Online Learning in Monkeys

Xiaojin Zhu[†], Michael Coen[†‡], Shelley Prudom[§], Ricki Colman[§], and Joseph Kemnitz[§]
[†]Department of Computer Sciences
[‡]Department of Biostatistics and Medical Informatics
[§]Wisconsin National Primate Research Center
University of Wisconsin, Madison, WI 53706, USA
{jerryzhu, mhcoen}@cs.wisc.edu, {sprudom, rcolman, kemnitz}@primate.wisc.edu

Abstract

We examine online learning in the context of the Wisconsin Card Sorting Task (WCST), a task for which the concept acquisition strategies for human and other primates are well documented. We describe a new WCST experiment in rhesus monkeys, comparing the monkeys' behaviors to that of online learning algorithms. Our expectation is that insights gained from this work and future research can lead to improved artificial learning systems.

WCST as Online Learning

Online learning (Blum 1998; Kivinen, Smola, and Williamson 2004) assumes that an adversary provides data sequentially to a learner. In particular, the adversary can change the target concept at any time (called a *concept drift*), unbeknownst to the learner. This is in sharp contrast to the stationary independent and identically-distributed (iid) data assumption in standard learning models. As such, online learning naturally models changing environments. An interesting question is how intelligent animals learn under similar conditions – the real world certainly has concept drifts, and experiences arrive sequentially. In order to study learning in natural and artificial systems, we connect the Wisconsin Card Sorting Task (WCST), a well-known neuropsychological test for humans (Milner 1963) and other primates (Nakahara et al. 2002), to online learning algorithms. While the WCST has been studied computationally before at the neurological level (Dehaene and Changeux 1991), we believe this study is the first to address the connection between online learning in primates and learning algorithms. Understanding how monkeys and people approach this game allows us to explore to what degree they employ similar strategies. Studying high level cross-species cognition provides insight into the origins and primitives of abstract reasoning, which is fundamental to modeling it computationally.

The WCST tests for cognitive "set shifting" (i.e. the readiness to abandon an old learned concept after a concept drift), which is difficult for human patients with pathology in the frontal lobes of the brain. We describe a simplified WCST for rhesus monkeys (*Macaca mulatta*), see figure below. In each trial, three objects (Circle (C), Star (S), Triangle (T)) appear simultaneously on screen, each with a different color (Red (R), Green (G), Blue (B)). The shape and color combination is randomly determined; all combinations will appear. The target concept is initially R. The monkey is rewarded with a food pellet for touching the red object regardless of its shape. In an important manipulation, once the monkey has had 10 consecutive correct trials, a concept drift happens where the target concept is changed to T without warning. The monkey must now touch triangular objects to receive pellets. After 10 consecutive touches of triangles the concept drifts again to B. After 10 consecutive touches of blue objects the concept drifts one last time to S. The task ends when the monkey makes 10 consecutive correct touches of stars. We model the WCST as an online game.

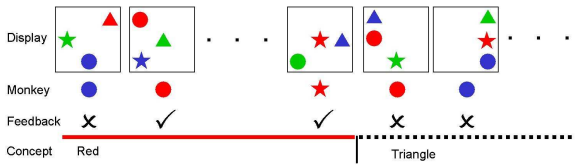

Let $x = (f_1 \dots f_d)$ be an object with d Boolean features. In our experiment the $d = 6$ features specify whether the object is red, green, blue, circle, star, and triangle, respectively. For example, a red star has the feature vector $x = (100010)$. A concept is a feature vector with exactly one 1. In each trial, the adversary may change the concept (which corresponds to a concept drift). The adversary then prepares $d/2$ objects $x_1, \dots, x_{d/2}$ and shows them to the learner. Each feature appears once in these objects. The learner picks one object x, and receives a correct/wrong feedback from the adversary depending on whether x fits the current concept. The game then repeats. The learner's goal is to minimize the number of mistakes. This is a special case of the disjunction learning problem. We present a simple online learning algorithm:

1. Initially $h = (1 \dots 1)$ (d ones). Repeat 2–4:
2. Randomly pick $x \in \{x_1 \dots x_{d/2}\}$ for which $h \wedge x \neq 0$
3. If x is correct, $h = h \wedge x$.
4. If x is wrong, $h = h \wedge \neg x$. If $h = 0$, reset $h = (1 \dots 1)$.

The learner starts with the hypothesis h that any feature can be the concept. Step 2 picks an object x (say "red star") that fits the hypotheses. If x is correct, the true concept must be either "red" or "star", thus the learner can remove all

other features from the hypothesis (step 3, $\wedge$ is elementwise AND). If x is wrong, the true concept can be neither "red" nor "star", and these features are removed (step 4). With changing concepts, the hypothesis can become empty–the learner detects a concept drift. It simply learns from scratch again. We can bound the number of mistakes:

Theorem 1. *For* any *input sequence with m concept drifts, our algorithm makes at most $(2m+1)(d-1)$ mistakes.*

We omit the proof due to space. Note there are online algorithms with tighter bounds (e.g., Winnow (Auer and Warmuth 1995) can achieve $O(m \log d)$), but the present algorithm is simpler for comparisons to natural learning.

