

## FROM BOXES AND ARROWS TO CONVERSATION AND NEGOTIATION

### Or how Research should be Amusing, Awful, and Artificial

#### ABSTRACT

The story of how a graduate student went from formalism to data, a brief tale of how engineering without tradition can lead thought in the right direction, and a mild caution of how intellectual skepticism is worth little without a corresponding dose of intellectual enthusiasm.



#### Jussi Karlgrén

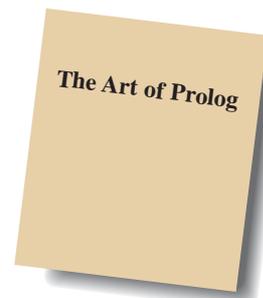
Jussi Karlgrén is a researcher at the Swedish Institute of Computer Science where he specializes in the study of text and interaction, especially as regards reader understanding of topicality and stylistic features of text. He holds a candidate degree in Computational Linguistics and Mathematics, a licentiate degree in Computer and Systems Sciences, and a doctorate in Computational Linguistics.

When I first came to DSV in the Fall of 1987 I was looking for diversion. I was busy finishing up my Bachelor's degree in Computational Linguistics and felt I needed classes to broaden my perspective.

This was at a time when DSV was within easy reach for curious students from other departments – located on campus at the university instead of in an industrial park, so I could simply walk over from my department to the other hallway and look for interesting things to think about. DSV had plenty to offer, and I signed up for a class called “Artificial intelligence programming” which, as it turned out, taught the basics of structured programming in high-level languages such as Prolog, Lisp, and Smalltalk. It was taught by Carl Gustav Jansson and several seasoned graduate students – most of whom today indeed have graduated and many of whom still can be found at the department or in its close vicinity.

Even today the sense of awe and possibility the classes instilled in me is easy to recollect. The course books are still some of my favourite volumes: “The Art of Prolog” and “The Structure of Computer Programs”. One with Asian artwork, the other with crudely drawn wizards on the cover. The contents were no less inspiring – they expressed the sense that everything was possible given the appropriate computational model. Wizards! Exotism! Fun!

This is where the subtitle above comes in: an apocryphal anecdote tells us what the then King of England (it is unclear which one – the story is attributed to various monarchs) said when first confronted with the design of St Pauls cathedral: “Amusing, Awful, and Artificial”. The king intended this as praise, meaning that the building was pleasantly thought-provoking, awe-inspiring,



Bildförslag: En av de nämnda kursböckerna. Jag har ännu inte fått boken från institutionen.

and artful. The words have since shifted in meaning (which is why the story is told today) but the combination of the three is what most of us in research strive for in what we do. We would do well if we more often and more explicitly thought along those lines. Linguistics is a methodical area of study.

Linguistics is mostly about applying formal models to arguably that most fluid of symptoms of human intellectual activity: verbal communicative behaviour. But linguists expend or used to expend enormous effort on debating the absolute merits of various different formalizations rather than working on applications of the same. At the time I am writing about, linguists in general, and leading researchers here at the department in Stockholm in particular, had already started turning to more data-oriented and less theory-oriented methodologies: corpus linguistics, focusing on the methodology of collecting data; typology, focusing on the differences and relations between languages rather than the specifics. But this reform had not yet percolated down into the undergraduate curriculum, nor made its way into basic textbooks. I had spent a long time trying to make sense of arguments of formalism rather than understanding people, how they speak, and how they think.

The field of artificial intelligence in the eighties was neither straight-laced nor crabby. Thinking about it, I can still feel that rush of boundless potential I experienced in taking those first classes. I still believe in those first realizations. Yes, we can model human cognition! Since I had already burnt my fingers on the fruitless representational debates of linguistics at the time I knew from the start that finding the perfect model counts for little in the end, whereas representing the salient aspects of behaviour does.

So far, this, of course, has been a common story, told and told again: a young and enthusiastic graduate student finds that many of the theories of the preceding generation of researchers is overloaded with ideology and not sensitive enough to the reality they model. That realization, in fact, drives much of research. The risk every disillusioned researcher runs in this situation is to tumble into unproductive agnosticism: collect data, gripe about dysfunctional models, and do nothing. In the words of a traditional linguist: “Viel Data und wenig Theorie”. Skepticism rampant leads nowhere!

