The Annals of Applied Statistics
2010, Vol. 4, No. 4, 1638–1641
DOI: 10.1214/10-AOAS381

REMEMBRANCE OF LEO BREIMAN

BY PETER BÜHLMANN

ETH Zürich

1. How I met Leo Breiman. In 1994, I came to Berkeley and was fortunate to stay there three years, first as a postdoctoral researcher and then as Neyman Visiting Assistant Professor. For me, this period was a unique opportunity to see other aspects and learn many more things about statistics: the Department of Statistics at Berkeley was much bigger and hence broader than my home at ETH Zürich and I enjoyed very much that the science was perhaps a bit more speculative.

As soon as I settled in the department, I tried to get in touch with the local faculty. Leo Breiman started a reading group on topics in machine learning and I didn't hesitate to participate together with other Ph.D. students. Leo spread a tremendous amount of enthusiasm, telling us about the vast opportunity we now had by taking advantage of computational power. Hearing his views and opinions and listening to his thoughts and ideas has been very exciting, stimulating and entertaining as well. This was my first occasion to get to know Leo. And there was, at least a bit, a vice-versa implication: now, Leo knew my name and who I am. Whenever we saw each other on the 4th floor in Evans Hall, I got a very gentle smile and "hello" from Leo. And in fact, this happened quite often: I often walked around while thinking about a problem, and it seemed to me, that Leo had a similar habit.

2. Witnessing three of Leo's fundamental contributions. I only got to know Leo Breiman in his late career. Nevertheless, between 1994 and 1997 when I was in Berkeley, I could witness Leo's exceptional creativity when he invented Bagging [Breiman (1996a)], gave fundamental explanations about Boosting [Breiman (1999)] and started to develop Random Forests [Breiman (2001)].

2.1. *Bagging.* I had the unique opportunity to listen to Leo Breiman when he presented Bagging during a seminar talk at UC Berkeley. I was puzzled and intrigued. At that time, I was working on the bootstrap and what Leo said didn't sound right to me: using the bootstrap language, he proposed to use $\hat{\theta}_{\mathrm{Bag}} = \mathbb{E}^*[\hat{\theta}^*]$, where $\hat{\theta}$ is the output of a "complex algorithm" based on the original observations and $\hat{\theta}^*$ denoting the analogue based on the bootstrap sample. Trivially,

$$\hat{\theta}_{\mathrm{Bag}} = \hat{\theta} + (\mathbb{E}^*[\hat{\theta}^*] - \hat{\theta}),$$

Received July 2010.
AMS 2000 subject classifications. Primary 62G08, 62G09; secondary 68T10.
Key words and phrases. Bagging, boosting, classification and regression trees, random forests.