

## A UPDATING FORMULA IN VARIANCE ESTIMATION

Here we show that the weighted standard deviation in equation 6 can be obtained by updating  $\sigma$  with formula 7 in Variance Estimation. For simplicity, the inputs to all the following functions are hidden as they are the same state-action pair  $(s, a)$ .

While being updated with equation 1, we have

$$Q_n = (1 - \alpha)^n Q_0 + \sum_{i=1}^n \alpha(1 - \alpha)^{n-i} G_i \quad (13)$$

$$= \alpha \sum_{i=1}^n (1 - \alpha)^{n-i} G_i \quad (14)$$

The equal sign in formula 14 is established by setting  $Q_0 = 0$ .

If we expand  $1/\alpha$  as a power series

$$\frac{1}{\alpha} = \frac{1 - (1 - \alpha)^n}{1 - (1 - \alpha)} \cdot \frac{1}{1 - (1 - \alpha)^n} \quad (15)$$

$$= \sum_{i=0}^{n-1} (1 - \alpha)^i \cdot \frac{1}{1 - (1 - \alpha)^n} \quad (16)$$

$$\approx \sum_{i=0}^{n-1} (1 - \alpha)^i \quad (17)$$

$$= \sum_{i=1}^n (1 - \alpha)^{n-i} \quad (18)$$

The approximate sign is used here to reflect that the second term of the multiplication approximates 1 as  $n \rightarrow \infty$  because  $\alpha \in (0, 1]$ . The value function can be re-written as an exponentially weighted sum of the return sequence by applying the above expansion of power series:

$$Q_n \approx \frac{\sum_{i=1}^n (1 - \alpha)^{n-i} G_i}{\sum_{i=1}^n (1 - \alpha)^{n-i}} \quad (19)$$

We denote  $A_n = \sum_{i=1}^n (1 - \alpha)^{n-i}$ , then

$$\sigma_{n+1}^2 = \frac{\sum_{i=1}^{n+1} (1 - \alpha)^{n+1-i} (G_i - Q_{n+1})^2}{A_{n+1}} \quad (20)$$

$$= \frac{\sum_{i=1}^n (1 - \alpha)^{n+1-i} (G_i - Q_{n+1})^2}{A_{n+1}} + \frac{(G_{n+1} - Q_{n+1})^2}{A_{n+1}} \quad (21)$$

$$= (1 - \alpha) \cdot \frac{\sum_{i=1}^n (1 - \alpha)^{n-i} (G_i - Q_{n+1})^2}{A_{n+1}} + \frac{(G_{n+1} - Q_{n+1})^2}{A_{n+1}} \quad (22)$$

$$= (1 - \alpha) \cdot \frac{\sum_{i=1}^n (1 - \alpha)^{n-i} (G_i - Q_n + Q_n - Q_{n+1})^2}{A_{n+1}} + \frac{(G_{n+1} - Q_{n+1})^2}{A_{n+1}} \quad (23)$$

$$= (1 - \alpha) \cdot \left[ \frac{\sum_{i=1}^n (1 - \alpha)^{n-i} (G_i - Q_n)^2}{A_{n+1}} + \frac{\sum_{i=1}^n (1 - \alpha)^{n-i} (Q_n - Q_{n+1})^2}{A_{n+1}} \right] + \frac{(G_{n+1} - Q_{n+1})^2}{A_{n+1}} \quad (24)$$

$$\approx (1 - \alpha) \cdot \frac{A_n}{A_{n+1}} \cdot [\sigma_n^2 + (Q_{n+1} - Q_n)^2 + 0] + \frac{(G_{n+1} - Q_{n+1})^2}{A_{n+1}} \quad (25)$$

$$\approx (1 - \alpha) [\sigma_n^2 + (Q_{n+1} - Q_n)^2] + \alpha(G_{n+1} - Q_{n+1})^2 \quad (26)$$

where the first approximate sign comes from formula 19 and the second one comes from formula 17.

Though the updating formula is biased as  $n$  doesn't approach infinity in practice, the bias is negligible. Because all the above approximations are rooted in formula 17 and  $(1 - \alpha)^n$  converges to 0 rapidly as  $n$  grows.

