

Stephen Downes

An Introduction to Connective Knowledge

2008

<https://doi.org/10.25969/mediarep/19789>

Veröffentlichungsversion / published version

Sammelbandbeitrag / collection article

Empfohlene Zitierung / Suggested Citation:

Downes, Stephen: An Introduction to Connective Knowledge. In: Theo Hug (Hg.): *Media, Knowledge & Education. Exploring new Spaces, Relations and Dynamics in Digital Media Ecologies*. Innsbruck: Innsbruck University Press 2008, S. 77–102. DOI: <https://doi.org/10.25969/mediarep/19789>.

Nutzungsbedingungen:

Dieser Text wird unter einer Deposit-Lizenz (Keine Weiterverbreitung - keine Bearbeitung) zur Verfügung gestellt. Gewährt wird ein nicht exklusives, nicht übertragbares, persönliches und beschränktes Recht auf Nutzung dieses Dokuments. Dieses Dokument ist ausschließlich für den persönlichen, nicht-kommerziellen Gebrauch bestimmt. Auf sämtlichen Kopien dieses Dokuments müssen alle Urheberrechtshinweise und sonstigen Hinweise auf gesetzlichen Schutz beibehalten werden. Sie dürfen dieses Dokument nicht in irgendeiner Weise abändern, noch dürfen Sie dieses Dokument für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen.

Mit der Verwendung dieses Dokuments erkennen Sie die Nutzungsbedingungen an.

Terms of use:

This document is made available under a Deposit License (No Redistribution - no modifications). We grant a non-exclusive, non-transferable, individual, and limited right for using this document. This document is solely intended for your personal, non-commercial use. All copies of this documents must retain all copyright information and other information regarding legal protection. You are not allowed to alter this document in any way, to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute, or otherwise use the document in public.

By using this particular document, you accept the conditions of use stated above.

An Introduction to Connective Knowledge

Stephen Downes

Abstracts

This paper provides an overview of connective knowledge. It is intended to be an introduction, expressed as non-technically as possible. It is intended to respond to writers like Chris Anderson (2005) who, like many other writers, describe connective forms of knowledge as probabilistic.

These forms of knowledge, created by means of interactions among connected entities, are outlined in publications such as James Surowiecki's (2005) *The Wisdom of Crowds*. They should not be thought of as probabilistic, but rather, as a distinct and new form of knowledge. The purpose of this paper is to introduce the reader to this new, connective, form of knowledge.

It should go without saying that these are my own thoughts, and this discussion should not therefore be considered an authoritative reference on the subject. Moreover, this is intended to be a brief overview, and not an academic treatise on the subject.

a. Types of Knowledge

You probably grew up learning that there are two major types of knowledge: qualitative and quantitative. These two types of knowledge have their origin in major schools of history and philosophy, the former in the works of the ancient Greeks, and the latter in Arabic and then later Renaissance philosophy.

Connective knowledge adds a third major category to this domain, knowledge that could be described as distributed, because it is spread across more than one entity. A property of one entity must lead to or become a property of another entity in order for them to be considered connected; the knowledge that results from such connections is connective knowledge.

This is more than just the existence of a relation between one entity and another; it implies interaction. A relation – such as 'taller than' or 'next to' – is a type of quality. It describes a property of the object in question, with reference to a second object. But the fact that I am, say, 'taller than' Fred tells us nothing about how Fred and I interact. That is something different.

This is why it is incorrect to represent distributed knowledge merely as a type of probabilistic knowledge. The logic of probability implies no connection between correlated events; it merely observes a distribution. A connected system may exhibit probabilistic characteristics, but it is not itself probabilistic.

Probabilistic knowledge is a type of quantitative knowledge. It is based on the counting of things (or events, or whatever) and of comparisons between one count and another (one needs only to read Carnap (1967) to see this clearly). A poll, for example, gives us probabilistic information; it tells us how many people would vote today, and by inference, would vote tomorrow. But the fact that Janet would vote one way, and I would vote one way, tells us nothing about how Janet and I interact.

Connective knowledge requires an interaction. More to the point, connective knowledge is knowledge *of* the interaction. If Janet votes a certain way because I told her to, an interaction has taken place and a connection has been established. The knowledge thus observed consists not in how Janet and I will vote, nor in how many of us will vote, but rather, in the observation that there is this type of connection between myself and Janet. This knowledge may have nothing to do with voting at all. Rather, we may recognize it as part of a larger pattern of interaction between the two of us.

b. Interpretation

What we 'know' about the world is irreducibly interpretive. That is to say, we do not through our senses and cognition obtain any sort of *direct* knowledge about the world, but rather, interpret the sensations we receive. This is true not only of connective knowledge, but of all three types of knowledge.

Consider qualities, for example. We take it as basic or atomic (see people like Ayer (1952) for example) that a statement like 'this apple is red' represents a pure and unadjusted fact. However, looking at this more closely tells us how much we have added to our original sensation in order to arrive at this fact:

First of all, the apple itself has no inherent colour. Colour is a property (specifically, the wavelength) of light reflecting off the apple. In different coloured light, the apple will appear to us differently – it appears white in red light, for example, or gray in diminished light. Yet we say the apple is 'red' – standardizing our colour descriptions to adapt to the natural light that surrounds us day to day.

Second, our perception of the apple as 'red' depends on our organizing light patterns in a certain way. When I was a child, the spectrum had six colours – red, orange, yellow, green, blue and purple. As an adult, I find that a seventh – indigo – has been added. It's not that a new colour came into existence when I was twenty, it's that our nomenclature changed. In a similar way, we can divide the colours of the spectrum in numerous ways: 'red', for example, can include shades as varied as 'crimson' and 'cherry'. Or '#ff0000'.

And third, when we say that 'the apple is red' we are drawing on our prior linguistic ability to use the words 'apple' and 'red' correctly and apply them to appropriate circumstances. Indeed, our prior knowledge often shapes our perceptions themselves: were you shown an apple in diminished light, so that all you could see was gray, and asked what colour it was, you would still respond 'red' because of your prior expectations about apples and redness.

Less intuitively so, but equally clearly, interpretation applies to quantitative knowledge as well. It is easy to say that a sentence like 'there are twenty schoolchildren in the yard' is a basic fact, but this all depends on how you classify schoolchildren. Suppose, unknown to us all, one of the children had just been expelled; is our statement now false? Not obviously so. Perhaps one of them is over sixteen – is this person still a child (and hence, a schoolchild)? It depends on your point of view.

Quantification is essentially the enumeration of members of a category or set. Consequently, it depends crucially on how that set is defined. But membership in a set, in turn, is (typically) based on the properties or qualities of the entities in question. So such membership is based on interpretation, and hence, so is counting.

One might be tempted to say that even though applied instances of counting are based on interpretation, mathematics itself is not. But in my view, this too would be mistaken. For one thing, as people such as Mill and Kitcher (1985) argue, the rules of mathematics depend on empirical verification for their importance: we say that one plus one is two, not out of some innate sense of goodness, but because when we put one sheep together with another, we observe that there are two. Nothing but our observations prevents us from saying that one plus one is three, and in some contexts such a statement makes perfect sense.

c. Emergence

Emergence is a hard concept, but at this point I can gloss it with a simple characterization: emergence is interpretation applied to connections.

There are two (equally valid) ways of thinking about this:

First, we may perceive an actual set of connections linking a group of entities as a distinct whole. For example, when one domino topples another, and so on, in turn, and we observe this from a distance, we may observe what appears to be a wave moving through the dominoes. The wave that we observe can be said to be an 'emergent phenomenon' – it is not a property of the dominoes themselves, or even of the falling of the dominoes, but of the connectedness of the falling – because one domino causes the next to fall, we see a wave.

