Things Do Things to Things: Social calculus, agentive grammar and the beginnings of language[1]

Martin Edwardes
King's College London
martin.edwardes@kcl.ac.uk

Contents

Abstract

It is over 60 years since Chomsky (1957) promised a mechanistically cognitive explanation of human language. Today, a mechanistic solution is being implemented in devices like Siri, but it does not work in a way Chomsky and Generativists envisaged.

In response, Chomsky has replaced the increasingly baroque Generativist language engine, Transformational Grammar, with the single cognitive capacity of recursion, or MERGE. He takes the view that this was an evolutionary macro-mutation – sudden, unexpected, and inexplicable. However, many linguists find this an insufficient explanation for the way modern human language emerged. They have, instead, been looking to the significance of consciousness, attention, social calculus, and self-reference to understand the evolution of human communication.

This paper builds upon this revived cognitive voice in linguistics. It proposes that an understanding of language, personification, and self- and other-modelling relies on social calculus: our capacity to represent our familiar conspecifics as animate entities, and to attribute personhood to them. From the cognitive advantages of anticipating the actions of the animate things around you, through the attribution of intention to those animate things, to

---

the recognition of the personhood creating those intentions (and, thorough that, recognition of our own personhood), language has relied on general socio-cognitive mechanisms to give it its form and function. There seems to be no special language-only lump in the brain, and no special language-only "lump" in our evolution.

The paper looks at how the socio-cognitive forms that underlie our social calculus have been recruited into human communication, defining the basic linguistic forms we use. It also considers the role that personification plays in the metalanguage of language, allowing us to model the communicative act itself as a form of social calculus.

Introduction

Linguistics has a problem. Generativism, one of its key theoretical approaches, takes the view that language is the product of a unique encoding and decoding system which is completely separate from meaning exchange. As Noam Chomsky, the leading Generativist, himself puts it:

> "The language faculty itself uses the conceptual resources that are available … {*ellipsis in original text*} It's a little hard to say what language is 'itself'. Does the English language include the word *gravitation*? We're somewhere in an area now where our linguistic capacities and our science-creating capacities are interacting. We don't understand either of these systems enough to know where to go on.
> …
> "The Sciences provide completely different kinds of ways of looking at the world, which are completely counter to common sense. In fact, common sense – at least in the advanced sciences – has been completely abandoned; it's not even a criterion any more. But that's a very modern, very special development, which holds [only in] {*square brackets in original text*} certain areas." (Chomsky, 2012, pp74-75).

To question whether the word *gravitation* is part of the English language is certainly counter to common sense; but does this rather arbitrary pronouncement show that linguistics is an advanced science where common sense can be abandoned if it makes the theory fit the evidence? Clearly not if we don't understand our linguistic capacities sufficiently to know where to go on.

In contrast to the Generativist position, which states that language evolved as a cognitive tool, many modern linguists take a Social and Communicative Approach to Language (SCAL). These linguists come from a range of theoretical positions (Cognitivism, Functionalism, Pragmatism, Integrationism, to name a few), which more closely match the general population's intuitions about language: that language is a system for communication which is all about the exchange of

meanings. What makes language special is the range of meanings that humans can negotiate between themselves, and the way that those negotiated meanings allow us to work together in complex joint enterprises.

As part of this SCAL, Clay Beckner et al describe language as a Complex Adaptive System, or CAS, saying:

> "This system is radically different from the static system of grammatical principles characteristic of the widely held generativist approach. Instead, language as a CAS of dynamic usage and its experience involves the following key features: (a) The system consists of multiple agents (the speakers in the speech community) interacting with one another. (b) The system is adaptive; that is, speakers' behavior is based on their past interactions, and current and past interactions together feed forward into future behavior. (c) A speaker's behavior is the consequence of competing factors ranging from perceptual mechanics to social motivations. (d) The structures of language emerge from interrelated patterns of experience, social interaction, and cognitive processes." (Beckner et al, 2009, p2).

