

OCR learned by Structured Output SVM

UCU Summer school

July 13, 2017

1 Introduction

We will consider a problem of recognizing a given name from handwritten characters. The inputs are binary images containing a line of text. All characters have a fixed width so that we do not have to solve the segmentation problem. A sample of the input images is shown in Figure 1. In Section 2 we will describe two linear classifiers for sequence recognition each of them using a different model of the sequences. Your task will be to learn parameters of these classifiers by Bundle Risk Minimization method and to evaluate their performance using data described in Section [?].

2 Structured output classifiers

The input is a binary image $\mathcal{I} = \{\mathcal{I}_1, \dots, \mathcal{I}_L\} \in \{0, 1\}^{16 \times 8 \cdot L}$ displaying a sequence of L handwritten characters. Each sub-image $\mathcal{I}_i \in \{0, 1\}^{16 \times 8}$, $i \in \{1, \dots, L\}$ is of fixed size hence the segmentation of the image L into characters is known. We represent the input image \mathcal{I} by a matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_L) \in \mathbb{R}^{d \times L}$, where d is the number of features. The i -th column $\mathbf{x}_i \in \mathbb{R}^d$ is a feature description of the i -th sub-image \mathcal{I}_i .

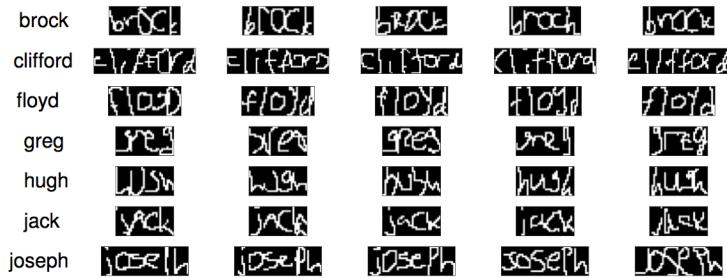


Figure 1: A sample of the input images for a selected sub-set of names

The task is to recognize a sequence of characters $y = (y_1, \dots, y_L) \in \mathcal{Y}^L$, $\mathcal{Y} = \{a, b, \dots, z\}$, depicted on the image \mathcal{I} using the features \mathbf{X} . Below we describe different classification strategies and a way how to evaluate their performance.

2.1 Independent linear multi-class classifier

Given a feature representation $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_L)$ of an input image \mathcal{I} , the predicted for each sub-image independently by a multi-class linear classifier $h: \mathcal{X} \rightarrow \mathcal{Y}$:

$$y^i = \operatorname{argmax}_{y \in \mathcal{Y}} \langle \mathbf{w}_y, \mathbf{x}^i \rangle, \forall i \in \{1, \dots, L\},$$

where the parameters are the character templates $\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{Y}|}) \in \mathbb{R}^{d \times |\mathcal{Y}|}$. You may consider adding biasses into \mathbf{w} . Note, the length L of the unknown sequence can be deduced from the width of the input image \mathcal{I} or the feature matrix \mathbf{X} , respectively.

2.2 Linear structured output classifier

A linear classifier over sequences $h: \mathcal{I}^L \rightarrow \mathcal{Y}^L$

$$(\hat{y}_1, \dots, \hat{y}_L) = \operatorname{argmax}_{(y_1, \dots, y_L) \in \mathcal{Y}^L} \sum_{i=1}^L q(x_i, y_i) + \sum_{i=1}^{L-1} g(y_i, y_{i+1}) \quad (1)$$

where $q(x^i, y^i) = \langle \mathbf{w}_{y^i}^q, x^i \rangle$ and $g(y^i, y^{i+1}) = \mathbf{w}_{y^i, y^{i+1}}^g$, where $\mathbf{w}^q = (\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{Y}|}) \in \mathbb{R}^{d \times |\mathcal{Y}|}$ and $\mathbf{w}^g \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$. In this class the pairwise function for classifier can be written in the matrix form $q(\mathbf{w}, \mathbf{x}) := \mathbf{W}^q \mathbf{x}$, where \mathbf{W}^q is vector $(\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{Y}|})$ stacked row-wise. Pairwise function is simply $g := \mathbf{W}^g$,

For begin short we denoted (1) as $\langle \mathbf{w}, \phi(x^i, y^i) \rangle$ in the lectures.

3 The task

In this task you have to implement surrogate losses

- for multi-class linear classifier:

$$\psi(x^i, y^i, \mathbf{w}) = \max_{y \in \mathcal{Y}} \{\llbracket y^i \neq y \rrbracket + \langle \mathbf{w}_y, x^i \rangle - \langle \mathbf{w}_{y^i}, x^i \rangle\}$$

- for structured-output classifier on the sequences:

$$\psi(x^i, y^i, \mathbf{w}) = \max_{y \in \mathcal{Y}} \sum_{k=1}^L \llbracket y_k^i \neq y_k \rrbracket + \langle \mathbf{w}, \phi(x^i, y) \rangle - \langle \mathbf{w}, \phi(x^i, y^i) \rangle,$$

for learning linear multi-class and structured output classifiers respectively using the Bundle Method for Risk Minimization (BMRM) solver. It accesses

manually specified surrogate loss as first order oracle. That means, to use BMRM you have to provide a python class with overloaded method `__call__` that takes numpy vector $\mathbf{w} \in \mathbb{R}^{n \times 1}$ and returns pair: (surrogate risk, subgradient) for given input \mathbf{w} . We provide you a python implementation of BMRM.

Useful hint for both tasks.

The function $F(\mathbf{w}) = \max_{i=1,\dots,n} (A\mathbf{w})$, where $A \in \mathbb{R}^{n \times n}$, has subgradient a_i^T , where a_i^T is the i -th row of the matrix on which maximum is achieved.