## B AN EMPIRICAL ANALYSIS

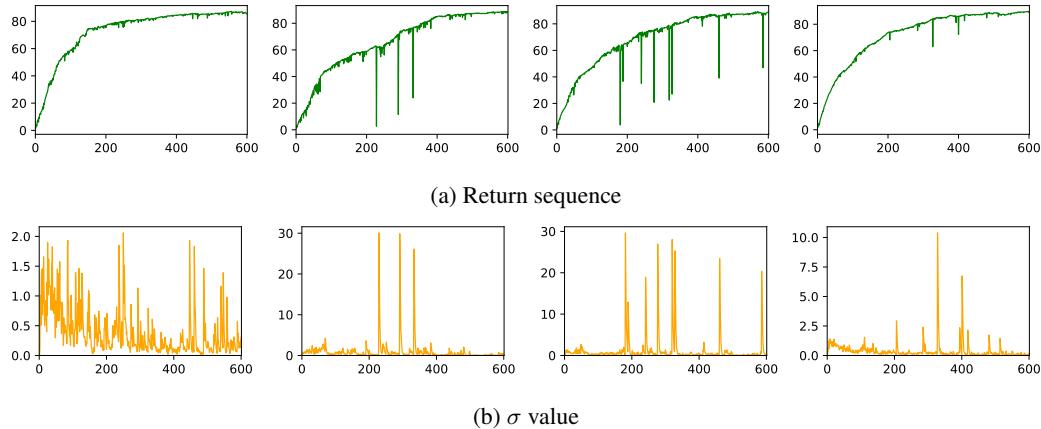


Figure 5: Raw return sequences (a) and corresponding  $\sigma$  values (b) over visitations.

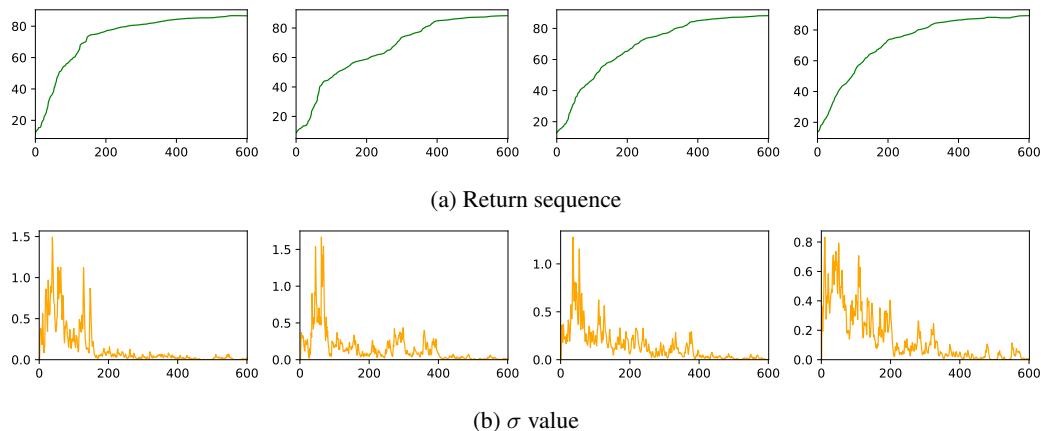


Figure 6: Smoothed return sequences (a) and corresponding  $\sigma$  values (b) over visitations.

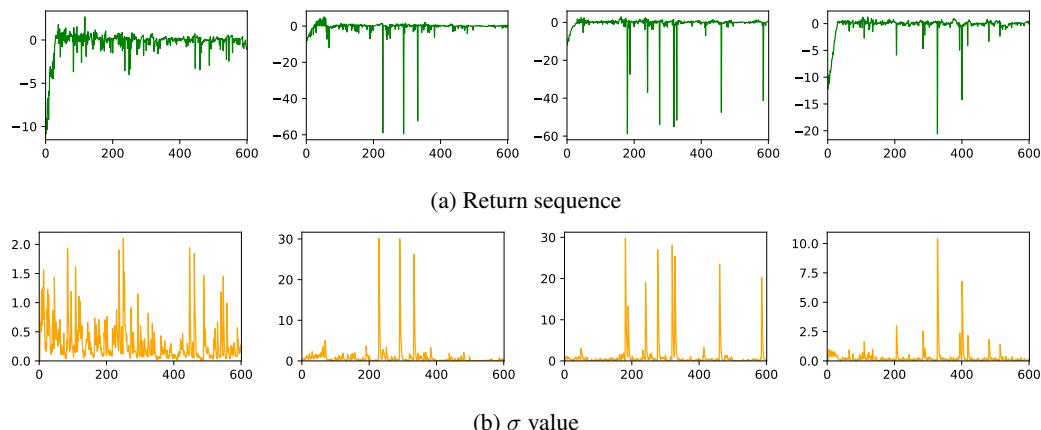


Figure 7: Residual return sequences (a) and corresponding  $\sigma$  values (b) over visitations.

