Online Supplement to "The Knowledge-Gradient Algorithm for Sequencing Experiments in Drug Discovery"

Diana M. Negoescu
Department of Management Science and Engineering,
Stanford University, Stanford, CA 94305, USA,
Peter I. Frazier
School of Operations Research and Information Engineering,
Cornell University, Ithaca, NY 14853, USA,
Warren B. Powell
Department of Operations Research and Financial Engineering
Princeton University, Princeton, NJ 08544, USA,
{negoescu@stanford.edu, pf98@cornell.edu, powell@princeton.edu}

In this section we present further numerical experiments using the Free-Wilson model, which complement and extend the results presented in Section 6.2.1.

Since the experiments in Section 6.2.1 showed that pure exploration is more competitive than OFAAT, the experiments we present here compare KGCB only to pure exploration. In these further numerical experiments, we randomly selected sets of compounds from the full set of 87120.

Figure 1 shows a sample path of 200 measurements on a set of randomly selected 1000 compounds, giving the opportunity cost at every iteration on the top panel and the true value of the compound that is measured at each step on the bottom panel. While pure exploration fails to find the best compound in the first 200 measurements, KGCB manages to do so in about 15 measurements (on this sample path), which is a significant improvement if we think about the amount of time and money each measurement requires. As the bottom panel shows, part of KGCB's success in this sample path was the fact that it measured two very good compounds early on, at steps 10 and 15, while pure exploration was measuring mediocre compounds for a long period of time, which prevented it from discovering the best compound quickly enough.

We also present the results from running the code on data sets of randomly selected 10000 compounds. Figure 2 shows four such sample paths of the new implementation. As the figure shows, the performance of the KGCB algorithm is encouraging, as it is usually able to find the best compound by the 75th measurement, and gets very close to finding it by the 25th measurement. In most sample paths, KGCB does at least as well as pure

exploration - the policy in which we choose compounds to measure completely at random, but update the belief in the same way as we do for the KGCB policy.

We are interested also in the typical performance of the KGCB policy. That is, if we are able to perform only one sample path (which is generally the case in practice), we are interested in knowing the probability that the KGCB policy performs better than an exploration policy.

To better get a sense of the distribution of the relative performance of KGCB versus pure exploration, we have made a box and whiskers plot of the differences between opportunity costs at every 10 measurements. As shown in Figure 3, except for the first measurement, all the other measurements have the lower quartile cutoff at or above 0, which suggests that the probability that KGCB is better than pure exploration, on any measurement of any sample path, is higher than 1/2.

Having tested the KGCB policy on data sets of 10000 compounds, we increased the size of the data sets further to 25000 compounds. Figure 4 shows four sample paths from running the KGCB and pure exploration policies on four randomly selected data sets of 25000 compounds. The rate of convergence for these plots is slower than for 10000 compounds, but the KGCB policy still manages to get reasonably close to the best compounds after about 50 measurements.

Figure 5 shows the mean opportunity cost for nine sample paths using nine different randomly chosen data sets of 25000 compounds. Although not as impressive as the mean plot for 10000 compounds, this plot also attests that there is value in using the KGCB policy as opposed to a pure exploration policy.

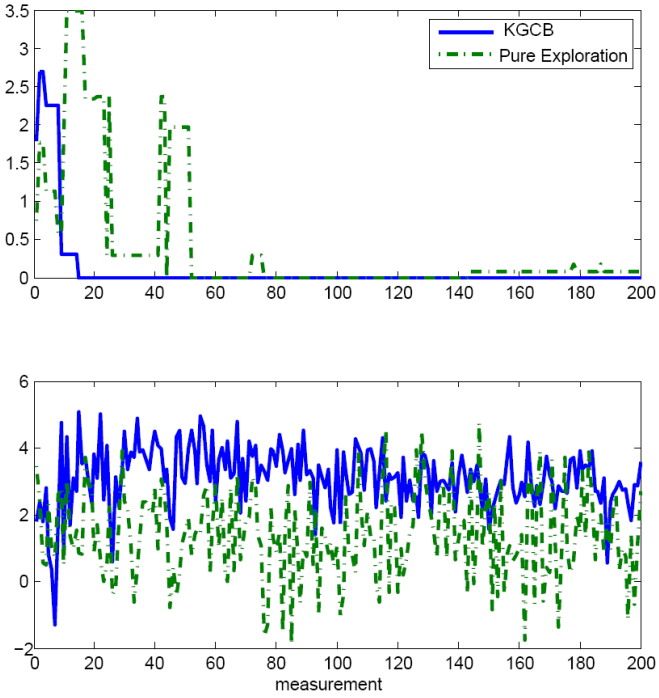

Figure 1: A sample path of 200 measurements on a data set of 1000 compounds showing opportunity cost (top) and the true values of the compounds tested at each time step (bottom).

