type questions to the evening panel here…

1. Hyperparameter optimization
    a. Can we say that fixed-space hyperopt is close to be solved? No :-)
    b. Random search has been used for parameter tuning for over 40 years in other communities (e.g., evolutionary computation) and only started to be used in ML 5 years ago. Do you know of any other examples like that we might still be missing out on?
    c. We have a handful of tools (SMAC, hyperopt, spearmint/whetlab, CMA-ES optunity, TPE, pyBO). How do we choose? Are we accumulating meta knowledge?
    d. Conditional hyperparameters still seems a tough problem, little research.
    e. In some papers (for instance deep learning) it is not always quite clear whether the contribution alone leads to much better results or if the authors were just able to set their hyperparameters better. Could reviewers insist of a proper hyperparameter optimization in order to make a fair comparison between different machine learning methods?
    f. We have nowadays many sophisticated hyperparameter optimization methods. However they also demand for a lot of computation power (GPUs, Parallel Computing on a CPU cluster) and time (Bayesian Optimization still needs many function evaluation especially in high dimensional spaces). Can we improve those methods such that they not only lead to good results but also allow users to run it on their laptop in a reasonable time?
2. What do you think about Rich's idea on running a data diff challenge?
3. Anything we can do on problem formalization? "We have data, what can we do with it?"
4. @Rich: focusing on some parts of the pipeline will be useful to focus work. However, the parts don't factorize perfectly: different ML methods want different data coding, normalization, calibration, etc. What's the best way of moving forward?
5. Is the human in the loop (of the development pipeline) possible/desirable?
6. Industry seems to be ahead of us, especially in accumulating multi-problem experience. Beside big IT companies which are obviously building up a lot of experience but perhaps on a smaller segment of type of problems, startups that claim to automatize ML also accumulate experience on the tail of the problem distribution (eg DataRobot has built around 30 million models). Do we have any contact with these startups? Any idea how to involve them in research? Are they making us obsolete?
7. At what point can we be confident enough in an AutoML system that we could hand it over to the average Excel user without fear of costly or, potentially, dangerous results? How much work is left to close the gap on getting there? 20 years.
8. @Schmidhuber about self-improving learning: Do you have comments about the concerns of the Future of Life Institute about the pitfalls of AI and the future dangers of it?