The SCAL has been productive in explaining many of the aspects of language that Generativism treats as inexplicable or mysterious – such as how and why language evolved. A consensus is now building that modern human language is the product of a series of social and communicative needs which have defined Homo as a clade and directed our evolution. As Sławomir Wacewicz describes it:

> "whereas literature in the 1990s tended to focus on the biological evolution of a genetically determined human language capacity, recent accounts acknowledge to a far greater extent the role of cultural evolution, which could not only have worked "on top" of the genetic basis, once it was brought about by biological evolution, but could have worked together with the latter process in a co-evolutionary feedback loop." (Wacewicz, 2016, p72).

Basically, increasingly complex socialisation requires increasingly complex negotiation, which requires increasingly complex communication; so individuals better able to handle increasing complexity in these three areas got more genes into the future than those less able. It is a classic evolutionary model, with fitness driving species development via selection between gene alleles; and it is unlike the Generativist model, where a mutation creates a capacity for complexity which meets a previously unknown need. In the Generativist model, the species would appear to be unaware of a need for complexity until they developed the capacity for it.

Yet for Chomsky, the mutation which produced recursion, the key cognitive difference between modern humans and other species, was not a product of social need:

> At the Alice V. and David H. Morris Symposium on the Evolution of Language held at Stony Brook University in October 2005 (and elsewhere), I quoted evolutionary biologists Salvador Luria and Francois Jacob, both Nobel Laureates, as expressing the view that communicative needs would not have provided "any great selective pressure to produce a system such as language" … If the rewiring of the brain that yielded recursive generation of hierarchically structured expressions took place in an individual, not a group (and there seems to be no coherent alternative), then interaction must have been a later phenomenon. Language would have evolved first as an internal object, a kind of "language of thought" (LOT), with externalisation (hence communication) an ancillary process. (Chomsky, 2008).

The question of how the "language of thought" differs from other cognition is not, and has not been, addressed.

Generativist theory and SCAL have produced very different models of language evolution, as table 1 shows.

	Generativist Theory	SCAL
What	Language is a monolithic cognitive system which is independent of all other cognitive and communicative systems. Its communicative use is incidental.	Language is a communicative tool which is segmented, differentiated and hierarchical, and which can be recursive. It is the product of a series of evolutionary events in cognition, communication and socialisation.
How	As a single genetic mutation which gave early humans the capacity for recursive cognition (Chomsky's MERGE). Once evolved, it was so useful that language-users quickly replaced non-language-users in the population.	As a series of genetic changes, each of which met a particular cognitive or communicative need. **PLUS** As a response to an increasingly sophisticated socialisation and enculturation.
Why	The mutation was random, so there is no reason why.	Because evolutionary pressures favoured co-operative and communicative traits in the population.
When	As a single event between 150kya and 40kya.	As a continuing process between 1mya and today.
Where	Unknown.	Initially Africa.
Who	Later Homo sapiens.	Homo heidelbergensis was likely first to use protolanguage, so it was genetically available to descendent species, and possibly culturally available to co-existing species.

Table 1: Comparison of Generativist and SCAL approaches to language evolution

What is social calculus?

Why does a paper about social calculus start with a discussion about a dispute in linguistics over the evolution of language? The answer is that, as a social and communicative tool, a key function of language is to allow us to share each other's social models (Dunbar, 1996). To do this, the sharing of social models requires a level of communicative complexity which corresponds to the cognitive complexity behind those social models; and, as those social models form a complex cognitive calculus – a social calculus – in each individual human brain (Bickerton, 2002), the complexity of our communication needs to match the complexity of our social calculus. This means that our understanding of the nature and evolution of social calculus is intimately linked to our understanding of the nature and evolution of language. Robin Dunbar has explored the sharing of social models and the complexities of social calculus in other primates as well as humans, and has shown that there seems to be a correlation between aspects of brain size and group size (Kudo & Dunbar, 2001); and there seems to be a correspondence between brain size and the levels of Theory of Mind that a species can compute (Dunbar, 2004, ch3).

