WORKSHEET

U12.1 Learning a behaviour

Learning, and the capstone · University · about 30 min

Name _____ Class _____ Date _____

What this lesson is about

What learning buys over a hand-written rule, what it costs, and the smallest honest example.

- What learning is, exactly
- Three settings you will meet
- When learning is worth it
- When it is not
- About `position()` in this module
- The smallest example there is

Questions

7 marks in all

1. Which three things must be present for there to be anything to learn? [1 mark]

Tick every answer that is true.

- ☐ **A** A family of candidate answers, the hypothesis class
- ☐ **B** A score: a loss to make small or a reward to make large
- ☐ **C** A search that finds a good candidate
- ☐ **D** A neural network
- ☐ **E** Thousands of labelled examples

2. The data sheet says 20 cm/s at command 100. The robot drives at command 60 for 4 s and travels 44.2 cm. What does this print? [1 mark]

```
nominal = 20.0 * 60 / 100.0 * 4.0
measured = 44.2
print(nominal, measured, round(measured / nominal, 3))
```

3. What is the name for the family of candidate answers a learner chooses between, such as all straight lines with a slope and an intercept? [1 mark]

4. A robot tries lateral offsets on short trials and keeps whichever scores best. Which setting is this? [1 mark]

- ☐ **A** Policy search: the data is trials and their scores
- ☐ **B** Supervised learning: the data is inputs with the right answers attached
- ☐ **C** System identification: the data is commands and what the robot did
- ☐ **D** None, because no model is fitted

5. Which of these is a poor use of learning? [1 mark]

- ☐ **A** Learning the inverse kinematics from data when the geometry gives them exactly
- ☐ **B** Measuring this robot's actual speed per unit of command
- ☐ **C** Refitting the drive model after changing to a new mat
- ☐ **D** Learning what 64 depth readings mean about which way is passable

6. In this module `position()` is available. What is the rule for using it? [1 mark]

- ☐ **A** The truth may inform the score, never the policy
- ☐ **B** It may be used anywhere, because it is the lab's measuring rig
- ☐ **C** It may be used by the policy but not by the score
- ☐ **D** It must not be used at all

7. The single gain from one 4 s run predicts the robot's travel better than the data sheet. What is still missing? [1 mark]

- ☐ **A** Any idea of how much to trust it: one run gives no spread
- ☐ **B** A loss function
- ☐ **C** A hypothesis class
- ☐ **D** A search

The task: the data sheet is not your robot

Print `nominal:` the travel the data sheet predicts at command 60 for 4 seconds, `measured:` what this robot actually did, and `gain:` the ratio.

```
from bugbot import *
connect()

CMD = 60
SECONDS = 4.0
V_MAX = 20.0
```

Plan your program here, then type it in and press Run.

Do it on the robot

www.bugbotlab.com/learn/u12-1-learning-a-behaviour/

The simulator checks it and tells you when it passes. Nothing to install, no account.

Challenges

1. Repeat the measurement five times and report the mean and the spread. How many runs before the mean stops moving?
2. Measure the gain at command 30 as well as at 60. Is it the same number?
3. Name one thing about this robot that a gain cannot capture, however carefully you measure it.