

A TOURIST IN THE AI-LAB BEFORE THE REVOLUTION

Walter Daelemans
*Professor of Computational Linguistics,
University of Antwerp, Belgium*

I first met Luc Steels at a symposium, I think in Antwerp in 1984, when I was working in Nijmegen for Gerard Kempen on a language technology project. My first and temporary job in Nijmegen would soon end, and I was about to be conscripted for military service in the military academy in Brussels. Luc and Gerard were both involved in an EU Esprit project on office automation (in the very First Framework program!). Characteristically for Luc, I was as good as hired after a short chat, and he invited me to join the AI-LAB in Brussels after my military service. I started in Brussels with a week's delay and a face littered with remnants of blisters because of the chickenpox which I had picked up in the army. I was 26, just married, and it was 1986, the period of glasnost and perestroika, and of the exploding Challenger. It was an exciting time to work at the AI-lab as well. The knowledge level work of the early Steels was at its pinnacle but the complex dynamics work (the "Artificial Life route to AI") was already budding. I distinctly remember Luc explaining neural networks on the blackboard (this was at the very start of the previous neural network revolution with the Rumelhart and McClelland volumes on Parallel Distributed Processing) and dismissing them unceremoniously in favor of systems based on complex dynamics. Pattie Maes, Viviane Jonckers, Kris Van Marcke, and Walter Van De Velde were finishing their PhDs on knowledge-based approaches; Jan Paredis, Jo De Cuyper, and Bernard Manderick were starting in the new paradigm.

For me, the excitement was also in large part due to the presence of Symbolics Lisp Machines (I have never again liked hardware as much as I liked that machine). I still own a Symbolics poster of that time (see illustration), with a LM on it and a rather corny slogan, operated by a grey-haired man with glasses. "Isn't it great you now finally look like that old guy", my daughter ironically said recently when she saw the poster.

Other attractions in the late eighties AI-LAB were an artist in residence (Peter Beyls) and an American graduate student (Ken Haase). The latter imported MIT graduate school working ethics into our laid back European style of working. He was first in, last out, always working, or at least doing stuff on the computer. Luc himself was an exponent of this U.S. style as well. At the time he was interviewed in a weekly magazine (Knack), and the interviewer took ample space to describe him as an "American style professor", in jeans and sneakers, so it must have been something special back then for a Belgian academic. As a matter of fact, not much later I was thrown out of the "University Club" myself for wearing jeans: times were indeed very different then. I liked the slightly nerdy atmosphere in the lab with Eric Wybouw, the system manager, "walling" (write all) "food-p" (shall we go for lunch?) over the local network, and Peter Strickx demonstrating hefty things with Lisp Machines (well, at least various games).

Realists work on the Future – now.



Luc's management style was certainly an inspiration for me: a merit-driven, open and flexible organisation, encouraging everyone to take initiative. It is something I aspired to but never quite succeeded in achieving in my own teams afterwards. While at the AI-LAB, I was allowed to finish the computational linguistics PhD I started in Nijmegen and defended in Leuven, and go to the University of Sussex for a three-month research visit. Not many research directors would allow so much freedom (or so much carelessness about how project money was spent). Regardless how large his group was, Luc always focused on his own research first and gave his people every opportunity to focus on their own. It is one of my regrets that because of this, Luc and I never worked closely together. Somehow, our research interests never converged.

I decided to be a computational linguist rather than an AI researcher, and at that time, Computational Linguistics (CL) was moving away from AI, and they now still are far removed organisationally, with largely disjunctive research communities. Interestingly, deep learning (the current hype cycle in neural network research) promises to bring the various subfields of AI together again.

My 1987 PhD was firmly rooted in the then dominant linguistic knowledge-based approach. I had even written part of the code in KRS, the frame-based language developed at the AI-LAB, in an attempt to do something slightly relevant for the group and not feel a complete tourist. In any case, if it hadn't been for the exposure to Luc's and other lab members' ideas about learning and complex dynamics, I would have continued in this paradigm. Now I found myself experimenting with genetic algorithms for learning phonological rules, and with statistical pattern matching approaches for syntactic analysis. In the late eighties! When I got a position as a lecturer in 1989 at Tilburg University, I became one of the first researchers in Europe doing statistical computational linguistics, just before what later would become known as the statistical revolution. Unfortunately, I lost important years still trying to sell my knowledge-based work. It did not help that I was working on Dutch phonology and morphology at a time when grammar-based syntactic models of English, rooted in unification logic, were all the rage. Ironically, in some of his latest work (Fluid Construction Grammar) Luc is actually looping back to the grammar-based language models I so successfully escaped from at the time.

If I had been less lazy and better educated (like Luc I am a "germanic philologist" by training, and I was wrestling with the mathematics of statistical methods at the time), I could have been one of the leaders of the statistical revolution. Now it took me until the mid nineties to, after the early experiments, seriously get up to speed with work on statistical natural language processing. I also founded ILK then, the Induction of Linguistic Knowledge research group, which despite my slow start, still was the first group in Europe explicitly dedicated to the machine learning paradigm in natural language processing. Thanks to the AI-LAB, I ended up an early adopter of the "new" approach. If it weren't for the inspiration of Luc and his creation of an environment where I came into contact early with statistical approaches, and where it was only natural to think revolutionary, I would not have created ILK against the hostility and indifference of the dominating grammar-based and formal semantics context in Computational Linguistics. A successful academic career would have been much less likely for me.

I haven't seen this innovative spirit in any of the research groups I was involved with afterwards. In general, Computational Linguistics researchers tend to work on what promises to bring most exposure and output (currently deep learning), which mostly ends up being the same thing everyone else (and especially the big labs in US) are working on. We also lack people looking at the big theoretical issues and at cognitive relevance the way Luc does. Currently, the field is dominated by engineers suffering from accuracy fetishism.

Of the more than 30 years that I have been in research, the merely 2.5 years at the AI-LAB in Brussels in the late eighties stand out in full color. I could have done much worse than start in this environment. The example set by Luc and my AI-LAB colleagues has had a lasting effect on how I think about what research should be, and what a research group should look like.

I wrote this text in one fell swoop in a hotel room in St. Petersburg trying to ignore the irritating sound of Russian snoring coming through the wall of the next room and keeping me awake. Yet another valuable lesson about research I learned from observing Luc: never waste a chance to work.