For the individual, as for the clade, social calculus begins as a private cognitive modelling of the relationships in the local group; it is not a shared public modelling so does not need the added complication of being communicable. In this respect, it follows the Generative principle that cognition precedes communication. Unlike the Generative "language of thought", however, social calculus is not of a different nature to other cognition: it is functionally specific but does not need to be viewed as cognitively novel. Indeed, the mechanisms of social calculus are easily mapped to other cognitive functions (things...do-things-to...things), which means that we cannot know whether social calculus was the first usage of this cognitive functionality, or a later exaptation from an earlier usage.

What we can say is that social calculus consists of two types of construct. The older type is the Relationship-A construct (Edwardes, 2014a): I am able to recall my relationship with you as an affective, or emotional, response. The nature of the response is specific to the individual being responded to, and it is usually stable, although not immutable: while my subliminal response to another individual tends to remain the same, if interpersonal circumstances change then my response can also change to meet the new circumstances. This Relationship-A construct is attentional (it can be stored in memory subliminally, but the individual must be aware of it if it is part of a social calculation), but does not need to be intentional (the individual does not generate the construct, only knows it). There is also no need to recognise the personhood of others in this construct: It is about my feelings towards A, not about A as an object.

The second type of construct in social calculus is A-Relationship-B (Edwardes, 2014b): I am able to recall your relationship with someone else. This construct differs from Relationship-A in a vital way: it is not affective. My emotional responses to A and B must not interfere with my understanding of the relationship between A and B: I may dislike both A and B, but I need to be able to honestly model their good relationship with each other if I am to accurately model their potential alliance against me. So, although I need to understand A's relationship with B without direct reference to my relationships with A and B, the relationship between A and B, when combined with my relationship with A, can indicate or influence my relationship with B (and vice versa). This means that A-Relationship-B calculus is attentional, like Relationship-A calculus; but it is also intentional – there is a level of awareness and conscious calculation involved; and it is dispassionate – my emotions don't come into my representation of the relationship.

A-Relationship-B calculus corresponds to Machiavellian Intelligence, a capacity that Andrew Whiten & Richard Byrne (1988) attributed to both modern chimpanzees and humans. It seems likely, therefore to have been present in our last common ancestor and to be, in evolutionary terms, an ancient capacity. Machiavellian Intelligence is a term that has been disparaged as inaccurate by both supporters of primates and supporters of Niccolo Machiavelli; but is too useful a term to be revised at this stage. It represents the fact that apes (including humans), and maybe other species, are able to model the relationships between others in their group, and use their knowledge about those relationships to navigate their social environment. On that basis, we need no explanation for the origin of A-Relationship-B modelling in the human clade: it was there before humans began.

How does awareness of social calculus become awareness of self?

As well as Machiavellian Intelligence, social calculus also corresponds, in many ways, to the theory of Theory of Mind (ToM): the idea that humans can model other individuals as having intentions, and therefore attentional minds. There has been some dispute over whether other apes have ToM, an issue that Josep Call & Michael Tomasello (2008) addressed in a 30-year review of the topic. They answered the question with a resounding "maybe"; it all depends on your definition of ToM. They said:

> In a broad construal of the phrase 'theory of mind', then, the answer to Premack and Woodruff's pregnant question of 30 years ago is a definite yes, chimpanzees do have a theory of mind. But chimpanzees probably do not understand others in terms of a fully human-like belief–desire psychology in which they appreciate that others have mental representations of the world that drive their actions even when those do not correspond to reality. And so in a more narrow definition of theory of mind as an understanding of false

beliefs, the answer to Premack and Woodruff's question might be no, they do not. (Call & Tomasello, 2008, p191).

In the terminology used here, we can represent the broad construal as Objective ToM and the narrow construal as Subjective ToM. Objective ToM means the individual is aware of others as objects with goals and targets (objects with objectives); but the individual does not need to empathise with them as other beings; it allows the individual to use social calculus in a Machiavellian way to enhance the individual's surviving and thriving. In contrast, subjective ToM means the individual is aware of others as intentional beings; the individual can empathise with other individuals and behave in ways which advantages them over the self. The line between THEM and ME is blurred, because ME has ceased to be a subliminal given and become merely a special form of THEM.

How has this happened? What special event has both raised the need for the individual to attentionally model themself, and given them the capacity to be dispassionate about themself? The answer lies in what happens when individuals start to share their social calculus models.