Second, we may perceive something as a distinct whole and interpret this as a set of connections. For example, when we look at the image of Richard Nixon on the television, we do not perceive the individual pixels, but rather, the image of a person. But our inference goes beyond merely the observation of the person; if asked, we would say that the appearances of the pixels are connected to each other, through the mechanism of having a common origin (Richard Nixon himself) and the mechanism of video broadcasting.

Emergence is fundamentally the result of interpretation. As mystics (and Spinoza) are fond of arguing, everything is connected. At a certain point, as the old saying goes, when a butterfly flaps its wings in China, the result is a thunderstorm in Halifax. But broadcasters in Halifax do not watch butterflies in China in order to predict the weather, because this connection will be of no use to them. Typically, they will look at more intermediate events, themselves emer-

gent properties, such as waves of air moving through the atmosphere (known locally as 'cold fronts').

In the same way, the observation of sets of connections between entities depends a great deal on what we already believe. That is why we see swans in clouds or faces on Mars when, manifestly, there are none. We have brought our prior knowledge of connected entities to bear on our interpretations of these phenomena. As Hume (1999) would say, our 'perception' of a causal relationship between two events is more a matter of 'custom and habit' than it is of observation.

d. Physicality

We generally think of knowledge as being about facts, and about facts in turn as being grounded in an independent reality, a physical reality. Consequently, it is natural for us to say, for example, that when we see that something is red, that there is a physical basis for that statement, that even if we bring some interpretation to bear, there is some physical fact of the matter that makes the apple red, and not blue.

Certainly, were we not to think of things this way, we would be hard pressed to say anything about anything. Physicality provides us with a substrate on which to hang our interpretations. It is, as Kant (1999) would say, a necessary condition for the possibility of perception. Physicality moreover offers us a means of sorting between what might be called 'correct' interpretations and 'misperceptions', between reality and a mirage.

All this may be the case, but nonetheless, there is nothing in our interpretations that is *inherently* based in physical reality, and hence, nothing that precludes our discussion of them without reference to this foundation. Indeed, this has been enormously useful in other domains. Despite, for example, the empirical basis of mathematics, it is much more productive and useful to refer to quantity without reference to the physical entities being counted, to (in other words) think of quantity in the abstract. The same is true of quality. Thinking of quality in the abstract leads to Aristotle's syllogisms (Kemmerling 2001) and is the basis of categorical reasoning.

Moreover, non-physical entities may have (or be attributed) properties that are themselves (on this theory) based in physical properties. In our ideas and dreams, we think of vivid colours and large numbers. And the ideas are transferable. Consider the concept of 'purple prose' (Stover 2007) – an expression which if applied literally is in all cases either meaningless or false, yet of significant utility and meaning.

What is to be learned from this? That the entities in the various categories of knowledge – be they properties or numbers – are themselves *not real*. When we talk about 'redness', we are not talking about something that has an independent, concrete existence in the world, but rather, in something that exists (insofar as it exists at all) only in our own minds. When we talk about the number 'four', we are not describing some Platonic entity (Dye 2003), but rather, nothing more than our own thoughts or sensations.

That does not make them less 'real'. Our perception of the colour 'red' is as real as any phenomenon in the world. It is merely to distinguish between the perception, which results from a complex of factors, from the physical entity, which ostensibly caused it.

In a similar manner, our interpretations of connections is distinct from the actual set of interactions that may exist in the world. Consider, for example, conspiracy theories – the postulation of a complex and inter-related set of people and events leading to the conclusion that someone is out to get you. Such theories, notoriously, have no basis in the physical world. But they may nonetheless be contemplated, and discussed, and passed along, as though they were real. And the experience of a conspiracy theory may be, to the perceiver, every bit as real to the person having the experience.

There is a tendency on the part of readers, whether of talking about crickets, or of Shirky (2003) talking about power laws, to represent connections as something 'natural' and 'real' that is simply 'out there' – as though what is said about networks of connections represents some immutable law of nature. Quite the converse is the case; our understanding of the existence of connections, and the nature of the networks they form, is something we bring to the table, an interpretation of what we think is salient.

e. Salience and Inference

Our knowledge consists of interpretations of perceptions, which are in themselves distinct from any physical reality that may have caused them. In this sense, one might say that these interpretations are 'constructed' – that is, they are the result of some mental or cognitive process, rather than something that comes delivered to us already assembled.

Inference is, broadly speaking, the manipulation of these bits of knowledge, in the abstract, to produce new bits of knowledge. In our mind, for example, we can postulate that if a red light is added to a yellow light, the result will be an orange light. Or that two sheep added to two sheep will result in four sheep. Often, subsequent perceptions will confirm such predictions, thus leading us to rely more greatly on the manipulations that resulted in them (and less greatly on manipulations that did not result in them, though the human mind is notoriously fickle in this regard).

All such inferences, however, are the result of a complex process of selecting what might be called the most 'salient' data. The counting of sheep, for example, is of utility only to people who own sheep (or are reading philosophy papers). Normally, during the course of our everyday lives, we have little need to count sheep, and so for the most part we ignore the actual number of sheep present to us at any given time. In a similar manner, when we perceive an orange light, we do not typically view it as a confirmation of the idea that red and yellow make orange. Unless we are visual artists we see it merely as an instance of 'orange'.

Our inferences, therefore, are based on salience, where *salience* may be thought of as the importance, relevance or vivacity of some property or perception. We 'pick out' those perceptions that will be of use to us, and disregard the rest. This is not often even a conscious process; it is based in part on innate reactions (such as jumping when we hear a loud sound) and largely on prior expectations. Our past knowledge has led us to recognize that something

that looks and sounds like a tiger is something we should pay attention to, and so our inference engine kicks into high gear.

In the same way, some connections are more salient than others. Think about your sense of place or location. It is centered on the city or town in which you are located, the streets spreading out from you in a pattern unique to your position. Change your location, and your map of the world changes with it; WalMart, which was once 'two blocks away', is now 'one block away'. Or consider your circle of friends: again, you are at the centre, with your closest associates at near proximity, with acquaintances more distance. Your friend, however, will count a different set of people as being most proximate, and others, including some you hold more close, as more distant.

Things become even more complex when considering the mind. We know that the mind is a massively connected set of neurons, but where is the point of view from which we regard these connections? While we can consider the bird's eye view in the abstract, and speak dispassionately about the hippocampus or the corpus callosum, we cannot adopt such a frame of reference with respect to our own thinking. And yet, it seems manifest that there *is* a point of view with which we regard our own mind; it is the essence of conscious thought, that we are aware of our mental processes at the same time we are having them.

Again, it is that which is most salient that comes to the fore here. You may have mental representations of hundreds or even thousands of people but, if you are enamoured, be thinking only about one. Your body consists of millions of nerve ends, but if you have a toothache, your attention is focused only on those few related to the tooth. In a similar manner, it is only your most active and your most consistent thoughts that intrude on your consciousness, and it is through the lens of those thoughts that you interpret phenomena (and through phenomena that you have those thoughts).

Inference is the observation of salient similarities among thoughts and perceptions. It is the recognition of common properties – qualities, quantities and connections – among varied perceptions, and the consequent drawing of connections between those entities, and between other properties of those entities. Seeing that two sheep and two sheep make four sheep, you are led (via the salience of quantity, and the newly found salience of cows) to contemplate the idea that two cows and two cows might make four cows.

f. Associationism

This process of inference has a history in philosophy under the heading of 'associationism', a type of reasoning associated with (until the advent of logical positivism) empirical philosophy and people such as Hume and Mill.

The central idea of associationism is this: two things that are relevantly similar become connected in the mind. This connection or association in turn allows knowledge about one to be inferred of the other. Thus, if we experience one tiger-like creature, and it tried to eat us, then if we see a relevantly similar tiger-like creature, we are led (as Hume would say, naturally and senselessly) to believe that it will try to eat us as well. Eventually, a complex of beliefs

about tiger-like creatures is formed, and some indeed become strong enough to allow us to contemplate a new (and dangerous) category of entity, given the name 'tiger'.

