# *Complexity Science and 21st Century Issues*

## *London School of Economics 25-26 March 2004*

**Dr Luc Steels**, University of Brussels (VUB AI Lab) and Sony Computer Science Laboratory, Paris.

### *"Language as a Complex Adaptive system"*

Natural language is an example of a complex system. Not only is it multilayered in terms of its words, phonetics, its syntax and its grammar, but it is a system acquired and preserved in cultural learning, something which has to be built and rebuilt by one generation after another. Most people use forty or fifty thousand words which have many meanings used in many different ways and although linguists sometimes pretend that there is a set grammar, history shows us different. In short language is a complex evolving system.

We can see examples of Latin words evolving into French such as 'scisse' (to know) becoming 'savoir' or 'posse' (to have power) turning into 'pouvoir'. But what also occurred was the introduction of the article 'le' as in 'le savoir' and 'le pouvoir'. Latin and old French do not have articles and Latin verbs were only prefaced by a pronoun for emphasis. The word 'das' (you give), for example, becomes in French 'tu donne'. So it is not just the words but the syntax and grammar which change.

But although language is a complex system we have no theories about how new meanings arise, though it's obvious that the things that people talked about, in say the Latin period, are very different from contemporary periods. If we think about all the concepts that have grown up around computers and telecommunications, such as 'the mouse', 'the Web', 'uploading', 'e-mail' and so on, we can appreciate how creative language is.

Then there is the question of how language changes in a population. How do new sounds and grammatical constructions enter the language and spread? How are our language and the way we view the world related? How does our language influence the way we view reality and vice versa? Language can shape the social network but the network also determines how language evolves. However if you open up any book on linguistics there are very few occasions on which these kinds of questions get raised. Most linguists describe the way in which language deals with phenomena and don't worry much about the evolutionary aspects and the question of origin. The suggestion here is that perhaps complex systems can help us in answering these questions in some of the following ways:

1. In complexity science there is a bag of concepts for understanding and building systems, which aim at simulating things like ecosystems or the stock market. There is a rich set of metaphors which come from the different sciences such as natural selection, self organisation and emergence.

2. Then there is also a bag of tools with which to analyse the behaviour of systems of which almost none have been used for language.

3. And thirdly a methodology has evolved in this field which I will call the 'methodology of artificial systems'. This is a bit different from simulations. In the case of simulations researchers use a model which they think is realistic and then they

run computer programs to see whether the model matches the real data and makes the same sort of predictions. Language systems are so complicated that building a model seems daunting, but what we can do is to build artificial systems that show certain phenomena and hope that they have some of the properties of the things that we are interested in. So in our case, we build systems with agents as real robots and try to see how they interact with each other via a certain kind of cognitive mechanism. It is the interaction of robots with each other and with their environment that leads to the invention of shared ways of conceptualising the world through negotiation. What we are trying to do is to produce artificial communication systems with language like properties. It's important to remember that language is used in new interactions with the world and with other people so in order to understand where new meanings come from we cannot just have a computer which sits on a desk we need to put our systems in environments that have rich interactions so that new meanings can arise.

Just to amuse you here are the kind of robots we have used before our experiments with language. It's an experiment from the early 90s inspired by Walter Grey who had the idea of building robots which get their power from a charging station for which they have to compete. There are boxes which the robots have to push against to make sure that there is enough energy in the charging station which they then go into for recharging. We can see one robot going into the charging station to recharge itself and then another. So it is like a little ecology where the behaviour of the robots becomes meaningful from the viewpoint of recharging and self sufficiency. Now one thing you see here is another robot also going into this charging station and actually pushing the other one out and people say 'take that robot out, he's bad 'which shows that these kind of experiments can be modelled like biological systems where animals have to work for food. We can see how the robots can learn to cope with the demands of their environment and see what kind of relationships develop.

