should be recallable in every relevant context. So when the paper gets shared it is "semanticized": uniquely identifying things of interest, and "connecting the dots" in ways

that the meaning of the connection is made explicit in our own Knowledge Graph, with the potential to be shared and collaborated upon.

MindGraph supports self-sovereign identity and ownership of your own data. Because the identity of a node is tied to the person/agent who creates it, the claims are made not on the basis of the false epistemology of single version of truth, each node is identified as belonging to its creator, so MindGraph is able to capture evolving and unreliable knowledge; recording conflicting views; expressing degrees of uncertainty about the various knowledge claims. Unlike the semantic web, MindGraph does not force you to resolve conflicts at the point of contribution, but can do so at the point of accessing nodes authored by different individuals or agents. We detail [Transitional Modelling](#) in MindGraph and other proposed enhancements to the Semantic Web via MindGraph in a paper for this year's International Semantic Web Conference. We show in that paper how MindGraph, anchored in the Semantic Web, can give rise to emergent, decentralised social knowledge networks. We also make the claim there also that MindGraph becomes an engine which grows co-evolving new knowledge through decentralised emergent peer to peer collaboration and how knowledge can be federated based on the [HyperKnowledge](#) Protocol. The Semantic Web deals with only notable things and was designed as providing the means for "[machine facilitated global knowledge exchanges](#)". MindGraph, in contrast, focuses on non-notable knowledge that is personal, and not (yet) established or generally accepted. MindGraph can be a bridge between the two worlds and enhance both at the same time. Though we may bemoan the Semantic web for its objectivist bias, but it must be acknowledged that our vision of Personal Knowledge Augmentation can only have been built upon that foundation, as it captures what is "agreed" upon, so that we can relate the new and controversial to something which is shared and assumed to be "understood".

MindGraph can greatly benefit the Semantic Web by offering a way to get ordinary humans in the loop not ontology experts. It can deliver us from the high cognitive load associated with the existing ways of contributing to the Semantic Web. All articulations as Linked Text in MindGraph provide human-created meaningful contexts for the notable entities they reference from the Semantic Web. As MindGraph by construction contains "semanticized" text, these texts can be mined by simple NLP algorithms to discover when two entities drawn from different sources are referring to the same thing. This provides a way of alleviating the problems associated with "heterogeneity of data" which machines cannot solve, because they lack real understanding of what the text is about. MindGraph supports what we call Augmented Authoring with 'Linked Text'. This is inspired by Engelbart's [Authorship Provisions in AUGMENT - 1984](#). It creates a graph based, high resolution addressing of text fragments, which in turn can be linked to other fragments and other nodes. 'Linked Text' uses HTML, as a HyperMedia format allows linking to other things via non-linear hyperlinks, as in a Wiki. These links can be to entities in the graph, and to other fragments which are semantically identified. All fragments in Linked Text are deeply addressable via hyperlinks from other fragments in other nodes. By construction these links are bi-directional. When a Linked Text fragment links to another fragment in another node, these two nodes get associated. For each node we can find all the other nodes which reference them, and this is crucial for spanning high resolution meaningful contexts, and bringing into view all related

nodes so that we can make semantic rearrangement and reuse of these fragments. This eliminates the repetitions and scattering of knowledge, and at the same time enhances the power of associative recall. MindGraph provides an extensible Augmented Linked Text editor which allows text to be authored as nested lists, and any element to be mapped on demand to a named narrative entity, enabling these fragments to be reused and become targets of deep meaningful links. It also allows the “semantification” of texts, allowing content to be richly linked to other entities, and other relevant fragments, or to form narrative trails via transclusion.

Google pioneered the use of Knowledge Graphs as mapping “strings to Things”. This enabled Google to support what can be best described as contextual recall i.e. giving answers based on the automatic availability of associated contexts. Although there is an API which allows you to map strings to “Things”, the API does not allow you to retrieve it’s context, or its immediate connected neighbourhood of nodes, nor can it find out anything about these connections. MindGraph, in contrast, not only enables contextual recall, one level at a time, or along specified connections, or traversal logic, it can also return the entire connected neighbourhood of nodes. MindGraph also provides rich visual direct manipulation tools, and an interoperability hub to create your own Contextual Recall patterns. These contexts are recalled in ways that are amenable to be presented through a rich variety of visualizations, for example hyperMaps (Voronoi Maps), which provide visually organised Gestalt views that are exploreable, and which enable deep associative complexes to be comprehend-able as gestalts. These maps emerge dynamically from the territory that they map. They give you what you have in a self revealing spontaneous order. This will often reveal to us more than we already know. They also act as serendipity engines helping us to discover previously unconsidered topics and authors, and opportunities to make new connections. They give us synoptic views of narrative trails as deeply rearrangeable, nested, dynamically controlled outlines. MindGraph supports the constructions of Narrative trails using Linked Text fragments. From these semantically addressable fragments MindGraph can weave linearised scholarly papers ready for on-line (self) publication. "Papers" produced by MindGraph can allow rapid comprehension and drilling down to relevant bits, within a self organising revealing context. This will reduce the need to read stuff from cover to cover, instead treating publications as a portal into our own world of relevant meaning, taking creative knowledge work to the point where the author left it at the point of sharing.

From existing papers, narrative trails and slideshows can be created through HyperCard-like user experience grounded in a MindGraph. We call this HyperMindCard. Using these cards we can construct presentations that give an overview of the topics. From these cards the scaffolding with which the original papers were erected can be accessed. In this way publications will always remain connected to the context of discovery from which they emerged so that the researcher or anybody else can pick up work on that topic exactly at the point where they left off. Put an end to the artificial separation of context of discovery and justification In the case of collaboration, a full audit trail of contributions can be reconstructed on demand, if this is permitted by the contributors. MindGraph supplies a kernel which via indwelling empowers us to augment our tacit awareness, and participate in the growth of human Extelligence. We are launching a subset of the above capabilities in a prototype called

TrailMarks, which empowers its users to mark the narrative trails they blaze across the Web, Thus these trails will never fade. It can reference all the relevant context from which they meaningfully emerge. It provides deep semantic links to everything they find, marking them as relevant, or if they have created themselves, giving users auto-associative recall of all the relevant contexts. We can share purposeful extracts as linear virtual documents, along with the scaffolding with which they were erected, in a form that facilitates peer to peer collaboration and complete immutable audit trails for each contribution. Through appropriate visualizations, TrailMarks enables users to comprehend at a glance what there is, and to make rapid and incisive choices on what to focus on next. Building on earlier advances it harnesses the capabilities of Google Search, WikiData, WorldBrain, and Hypothesis. Trailmarks is constructed as a decentralised knowledge interoperability hub, and two way interoperability with Google Docs, WordBrain and Hypothesis is in development.

The root of the word intelligence is inter (among) and legere (choose) i.e. the power of the mind to identify and connect items. The power of a graph is its ability to identify items of interest and connect them in meaningful ways. When a new connection is created it makes two neighbourhoods of associative complexes adjacent. The result is a rearrangement of the graph, which can be presented as a new experience that can lead to new insights. As our comprehension (the Latin etymology derives from grasping) becomes deeper and more extensive, our knowledge grows. The epistemology which underlies MindGraph could be described as an Epistemology of Adjacency. It relies upon our power to choose and connect elements from our awareness, construct patterns, and augment our tacit awareness. Through time binding it is our sacred duty to pass on as our legacy not only that which we have acquired, but along with our additions to the common record the scaffolding with which those additions were erected. Human civilization (in halting steps) emerges and grows via the refinement of our efforts to register and create meaning in our unending quest to realise the good.