The sharing of social calculus models is not something that happens because it can, there are considerable obstacles to be overcome: the Sender's Dilemma (why should I give away valuable information?) and the Receiver's Dilemma (why should I believe you when you could be lying?) both need to be addressed; the transmission mechanisms need to be established (just because we share a cognitive mechanism for social calculus does not mean that we share a communicative mechanism able to convey social calculus); the mechanisms for negotiation towards meaning need to be defined (sharing social calculus is sharing opinions not sharing facts, and opinions need to be negotiated); and a whole series of other issues need to be discussed. This, though, is a problem for another paper, here I will concentrate on what happens after the sharing has begun.

When individuals started sharing their social calculus, they would have initially shared the A-Relationship-B models they have built up through observation. Many of these will already have been observed by the individual receiving the model, so they refute or reinforce the receiver's own social calculus. There is, however, one set of models which will be completely new for the receiver: the models which the sharer has observed in which the receiver is a protagonist. Or, to put it another way, what happens when a receiver receives an A-Relationship-B construct in which they are A or B?

While a system of social calculus remains inside a single head, it can rely on two key assumptions. The first is that all the modelled objects are third-person: there is no YOU to be considered, only lots of THEMs. The second assumption is that the self is the context in which the modelling happens. This context is a fixed and inactive background to the modelling, not an active component in the model; so there is no ME to be considered. When the social calculus is communicated, these two assumptions can no longer be made. I may share an A-Relationship-B model with A or B – which means that, on some level, I can recognise that one of the third-person labels in the shared model is both a component in the message and a component around the message. This is mind-blowing for the sharer but, essentially, innocuous: one of the THEMs in this exchange is a YOU – but a YOU is only a privileged THEM.

It gets more interesting when considering what the receiver needs to do when receiving an A-Relationship-B model in which they are A or B. The first time this happens, they will have no model in their social calculus to represent themself; they have to create a third-person representation of themself. This is not itself problematic, every time a new person joins the group they need to be modelled as a new node in every individual's social calculus; but in this case, the "new" group member is new only to the receiver, and it is intimate to the receiver in the way no other THEM is. It requires subjective ToM on a new level: shared social awareness has generated self-awareness.

Modelling the self and self-modelling

The emergence of self-awareness out of social calculus raises important issues about what a self is. It seems to be less a baseline from which an inevitable personality emerges, and more a mutative response to a social context. Evidence in support of this view is beginning to accumulate. For instance, Nasrine Hazem et al (2018) show that engagement with another individual enhances awareness of the physical self; and Sandra Weltzien et al (2018) showed that children aged 7 to 8 are easily primed to selfish or prosocial behaviour by focussing their attention on either themselves or their friends. The self in the social calculus mix is not the product of awareness of a predefined self, it is the product of awareness of models of the self – which are not necessarily models informed by self-knowledge.

It seems that the homogenous self that we believe ourself to be is actually a composite of many different offered models. The composition is further complicated by a series of other modelled selves, derived from the offered models but further adjusted by the experiences of the self themself. Together, this range of modelled selves is identified here as the Seven-Selves Modelling Hypothesis (SSMH), which attempts to provide a comprehensive description of the human