Here's another experiment which illustrates a doggy version of the Turing test. You will remember the Alan Turing test in which an observer is outside a room in which (unbeknown to the observer) is a computer? The observer can ask any questions he or she likes. If after some lengthy interrogation with searching questions the observer cannot tell whether it's a computer in the room or a real person then artificial intelligence is possible. Well this is not a good test but here we have a robot dog and a real dog and a piece of meat in the middle. You can see the real dog has no doubts about the reality of the robot dog.

We try to do a similar sort of thing for language by building artificial systems and putting them into partially natural situations to see what kind of behaviour comes out. Typically we've settled on what I call the 'language game approach' which has been inspired by Wittgenstein and other language philosophers and assumes that words get their meaning in use. A language game is also something which involves an interaction with the world which is verbal though of course it need not be. I mean if I ask the person in front of me if he wants some water and he says 'yes', then I give him some water. So a language game is about action in reality. I am doing something and I'm using language as part of the interaction.

Each robot in this game has mechanisms for sound generation and certain learning mechanisms and we try to see whether we can start to build up artificial languages in terms of establishing an ontology. By 'ontology' I mean the concepts and objects in a loose sense which are used to express meaning. Here's a little experiment which is not yet about this kind of meaning. We're just interested in use of the sound here so we equipped each robot with a voice simulator and a repertoire of sounds that were similar to sounds in natural language. We also modelled the perception of the sounds in each robot by putting in a neural network which learnt to

categorise and thus recognise and later reproduce.

The robots play a little game in which there is first some calling of each other but because it's a group they have first to get the attention of another robot so one starts an imitation game. The agents start by making completely random sounds, but by categorisation they start imitating and attracting the attention of each other. This is a very simple and amusing illustration but behind it is some very sophisticated machinery. There are neural networks that get initialised and these categorise the sound and control the articulator. So we start with networks that are randomly initialised and progressively attractors form in the landscape of sounds. This already gives a hint for people that are into complex systems that we are using many of the ideas that are current in that discipline. Sounds then are viewed as attractors in a potential sound landscape.

The other thing we are doing in our language games is getting real robots to describe things in the robot world. Understanding the origin of meaning is important in a number of fields including economics and specifically for understanding innovation through conceptual change. I think this is particularly relevant for cognitive science because there is a revolution going on in linguistics and semiotics to bring in the ideas of complex systems

What interests me here is the impact of information technology and I believe that if you can make a theory of the sort that I would like, then this would have an impact on many other fields of study. In the 70s a system was built at M.I.T. based on Locke's empirical theory that the mind starts as a 'clean slate' and concepts about the world are acquired through experiencing instances of qualities such as 'red'. You could have a dialogue with the system and say things like 'pick up a red box' and you'd see a movement on the screen and the red box would be picked up. You could also ask sophisticated questions such as 'what does the box contain?' At the time this came out everybody was amazed because it seemed that there was language processing with reasoning and maybe the artificial intelligence test had succeeded. But very soon there was a criticism by John Searle who introduced the metaphor of the 'Chinese room' in which a computer replies to questions in Chinese. Searle asserted that what the program was doing in the machine was mere symbol processing and said: "I could sit in that room and go through all their motions without knowing any Chinese". Searle's thought experiment was ultimately a metaphor against the strong A.I. argument that computers could be given intentional states.

Searle's argument was only partially right. Mere symbol processing and a fixed grammar will not lead to artificial intelligence, but it assumes that computers do not have any contact with the real world so they couldn't possibly learn about it. There were in fact systems already built at the time like 'Shakey' built by Neilson and his group at the Stamford Research Institute which were already grounded in the world. You could say to a robot: 'go to the door' and it would actually go to the door, but Searle is right about symbolic representations programmed in completely and not autonomous in that they were developed by the agent itself. If you use a search engine on the Internet by asking a question the Internet knows nothing about the subject but it will come back with an answer and the answer may depend upon many subtle connections. So the question is what kind of meaning does the Internet come up with and how can information systems autonomously create meaning which is grounded in the world?