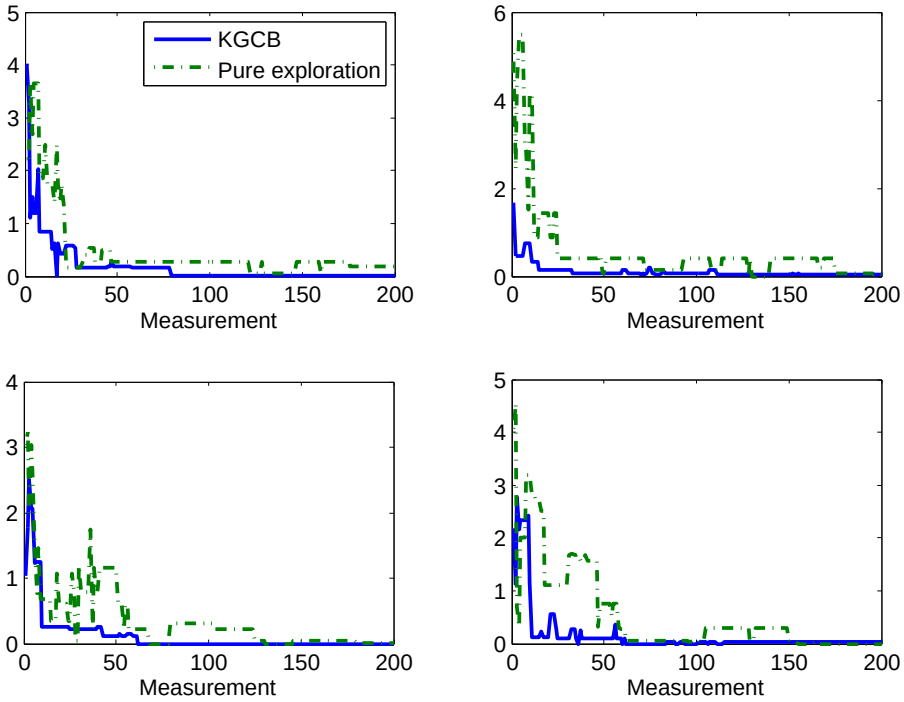

Figure 2: Four sample paths using data sets of 10000 compounds and a noise standard deviation of 0.38.

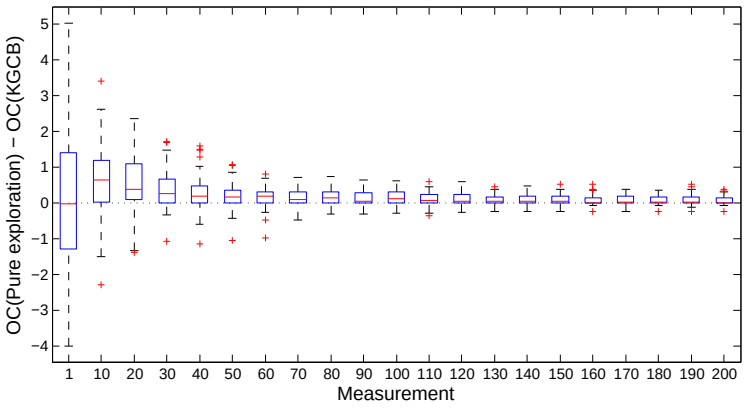

Figure 3: Distribution of difference between opportunity costs between pure exploration and KGCB using 75 sample paths of 10000 compounds each and a noise standard deviation of 0.38.

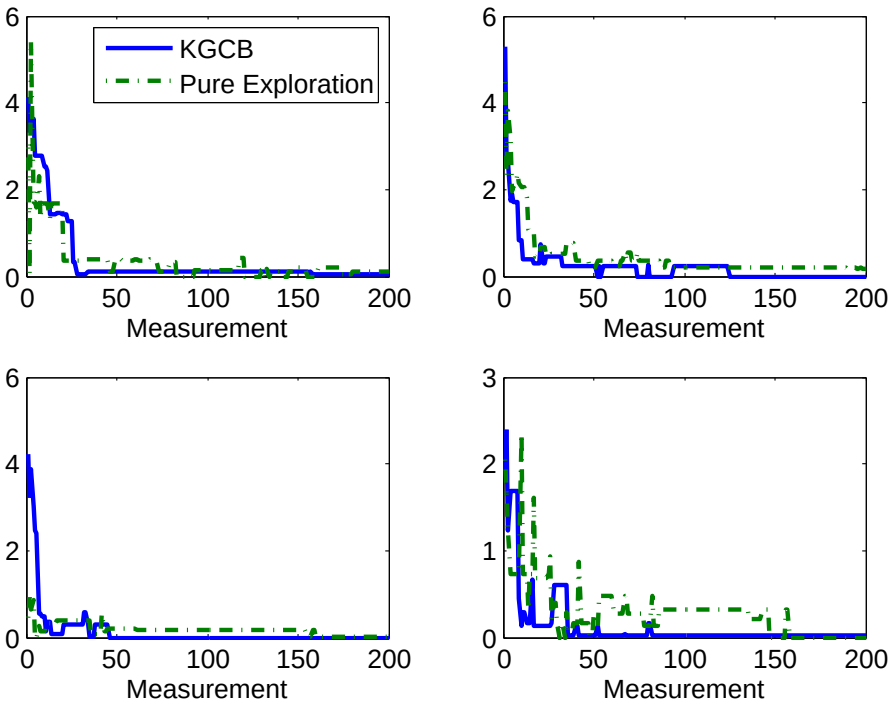

Figure 4: Four sample paths suing data sets of 25000 compounds and a noise standard deviation of 0.38.

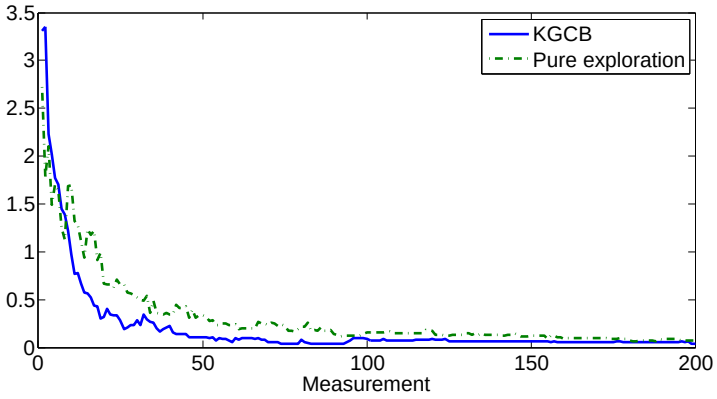

Figure 5: Average over nine runs of sample paths using data sets of 25000 compounds and a noise standard deviation of 0.38.