Selecting only the "important" samples for Experience Replay makes the batch less diverse and worsens forgetting in the model
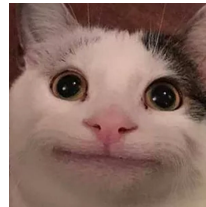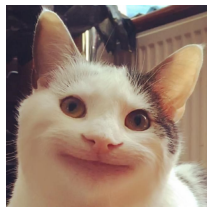
BATCH SAMPLING FOR EXPERIENCE REPLAY

Andrii Krutsylo

Experience Replay

1. Initialize the memory buffer to store past experiences.

2. Train the model on new tasks while storing some of the samples in the memory.

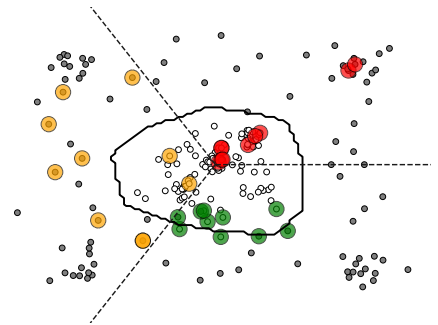3. Periodically sample from the memory buffer and train the model using the replay batch.

Replay Batch

- Maximally interfered samples after model update (MIR, MGIR, GMIR)

- Best able to preserve latent decision boundaries (ASER)

- Most conflicting with the current input data (AdaER)

Problem

Similar samples (in the feature space of the model) would have similar scores.

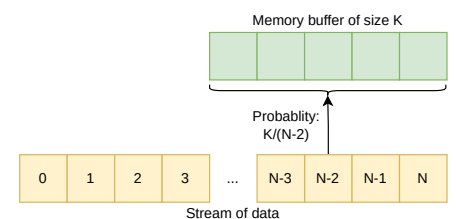Sample A   Sample B

Solutions

- Prioritize diversity (GRASP)

- Include diversity in the importance score of the sample

- Evaluate the batch (~263 trillion possible batches with a memory buffer size of 1000 for a batch size of 10)
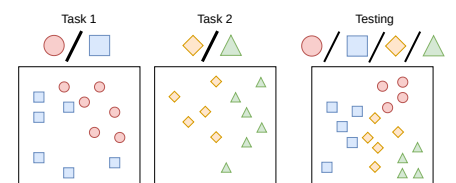
Batch properties

- **Representative**
- **Diverse**
- **Important**

Reservoir sampling to the memory buffer

Memory buffer of size K

Probablity: K/(N-2)

| 0 | 1 | 2 | 3 | ... | N-3 | N-2 | N-1 | N |

Stream of data

Continual learning

Task 1   Task 2   Testing