Sutton & Barto (2018) prove that the policy and value functions monotonically improve over sweeps in policy iteration. In this simple form, return sequences for individual state-action pairs are monotonically increasing and usually converge fast. However, the same statement is not true when the policy is updated with Monte Carlo settings where the returns are random samples. In this situation, fluctuations in returns are inevitable due to the stochastic property; an incremental updating scheme is adopted to stabilize learning as well as avoid the excessive impact of malicious samples.

Here we illustrate the characteristics of the return sequences and analyze how variance guides exploration. The raw return sequences shown in Figure 5a(a) are randomly sampled from frequently visited state-action pairs in the Cartpole balancing problem and truncated to be of the same length. To disentangle the impact of the convergence trend from transient fluctuations, smoothed versions of those sequences are extracted from the raw sequences (shown in Figure 6a), while the residuals are shown in Figure 7a. Then we apply our Variance Estimation method on each of the sequences independently and show the  $\sigma$  values over visits below the sequences.

While  $\sigma$  values are integrated into the Q values used by our exploration policy (shown in Equation 5), a greater  $\sigma$  value usually results in an increase in exploration budget. Meanwhile, the nature of weighted sum balances the importance of learned Q value and the history of its change which is captured by  $\sigma$ .

In Figure 6b, we observe that the  $\sigma$  value is greater when the return sequence changes at a faster rate. As the sequence converges, the  $\sigma$  value approaches zero. An interesting but unobvious observation is that the  $\sigma$  value spikes when there is change in the convergence rate of the return sequence. Since variance is essentially a Euclidean distance metric, it is capable of capturing second-order information.

Sequences in Figure 7a isolate the impacts of transient fluctuations from the overall trend. We observe that whenever there is an excessive fluctuation appears,  $\sigma$  value spikes to a high magnitude to demand immediate exploration. Once the return goes back to its normal range, the  $\sigma$  value decreases simultaneously. Those quick responses are useful since excessive fluctuations are harmful to the estimation of Q values. A timely investigation of exploration budget eliminates this negative impact before it propagates to more states. Meanwhile, frequent fluctuations beget an increase in  $\sigma$  value and result in more exploration to determine its value.

In conclusion, the exploration policy on our constructed upper bound effectively allocates exploration budget to accelerate convergence in important states as well as alleviate the impact of fluctuations.

## C HYPERPARAMETERS

Table 2: Atari DQN Hyperparameters

Hyperparameter	Value	Description
$c_{\text{V-DQN}}$	0.1	Weighting factor of $\sigma$ -stream in exploration policy in V-DQN
$c_{\text{TD-DQN}}$	0.1	Weighting factor of $\sigma$ -stream in exploration policy in TD-DQN
mini-batch size	32	Size of mini-batch sample for gradient step
replay buffer size	1M	Maximum number of transitions stored in the replay buffer
initial replay buffer size	50K	Number of transitions stored in the replay buffer before optimization starts
optimization frequency	4	Number of actions the agent takes between successive network optimization steps
update frequency	30000	Number of steps between consecutive target updates
$\epsilon_{\text{init}}$	1.00	Initial exploration rate of $\epsilon$ -greedy method
$\epsilon_{\text{final}}$	0.01	Final exploration rate of $\epsilon$ -greedy method
$N_{\epsilon}$	1M	Number of actions that the exploration rate of $\epsilon$ -greedy method decays from initial value to final value
$\alpha$	0.0000625	Adam optimizer learning rate
$\epsilon_{\text{ADAM}}$	0.00015	Adam optimizer parameter
evaluation frequency	250K	Number of actions between successive evaluation runs
evaluation length	125K	Number of actions per evaluation run
evaluation episode length	27K	Maximum number of action in an episode in evaluation runs
max no-op	30	Maximum number of no-op actions before the episode starts