What I would now like to do is to make some comments about language and in particular verbal communication. So this is about the ideas that we are using to build robotic systems and I'll give you some details on the experiments that we have been doing and then the conclusions. First some ideas about communication. Suppose

that you're going with a friend to the movies and you decide to take the bus. You're walking and suddenly there's a bus coming. One of you starts running and the other follows. It could also be that one of the two says 'the bus' and both start running. What is going on here is interaction between agents in a particular situation and particular context. That's something we can never forget about and it's where we have to start in our robotic experiments.

There has to be a situation in which there is a sharing of intention and a common ground before anything in terms of communication can happen. Of course in the case of going to the movies the intention is to catch the bus and it results in both people running. But there are meanings between what is said and what is done that are not implicit though it's clear that the connection between what is said and what is done somehow has to go through meanings.

For the purposes of our research we adopted a very pragmatic view that meanings are distinctions that are relevant for the interaction between the agent and the environment or between agents among themselves. So they're not Platonic entities or anything like that, but features of the environment relevant for the interaction. For example, if you standing in front of a traffic light the distinction between red and green is pretty important. But in order to communicate meanings have to be recognised. Meanings have to be put into categories. Some neural networks do categorisation and it's only when meanings can be cognitively recognised that they can start to play a role in explicit communication. Typically in logic, for example, predicates denote categories though I want to avoid confusion between meanings, categories and predicates because not all features of the environment are meaningful or explicitly recognised by our cognitive system.

The second thing is that when we want to communicate something we only make selected meaning explicit. There may be lots of things floating around in our shared or common situation, like wanting to go to the movie. We both might know we need to take the bus because we are late but none of this is being said. The only thing which is said is 'the bus'. So that's very important and the reason why it's so difficult for computers to deal with language is because language only gives a hint about what we want to say or how we want someone else to act. What I would like to say about that, is that language is not a code. Unfortunately people in complex systems who are interested in language often assume that it is a code. I mean by that, that in a code all the information is in the message whereas in language there is only partial information. Also, in a code the conventions are fixed, whereas in a representational system it is much more flexible. Language is highly contextualised, whereas say, drawings are much less contextualised and for a code you don't need to know who is sending the message. You don't need to know when it was sent, whether the person was feeling well, who else was there etc. etc. You just need the coding mechanism and the roles and you can do it, whereas in the case of representations in general, you need inference, you need common grounds. All of which are necessary to figure out what is going on.

This is a drawing by a four year old child. What do you think it is? Child play? A garden maybe? On the train maybe? It is a bus. And this is representation and in a child's mind a bus is big so the whole page is filled. A bus has many wheels and many windows and this is the conductor taking the money. And, interestingly you notice the different perspectives that are put in this picture. There is perspective from the side. There's perspective from the top and a different perspective with all the wheels showing. So the perspective is both inside and out. And this is the conductor of the bus. This is a typical example of a representation. You select things that are important to you and the conductor of this bus is clearly important to this child. And you see a

hand sticking out to take the money. So the whole thing is about seeing things from different perspectives. Of course conventions are invented like many wheels, and humans have this incredible ability to invent representations Even though a lot is conventionalised children are known to be very creative in language.
The next observation is that there are important differences in what meanings can or should be made explicit in different languages. This is something you start to appreciate if you know many languages. I'll give just two examples of this and the first is called 'time'. You might think that time is something that we all share and is universal and everybody thinks about it in the same way. But if you start looking at languages there are enormous differences in how cultures conceptualise time and also what they make explicit about time. Take tenses in English: you say 'he came', 'he comes', and 'he will come' as past, present, future. You also have questions about where you are in a particular event: like ' is it finished?' or ' Is it still going on?' or 'what will it do?' These are things that English brings out, but in Chinese you don't have explicit tenses. In Chinese you have one word for say 'go' and if you want to express the time you do it with other words though it's not made explicit or obligatory through the verb.

