CS103 Winter 2019 – Assignment 1

Due: January 31, 2019

Notation. Let G be a group acting on a set X, *i.e.*, we have a group action $\cdot : G \times X \to X$ satisfying:

1. $e \cdot x = x$ for all $x \in X$, where $e \in G$ is the group identity;
2. $(gh) \cdot x = g \cdot (h \cdot x)$ for all $g, h \in G$ and $x \in X$.

Given $x \in X$, we define the orbit $\mathrm{Orb}(x)$ of x as

$$\mathrm{Orb}(x) = \{g \cdot x : g \in G\},$$

that is, $\mathrm{Orb}(x)$ is the set of elements where x can be mapped through the action of the group. The set of all orbits is denoted by $X/G := \{\mathrm{Orb}(x) : x \in X\}$. The stabilizer $\mathrm{Stab}(x)$ of x is the set

$$\mathrm{Stab}_G(x) = \{g \in G : g \cdot x = x\},$$

that is, $\mathrm{Stab}(x)$ is the set of group elements that leave x fixed.

Exercise 1 (5 points). Consider the relation $x \sim_G y$ if $y \in \mathrm{Orb}(x)$, *i.e.*, if y is in the orbit of x. Prove that $\sim_G$ is an equivalence relation on X, that is, prove that:

1. $x \sim_G x$ for all $x \in X$ (reflexive);
2. If $x \sim_G y$, then $y \sim_G x$ (symmetric);
3. If $x \sim_G y$ and $y \sim_G z$, then $x \sim_G z$ (transitive).

Exercise 2 (5 points). Let $G = \{e, r, r^2, r^3\}$ be the group of rotations in the plane that are multiple of 90 degrees. Let X be the set of all unit squares whose sides are colored in red and/or blue. Notice that G acts naturally on X by rotating the colored squares by 0, 90, 180 or 270 degrees. How many distinct orbits does this action have (*i.e.*, what is the cardinality of X/G)? This is the same as counting the ways we can color a square ignoring colorations that differ only by a rotation.

Exercise 3 (10 points). **(a)** Let $X = \{f : \mathbb{R}^2 \to [0, 1]\}$ be the set of all black and white images, and let $G = \mathbb{R}^2$ be the group of translations of the plane. The translation group G acts on X by translating the images, that is, given a translation $g := (u, v) \in G$ and an image $f(x, y)$, we can define the translated image $(g \cdot f)(x, y) := f(x + u, y + v)$. For each of the following images f, describe $\mathrm{Stab}_G(f)$, that is, the set of translations that leave the image f unaltered:

1. $f(x, y) = 0$ for all $(x, y) \in \mathbb{R}^2$, that is f is the the uniformly black image;
2. $f(x, y) = 1$ if $x < 0$, and 0 otherwise, that is, $f(x, y)$ is the image that is white in the semi-plane $x < 0$;
3. $f(x, y) = 1$ if $x \cdot y > 0$ and 0 otherwise.

(b) Recall that an equivariant reference frame detector is a function $R : X \to G$ such that $R(g \cdot f) = g \cdot R(f)$ for any $f \in X$ and $g \in G$. Given the above results, is it always possible to construct an equivariant frame detector? Explain.

Exercise 4 (10 points). In a machine learning problem, we are given as input a K-dimensional vectors $x = (x_0, \dots, x_{K-1}) \in \mathbb{R}^K$, and we want to regress the associated binary label $y \in \{0, 1\}$.

We know that the label y is invariant to cyclical permutation of the sequence. That is, let $G = \{0, 1, \dots, K-1\}$ be the group of integers with addition modulo K. Then, $t \in G$ acts on the vectors in $\mathbb{R}^K$ by cyclically shifting the elements of the sequence

$$t \cdot x = (x_t, x_{t+1}, \dots, x_{K-1}, x_0, \dots, x_{t-1}).$$

We know that x and $t \cdot x$ always have the same label y.

(a) Describe all the linear representations $\phi : \mathbb{R}^K \to \mathbb{R}^K$ that are equivariant to the action of G, assuming that G acts on both domain and codomain by cyclically shifting the elements of the vector.

(b) Describe all the linear representations $\phi : \mathbb{R}^K \to \mathbb{R}$ that are invariant to the action of G.

Coding assignment (optional, 15 points). On the course website (under the Assignments section) you will find a dataset satisfying the conditions of Exercise 4.

(a) Using this dataset, train a deep fully connected network to regress the binary label y given the input vector $x \in \mathbb{R}^K$. Try changing the number of samples used to train and plot the testing error of this network as a function of the number of training samples.

(b) Now design a network for this problem using only equivariant and invariant representations [Hint: try to reuse the convolutional layers already implemented in all main deep learning libraries]. Again, plot the testing error of this network as a function of the number of training samples.

(c) Which one of the two networks performs better? Which one is more data efficient (*i.e.,* obtains lower errors with less data)?