***“In the words of a traditional linguist: ‘Viel Data und wenig Theorie’.***

### **From DSV to SISU**

I continued studying at DSV taking classes such as “Knowledge representation and reasoning”. After gaining my degree I signed up as a graduate student at the department in Carl Gustaf Jansson’s research lab (which was more difficult than it sounds – Anne Marie Philipson spent a fair amount of effort to help me dodge the bureaucratic obstacles). I almost immediately found employment at SISU, a research institute closely associated with DSV.

SISU busied itself mainly with applications of conceptual

modelling, building models of human knowledge using a generalized entity–relationship framework. The approach taken at SISU was to build models that could be put to use. Their exact form and principles were of less consequence than their quality – and this was best assured through judicious editing and non-trivial intellectual investment on the part of the knowledge engineer. The project I worked in was a project at IBM, to build a written English-language dialogue interface for a relational database. To make it work, it needed a fairly complete model of the domain it was to be employed in. The amount of human knowledge engineering put into constructing conceptual models at SISU, while quite rich, practical, and useful, was neither Artificial, Amusing, nor Awful for those of us who, like me, were more interested in modeling conceptual behaviour than in representation. I was looking for mechanisms to model automatic acquisition and maintenance of conceptual structures. Many of the graduate students at DSV seemed to work with exciting machine learning experiments – much more Artificial, Amusing, and Awful as far as I could tell.

### How find The Optimal Representation?

So I gave up the job at SISU, better to study the behaviour of human communication, and returned to DSV. Back at the department – still located in its less than stylish but intellectually stimulating environment on campus – I spent hours on end in the cafeteria, in Jacob Palme's classes and seminars on computer-mediated human communication, and on the Usenet discussing Chinese rooms (Searle, 1980; Harnad, 1989), human communicative abilities and human cognition. Is adequate communicative behaviour a sufficient condition to postulate cognitive realism? If something passes the Turing test ▶ – is that something sentient? And is it not sentient if it fails the test? Computer science is not equipped to answer those questions without reaching out to other areas of study. Fishing about, I found parallel distributed programming and connexionism, which seemed to be the answer to many of my doubts (Rumelhart and McClelland, 1986). Here I found a representation intended to model behaviour rather than formal characteristics! I went to a Summer school to further study neural networks and connectionist representations where details of the representation were carefully and anxiously scrutinized by leading proponents of the field with regard to their orthodoxy of representation. That summer school cured me of that remedy, and in consequence of all remedies. This is one of the major lessons I learnt in my early graduate years: I ceased to search for The Optimal Representation and have since then used numerous different representations and conceptualizations of data in the projects I have worked on.

My main project in the late eighties concerned recommender systems (Karlgrén, 1990). I built prototypes to use algebraic



A group of "natural language" researchers, among them Jussi Karlgrén worked at the IBM laboratory on the island of Lidingö outside Stockholm. The information pamphlet about the project is from 1988.

### ▶ The Turing test

was introduced by Alan M. Turing (1912-1954) in his 1950 article "Computing Machinery and Intelligence" (Mind, Vol. 59, No. 236, pp. 433-460)

The Turing Test is meant to determine if a computer program has intelligence.

analysis to extract user preferences from observations on user information access behaviour: the objective was to build a system which would recommend users a document based on both their individual reading behaviour and on the reading behaviour of other users like them. The behaviour I wanted to model is that of a competent librarian: "If you like that book, you might like this one." As it turned out, data to get the system in the air were difficult to come by. Even today, with the WWW in general use, bootstrapping is a major task for the productive deployment of recommender systems; then, obtaining data, any data, was a real challenge. And in fact, data collection has been a major bottleneck for any fruitful representation and modeling of intellectual behaviour: it is difficult to obtain realistic amounts of data to prove or disprove the models in question. In the end, a few years later, I ended up using .newsrc files extracted from my colleagues home directories (Karlgrén, 1994 a) which rendered me a severe dressing down by reviewers who felt I had been disrespectful of my colleagues' integrity. This project used simple raw data, manipulated them minimally, and extracted knowledge from them in the simplest manner conceivable. Today, numerous commercial implementations operate using these and similar algorithms. An interesting facet of intelligent behaviour can be modeled with simple tools – if the data are there.