In languages in South American you actually have a much more refined system of tense; in that the distinctions made about the past and future are much more detailed. In some languages you always need to say where you got the information from. So you would say something like: 'I saw it myself' or 'I heard it from somebody' or 'I believe it but I'm not sure'. And in that language it's obligatory to make that distinction.

There are also differences in the way the relationship between an object and an event is expressed. Whereas In German a relationship is expressed by case endings, in English it is by word order and preposition. This is interesting because English used to have a declension system which subsequently disappeared, making word order and preposition use more obligatory.

I had the good fortune a few weeks ago to be in Kenya and the Masai language is fascinating from many anthropological viewpoints. They have a phrase which means: 'the woman will get milk for her husband with a cup'. What they do is repeat the word 'milk' a couple of times in the sentence to express the roles of a woman, a husband and a cup..

There are not only differences between languages but there are also differences between how individuals in the same language group express themselves. And finally there is the spread of words and constructions and biologists will note the similarity to the spread of viruses in a population

OK, with that in mind what about change? New words may pop up or old words can gain new meaning and be used to express something not expressed before. Word order may initially be very free and become stricter (syntax). Occasionally a whole construction will be reused for something else which is new. This occurs typically by what is known as semantic bleaching where some of the properties of the original construction are lost and some new ones made. But there's also syntactic bleaching where there is construction change.

Here's an example in English where there is both. In English the word 'will' used to be a verb of volition so it was possible to say 'I will a book' meaning 'I desire a book' which has the French equivalent 'Je veux un livre' (I want a book.). What has happened in some constructions is that 'will' has, by semantic bleaching, lost some of its meaning of volition and become an expression of future intention. So, '1 will do this' or 'I want to do this' still have a sense of volition. But 'it will rain tomorrow' contains no sense of volition at all and you do not say in French: 'il veut pleuvoir

demain'. So what has happened is that the word moved from a main verb to a future tense auxiliary and that's how grammar evolves. Many examples have been recorded by linguists and many seem to a large extent to be universal though we don't have any models of the evolutionary process.

Another interesting example is the development of negation. In old French ‘I am not going’ is 'Je ne vais ’ and the 'ne' comes from the Latin 'non'. But at some point people didn't find that was a strong enough expression of negation so they added ‘un pas’ as in ‘Je ne vais un pas' meaning ‘I’ll not go one step'. There were also competing expressions such as ‘Je ne vais une goutte’ meaning ‘I’ll not go a drop’. However at some point 'pas' won the day and changed by syntactic bleaching from a noun to a participle and lost the article ‘un’. So we have 'Je ne vais pas' which literally means 'I not go step' or 'Je ne veux pas' (I not want step) though there’s no thought of a step in the expression. And actually in modern French even the 'ne' is disappearing so you say 'Je veux pas'' which literally means 'I want step' but it’s used as a negation.

What is happening in this general evolution is that agents are inventing new categories, inventing new language constructs and re-using existing constructs for new purposes. And some of these then propagate in the population if they fulfil a need. So this is language as a complex adaptive system, where we have an ecology of meanings or constructs, which are in competition with each other. If we have a large population and a word is invented somewhere for something and another word is invented for the same thing somewhere else, there will be competition and maybe one will dominate.

In a large population we also get more variety of expression and selection is based on a number of forces. One is to maximise the expressive power. We want to be understood. We want to say what we want to, say and if we cannot find a ready-made expression then we invent something and we do that all the time. Language is not a fixed system. It is something that is created and recreated

In maximising communication we want to be precise, but not more precise than we have to be. In doing so we seek to minimise the effort both from the viewpoint of the speaker and from the viewpoint of the listener. We seek to minimise the effort of the listener because if he or she has to make lots of inferences in order to understand then he or she may get bored and that would decrease communicative success. We also need to minimise our effort because there may be very little time to say something in a changing conversation. So these are forces acting and these skills have to be learnt.

