AI Magazine Volume 29 Number 2 (2008) (© AAAI)

Learning from Noise

Martin Eric Mueller

A graduate student with no AI or machine-learning education was given the problem of training a classifier that predicts a radio-frequency identification (RFID) location using field strength data from several antennas. Because the data consisted of long records of real values, the student was advised to use artificial neural networks. After several weeks of producing random classifiers, the student showed up at my office and asked whether I could help. It always seems a good idea to analyze the data first, so we constructed a primitive visualization: signal strength of four antennae over time. The graphs looked like we'd glued a pen on a dog's tail while showing him a juicy T-bone steak. I suggested we add a few functions, such as pair-wise difference, mean, deviation, and so on—just to get a feel for the data.

The image did not change in general; it was like having the dog run over the same picture several times. So I suggested we sort the data points by the actual target location and then see whether the plot changed at all. It did not. Same dog, same steak—this time the dog was just jumping back and forth instead of walking from left to right. In other words, the data simply didn't include any information that could be used to extract knowledge.

It turns out the data had been collected in a building with steel girders whose reflections had reduced it to white noise. The experimenters hired the student because they were unable to learn a Bayesian classifier. People often show a behavior known as confirmation bias: if something doesn't work as expected, try again—just a little bit harder. In this case, it led to a totally unnecessary attempt to extract information from nothing.

Martin Eric Mueller is assistant professor in the Department of Computer Science at the University of Augsburg. Since 1997, he has been researching in machine learning and its applications to adaptive web search and user modeling (for which he earned his Ph.D. in 2001 from the University of Osnabrück) and, since then, to human-computer interaction, context-aware systems, and cognitive robotics.