Experiments in Monkeys and Computers

Monkey behaviors in the WCST. This experiment was part of a cognitive survey of aged rhesus macaques in a dietary restriction study at the Wisconsin National Primate Research Center. Seven monkeys completed the WCST using a touch screen apparatus, which they learned to operate through a series of acclimation tasks. Animals appear to use different strategies in the WCST and appear to advance in stages. First they may have difficulty "unlearning" the previous concept and commit large numbers of errors, mostly perseverative with the occasional pellet motivation to continue to perform. They then seem to realize the possibility of a new concept and will often begin to have 2 to 4 correct responses in a row. We make the following observations.

1. *Monkeys adapt to concept drifts slowly.* As Table 1 m-columns show, the monkeys took about 300 trials to abandon an old concept and succeed 10 trials in a row on the new concept. This is slow, but clearly better than random guessing, which needs around 8.8×10^4 trials. Paired t-tests at 0.05 level do not reveal significant differences between the number of trials for the four concepts: the monkeys do not seem to get better at concept drifts.

	m-trials	m-errors	m-presv	c-errors	c-presv
R	425±254	242±187	-	4±2	-
T	249±95	113±43	89±23	4±3	3±3
B	437±301	247±193	186±109	3±2	2±1
S	279±101	132±48	94±30	2±1	1±0

Table 1: Trials, errors, and perseverative errors for each concept in monkey- and computer-experiments (mean±std).

2. *Perseverative error dominates.* On average a monkey made 732 errors in all four concepts. A perseverative error is a mistake which would have been correct under the immediately preceding concept. Perseverative errors account for the majority (369 out of 490, 75%) of errors a monkey made in T,B,S. Perseveration can span multiple concepts: for example, the monkeys select many red objects not only in T, but also in the beginning part of B.

3. *The monkeys do not slow down after a concept drift.* We do not observe statistically significant differences in response time before and after the concept drifts. It seems they do not pause and puzzle over the "food reward malfunction" after a concept drift.

Online learning algorithm performance. We subject our online learning algorithm to the same data sequence received by each monkey individually. The results strongly correspond to what would be expected from the subset principle and the notion of identification in the limit. This is illustrated in Table 1's c-columns, where the mistakes happen very soon after each concept drift, and the algorithm quickly adjusts to the new concept, making no further mistake until the next concept drift. The worst data sequence produces 13 errors in all four concepts, as bounded by Theorem 1 ($m = 3, d = 6$, bound=35).

A slow, stubborn algorithm. We make two modifications: i) "slow": In each trial, skip learning (steps 3, 4) with probability α. This assumes that it takes many tries to modify h. ii) "stubborn": when h becomes 0 in step 4, insist on the old, incorrect h with probability β, otherwise reset $h = (1 \ldots 1)$. By setting α, β judiciously, we can simulate the imperfect monkey responses. For example, when $\alpha = 0.93, \beta = 0.96$ the algorithm makes on average 563 errors, with 312 perservative errors (67%) out of 469 in T,B,S.

Conclusions and Future Work

We presented a Mistake Bound analysis of the WCST. We performed WCST experiments in monkeys and compared them to online learning algorithms. We suspect the primary cause for the "inefficiency" in the monkeys' responses is due to the fact that they must first acquire the hypothesis space for the game, whereas the algorithm's search space is predefined. Thus, the monkeys do not initially even know what "game" they are playing. It is quite plausible that many of the observed phenomena with the monkey population are consequences of their acquiring the inductive bias of the WCST and do not actually reflect the skill a more trained monkey would display. This perhaps explains the absence of any delay following a concept shift, as the monkeys may not yet be actually playing the WCST as we understand it. Exploring this possibility will be the subject of future work.

Acknowledgment We thank Timothy Rogers for WCST discussions. Research supported in part by the Wisconsin Alumni Research Foundation, NIA grant P01 AG11915, NCRR grant P51 RR000167, and the University of Wisconsin School of Medicine and Public Health.

References

Auer, P., and Warmuth, M. 1995. Tracking the best disjunction. In *The proceedings of the 36th annual symposium on foundations of computer science*, 312–321.

Blum, A. 1998. On-line algorithms in machine learning. In Fiat, and Woeginger., eds., *Online Algorithms: the State of the Art*.

Dehaene, S., and Changeux, J.-P. 1991. The wisconsin card sorting test: Theoretical analysis and modeling in a neuronal network. *Cerebral Cortex* 1:62–79.

Kivinen, J.; Smola, A. J.; and Williamson, R. C. 2004. Online learning with kernels. *IEEE Transactions on Signal Processing* 52(8):2165–2176.

Milner, B. 1963. Effect of different brain lesions on card sorting. *Archives of Neurology* 9:90–100.

Nakahara, K.; Hayashi, T.; Konishi, S.; and Miyashita, Y. 2002. Functional MRI of macaque monkeys performing a cognitive set-shifting task. *Science* 295(5559):1532–1536.