One of the experiments we did on a large scale was the talking heads experiment. If you try to use a number of autonomous robots that move around things get very complicated so we used a couple of ‘pan-tilt’ cameras as agents in a room with a whiteboard on which could be pasted various figures to change the environment. These are our two agents and they switch roles as speaker and listener. We wanted a simple as possible language game that had the minimum ingredients to investigate the origins of concepts expressed in a lexicon. Initially the agents do not have a way of conceptualising their world which had to be developed in the course of the game.

When the game starts the speaker agent pans the camera and selects an object. Since the cameras don’t have arms the listener agent has to judge the direction by the camera position. It’s not very exact but good enough. When the speaker has selected an object he wants to draw the listener’s attention to it and the listener will then respond. So this is the basis of the game. It’s the same kind of interaction as if I say, ‘where is the water?’ and you say 'there'.

Now if I say something like: 'googlagoo' you won't know what it is. But if I

point and say 'googlagoo' then you begin to and this is the game they play. One of the agents wants something and has an object in mind and wants the listener to pay attention to it. Now the other agent does not know the meaning of the word and can only begin to learn it by interactions with the world. The speaker picks out an object in the world say a black box on a table but cannot immediately convey the meaning of the word. There's no telepathy and no sharing of concepts before they start so they can only infer possible meanings by referring to the world. This is of key importance. What does 'googlagoo' refer to in the environment? There are a number of possible referential meanings. So the meaning has to be negotiated. The speaker may not succeed for quite a while until maybe there's another black box not on a table so the listener can begin to know that the table is not included. So this is the kind of dynamics we have in the games.

In order to do a large number of experiments we built a system where agents were allowed to travel through the internet from one location to another playing different games. In this way we had a population of a few thousand agents who played for 3 or 4 months in places around the world. But we needed to put into the system all the processes that go on in both the speaker and the listener for formulating the expression and understanding it and that wasn't simple. Apart from machine vision, we had to structure the object events giving different ways of conceptualising things, mechanisms for grammar and everything to enable the building of words and structure. So I don't don't want to pretend that we started from nothing, the agents are very complex from the start

When a speaker of a language selects a topic it is related to the kinds of social interactions that are known. Conceptualisation means finding a category meaning which is distinctive i.e. true for the topic but not for any of the other objects in the environment. So we are assuming a principle of relevance here. In order to communicate we want to find things that are relevant for the goal of communication and we need to be precise. If I wish to pick out a particular cup when there are two cups in front of me then I have in some way to individuate the cup I want. What the speaker aims to do is to find a conceptualisation of the world that not only picks out a topic once but will do so again if necessary.

Conceptualisation is translated into language via a verbalisation process. I don't have time to say very much about this two way process but the repertoire that an agent builds is a set of meaning structures and a set of forms that are related to them. It's an associative memory of meanings and forms that can be used in both directions. That means if you want to express a meaning you use this memory to find a form and the form can be as simple as a word or it can be a grammatical construction of some form. If you want to go in the other direction as a listener you are looking up in the associative memory and bits of a form which relate to the meaning. So the agents don't have just one form for one meaning or one meaning for one form, but a memory which has multiple relations between form and meaning. There's also a kind of evaluation process or score for each of them. Again this is like a little ecology where there is competition for the expression of a certain meaning or a competition between the meanings of different forms. What happens is that every time that the agents are speaking after the game they are adapting these scores. If there's a certain meaning that is expressed and a certain form used and it is successful then the score of that will go up and the score of competitors will go down. There's a kind of lateral competition and conversely if there's a form used here and you used a certain meaning successfully to pass and interpret that form then the score will go up and the score of competitors will go down. So there's a positive feedback loop; the more a word is used the more successful it will be and hence the more it will be used as the agents

increase the score and decrease the score of competitors. After a while you see organisation in the system as words form a coherent set of conventions.

