Dopamine

A Research Framework for Deep Reinforcement Learning

Pablo Samuel Castro @pcastr 5 señor swesearcher





Where I'm from





Core team







Carles Gelada Subhodeep Moitra



Saurabh Kumar

Pablo Samuel Castro

Marc G. Bellemare

And many other contributors.





Dopamine is a research framework for fast prototyping of reinforcement learning algorithms.

It aims to fill the need for a small, easily grokked codebase in which users can freely experiment with wild ideas [speculative research].





"Double DQN leads to more conservative policies in fast-paced Atari games"



"Double DQN leads to more conservative policies in fast-paced Atari games"





"Double DQN leads to more conservative policies in fast-paced Atari games"





Launch stats

Launched on August 27th to a very a positive reception



In a new blog post, Google Brain team researchers @pcastr & @marcgbellemare share a new @TensorFlow-based reinforcement learning framework that aims to provide flexibility, stability, and reproducibility for new and experienced RL researchers alike.

Following

V



1:11 PM - 27 Aug 2018





Launch stats





Launch stats





Code Stats

- 12 python files (excluding tests)
 - A little over 2000 lines in total
- 98% code coverage
- Initial offering:
 - Atari environment (ALE)
 - 4 agents: DQN, C51, Rainbow, IQN
- Tensorboard integration



Code Design





Code details

- Code is well-documented
- We ran 5 independent runs for all 4 agents on all 60 games
- We provide:
 - TensorFlow checkpoints for each of these runs
 - pickle files to easily visualize in colab (or anywhere)
 - Tensorboard event files
 - JSON files with data for plotting
- Colabs for extra documentation and instruction



DQN gin-config

```
1 \text{ DONAgent.gamma} = 0.99
2 DQNAgent.update horizon = 1
3 DONAgent.min_replay_history = 20000 # agent steps
4 DONAgent.update period = 4
5 DONAgent.target update period = 8000 # agent steps
6 DQNAgent.epsilon_train = 0.01
7 DQNAgent.epsilon_eval = 0.001
& DQNAgent.epsilon_decay_period = 250000 # agent steps
9 DQNAgent.tf_device = '/qpu:0' # use '/cpu:*' for non-GPU version
10 DQNAgent.optimizer = @tf.train.RMSPropOptimizer()
11
12 tf.train.RMSPropOptimizer.learning_rate = 0.00025
13 tf.train.RMSPropOptimizer.decay = 0.95
14 tf.train.RMSPropOptimizer.momentum = 0.0
15 tf.train.RMSPropOptimizer.epsilon = 0.00001
16 tf.train.RMSPropOptimizer.centered = True
17
18 Runner.game_name = 'Pong'
19 Runner.sticky actions = True
20 Runner.num iterations = 200
21 Runner.training_steps = 250000 # agent steps
22 Runner.evaluation_steps = 125000 # agent steps
23 Runner.max_steps_per_episode = 27000 # agent steps
24
25 WrappedReplayBuffer.replay_capacity = 1000000
26 WrappedReplayBuffer.batch_size = 32
```



Comparison with published settings





Episode ends: LifeLoss vs GameOver

AtariPreprocessing.terminal_on_life_loss = True



<u>Machado et al., Revisiting the Arcade Learning Environment: Evaluation Protocols and</u> <u>Open Problems for General Agents, Journal of Artificial Intelligence Research, 2018.</u>



Train vs Eval curves





Sticky vs non-sticky actions



Runner.sticky_actions = False



New agent based on DQN

@title Create an agent based on DON, but choosing actions randomly. LOG PATH = os.path.join(BASE PATH, 'random dqn', GAME) class MyRandomDQNAgent(dqn agent.DQNAgent): def init (self, sess, num actions): """This maintains all the DQN default argument values.""" super(MyRandomDQNAgent, self). init (sess, num actions) def step(self, reward, observation): """Calls the step function of the parent class, but returns a random action. = super(MyRandomDQNAgent, self).step(reward, observation) return np.random.randint(self.num actions) def create random dqn agent(sess, environment, summary writer=None): """The Runner class will expect a function of this type to create an agent.""" return MyRandomDQNAgent(sess, num actions=environment.action space.n) # Create the runner class with this agent. We use very small numbers of steps # to terminate quickly, as this is mostly meant for demonstrating how one can # use the framework. We also explicitly terminate after 110 iterations (instead # of the standard 200) to demonstrate the plotting of partial runs. random dgn runner = run experiment.Runner(LOG PATH, create random dgn agent, game name=GAME, num iterations=200, training steps=10, evaluation steps=10, max steps per episode=100)



New agent from scratch

```
class StickyAgent(object):
  """This agent randomly selects an action and sticks to it. It will change
  actions with probability switch prob."""
  def init (self, sess, num actions, switch prob=0.1):
    self. sess = sess
    self. num actions = num actions
    self. switch prob = switch prob
    self. last action = np.random.randint(num actions)
    self.eval mode = False
  def choose action(self):
   if np.random.random() <= self. switch prob:</pre>
      self. last action = np.random.randint(self. num actions)
    return self. last action
  def bundle and checkpoint(self, unused checkpoint dir, unused iteration):
    pass
  def unbundle(self, unused checkpoint dir, unused checkpoint version,
              unused data):
    pass
  def begin episode(self, unused observation):
   return self. choose action()
  def end episode(self, unused reward):
   pass
  def step(self, reward, observation):
    return self. choose action()
def create sticky agent(sess, environment, summary writer=None):
  """The Runner class will expect a function of this type to create an agent."""
  return StickyAgent(sess, num actions=environment.action space.n,
                    switch prob=0.2)
sticky runner = run experiment.Runner(LOG PATH,
                                      create sticky agent,
                                      game name=GAME.
                                      num iterations=200,
                                      training steps=10,
                                      evaluation steps=10,
                                      max steps per episode=100)
```



Baselines plots



spaceinvaders •

spaceinvaders



Save as SVG Save as PNG View Source Open in Vega Editor



Common requests

- Non-Atari gym environments (Cartpole, Acrobot, etc.)
 - \circ very soon!
- Policy gradient methods
- Distributed training
- Visualizations



¡Gracias! github.com/google/dopamine

Pablo Samuel Castro @pcastr 5