### Research at SICS

My experiences with conceptual modeling made me think about how to transmit the fairly stable representations to the more flexible counterpart in conversation – how to teach the user the representation the system purported to have learnt from users. I went back to the project I had worked on with SISU (which in meanwhile had completed its system development phase without my help and were shipping systems to delighted users) and studied how users mirrored the linguistic usage of systems. Under Carl Gustaf Jansson's supervision, this gave me my licentiate degree and a platform for several years of continued research at SICS, another research institute in the department's vicinity, where I still work (Karlgrén, 1992). I continued to study how a system might be built to express itself through superficially spurious and redundant conversational moves and thereby enrich the interaction (Karlgrén, 1994 b). The idea was to use a stable, fixed, and unchangeable conceptual model as a base, and to build an interaction module which traversed it in cases where user input was ambiguous or inexact, mumbling about related content: "Salaries for consultants? Consultants have costs. Employees have salaries." Here, even a simple symbolic interaction module with a static representation could produce something which gave a much more flexible interaction model. Again, an interesting facet of intelligent communicative behaviour can be modeled with simple tools – if the knowledge of counterpart adaptivity is there.

### The trap of doubting everything and building nothing

So these were the more important lessons I have learned during my time at DSV. Most of the pieces of the puzzle were there all along. It has taken me a while to compose a picture out of them. Some things I know now: First, and this is a personal lesson I have chosen to learn, knowledge, whether in human or machine, is to be made and created on the fly. Second, no representation is better than its use. It took a number of false starts to learn this. I am not alone in this insight: several other contributions in this volume will bear witness to the same intellectual progression. Third, and this is the most general one, and one I would like to have other young researchers remember: we cannot do sensible things if we do not try for awful, amusing and artificial! While we should be wary of representational conviction, spending time discarding representation after representation is not enough. One should not tumble into the trap of doubting everything and building nothing. There is nothing amusing, awful, nor artificial in critique alone, deconstruction is not a sufficient goal for research, and skepticism needs to be tempered with enthusiasm to lead anywhere.

The road from representational pettiness via the unproductive reaches of viel Data und wenig Theorie to a theory of data, its use and the behaviour of its users is long and no less laborious than the work involved in proving or disproving theories. But it is more productive! And turning back to the theme of this present volume: how did we get here? The group at DSV running research projects on representation and reasoning proceeded seamlessly from one false start to another – typically for computer science, with no preset intellectual compass founded in old schools of thought or solid convictions, luckily the then young researchers never got bogged down into limiting and constraining tradition. The short and unfettered intellectual direction coupled with the enthusiasm for experimental engineering that is computer science is to thank for much!

***“There is nothing amusing, awful, nor artificial in critique alone, deconstruction is not a sufficient goal for research, and skepticism needs to be tempered with enthusiasm to lead anywhere.”***

### References

- Harnad, Stevan. 1989. Minds, Machines and Searle. Journal of Theoretical and Experimental Artificial Intelligence 1: 5-25.
- Karlgren, Jussi. 1990. An Algebra for Recommendations: Using Reader Data as a Basis for Measuring Document Proximity. SYSLAB Technical Report 179. Stockholm University, Stockholm, Sweden.
- Karlgren, Jussi. 1992. The Interaction of Discourse Modality and User Expectations in Human-Computer Dialog. Licentiate thesis, department of Computer and Systems Sciences, Stockholm University
- Karlgren, Jussi. 1994 b. Mumbling – User-driven Cooperative Interaction. SICS Technical Report T94:01. Swedish Institute of Computer Science, Stockholm, Sweden.

**JUSSI KARLGREN**

---

Karlgren, Jussi. 1994 a. Newsgroup Clustering Based On User Behavior – A Recommendation Algebra. SICS Technical Report T94:04. SICS, Stockholm, Sweden.

Rumelhart, David E.; McClelland, James L. Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Psychological and Biological Models. Cambridge, Massachusetts: MIT Press. 1986.

Searle, John. 1980a. Minds, Brains, and Programs. Behavioral and Brain Sciences 3: 417-424.