capacity for self-modelling. A full representation of an individual's self-modelling would be something like figure 1.
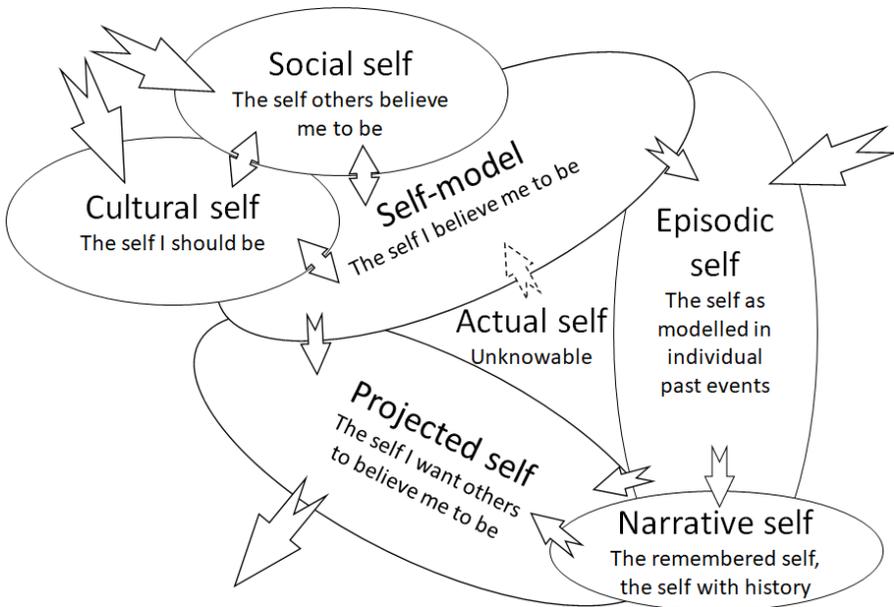

Figure 1: The Seven-Selves Modelling Hypothesis (SSMH)

In figure 1 we see the models of the self offered by others (the social selves) being adjusted by another externally presented model, the ideal of selfhood (the cultural self). Together, these two sets of models provide the material to generate a single integrated self-model. This self-model changes rapidly over time, however: the singularity and integration of the self-model are time-constrained. The self-model is also affected subliminally by the limits the actual self imposes; but these subliminal limits cannot be interrogated by the conscious self-model, so will not be addressed here.

The next cognitive function that affects awareness of selfhood is memory: we remember events by recalling what we currently know about the nature of the event, the roles of others in that event, and our own role in the event. Because the self-model mutates over time, the self-model of today cannot faithfully represent the self-model of yesterday; so to properly assess a memory of an event, the individual has to be aware of, or construct, a previous self-model in relation to the event. Each memory therefore has a reference to a modelled memory self, or episodic self; but each episodic self is not a simple memory of the self at the time of the event, it is moderated by the current self-model and by the recollections of

others (Gardiner, 2001). Robert Numan (2015) describes this as a comparator circuit between the hippocampus (the memory component) and the pre-frontal cortex (the intentional self component).

Linked to the episodic self is the narrative self. This is a meta-memory function, tying together remembered selves with the current self-model, and thus giving the illusion of a continuous self through time. Although the current self-model may have little correspondence to the episodic self-model, we assume a series of mutative selves linking the two selves together. Galen Strawson (2004) argues that the narrative self is not just an illusion, it is not even a necessary component of human cognition; he cites Michel de Montaigne and others as examples of lives lived well but without a narrative component. However, we live in a culture which is based upon the assumption of a narrative self. We have words like "responsibility" and "duty", which relate the actions of our past selves to our present self, and our legal system is predicated on our ownership of our past actions: even if we do not accept that the consequences of the actions of our past selves affect our present self, the people around us will do so.

This leaves the projected self, the image we try to present to the world. This is generated from all the other selves, and is therefore quite difficult to define: different emphases by different individuals create differently oriented projected selves. The projected self is also the self where social deception occurs, from the simple omission or downplaying of key characteristics to the projection of a completely fictitious self. The former is something we all indulge in; the latter is the territory of the narcissist, the sociopath and the confidence trickster. In a simple Darwinian World we would expect a constant arms race between the deceptive projected self and detection of deception; but in humans, living in stable groups of known individuals who are sharing social calculus models, reputation is a powerful mitigator of deception. It is only strangers which have the power to deceive effectively, which may explain our tendency towards chauvinism against out-group individuals. Instead of an arms race between deception and detection, we seem to have a default trust of small in-group deceptions in the projected selves of others, and a default distrust of all out-group projected selves. Our weakness, a willingness to reclassify an out-group individual as in-group based on short acquaintance, is where we are most easily deceived.