Various types of associationism exist, from association of impressions postulated by Hume to the similarity of phenomena described by Tversky (2003). Two major types of associationism are relevant to us here:

The first is simple associationism, sometimes known as 'Hebbian associationism' (Haselager 1999), which is postulated to be (and probably is) foundational in the forming of neural connections in the mind (its applicability to the world outside the mind is much less evident). The principle, specifically, is that if two neurons fire at the same time, a connection will tend to be formed between them. This is, of course, an 'all else being equal' hypothesis: the neurons have to be the sort of neurons that can form connections, there needs to be some sort of proximity between them, and they need to be (computationally and physically) compatible with each other. A lot like a love story.

The second may be classified under the (inaccurate) heading of Boltzmann associationism (Hinton & Sejnowski 1989). Derived from the idea of the Boltzmann machine, this sort of associationism is an expression of (something like) thermodynamic forces. Think of it as the network attempting to settle into a 'balanced' or 'harmonious' state. The idea behind Boltzmann associationism is that a certain amount of energy applied to a system will create a certain amount of kinetics – in other words, your brain goes on thinking even though it's not receiving input. In the absence of external influences to cause Hebbian connections, the brain settles into a (thermodynamically) stable configuration.

Whether such modes of associationism, or any other method of connection-forming, is at work within any particular system, is a question for empirical observation. Probably, in any given system, it will be a combination. And as before, in addition to specific connection-building mechanisms, there will be a requirement for enabling factors, such as proximity, and corrective factors, such as is described by back propagation (Rumelhart, Hinton & Williams 1989).

We understand similarity well enough with respect to quality and quantity. Things can be more or less alike – large, round and orange, say. And we can see how though this similarity how an association can be formed – our perception of (what we interpret to be) two different orange phenomena leads us to draw an association between them. Quantities, as well, are associated: we have never experienced a rainfall of six inches of milk, but we can easily imagine what it would be like, based on our experiences with six inches of water.

In the case of connections, the concept of similarity is less intuitive, but breaks into two major categories:

First, we can say that two entities are connectively similar if they share connections with the same set of entities. For example, Michael and I may be connectively similar, even if we have never met, if we share the same group of friends. Of course, such a similarity makes it more likely that a connection would form between us: but it is important to note the directionality here. The similarity precedes the connection.

Second, we can say that two entities are connectively similar if they share similar sets of connections. For example, Paul and Michelle may be political activists, but working for different political parties. In such a case, they will share the same *types* of connections, but with different sets of friends. Such sets of connections are (more or less) *isomorphic*. It is worth noting that this isomorphism will tend to lead to a connection between the two groups (political parties tend to interact with other political parties, but much less so with hockey teams) which in turn again leads to connections between the members.

g. Distribution

At this point we reach a central concept of connective knowledge, that of *distributed* knowledge.

In the previous section we looked, a bit glibly, at the possibility of political parties interacting with each other. And this is a concept we can intuitively grasp; we see it every day in political debates, in the legislature, and as represented in political polls and newspaper articles.

But a political party is not (*per se*) a self-contained entity: it is an assemblage of individual people where these people are connected through some sort of common process (usually but not always involving a commonality of belief and participation in a membership process, such as signing a membership card and paying five dollars, along with an organized and often guided set of interactions between the members, such as are evidenced through a primary process or political convention).

The political party is a distributed entity. What is important to note is that it is *more* than merely a collection of associated or even similar people. A group of people, even if they all hold the same beliefs, and even if they all know each other, does not constitute a political party. Nor is it a question of quantity: a group of five people may constitute a (very unpopular) political party, while a very large group may not have any political existence at all. What makes a political party (and similar entities, such as corporations, hockey teams and university faculties) is the set of *connections* between its members, the existence of which is often manifest and recognized with special documents and legal standing.

It's a nebulous concept. The political party does not exist, is not contained, in any of its members, nor is it a mere aggregation of the properties or number of its members, but it would not exist without its members. The existence of the political party is *distributed* – there is no single place it could be said to be, but many places in which its existence could be said to be manifest. Each member forms a part of the political party, but they are not a miniature version of the party as a whole. The properties of the party are separate and distinct from those of the members.

We have here once again reached the concept of emergence, but from a different direction. Any property the political party may have is an emergent property. Consequently, it is a property that exists (in our minds) solely by virtue of it having been *recognized or interpreted* as such (which is why we have a formal process of 'recognizing' political parties). And yet, while this property depends on the constituent members, it is not in turn a property of the members (Davidson (1980) calls this supervenience). The emergent properties of a distributed

entity exist solely as a consequence of the *organization* of its parts, and not its membership, and specifically, from the fact that these parts are connected in a certain recognizable way.

Strictly speaking, every entity in the world is a distributed entity (save, perhaps, indivisible subatomic particles – and (in my view) these may exist only by virtue of a reverse distribution, consisting entirely of entities that are *larger* than they are, much like a point in a moire pattern – but this is *very* speculative). Every entity is composed of additional entities, and the properties of the entity in question are not all mere reflections of the smaller entities, but rather, unique properties, that come into existence because of the organization of those entities. Thus the same collection of carbon atoms may result in very soft charcoal or a very hard diamond.

When we speak of one of those properties, therefore – say, the hardness of a diamond – there is no *place* that we can point to where this property is located. There is no specific instance of the hardness of the diamond, save in our perception and interpretation that carbon atoms, when organized *this* way, are what we call 'hard'. The property of being hard, in addition to being distributed across the carbon atoms that constitute a diamond, in addition exists only as a result of our perception of it. Strictly speaking, were there no perceivers to recognize diamonds as being 'hard', there would be no 'hardness' for diamonds to have.

h. Meaning

Above, we discussed the possibility of considering properties separated from the physical entities that are instances of them. Thus, for example, we can think of 'red' without thinking of a 'red thing'. At the time, we said that it does not follow that there is any specific entity such as 'redness'. But now we have to ask, in what does 'redness' consist. Because there is a sense in which 'redness' is real: it is something we all understand, a concept that is useful in our daily lives. The sentence, "This photo needs more redness" is not something we would immediately dismiss as nonsense.

The concept of 'redness' is an example of distributed meaning. There is no particular place we could point to where the 'meaning' of 'redness' is located. Indeed, that we have a concept such as 'redness' in our minds is in itself only something that we could know through interpretation of the myriad patterns presented in our consciousness and our behaviour. No doubt we have numerous other similar concepts, however, because they are not salient – because they never play a role in higher order cognitive behaviour – we do not recognize them. We are, in a sense, blind to them, until through some process (such as a Rorschach test) they are searched for and observed.

In a sense, having the concept of 'redness' in our own mind is similar to having 'liberal' as a description of a political party. Low-level subsymbolic concepts exist in our minds – collections of connected neurons that themselves do not have meaning we would recognize, but which in combination eventually form higher-order structures that do correspond with the meanings of words (or melodies, or icons, etc), such as 'redness'. Saying that we have the concept of 'redness' in our mind is to pick out a particularly salient set of collections of connected neurons.

We can understand intuitively how the meaning of a word is distributed in this way if we reflect on the meaning of a specific word. For example, consider the word 'Paris'. We would at first blush take this word to refer to – and be – something concrete and definite, a city in north-central France. But the use of the word 'Paris' conjures different associations for different people. For example, 'city', 'France' and 'Eiffel Tower'. And some people think of plaster, other people think of Hilton, other people think of the left bank, other people think of Kurt Vonnegut.

But more: when we say that the meaning of the word 'Paris' is distributed, what we mean in addition is that the meaning of the word is *constituted* in part out of the same elements that constitute the meanings of these other words. We might say (loosely) that the connection between subsymbolic entities A,B, and C constitute the meaning of 'Paris', while the connection between B,C and D gives us the meaning corresponding to 'plaster' (obviously this is a vast oversimplification). When the meanings of words are distributed, the basis of their meanings – the smaller subsymbolic entities that make up the meanings – are *intermingled*. In a certain sense, you *can't* understand what 'Paris' means unless you at the same time understand what a set of other words, and indeed, other concepts (such as 'naming') mean.

