

Zen and the Art of AI Research

Frank van Harmelen

A student went to a famous Zen teacher and said earnestly,
“I am devoted to reach enlightenment. How long will it take me to reach it?”

The teacher’s reply was casual, “Ten years”.

Impatiently, the student answered, “But I want to reach it faster than that. I will work very hard. I will practice everyday, ten or more hours a day if I have to. How long will it take then?”

The teacher thought for a moment,
“20 years”.

Like everybody, I have learned many things by being a Ph.D. student under Alan’s supervision. I learned about Constructive Logic, the Proofs as Programs paradigm, proofs by induction, heuristics for theorem provers, and a whole lot more.

But these are lessons that I could have learned in many other places, and from many other supervisors, or even from books and papers. Instead, the really valuable lessons were the unspoken ones.

It’s like studying with a Zen Master. Many skills and much knowledge can be made explicit, like the lessons that a Zen Master writes down in his books, or tells you in his speeches.

But you cannot learn the real Zen in this way. The real knowledge can only be acquired by being with the Master, by seeing how he practices Zen, and by internalising his skills by absorption.

Below are some of the lessons I have learned from the Master, as well as some of the lessons I am still struggling with.

Lessons learned from the Master by this student

- how to write an AI paper
- how to present an AI paper
- doing research in a group means not being afraid to show your mistakes
- this means that you must make yourself vulnerable

- putting early ideas to paper is productive
- running a blue note system is a good way to encourage this
- PhD supervision is a structured activity, with monthly meetings, a multi-year plan, etc
- it's worthwhile to discuss methodology explicitly, including worries about methodological problems in the field as a whole
- don't worry when your colleagues sometimes give you funny looks for worrying about methodology
- the "soft", conceptual problems in AI are sometimes (often? always?) the hardest ones
- therefore, one shouldn't only hide in the "easy", formal problems.
- it pays to do structured personal time management, for example by allocating fixed hours of the day for certain tasks
- breadth of knowledge is equally important as depth of knowledge
- therefore, always go to seminars and invited lectures, even when they are not directly related to your own research. They may be at some point.
- don't be afraid to ask questions during talks and lectures
- the goal of a talk is to explain some content matter to the audience, not to intimidate them with the intelligence or skills of the speaker
- informal activities and get-togethers as essential for a research group
- doing services for the community are an essential part of being a researcher
- scientists should take their social responsibility seriously
- therefore, scientists should be concerned with the applications of their science and its role in society as a whole

Lessons not yet mastered by this student

- how to run a research group over 25 years and maintain consistent high quality
- how to run a research group over 25 years and maintain a coherent focus