## D EXPERIMENTAL RESULTS ON ATARI GAMES

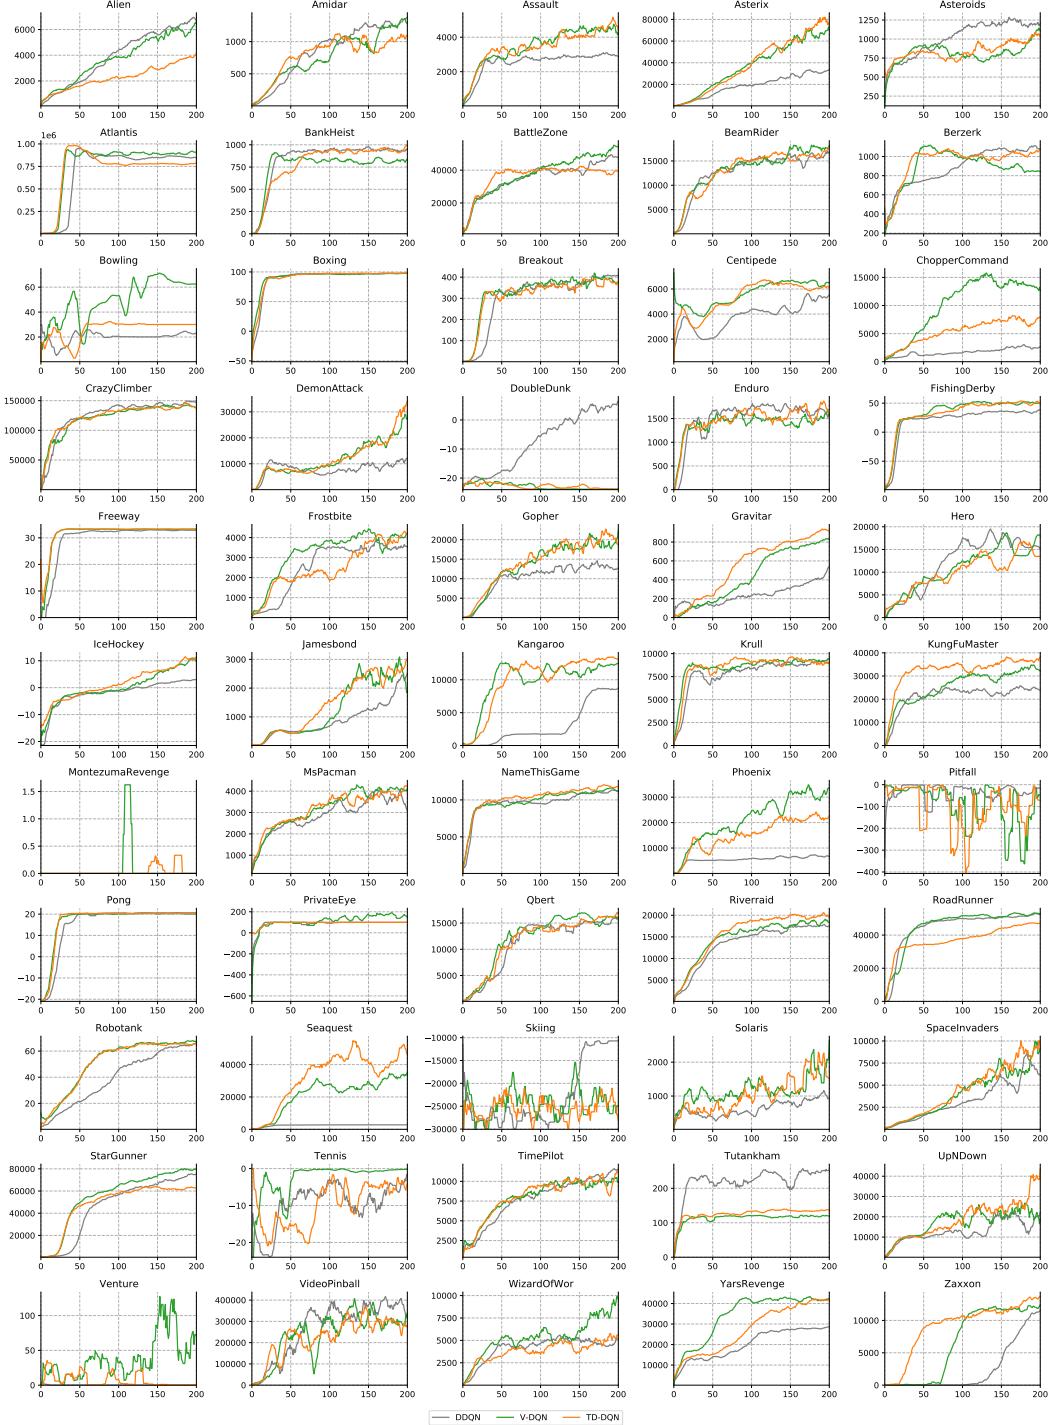


Figure 8: Training curve on Atari games from a single training run each. Episodes start with up to 30 no-op actions. Each data point is an average of episode rewards from 500K frames of evaluation runs, and smoothed over 10 data points.