This may seem like a hard, even impossible, concept, but it is one that we work with and manage every day. One might ask, for example, “where is Edmonton?” The answer to that question does not exist as some sort of determinate, singular entity; it is mixed in with a variety of other concepts. “Edmonton is in Alberta,” we might answer, which draws the concept of 'Alberta' into our understanding. “Edmonton is in the Palliser Triangle,” a geographer might say, which in turn draws in another set of associations as part of the answer. Edmonton is at latitude 52 north, a cartographer might respond, involving in our understanding the nature and employment of Cartesian geometry. The answer to the question 'where is Edmonton' and the meaning of the word 'redness' are of a similar nature, entrenched in a complex and interwoven networks of other meanings.

i. Shared Meaning

From the writings of people like Wittgenstein we get the idea that meanings, in the truest sense, exist only when they are shared by a community of speakers. Wittgenstein (1999) even went so far as to say there could be no private language, that meaning is possible only if it is shared publicly.

This strikes many people as wrong because they think of meaning as reference or (following Kripke (1980)) following a necessary order of things in the world. The Tarski (1944) definition of truth – “‘Snow is white’ is true if and only if snow is white” strikes an intuitive chord with people, as it establishes an observable empirical basis in the meanings of words.

And indeed, it is our common experience of an independently existing physical world that also leads us to such intuitions. Never mind old folk tales like “The Eskimos have 22 words for 'snow,’” the fact remains that when an Inuit says 'snow' and when a Brazilian says 'snow' they mean the same physical entity, specifically, crystalline H₂O.

And yet – an Inuit would say 'snow' in Inuktitut, and a Brazilian in Portuguese, and the words in these two languages are different, and reflect different interpretations of reality. Languages are not isomorphic (Chomsky notwithstanding). The basis of English structure, for example, lies in the distinction between myself and the other, while in the French it is myself, my body, and the other. Neither is factually incorrect; snow is 'white' in each instance, and yet meaning diverges (or may diverge; as Quine (1964) says, we can't know for sure).

Meaning, for Wittgenstein, is established in the *act* of communicating. From the perspective of the current discussion, we would say something like this: the shared meaning of the word 'Paris' is an emergent property of the set of specific interactions between people involving the use of the word 'Paris' or of words associated with the word 'Paris'. Or as Wittgenstein said it, "Meaning is use."

It is important at this juncture to understand that this account of meaning does not contradict, nor even compete with, the account of meaning given above. Just as we can examine two different people to find different meanings of the word 'Paris', so also can we examine two distinct types of entity – a person and a society – in order to understand its meaning. Because there is no single and distinct entity which the meaning of the word 'Paris' must be. What connections are salient, what entities are salient, in our determination of the meaning of the word is a matter of context, a matter of interpretation.

When Wittgenstein says that there can be no private language he is, strictly speaking, wrong. I have numerous private words (which I won't share here, for otherwise they wouldn't be private, and I wouldn't have an example any more) and could in principle have a private language. Because having a language is not a case of *knowing* the language, as Wittgenstein (on some interpretations (Kripke 1982)) argues. Having a language is being *organized* in a certain way. This organization is the 'black box' that gives us, as Ryle (1949) would say, "dispositions" to behave in certain ways, to (for example) utter the word 'Paris' when presented with a certain phenomena.

Indeed, to turn this around, 'knowing' anything is of a similar nature. To 'know' something is not to be possessed of a certain fact. There is no 'instance' of a piece of knowledge in our head. To 'know' is to be organized in a certain way, to have, if you will, a certain regularly occurring pattern of neural activity (and consequently, disposition to behave). Knowledge is, as Hume said, a 'habit of the mind'.

Indeed, if speaking a language, using a language, required 'knowing' a language (in the cognitive sense), then a child would not be able to speak a language, for a child employs linguistic constructions that he or she could not possibly identify or name (as a student of French, it is very frustrating to see a six year-old exercise more capacity in the language than I can). This is the sort of phenomenon that was perplexing to Chomsky (1980): how could someone speak a language without the mental capacity to 'know' it? But this is not sufficient reason to suppose Chomsky's syntactic structures are innate; it makes as much (if not more) sense to believe that they are (subsymbolic) organizations of neural connections.

None of this, though, should be interpreted to mean that language is *merely* a mental phenomenon. We remarked above that the meaning of the word 'Paris' could be understood both

from a personal and social point of view. But additionally, it should now be noted, that the personal and the social do not operate independently of each other. It is, after all, no coincidence that children grow up speaking the same language as their parents. The experience of linguistic elements as perceptions leads to the formation of linguistic elements as neural and mental structures, and the interaction of these back and forth lead to their being associated, and over time, more similar. Use of the language influences the speaker; use of the language influences the language.

j. Organization

It may seem odd at this juncture to speak of a language as a social phenomenon, and a language as a mental phenomenon, in much the same terms, and indeed even, interchangeable.

But it is not odd, nor even unintuitive, when it is recognized that meaning, both socially and neurally, have the same origin: *meaning is an emergent phenomenon*, arising from the connections between underlying entities. Socially, the underlying entities are speakers of the language, while mentally, the underlying entities are neurons and subsymbolic neural structures.

How could these be the same? One might ask. But that's a bit like asking how a neural cell and a Popsicle could both be pink. Or a bit like asking how there could at the same time be a thousand neural cells in a layer and a thousand people in a market. What makes language, both social and personal, similar is that both are derived from the same set of principles. And, indeed, it is due to their following the same principles that makes language possible at all! If we could not in society replicate the same sort of things that happen in our own minds, there would be no means by which we could communicate at all. Consider rabbits, who have active (though rabbit-like) mental lives: without the capacity to share meaning through networks of organized interaction, they are utterly unable to form a language.

The principles of organized networks of connections have received much attention in recent years, and deservedly so. We understand a great deal about how such networks work and about their properties. Conceptually, they have been studied under the heading of graph theory. Concrete instances of networks have been studied in the words of Watts (2003) and Barabasi (2002), among others. Computationally, networks have been the locus of investigation by people like Minsky and Papert (1987), Rumelhart and McClelland (1989). Social networks, and social networking software, have become a minor industry. And, of course, the internet itself has given us a large scale network to study up close and in detail.

Most work (to my observation) has been centered in two major areas: first, the properties of different types of networks (for example, random networks, loosely coupled networks, etc), and secondly, properties of the propagation of information through networks (as instanced in, for example, the 'six degrees' phenomenon). Additionally, though the investigation of dynamic networks, it has been shown how networks can grow naturally, with no intent or design, on the basis of a few very simple principles. Observation of these phenomena have explained such things as power laws, which describe disparate numbers of connections between nodes

in the networks, and cascade phenomena, in the process of examining the propagation of ideas and diseases through a society (or through a human body).

Much less has been said about what is probably the most important implication of this work: if a human mind can come to 'know', and if a human mind is, essentially, a network, then any network can come to 'know', and for that matter, so can a society. Just as the meaning of a word can be both personally based and culturally based, so also can knowledge itself be both personally and culturally based. Moreover, because we know that people can learn, we can now also that societies can learn, and conversely, through the study of how a society can learn, we can understand more deeply how a person can learn.

k. Social Knowledge

Social knowledge is to a society what personal knowledge is to a person. It is a result of the connections between the individual members of society, resident in no single one of them, but rather a property of the society working as a whole. Numerous instances of such connections occur; where certain of those connections become salient, and are frequently activated through use, they are recognized as forming a distinct entity, producing a distinct type of knowledge.