How self-modelling defines Homo sapiens

If self-modelling is a product of shared social calculus, and only humans share social calculus, then we would expect only humans to have effective self-modelling mechanisms. However, the cognitive capacity for social calculus is unlikely to be limited to the Homo clade, and it may be quite widespread in nature; it would be a

useful tool for any species which lives in large, complexly organised social groups. If the sharing of social calculus is the key to self-awareness, then we would expect to see it present in some form in any animal which has been deliberately exposed to a satisfactory level of human language. Without going into detail here, we do see some intriguing aspects of self-awareness in the Yerkes bonobo group, who communicate with humans using a keyboard of symbols (Segerdahl et al, 2005, ch3); also in the Washington chimpanzee group, who communicate with humans using a version of American Sign Language (Fouts & Mills, 1997, ch12); and in Alex the grey parrot, who communicated using human speech sounds (Pepperberg, 1999, ch11).

A second way that self-modelling defines us as a species is in our capacity for joint enterprise. Joint enterprise represents our capacity to work together to produce outcomes which individuals could not produce by themselves. Humans can specialise in particular sub-tasks of a joint enterprise – and not just in an ad hoc way, an individual can become expert in one small part of the enterprise without needing to understand the contributions of others. They do not even need to know the purpose of the joint enterprise to contribute effectively. This specialisation is possible because the individual can model themself to themself and to others as an expert, and seemingly live a happy and fulfilled life within that specialised role. In fact, we tend to state our social roles quite early on in a conversation with a stranger: questions about work, school, and hobbies all figure in lists of useful conversation openers.

In terms of the outcomes of the joint enterprises, what humans achieve is several orders of magnitude greater than any other social animal. We have colonised every environment on Earth except the deep ocean, we have sent objects and people beyond the limits of this planet, and we have affected the environment of the planet so radically that we are in danger of changing it fundamentally. All of this is possible because we can model ourselves and each other not just as individuals but as specialists.

A third species-defining outcome of self-modelling is altruistic sacrifice. This does not just refer to the heroic extreme of self-sacrifice, it covers the everyday small sacrifices we make to keep society jogging along. The polite "after you", the coin thrown into a musician's hat, even the friendly word to the barista, all of these are sacrifices which make the biggest human joint enterprise – human society – work. They rely ultimately on the value of establishing a reputation as a good citizen, offering a projected self which demonstrably corresponds closely to the cultural self; but proximately, the sacrifices may be unobserved by anyone able to usefully broadcast the story of my generosity, especially in modern urban societies.

However, the fact that I have a third-party model of my self may be what makes this work. When I see someone else being generous it raises their reputation in my social calculus – they go up in my esteem; when I see myself being generous it raises the reputation of my self-model in my social calculus – I go up in my esteem. Nathaniel Branden (1994) said that "self-esteem is the reputation we acquire with ourselves"; the SSMH gives us the mechanism by which we can acquire a reputation with ourselves.

The route to self-modelling

How did we become the socially calculating, self-aware species we are? This paper has looked at some of the mechanisms, but it has not indicated a timescale or a developmental map. It is probably still too early in our study of early Homo to tie anything cognitive or cultural to a developmental calendar, but at least the capacities enabling social calculus to be shared, and self-modelling to begin, can be reasonably described.

The sharing of social calculus models would have required the pre-existence of several things. Foremost is the cognitive existence of social calculus itself, which has been discussed above. This relies on several capacities, the first of which is the presence of large social groups; these, in turn generate a genetic trend towards larger brains to handle the social relationships involved in the larger group. Second is a willingness to work together in joint enterprises – not the complex organisations we see today, but simpler groupings in which specialisations began to appear and be valued. The individual who can make good throwing stones and the individual good at throwing had to work together and share together in a joint hunting enterprise if they were to maximise the fitness of their individual skills.

The third pre-existent capacity required for sharing social calculus would have been a system of voluntary communication: we must have been able to exchange information in some way for early socialisation and enculturation to spread through a population. This communication system would have emerged from an earlier signalling system, where meanings did not need to be negotiated. Signals involve the production of vocal responses to environmental events and require no negotiation about their meaning; a communication system involves the intentional offering of information by one individual to another, and it requires negotiation towards meaning by both parties to reach understanding. What early humans were negotiating did not have to be social knowledge – in fact, it was probably environmental knowledge, as we see in modern primate signalling: warning signals, food indicators, locative vocalisations, attention-getting signals, and emotional displays.

