Erratum for section 8.3.7 in Rasmussen and Williams, Gaussian Processes for Machine Learning (MIT Press, 2006)

In section 8.3.7 we present a comparison of 4 methods (SR, SR, PP, BCM) on the SARCOS robot arm inverse dynamics data. Unfortunately there was an error in the scripts that meant that the noise variance was added in twice when computing the predictive variance for the PP runs; this affected the PP results for MSLL (but not SMSE). In this erratum we present corrected results for Table 8.1 (corrected values in bold) and Figure 8.1(b).

Method	m	SMSE	MSLL	mean runtime (s)
SD	256	0.0813 ± 0.0198	-1.4291 ± 0.0558	0.8
	512	0.0532 ± 0.0046	-1.5834 ± 0.0319	2.1
	1024	0.0398 ± 0.0036	-1.7149 ± 0.0293	6.5
	2048	0.0290 ± 0.0013	-1.8611 ± 0.0204	25.0
	4096	0.0200 ± 0.0008	-2.0241 ± 0.0151	100.7
SR	256	0.0351 ± 0.0036	-1.6088 ± 0.0984	11.0
	512	0.0259 ± 0.0014	-1.8185 ± 0.0357	27.0
	1024	0.0193 ± 0.0008	-1.9728 ± 0.0207	79.5
	2048	0.0150 ± 0.0005	-2.1126 ± 0.0185	284.8
	4096	0.0110 ± 0.0004	-2.2474 ± 0.0204	927.6
PP	256	0.0351 ± 0.0036	**-1.6940 ± 0.0528**	17.3
	512	0.0259 ± 0.0014	**-1.8423 ± 0.0286**	41.4
	1024	0.0193 ± 0.0008	**-1.9823 ± 0.0233**	95.1
	2048	0.0150 ± 0.0005	**-2.1125 ± 0.0202**	354.2
	4096	0.0110 ± 0.0004	**-2.2399 ± 0.0160**	964.5
BCM	256	0.0314 ± 0.0046	-1.7066 ± 0.0550	506.4
	512	0.0281 ± 0.0055	-1.7807 ± 0.0820	660.5
	1024	0.0180 ± 0.0010	-2.0081 ± 0.0321	1043.2
	2048	0.0136 ± 0.0007	-2.1364 ± 0.0266	1920.7

Table 8.1: Test results on the inverse dynamics problem for a number of different methods. Ten repetitions were used, the mean loss is shown ± one standard deviation.

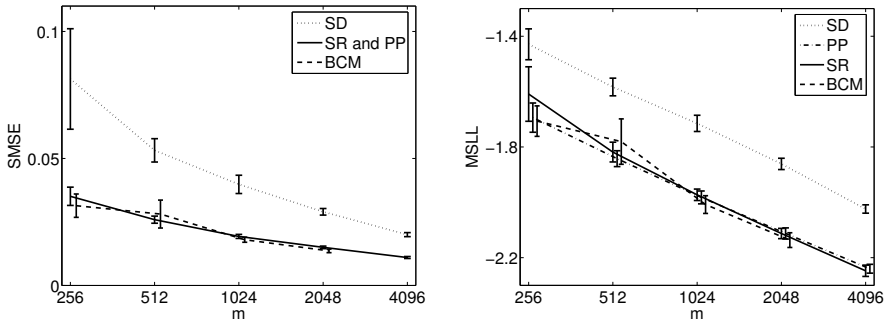

(a) (b)

Figure 8.1: Panel(a): plot of SMSE against m. Panel(b) shows the MSLL for the four methods. The error bars denote one standard deviation. For clarity in (a) the BCM results are slightly displaced horizontally w.r.t. the SR results, and in (b) both the PP and BCM results are slightly displaced horizontally w.r.t. the SR results.

The PP method seems to have slightly better MSLL performance than SR for small m (256 and 512).

Notice now that there is not much difference in performance between the SR, PP and BCM methods for various sizes of m on this problem, and that they all outperform the SD method. However, as mentioned in the original text, it may make most sense to compare performance against runtime, in which case SD for $m = 4096$ is competitive with SR for $m = 1024$.

Chris Williams
20 May 2006