Correction to results appearing in paper
"Maintaining Attentional Capacity in a Social Robot"

Yuval Marom

April 20, 2000

The results presented in the paper were erroneous due to a fault in the software, more specifically, in the part responsible for resetting the neural network between runs. As a result the network was not reset properly resulting in a violation of the assumption of independence between runs.

New features revealed by corrected data

After amending the problem, new features appeared in the data. In the context of this paper, these features are in terms of the energy acquired during recall. The table below summarises these new results.

- large variance in the low and middle threshold regions;
- two peaks, one in each of these regions, at thresholds of around 0.15 and 0.3 (as opposed to just one peak in the middle region, as reported in paper);
- these peaks are higher than the energy acquired by hand-crafted behaviour, and this is due to localisation (see explanations below).

energy ($\times 10^7$)	mean	std. dev.	max
th < 0.2	2.704	0.205	3.198
$0.2 \leq$ th ≤ 0.4	2.681	0.339	3.331
th > 0.4	2.354	0.080	2.558
attention strategy	2.865	0.257	3.270
random behaviour	2.328	0.056	2.427
hand-crafted behaviour	2.905	0.080	3.089

Table 1 (revised): Energy acquired during recall stage (see Table 1 in paper for more details).

Explanations for the new features

The new findings have encouraged us to examine each of the learning and testing/recall stages separately, and more closely. We repeated the experiments reported in the paper, but did not use the same threshold at learning and recall. We ran learning episodes of each threshold, and for each episode, we went through the entire threshold space at the recall stage. In order to eliminate bias of a particular learning episode, we repeated an instance of the same learning threshold many times. We found that the level of attentiveness at the learning stage should not necessarily be the same at recall!

Learning stage. At low thresholds learning is inconsistent. Low selectivity results in swamping of the neural network with undesired stimuli[1], sometimes noisy. At high thresholds the agent is much more selective, and the network is exposed to more consistent patterns.

Recall stage. At low recall thresholds, the agent tends to stay around bright regions, because it is more attentive and hence sensitive to slight changes in light intensity. This emergent localisation explains the high peak in the low threshold region[2]. However, if the threshold is too low, the agent may even localise around noisy parts of the environment, hence not acquire high energies. As the threshold increases, the emergent behaviour is less of localisation and more of *exploration*: the agent is less attentive to changes around it, and is more willing to move forward to new locations. As the threshold gets too high (> 0.4), the agent is not attentive enough, and exploration is at the expense of exploitation — critical stimuli are ignored.

In addition, when 'bad' cases of learning are identified (according to a desired ratio of learned components[3]) and removed, the variation in the results is much smaller.

Implications to Attentional Effort

It now seems that the attention allocation policy (see Section 2.2 in paper) needs to be different for learning and recall. It seems from the new data that the need to pay attention is greater when the agent is on its own than when guided through the environment by a capable supervisor. Less attentive means more selective, and it is especially important to be selective when learning by following, because the following is not precise: firstly, there is a small time lag, and secondly the path taken by the learner is not exactly the same as the teacher. The agent needs to be sure that what it is about to 'learn' is significant enough for its purposes, as discussed in the paper.

We have therefore repeated the experiments involving attentional strategy, with a low attentional effort maintained at the learning stage (0.35 to 0.4), and a higher effort maintained at the recall stage (0.55 to 0.7). The acquired energy has mean=2.933, SD=0.174, and max=3.681 (all $\times 10^7$).

We also looked at the proportion of 'bad' cases (see above), p, and tested whether it is significantly less than 0.5. When using a single policy, $p = 0.42$; when using a separate policy for learning and recall $p = 0.16$; only the latter is significant (at both 5% and 1% significance levels).

[1]For example, a light source on the left is also perceived on the right.

[2]The emergent hand-crafted behaviour is to move away from a bright source once going past it, and to go looking at other locations in the environment.

[3]Ratio of left : right : straight moves