However, environmental knowledge is usually shared by calling attention to the environmental event: a simple deictic signal like "look" or "there" begins the negotiation towards meaning. This means that no structural, or grammatical, complexity is required in a communication system sharing environmental knowledge: the vervet monkey snake warning is just a chutter (Cheyney & Seyfarth, 1980); the chimpanzee food call is just a grunt (Schel et al, 2013); the gibbon "I am here" call is just a hoo (Clarke et al, 2015). In each of these cases the call can often be modulated to more particularly reference the type of snake or food, or the individual calling, which is impressively subtle; but the sounds all rely on extrinsic meaning (the meaning is out there in the world), they have no context-free generalised reference, or intrinsic meaning, as the English words "look" and "there" have.

Fortunately, the sharing of social calculus carries us over the divide between extrinsic and intrinsic reference. The A-Relationship-B construct requires A, B, and the relationship to be individually meaningful: the particular grunts that represent A and B must mean – or be negotiated to mean – A and B to the sender and receiver, and the relationship grunt must mean that particular relationship; but the form A-Relationship-B is a framework into which any number of individual-representing grunts can be inserted, and any number of relationships can link them. In addition, sharing of social calculus requires an open-ended communication system: as new individuals join the group, new representing grunts need to be generated and negotiated into meaning; and as new types of relationship develop, new grunts to represent the relationships must be agreed.
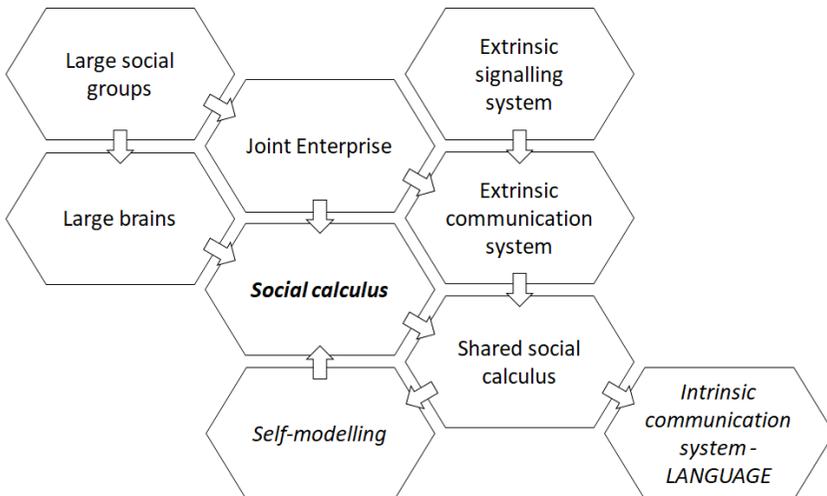

Figure 2: The Route to Self-modelling

After all that, the appearance of self-modelling would seem to be a side effect of shared social calculus; and that may indeed be what it is. That does not mean, however, that it is a small thing: it is the fly swallowed by the old lady, which set in train an increasingly complex and onerous system of ingestion. Self-sacrifice, religion, fiction, Twitter, Bitcoin … unlike the old lady and the horse, and thanks to our capacity to model ourselves, we can swallow anything.

References

Clay Beckner, Nick C. Ellis, Richard Blythe, John Holland, Joan Bybee, Jinyun Ke, Morten H. Christiansen, Diane Larsen-Freeman, William Croft & Tom Schoenemann (2009). Language Is a Complex Adaptive System: Position Paper, The "Five Graces Group". In *Language Learning* 59:Suppl.1, pp1-26.

Derek Bickerton (2002). How protolanguage became language. In Chris Knight, Michael Studdert-Kennedy & James. R. Hurford (eds.), *The Evolutionary Emergence of Language*. Cambridge University Press: Cambridge, UK, pp264-284.

Nathaniel Branden (1994). *The Six Pillars of Self-Esteem*. Bantam Books: New York, NY, USA.

Josep Call & Michael Tomasello (2008). Does the chimpanzee have a theory of mind? 30 years later. In *Trends in Cognitive Sciences* 12:5, pp187-192.