I'll just show you some data from the talking heads experiment. This graph shows the frequencies with which different words are used in the population for the same meaning. You see that in the beginning there is a lot of competition to be able to express a meaning in a number of games; 5000, 10000 etc. And this is the frequency of the words use in the population. And you can see that after a while there is one word that will dominate, not because we program it but due to the positive feedback. This is a graph that is showing the different meanings for the same word And again you see that typically there are many possible meanings and this goes back to the black box example i.e. you don't know and you interpret as well as you can and guess the possible meaning. You assume a meaning, though of course more than one meaning is compatible with the situation. In this experiment you have 311 different meanings floating in the population but by a process of maximal communication success these will decrease. You can also see how long it takes the agents to boot strap from scratch to a lexicon with different population sizes. The more agents you have the longer it will take before they reach coherence but probably the richer the language.

I'll also just show this one which is a bit like the traffic simulations - simulating a whole city with traffic and energy. Here we can have populations of agents and we can change population flow, take some agents out, put some agents in and see whether the system remains stable and see what the conditions are. We also introduce noise in the perception of words or errors in production of a word and see what impact it has on the stability of the communication system. All the sorts of games that people used to play with chemical reactions we can now do with language systems.

In another game we explore how quickly a lexicon builds and how stable it is. We started with a closed system which meant a group of agents who very quickly got 100% success in communication. Then we changed one agent every hundred games and found it had no impact because the new agent quickly learned the language which was already in the population. Here we introduced a word that was not always pronounced or heard in the right way and started getting change in the form of the word until one form predominated and became the new word

I'll close with one more example of an experiment between humans and robots because people say it's all well and good that robots speak their own language but we don't understand it. This is again through the eye of a robot. This is a ball. Now you might think a ball is an easy thing to recognise but through the eye of a robot, a ball is very complicated. I mean if you look at a number of images of a ball, you can see that almost none of them are a nice sphere. Just doing segmentation is extremely difficult and if you blow up to the pixels how many do you need?

Learning a language is very much learning the concepts that are used in the culture and we use a memory based approach in which the eye is given different views of the object in different contexts so a visual memory is built up.

In this experiment the robot is learning the concept of ball in a 'look game' at the same time as it is learning the word. In it the experimenter is saying 'look ball' and there are two things. One is the degree to which attention is shared. In other words how can you know what the robot is looking at and how can the robot know what you are bringing attention to? The second thing is that the robot has already learned the word for smile and so the first time it sees the ball it thinks 'this is smiley' because that's the closest thing it knows. So then the feedback is 'no, this is not Smiley it is 'ball'. And eventually the robot says 'ball'

There are a lot of things to be done particularly in the area of grammar which is my main focus at the moment and there are enormous opportunities for theoretical work in complex systems because the mathematical foundations for all this are almost completely lacking.

Questioner 1: Can you make a robot learn abstract concepts?

Answer 1: Well, we are focusing at the moment on visual categories and body interaction concepts. One thing that happens in language is that these more abstract notions are typically by metaphor and analogy, extracted from more bodily sensory/motor interaction. But of course if you want to do something like 'happiness' then the first question is; what does it mean for a robot to be happy. We cannot pretend that the robot is happy just by putting in a symbol for a smiley face.

Questioner 2: When you picked up the box and you said 'what do you call it?' you would be open to descriptions such as ‘rectangle’, which is a geometric notion so how would you deal with that?

Answer 2: Well this all depends on the agent but the most important criterion is that the meaning is something that is relevant and distinctive about that particular object and no other. So it's a way to zoom in. It's a bit like moving to a new house. You create categories which you impose on the world and then you use them, so it’s not that you have a million examples of boxes or things like that You create the distinction as you need it and then you name it and then the listener picks it up as well. So it's a very much the idea of language and conceptualisation as social constructions. Of course they are related to or grounded in the world so they are not arbitrary but imposed.