As an example, consider the knowledge of 'how to fly a person from England to Canada in a 747'. No single person possesses this knowledge, because it is the result of combining numerous instances of personal knowledge – from how to make tires to how to navigate a 747 to how to execute a landing while keeping the airplane intact. What makes these individual bits of knowledge combine to form an instance of social knowledge is that they are connected; knowing how to land an aircraft depends on, and makes sense, only in the context of knowing how to fly an aircraft, or to build an aircraft.

Though many instances of social knowledge go unobserved and unremarked, numerous examples may be adduced. For example, the knowledge of 'the value of wheat' at a given time is a type of social knowledge; it is the knowledge that results through the connections of millions of wheat buyers and wheat sellers in a marketplace. No individual has a grasp of 'the value of wheat' – they each make decisions to buy or to sell based on their own individual knowledge and needs. It is true that there is a 'market value' of wheat – but again, this is an *interpretation* of that social knowledge – not all instances of wheat-trading are taken into account, only those expressed in financial terms, and not all wheat-traders are considered (the child receiving wheat from her mother, for example).

Smith's 'invisible hand of the marketplace' is but *one* way of looking at particular types of social knowledge, specifically, those that may be expressed quantitatively, and on the basis of quantitative reasoning. Wheat may be valued non-quantitatively – by its taste, for example. Consider how society values chocolate, in comparison. The 'value of wheat', looked at from a connective perspective, is a consideration of the interaction between all statements concerning 'value' and all statements concerning 'wheat', and an interpretation of those statements. That we today express the value of wheat in economic terms says as much about the salience of financial value in today's society as it says about wheat.

Social knowledge has recently attained recognition (and value) under the heading of Surowiecki's (2005) 'wisdom of crowds'. But it is worth noting that many of Surowiecki's examples are cases where individual guesses "aggregated and then averaged." While Surowiecki stresses (correctly) the autonomy of those guesses, he does not so stress the equally important fact that those guesses are not independent events – they are connected, in some key way, to each other (for example, the people guessing the temperature of a room have also the property of being in the same room; those estimating the weight of objects all see the same objects, and in the same way).

Social knowledge is not *merely* the aggregation and averaging of individual knowledge (as if there could be such a thing – consider how in guessing weights we use a medium as average, while in electing leaders we use a mode as average). That is why such aggregation is not necessarily reliable – an aggregation that is considered independently of the connections between entities is like a count that is considered independently of the membership of a set. Consider, for example, counting sheep without worrying about whether what is being counted is a sheep. It can work sometimes – in sheep-filled rooms, for example. But more often, it will mislead.

I. Power Laws and Inequalities

Much of the work in networks has been on what are called 'scale-free' networks. In a scale-free network (as people like Barabasi have shown) some entities in the network have a much higher degree of connectedness than others. True, in any, there will be a certain variance in distribution, but in a scale free network this variance can be extreme. Consider, for example, a network like the internet, where some sites, such as Google, have millions of visitors, while other sites have only one or even none.

A network of this sort forms through a *dynamic process* where the presence of one entity leads others to connect to it. For example, consider the act of creating links on a web page. In order to create a useful link, it is necessary to connect to a site that already exists. This means that, all other things being equal, a site that was created first will obtain the most links, because it will have been a candidate for linkage for all subsequent websites, while a site that was created last will have the fewest links, because it has never been a candidate for links.

This effect can be magnified when *preferential attraction* is considered. For when creating a link on a web page, a designer wants not merely to link to a random page, but to a *good* page. But how does one judge what counts as a good page? One way is to look at what other people are linking to. The probability that the first page created will be found is greater than that for any other page, which means that the first page will obtain even more links that it would receive through random chance. With this and similar drivers, some websites obtain millions more links than others.

What's interesting is that though a similar process leads to the formation of scale-free networks in other areas, not in all cases is such an extreme inequality reached. What happens is that in some cases a structural upper limit is reached. Consider, as Barabasi does, the cases of airports and the power grid. Both are developed according to similar principles (airlines want

to land flights, for example, where other airlines land flights). And, not unexpectedly, a power-law distribution occurs. But there is an upper limit to the number of aircraft that can land in a single airport, and consequently, a limit to the size of the inequality that can occur.

Various writers (for example Shirky) write and speak as though the power law were an artifact of nature, something that develops of its own accord. And because it is natural, and because such systems produce knowledge (we will return to this point), it is argued that it would be a mistake to interfere with the network structure. This argument is remarkably similar to the argument posed by the beneficiaries of a similar inequality in financial markets. The rich get richer, benefiting from an unequal allocation of resources, but efforts to change this constitute 'interference' in a 'natural phenomenon', the invisible hand of the marketplace, intelligently allocating resources and determining priorities.

This may be true, if we think of networks as natural systems. But the absence of limits to the growth in the connectivity of some nodes should alert us that there is something else going on as well. And it is this: the networks we describe, and in some cases build (or through legislation, protect), are *interpretations* of the multifarious connections that exist in an environment or in a society. They depend, essentially, on a *point of view*. And, arguably, the inequalities of links on the web or money in society represent the prevalence of one point of view, or some points of view, over others. But to understand how this could be so, we need to look at networks, not as physical systems, but as semantical constructs, where the organization of links is determined as much by similarity and salience than by raw, epistemologically neutral, forces of nature.

m. Knowledge

What does it mean, even to say that a sentence has semantical import? To say, similarly, that we 'know' something? As suggested above, most of us remain committed to something like a Tarski semantics: we know something just in case what we know happens to be true. But of course, this fails to tell the whole story. The knowledge needs to be, in some way, in our mind (or in our society); it needs to be a 'belief'. And (so goes the argument) it needs to be in some way justified, through a process of verification, or at the very least, says Popper (1963), through the absence of falsification.

This view has its difficulties, as the Gettier (1963) counterexamples suggest. But (in my view) its most significant difficulties emerge when we try to articulate what it is that we know. Consider, for example, 'snow is white'. Sure, one could check some snow in order to determine that it is white, but only of one first understood what is meant by 'snow' and 'white' (not to mention, as Clinton taught us, 'is'). But as discussed above, that constitutes *the* meaning of, say, 'snow', is far from clear. there is no such single entity. What it means is a matter of interpretation. So, for example, does enumerating what constitutes instance of snow. Does 'yellow snow' count? Does snow produced by artificial ice machines count?

The behaviourist response to such dilemmas is to define 'knowing' that snow is white as a disposition to utter the word 'white' when presented with the question, 'what colour is snow'. And while we most certainly employ such tactics in the evaluation of knowledge (measuring

responses is, after all, the basis of testing and examinations), it remains unsatisfactory, because we need to know what puts the disposition to say 'white' into a student's mind in the first place. Is it the whiteness of snow? Is it the memorization of the sentence 'snow is white'? Is it a comprehensive understanding of the process of crystallizing H₂O?

From the discussion above, it should be clear that on the account being given here, to 'know' that 'snow is white' is to be organized in a certain way (one that is evidenced by uttering 'snow' when asked). To be organized in such a way as to have neural and mental structures corresponding to the words 'snow', 'is' and 'white', where those structures are such that the concept 'snow' is closely associated with (in certain contexts) the concept 'white' (obviously this is a gloss). Knowing that 'snow is white' is therefore being organized in a certain way, but not in some particular way (we couldn't examine one's neural organization and be able to say whether the person knows that snow is white).

This is a very different model of what it means to 'know' – for one thing, because it is based on organization and connectedness in the brain, the concept of justification and even of belief are nowhere present. What we 'know' is, if you will, a natural development that occurs in the mind, other things being equal, when presented with certain sets of phenomena; present the learner with different phenomena and they will learn different things. Like the Portuguese word for 'snow', for example. And whether something counts as 'knowledge' rather than, say, 'belief' or 'speculation', depends less on the state of the world, and more on the strength or degree of connectedness between the entities. To 'know' something is to not be able to not know. It's like finding Waldo, or looking at an abstract image. There may be a time when we don't know where Waldo is, or what the image represents, but once we have an interpretation, it is not possible to look without seeing Waldo, without seeing the image.