Noam Chomsky (1957). *Syntactic Structures.* Mouton de Gruyter: Berlin, Germany.

Noam Chomsky (2008). Human nature and the origins of language. In *Radical Anthropology* 2, pp19-23.

Noam Chomsky (2012). *The Science of Language: Interviews with James McGilvray*. Cambridge University Press: Cambridge, UK.

Esther Clarke, Ulrich H Reichard & Klaus Zuberbühler (2015). Context-specific close-range "hoo" calls in wild gibbons (Hylobates lar). In *BMC Evolutionary Biology* 15:56.

Robin I.M. Dunbar (1996). *Grooming, Gossip and the Evolution of Language*. Faber & Faber Ltd: London, UK.

Robin I.M. Dunbar (2004). *The Human Story: a new history of mankind's evolution*. Faber & Faber Ltd: London, UK.

Martin Edwardes (2014a). What were we talking about? Exchanging social models as a route to language. In Erica A. Cartmill, Seán Roberts, Heidi Lyn & Hannah Cornish (eds.), *The Evolution of Language: Proceedings of the 10th International Conference (EVOLANG10)*. World Scientific: Singapore, pp58-65.

Martin Edwardes (2014b). Awareness of self and awareness of selfness: why the capacity to self-model represents a novel level of cognition in humans. In G. Rundblad, A. Tytus, O. Knapton & C. Tang (eds.), *Selected Papers from the 4th UK Cognitive Linguistics Conference*, pp68-83.

Roger Fouts with Stephen Tukel Mills (1997). *Next of Kin: My conversations with chimpanzees*. Avon Books Inc: New York, USA.

John M. Gardiner (2001). Episodic memory and autonoetic consciousness: a first-person approach. In *Philosophical Transactions of the Royal Society London B* 356, pp1351-1361.

Nesrine Hazem, Morgan Beaurenaut, Nathalie George & Laurence Conty (2018). Social Contact Enhances Bodily Self-Awareness. In *Nature Scientific Reports* 8:4195.

Hiroko Kudo & Robin I. M. Dunbar (2001). Neocortex size and social network size in primates. In *Animal Behaviour* 62:4, pp711-722.

Robert Numan (2015). A Prefrontal-Hippocampal Comparator for Goal-Directed Behavior: The Intentional Self and Episodic Memory. In *Frontiers in Behavioral Neuroscience* 9:323.

Irene Maxine Pepperberg (1999). *The Alex Studies: Cognitive and communicative abilities of grey parrots*. Harvard University Press: Cambridge, MA, USA.

Anne Marijke Schel, Zarin Machanda, Simon W. Townsend, Klaus Zuberbühler & Katie E. Slocombe (2013). Chimpanzee food calls are directed at specific individuals. In *Animal Behaviour* 86:5, pp955-965.

Par Segerdahl, William Fields & Sue Savage-Rumbaugh (2005). *Kanzi's Primal Language: The cultural initiation of primates into language*. Palgrave Macmillan: Basingstoke, Hants, UK.

Robert M. Seyfarth, Dorothy L. Cheney & Peter Marler (1980). Vervet monkey alarm calls: Semantic communication in a free-ranging primate. In *Animal Behaviour* 28, pp1070-1094.

Galen Strawson (2004). Not every life is a narrative: A fallacy of our age. In *Times Literary Supplement*, 15 October 2004, pp13-15.

Sławomir Wacewicz (2016). A contemporary look at language origins. In *Avant* VI:2, pp68-81.

Sandra Weltzien, Lauren E. Marsh & Bruce Hood (2018). Thinking of me: Self-focus reduces sharing and helping in seven- to eight-year-olds. In *PLoS ONE* 13:1, e0189752.

Andrew Whiten & Richard W. Byrne (1988). Taking Machiavellian intelligence apart. In Richard W. Byrne & Andrew Whiten (eds.), *Machiavellian intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Oxford: Oxford University Press, pp50-65.

For more about social calculus, agentive grammar
and the beginnings of language, see:
http://martinedwardes.webplus.net/eaorc/eaorc_languageroute.html