Game	DDQN	V-DQN	TD-DQN
Alien	123%	<b>130%</b>	76%
Amidar	<b>99%</b>	97%	97%
Assault	764%	<b>1659%</b>	1324%
Asterix	537%	1280%	<b>1490%</b>
Asteroids	<b>2%</b>	1%	1%
Atlantis	7445%	7334%	<b>7714%</b>
BankHeist	165%	158%	<b>170%</b>
BattleZone	168%	<b>199%</b>	149%
BeamRider	143%	153%	<b>156%</b>
Berzerk	<b>51%</b>	49%	49%
Bowling	7%	<b>52%</b>	18%
Boxing	902%	903%	<b>904%</b>
Breakout	1764%	<b>1851%</b>	1644%
Centipede	57%	64%	<b>64%</b>
ChopperCommand	39%	<b>237%</b>	111%
CrazyClimber	639%	<b>658%</b>	646%
DemonAttack	563%	1415%	<b>1700%</b>
DoubleDunk	<b>1568%</b>	-145%	-90%
Enduro	267%	<b>284%</b>	281%
FishingDerby	151%	<b>164%</b>	164%
Freeway	130%	<b>131%</b>	131%
Frostbite	108%	<b>116%</b>	112%
Gopher	1071%	<b>1689%</b>	1665%
Gravitar	14%	22%	<b>25%</b>
Hero	<b>83%</b>	81%	82%
IceHockey	139%	227%	<b>253%</b>
Jamesbond	890%	<b>1315%</b>	1071%
Kangaroo	354%	529%	<b>552%</b>
Krull	909%	917%	<b>931%</b>
KungFuMaster	150%	199%	<b>214%</b>
MontezumaRevenge	-1%	<b>-0%</b>	-1%
MsPacman	32%	<b>33%</b>	29%
NameThisGame	207%	210%	<b>218%</b>
Phoenix	134%	<b>852%</b>	531%
Pitfall	5%	<b>5%</b>	5%
Pong	<b>116%</b>	115%	115%
PrivateEye	-1%	<b>-1%</b>	-1%
Qbert	146%	150%	<b>160%</b>
Riverraid	136%	152%	<b>156%</b>
RoadRunner	826%	<b>827%</b>	724%
Robotank	1006%	<b>1071%</b>	1051%
Sequest	6%	122%	<b>226%</b>
Skiing	<b>43%</b>	37%	35%
Solaris	-5%	<b>50%</b>	40%
SpaceInvaders	889%	<b>1527%</b>	1318%
StarGunner	922%	<b>984%</b>	760%
Tennis	<b>197%</b>	147%	146%
TimePilot	<b>387%</b>	365%	366%
Tutankham	<b>215%</b>	108%	109%
UpNDown	366%	516%	<b>707%</b>
Venture	-1%	<b>36%</b>	18%
VideoPinball	425%	<b>425%</b>	425%
WizardOfWor	168%	<b>309%</b>	206%
YarsRevenge	61%	<b>98%</b>	95%
Zaxxon	151%	175%	<b>183%</b>

Table 3: Normalized scores.