No wonder Dreyfus and Dreyfus (2005) talk about 'levels' of knowledge, up to and including an almost intuitive 'expert' knowledge. As a particular organization, a particular set of connections, between neural structures is strengthened, as this structure becomes embedded in more and more of our other concepts and other knowledge, it changes its nature, changing from something that needs to be triggered by cue or association (or mental effort) into something that is natural as other things we 'know' deeply, like how to breathe, and how to walk, structures entrenched through years, decades, or successful practice. Contrast this to a cognitivist model of knowledge, where once justification is presented, something is 'known', and cannot in later life be 'more known'.

n. Public Knowledge

'Public knowledge' is the explicit representation of social knowledge in language or some other concrete form. Public knowledge is what most people think of as 'knowledge' *per se*, it is what we attempt to teach our children, it is what is embodied in a canon and passed on to successive generations.

There are things known only by myself (think again of Wittgenstein's private language argument), such as who I like and why, or where I last stubbed my toe, that society either cannot or has no desire to come to know as a part of social knowledge. Such knowledge, personal

knowledge, does not externalize, because there is either no need or no mechanism with which to place it in the public domain.

Knowledge that is, for example, subsymbolic defies communication (it is not impossible to communicate, though – consider a shrug, a sigh, a knowing look). In order for private knowledge to become public knowledge, it must have some means of connecting with everything else that is considered public knowledge – through commonly understood utterances or actions.

But the mere communication of private knowledge in the public domain does not thereby convert it to public knowledge. It must be interpreted as such, recognized as such, in the public domain. In order for this to happen, the set of utterances ('Paris is the capital of France', say) must form a part of the communications, of the interactions, in the social network as a whole. Then this pattern of communication must in turn be *recognized* by some perceiver (or group of perceivers) as constituting a relevant underlying organization of communication informing (say) the behaviour of a society as a whole. Merely saying 'Paris is the capital of France' doesn't make it so; many other people must say it, and even then, the mere public utterance doesn't make it so; it be recognized as a constituent element of the body of knowledge possessed by a society.

It becomes evident that one's demonstration of having acquired 'knowledge' is very different in the case of public knowledge than it is for private knowledge, even when the instance known is the same. Knowing privately that 'Paris is the capital of France' may consist merely of writing the appropriate word on a piece of paper, but knowing the same thing publicly involves a complex of interactions and behaviours, consisting essentially of immersion (becoming a part of, and entity within the organization) in the knowing community, so that utterances of the word 'Paris' reflect, and are seen to reflect, an instance of the (generally recognized fact that) 'Paris is the capital of France'.

Knowing publicly is, as Kuhn (1970) said, knowing 'how to solve the problems at the end of the chapter'. It involves being able not only to produce specific behaviours, but in providing evidence of sharing in the same network of associations and meanings as others in the community, sharing a language, methodologies, riverbed assumptions. Failure to personally know something creates only a personal risk – one might travel to Lieges looking for the French parliament instead of to Paris. Failure to know publicly carries a greater risk: that of not being considered to be a part of the knowing community, of being, therefore, excluded from its interactions, and of being misunderstood when attempting to communicate.

This is why writers such as Wenger (1999) find such importance in communities of practice, and more, see such involvement as a process of (as he says) personal becoming. Interaction in a community of practice is to a significant degree an alignment of (certain parts of) one's personal knowledge with public knowledge – immersion produces a salience of certain utterances, certain practices, and thus promotes the development of corresponding (but probably not isomorphic) structures in the mind. It exposes a person to instances of knowledge statements and practices which, if they are sufficiently similar to preexisting organizations of neural and mental structures, increase, through association, their strength and importance.

Personal knowledge is distinct from public knowledge, but the two go hand in hand, and a person who is considered 'highly learned' is one who has internalized, to an expert degree, a great deal of public knowledge.

o. Knowing and Not Knowing (Cascade Phenomena)

But on what do we base public knowledge? What is the process of interpretation and recognition by which we, say, accept the theory of gravity and reject stories about flying saucers? What makes some knowledge part of 'social knowledge' and other knowledge (merely?) personal knowledge? Why would a community accept some things as 'known' and not others?

Knowledge is a network phenomenon, to 'know' something is to be organized in a certain way, to exhibit patterns of connectivity. To 'learn' is to acquire certain patterns. This is as true for a community as it is for an individual. But it should be self-evident that mere organization is not the *only* determinate of what constitutes, if you will, 'good' knowledge as opposed to 'bad' (or 'false') knowledge.

Consider public knowledge. People form themselves into communities, develop common language and social bonds, and then proceed to invade Europe, or commit mass suicide, or starve themselves to death. Nor is personal knowledge any reliable counterbalance to this. People seem to be as inclined to internalize the dysfunctional as the utile, the self-destructive as the empowering.

These are examples of cascade phenomena. A signal propagates from one entity in the network to the next to the next, seemingly without restraint, until it reaches every entity in the network. Such phenomena exist in the natural world as well: the sweep of the plague through medieval society, the failure of one hydro plant after another in a blackout, the bubbles in the stock market. Cascade phenomena are in one sense difficult to explain, and in another sense deceptively simple.

The sense in which they are simple to explain is mathematical. If a signal has more than an even chance of being propagated from one entity in the network to the next, and if the network is fully connected, then the signal will eventually propagate to every entity in the network. The speed at which this process occurs is a property of the connectivity of the network. In (certain) random and scale free networks, it takes very few connections to jump from one side of the network to the other. Cascade phenomena sweep through densely connected networks very rapidly.

The sense in which they are hard to explain is related to the question of why they exist at all. Given the destructive nature of cascade phenomena, it would make more sense to leave entities in the network unconnected (much like Newton escaped the plague by isolating himself). Terminating all the connections would prevent cascade phenomena. However, it would also prevent any possibility of human knowledge, any possibility of a knowing society.

p. Structure and Process

Nothing guarantees truth. It is tempting to suppose that we could easily restrain the excesses of cascading communities through a simple application of qualitative or quantitative knowledge obtained through other domains, but in practice we gain no increased certainty or security.

Consider, for example, qualitative knowledge. We are as apt to be misled by the information given by our senses as by any wayward community. Descartes (1996) records simple examples, such as mirages, or the bending of a stick in water, to make the point. Today's science can point to much deeper scepticism. Perception itself consists of selective filtering and interpretation. The mind supplies sensations that are not there. Even a cautiously aware and reflective perceiver can be misled.

Quantitative knowledge, the cathedral of the twentieth century, fares no better. Though errors in counting are rare, it is a fragile a process. *What* we count is as important as how we count, and on this, quantitative reasoning is silent. We can measure grades, but are grades the measure of learning? We can measure economic growth, but is an increase in the circulation of money a measure of progress? We can easily mislead ourselves with statistics, as Huff (1993) shows, and in more esoteric realms, such as probability, our intuitions can be exactly wrong.

In the realms of observation and mathematics, we compensate for these weaknesses by recognizing that a single point of view is insufficient; we distribute what constitutes an 'observation' through a process of description and verification. If one person says he saw a zombie, we take such a claim sceptically; if a hundred people say they saw zombies, we take it more seriously, and if a process is described whereby anyone who is interested can see a zombie for themselves, the observation is accepted.

Even then, we demonstrated caution though an explicit recognition that in the process of seeing we are interpreting. An observation of a certain phenomenon may be labelled the observation of 'zombies', but we consider alternative explanations. This is aided by ensuring that the observers of the phenomena have different sets of prior experiences, different world views, different ways they could interpret the phenomenon. Having every member of a religious sect report seeing zombies is less reliable than having members of different sects, scientists and sceptics report the same thing.