Game	Random	Human	DDQN	V-DQN	TD-DQN
Alien	128.3	6371.3	7807.3	<b>8236.9</b>	4895.4
Amidar	11.8	1540.4	<b>1521.6</b>	1495.0	1494.0
Assault	166.9	628.9	3697.5	<b>7829.8</b>	6284.9
Asterix	164.5	7536.0	39782.0	94509.1	<b>110013.6</b>
Asteroids	871.3	36517.3	<b>1464.3</b>	1317.6	1278.0
Atlantis	13463.0	26575.0	989675.0	975100.0	<b>1024975.0</b>
BankHeist	21.7	644.5	1050.1	1007.7	<b>1082.3</b>
BattleZone	3560.0	33030.0	53153.8	<b>62133.3</b>	47460.0
BeamRider	254.6	14961.0	21296.0	22765.7	<b>23234.0</b>
Berzerk	196.1	2237.5	<b>1228.1</b>	1205.0	1190.7
Bowling	35.2	146.5	42.7	<b>93.6</b>	54.7
Boxing	-1.5	9.6	98.6	98.7	<b>98.8</b>
Breakout	1.6	27.9	465.5	<b>488.3</b>	434.1
Centipede	1925.5	10321.9	6695.1	7271.9	<b>7282.8</b>
ChopperCommand	644.0	8930.0	3900.0	<b>20289.7</b>	9835.0
CrazyClimber	9337.0	32667.0	158346.2	<b>162828.0</b>	159952.0
DemonAttack	208.3	3442.8	18418.2	45977.5	<b>55200.6</b>
DoubleDunk	-16.0	-14.4	<b>9.1</b>	-18.3	-17.4
Enduro	-81.8	740.2	2113.5	<b>2250.2</b>	2225.0
FishingDerby	-77.1	5.1	46.6	<b>57.7</b>	57.7
Freeway	0.1	25.6	33.2	<b>33.6</b>	33.6
Frostbite	66.4	4202.8	4516.2	<b>4883.8</b>	4702.4
Gopher	250.0	2311.0	22331.4	<b>35061.4</b>	34555.7
Gravitar	245.5	3116.0	637.5	869.5	<b>976.1</b>
Hero	1580.3	25839.4	<b>21606.6</b>	21244.0	21476.5
IceHockey	-9.7	0.5	4.5	13.5	<b>16.1</b>
Jamesbond	33.5	368.5	3014.2	<b>4438.8</b>	3622.5
Kangaroo	100.0	2739.0	9450.0	14057.6	<b>14679.4</b>
Krull	1151.9	2109.1	9855.5	9930.0	<b>10065.0</b>
KungFuMaster	304.0	20786.8	31070.7	40984.4	<b>44177.4</b>
MontezumaRevenge	25.0	4182.0	0.0	<b>9.1</b>	3.3
MsPacman	197.8	15375.0	5038.5	<b>5246.7</b>	4585.7
NameThisGame	1747.8	6796.0	12181.1	12359.4	<b>12774.7</b>
Phoenix	1134.4	6686.2	8568.6	<b>48424.7</b>	30623.9
Pitfall	-348.8	5998.9	0.0	<b>0.0</b>	0.0
Pong	-18.0	15.5	<b>20.8</b>	20.6	20.6
PrivateEye	662.8	64169.1	100.0	<b>200.0</b>	100.0
Obert	183.0	12085.0	17551.4	18051.7	<b>19250.0</b>
Riverraid	588.3	14382.2	19322.9	21545.0	<b>22073.8</b>
RoadRunner	200.0	6878.0	55381.7	<b>55437.1</b>	48526.7
Robotank	2.4	8.9	67.8	<b>72.0</b>	70.7
Seaquest	215.5	40425.8	2789.8	49230.8	<b>91277.1</b>
Skiing	-15287.4	-3686.6	<b>-10314.1</b>	-10990.9	-11215.2
Solaris	2047.2	11032.6	1572.0	<b>6497.6</b>	5615.0
SpaceInvaders	182.6	1464.9	11580.8	<b>19757.7</b>	17086.2
StarGunner	697.0	9528.0	82076.7	<b>87577.8</b>	67768.6
Tennis	-21.4	-6.7	<b>7.5</b>	0.2	0.0
TimePilot	3273.0	5650.0	<b>12460.7</b>	11941.4	11975.0
Tutankham	12.7	138.3	<b>283.0</b>	147.8	150.1
UpNDown	707.2	9896.1	34346.4	48118.3	<b>65673.9</b>
Venture	18.0	1039.0	9.6	<b>382.9</b>	197.1
VideoPinball	20452.0	15641.1	584388.2	<b>632013.8</b>	631348.0
WizardOfWor	804.0	4556.0	7115.1	<b>12388.1</b>	8551.7
YarsRevenge	1476.9	47135.2	29332.9	<b>46319.6</b>	44894.4
Zaxxon	475.0	8443.0	12488.4	14409.8	<b>15028.3</b>

Table 4: Raw Scores.