In quantitative reasoning, we take care to ensure that, in our measurements, we are measuring the same thing. Through processes such as double-blind experimentation, we additionally take care to ensure that our expectations do not influence the count. In statistical reasoning, we take care to ensure that we have a sufficiently random and representative sample, in order to ensure that we are measuring one phenomenon, and not a different, unexpected phenomenon. In both we employ what Carnap (1967) called the requirement of the total evidence: we peer at something from all angles, all viewpoints, and if everybody (or the preponderance of observers) conclude that it's a duck, then it's a duck.

q. Reliable Networks

Connective knowledge is supported through similar mechanisms. It is important to recognize that a structure of connections is, at its heart, *artificial*, an *interpretation* of any reality there may be, and moreover, that our observations of emergent phenomena themselves as fragile and questionable as observations and measurements – these days, maybe more so, because we do not have a sound science of network semantics.

Where structures of connections (ie., networks) differ from sets of observations or measurements is that there is in principle no external entity to which we can appeal in order to check our understanding. In a networked society, every person is a member of the network, and all things being equal, there is not some other networked society against which we can test our conclusions (prior to the days of global communications, societies did test themselves one against the other, but unfortunately through war and other conflict, a solution that was worse than the problem and which clouded their ability to interpret connections in a rational and dispassionate way).

We have already seen that there are different types of networks – different ways sets of connections between entities can be generated and organized. Where the mechanisms that support knowledge in other realms come into play in the world of networks is that these mechanisms become *properties* of the networks we rely upon to generate and contain knowledge.

In a network, a cascade phenomenon is akin to jumping to a conclusion about an observation. It is, in a sense, a rash and unthinking response to whatever phenomenon prompted it. The mechanisms that push a stock market into a bubble are akin to a person being convinced by looking at the same thing over and over again. A network in the throes of a cascade needs the internal equivalent to a 'second set of eyes' to act as the bearer of sober second thought.

This capacity is crucially dependent on the structure of the network. Just as a network with no connections has no capacity to generate knowledge, a fully connected network has no defense against jumping to conclusions. What is needed is to attain a middle point, where full connectivity is achieved, but where impulses in the network ebb and flow, where impulses generated by phenomena are checked against not one but a multitude of competing and even contradictory impulses.

This is what the human mind does naturally. It is constructed in such a way that no single impulse is able to overwhelm the network. A perception must be filtered through layers of intermediate (and (anthropomorphically) sceptical) neurons before influencing the formation of a concept. For every organization of neurons that achieves an active state, there are countless alternative organizations ready to be activated by the same, or slightly different, phenomena (think of how even a seed of doubt can destabilize your certainty about something).

Knowledge in the mind is not a matter of mere numbers of neurons being activated by a certain phenomenon; it is an ocean of competing and conflicting possible organizations, each ebbing and subsiding with any new input (or even upon reflection). In such a diverse and

demanding environment only patterns of organization genuinely successful in some important manner achieve salience, and even fewer become so important we cannot let them go.

r. Network Structure

It is with these considerations that we return to the consideration of scale-free networks.

As mentioned above, a scale free network is characterized by a small number of entities is numerous connections, and a large number of entities with much fewer connections. It is worth noting that such networks are very tightly connected – in a scale free network a piece of information can reach an entire network very quickly.

While the human brain exhibits some scale-free properties, it is nonetheless not as imbalanced as even things like the economic system or the World Wide Web. Some neurons (or neural clusters) play important and central roles in the brain, but they are not millions of times more connected than most of the others. The brain is densely connected, but the connections are more equitably distributed.

This is no doubt a result of the physical limitations of neurons. But even more importantly, reducing the scale of the inequality between neurons also slows the propagation of impulses through the brain. It allows sub-organizations to develop – the alternative interpretations we can experience when observing a Gestalt phenomenon, for example. Were the structure of human thought to be replicated at the social level, what we would see is essentially a community of communities – the part of us (society) that likes knitting, the part of us that is a hedonist, the part of us that enjoys a good novel.

Networks that exhibit extreme power law distributions are unstable. Because, though the mechanism of highly connected nodes, a single impulse can be broadcast and accepted by the entire network all at once, there is no constraint should the impulse prove to be destructive or dysfunctional. The extremes in human social behaviour, wrought on a smaller scale by chieftains and kings, and on a global scale by mass media, should serve as ample evidence of this. With nothing to counteract an irrational impulse, the characteristic of the one becomes the characteristic of the whole, and the society spirals into self-destruction.

Chieftains, kings and broadcast media are inventions. They are ways we represent, in physical form, the set of connections we perceive to be extant in a society. But as interpretations of a complex set of connections, they are subject to individual points of view, prior conceptions and prejudice. As Rousseau (1968) observed, when the mechanisms of the whole are put into the hands of the few, the very nature of the whole is interpreted in such a way as to serve the needs of the few.

In order therefore to successfully counterbalance the tendency toward a cascade phenomenon in the realm of public knowledge, the excesses made possible by an unrestrained scale-free network need to be counterbalanced through either one of two mechanisms: either a reduction in the number of connections afforded by the very few, or an increase in the density of the local network for individual entities. Either of these approaches may be characterized under the same heading: *the fostering of diversity*.

For, indeed, the mechanism for attaining the reliability of connective knowledge is fundamentally the same as that of attaining reliability in other areas; the promotion of diversity, through the *empowering of individual entities*, and the reduction in the influence of well-connected entities, is essentially a way of creating extra sets of eyes within the network.

s. Truth

Recently a series of discussions took place regarding the relative 'truth' of entries in Wikipedia, a collection of articles created through a process of collective authoring, and *Encyclopaedia Britannica*, a collection of articles about similar topics written by a series of experts (Giles 2005).

Such discussions are difficult to resolve because, as we have seen, what constitutes the 'truth' of the matter is very much a matter of interpretation. Truth, as commonly conceived, is said to be based on facts (and mediated through 'truth-preserving' inference), but if even the simplest observation depends to a great degree on interpretation, then the foundation of truth itself is equally suspect.

And yet this post-modernist attitude to knowledge is difficult to reconcile with our intuitions. We do rely on facts, there is knowledge, and what counts as knowledge has the virtue of being true. And when a body of work such as Wikipedia is examined, some statements are regarded as, and universally acknowledged to be, true, while others (happily a much smaller set) are found to be 'not true'. This, indeed, was the basis on which the *Nature* comparison of the two encyclopedias was based.

What distinguishes Wikipedia from Britannica is not so much the account of truth it embraces as the process through which it arrives at truth. Wikipedia, much more so than Britannica, represents an instance of connective knowledge – it is an attempt to capture, as public knowledge, what can be observed via the interactions of numerous instances of private knowledge.

It should be clear and obvious at the outset that this is not some process whereby individual points of view are aggregated and averaged – such mechanisms are more evident in entities such as Google and Technorati and Digg. Rather, Wikipedia, through iterations of successive editing, captures the output of *interactions* between instances of private knowledge. The majority, typically, does not rule on Wikipedia; what matters is what is produced through the interaction.

In the case of Britannica, the same is the case. The authors, as experts, are typically those immersed in a knowledge community, who have in turn internalized the knowledge (both social and public) possessed by that community. The expert serves as a dedicated *interpreter* of that knowledge, an interpretation that is additionally subject to subsequent interactions with proof-readers and editors.

A priori, each approach has an equally plausible claim to being an effective (and reliable) generator of knowledge, which raises the question of how we will resolve the truth of the matter when (inevitably) there exists a point at which one encyclopedia says a statement is true and the other says the opposite.

Truth, in such a case, will come to depend not so much on the facts of the matter, but rather, through an examination of the process through which various types of knowledge are accumulated and interpreted. Just as the reliability of an observation report depends on how the observation is made, so to will the proclamations of connected communities of knowers.

t. Knowing Networks

Arguably, the following criteria will determine the difference:

First, *diversity*. Did the process involve the widest possible spectrum of points of view? Did people who interpret the matter one way, and from one set of background assumptions, interact with people who approach the matter from a different perspective?

Second, and related, *autonomy*. Were the individual knowers contributing to the interaction of their own accord, according to their own knowledge, values and decisions, or were they acting at the behest of some external agency seeking to magnify a certain point of view through quantity rather than reason and reflection?

Third, *interactivity*. Is the knowledge being produced the product of an interaction between the members, or is it a (mere) aggregation of the members' perspectives? A *different* type of knowledge is produced one way as opposed to the other. Just as the human mind does not determine what is seen in front of it by merely counting pixels, nor either does a process intended to create public knowledge.

Fourth, and again related, *openness*. Is there a mechanism that allows a given perspective to be entered into the system, to be heard and interacted with by others?

It is based on these criteria that we arrive at an account of a knowing network. The scale-free networks contemplated above constitute instances in which these criteria are violated: by concentrating the flow of knowledge through central and highly connected nodes, they reduce diversity and reduce interactivity. Even where such networks are open and allow autonomy (and they are often not), the members of such networks are constrained: only certain perspectives are presented to them for consideration, and only certain perspectives will be passed to the remainder of the network (namely, in both cases, the perspectives of those occupying the highly connected nodes).

Even where such networks are open and allow autonomy (and they are often not), the members of such networks are constrained: only certain perspectives are presented to them for consideration, and only certain perspectives will be passed to the remainder of the network (namely, in both cases, the perspectives of those occupying the highly connected nodes).

u. Remnants

This new knowledge is not inherently any more reliable than the old. A community that limits its diversity, that becomes closed, is as liable to err as a person who refuses to look around, refuses to take measure. A person, exposed only to limited points of view, with limited opportunities to interact, will be similarly bereft of insight.

It is, after all, a form of knowledge we have had all along, just as we have always have qualities, always had quantities.

Connective knowledge is no magic pill, no simple route to reliability. As the examples mentioned above (*part o*) demonstrate, a knowledge-forming community can be easily misled or deluded, just as a person can suffer from delusions and misunderstandings.

Indeed, if anything, the sort of knowledge described here is perhaps even more liable to error, because it is so much more clearly dependent on interpretation. Knowledge derived from a pattern may be formed from a partial pattern; the perceiving mind fills in the gaps of perception. From these gaps spring the seeds of error.

Moreover, as we enter the connected age, we live with remnants of the previous eras, years when connectivity in society was limited, control over perspective maintained by the beneficiaries of scale-free communications networks. History is replete with examples of the mind of one man, or one group in power, distorting the mechanisms of media to their own ends.

The examples range from very large to very small, from the rise of totalitarianism to the propagation of genocide to gender stereotypes, mass media marketing, and propaganda. Practitioners vary from dictators to slave owners to misogynists. The history of repression walks hand in hand with the history of the distortion of connective knowledge.

The purpose of this paper is not to provide truth, but to point the way toward the correction of these errors, both in ourselves and in our society. To show that, through attention to the underlying framework informing social and public knowledge, we can find a new renaissance, not perfection, but perhaps, a world less filled with ignorance and superstition.

Freedom begins with living free, in sharing freely, in celebrating each other, and in letting others, too, to live free. Freedom begins when we understand of our own biases and our own prejudices; by embracing autonomy and diversity, interaction and openness, we break through the darkness, into the light.

References

- Anderson, C. (2005) *The Probabilistic Age. The Long Tail*. Available from: http://www.thelongtail.com/the_long_tail/2005/12/the_probabilist.html [Accessed 07 December 2007]
- Ayer, A.J. (1952) *Language, Truth and Logic*. London: Gollancz. Dover Publications; 2nd edition.
- Barabasi, A.-L. 2002. *Linked: The New Science of Networks*. Perseus Books.
- Carnap, R. (1967) *Logical foundations of probability*. University of Chicago Press, 2nd edition.
- Chomsky, N. (1980) *Rules and Representations*. Basil Blackwell.
- Davidson, D. (1980) *Actions and Events*. Clarendon Press.

- Descartes, R. 1996. *Meditations on First Philosophy*. Revised Edition. John Cottingham, Ed. Cambridge University Press
- Dreyfus, H. & Dreyfus, S. (2005) Expertise in real world contexts. *Organization Studies*, 26(5), pp. 779–792.
- Dye, J. (2003) *Platonic Forms*. Available from: <<http://www.niu.edu/~jdye/forms.html>> [Accessed 07 December 2007]
- Gettier, E.L. (1963) Is Justified True Belief Knowledge? *Analysis* 23 pp. 121–123.
- Giles, J. (2005) Internet encyclopedias go head to head. *Nature* 438 900-901. Available from: <<http://www.nature.com/nature/journal/v438/n7070/full/438900a.html>> [Accessed 07 December 2007]
- Haselager, P. (1999) Neurodynamics and the Revival of Associationism in Cognitive Science. In: A. Riegler, M. Peschl, & A. Von Stein (Eds.), *Understanding Representation in the Cognitive Sciences*. Kluwer Academic/Plenum Publication. Available from: <<http://www.nici.kun.nl/~haselag/publications/NeuroAssoc99.pdf>> [Accessed 07 December 2007]
- Hinton, G.E. & Sejnowski, T.J. (1989) Learning and Relearning in Boltzmann Machines. In: David E. Rumelhart and James L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Volume 1. MIT Press.
- Huff, D. (1993) *How to Lie With Statistics*. W.W. Norton.
- Hume, D. (1999) *An Enquiry Concerning Human Understanding*. Tom L. Beauchamp, ed. Oxford.
- Kant, I. (1999) *Critique of Pure Reason*. Paul Guyer and Allen W. Wood, eds. Cambridge University Press.
- Kemmerling, G. (2001) Categorical Syllogisms. *Philosophy Pages*. Available from: <<http://www.philosophypages.com/lg/e08a.htm>> [Accessed 07 December 2007]
- Kitcher, P. (1985) *The Nature of Mathematical Knowledge*. Oxford University Press.
- Kripke, S. (1980) *Naming and Necessity*. Harvard University Press.
- Kripke, S. (1982) *Wittgenstein on Rules and Private Language*. Harvard University Press.
- Kuhn, T. (1970) *The Structure of Scientific Revolutions*. University of Chicago Press.
- Minsky, M. & Papert, S.A. (1987) *Perceptrons – Expanded Edition: An Introduction to Computational Geometry*. MIT Press.
- Popper, K. (1963) *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge.
- Quine, W.V.O. (1964) *Word and Object*. MIT Press.
- Jean-Jacques Rousseau. 1968. *The Social Contract*. Maurice Cranston, trans. Penguin Classics.
- Rumelhart, D.E. & McClelland, J.L. (1989) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Volume 1. MIT Press.

- Rumelhart, D.E., Hinton, G.E. & Williams, R.J. (1989) Learning Internal Representations by Error Propagation. In: David E. Rumelhart and James L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Volume 1. MIT Press.
- Ryle, G. (1949) *The Concept of Mind*. University of Chicago Press.
- Shirky, C. (2003) Power Laws, Weblogs, and Inequality. *Clay Shirky's Writings About the Internet*, Available from: <http://www.shirky.com/writings/powerlaw_weblog.html> [Accessed 07 December 2007]
- Stover, D. (2007) The Purple Prose Eater. Available from: <<http://www.debstover.com/purple.html>> [Accessed 07 December 2007]
- Surowiecki, J. (2005) *The Wisdom of Crowds*. Anchor
- Tarski, A. (1944) The Semantic Conception of Truth and the Foundations of Semantics. *Philosophy and Phenomenological Research* 4.
- Tversky, A. (2003) *Preference, Belief, and Similarity: Selected Writings*. MIT Press.
- Watts, D.J. (2003) *Six Degrees: The Science of a Connected Age*. W.W. Norton & Company.
- Wenger, E. (1999) *Communities of Practice: Learning, Meaning and Identity*. Cambridge University Press.
- Wittgenstein, L. (1999) *Philosophical Investigations*. Sections 244–271. Prentice-Hall, 3rd Edition.