Underwater Mine Classification with Imperfect Labels

David P. Williams

williams@nurc.nato.int

NATO Undersea Research Centre • Viale San Bartolomeo 400 • 19126 La Spezia (SP) • Italy

Abstract

A new algorithm for performing classification with imperfectly labeled data is presented. The proposed approach is motivated by the insight that the average prediction of a group of sufficiently informed people is often more accurate than the prediction of any one supposed expert. This idea that the "wisdom of crowds" can outperform a single expert is implemented by drawing sets of labels as samples from a Bernoulli distribution with a specified labeling error rate. Additionally, ideas from multiple imputation are exploited to provide a principled way for determining an appropriate number of label sampling rounds to consider. The approach is demonstrated in the context of an underwater mine classification application on real synthetic aperture sonar data collected at sea, with promising results.

Objective: Classify underwater objects as targets (*i.e.*, mines) or clutter (*e.g.*, rocks).

Training data collection:
- Deploy known targets at a site.
- Collect sonar imagery using an autonomous underwater vehicle (AUV).
- Manually label objects based on target-deployment knowledge.

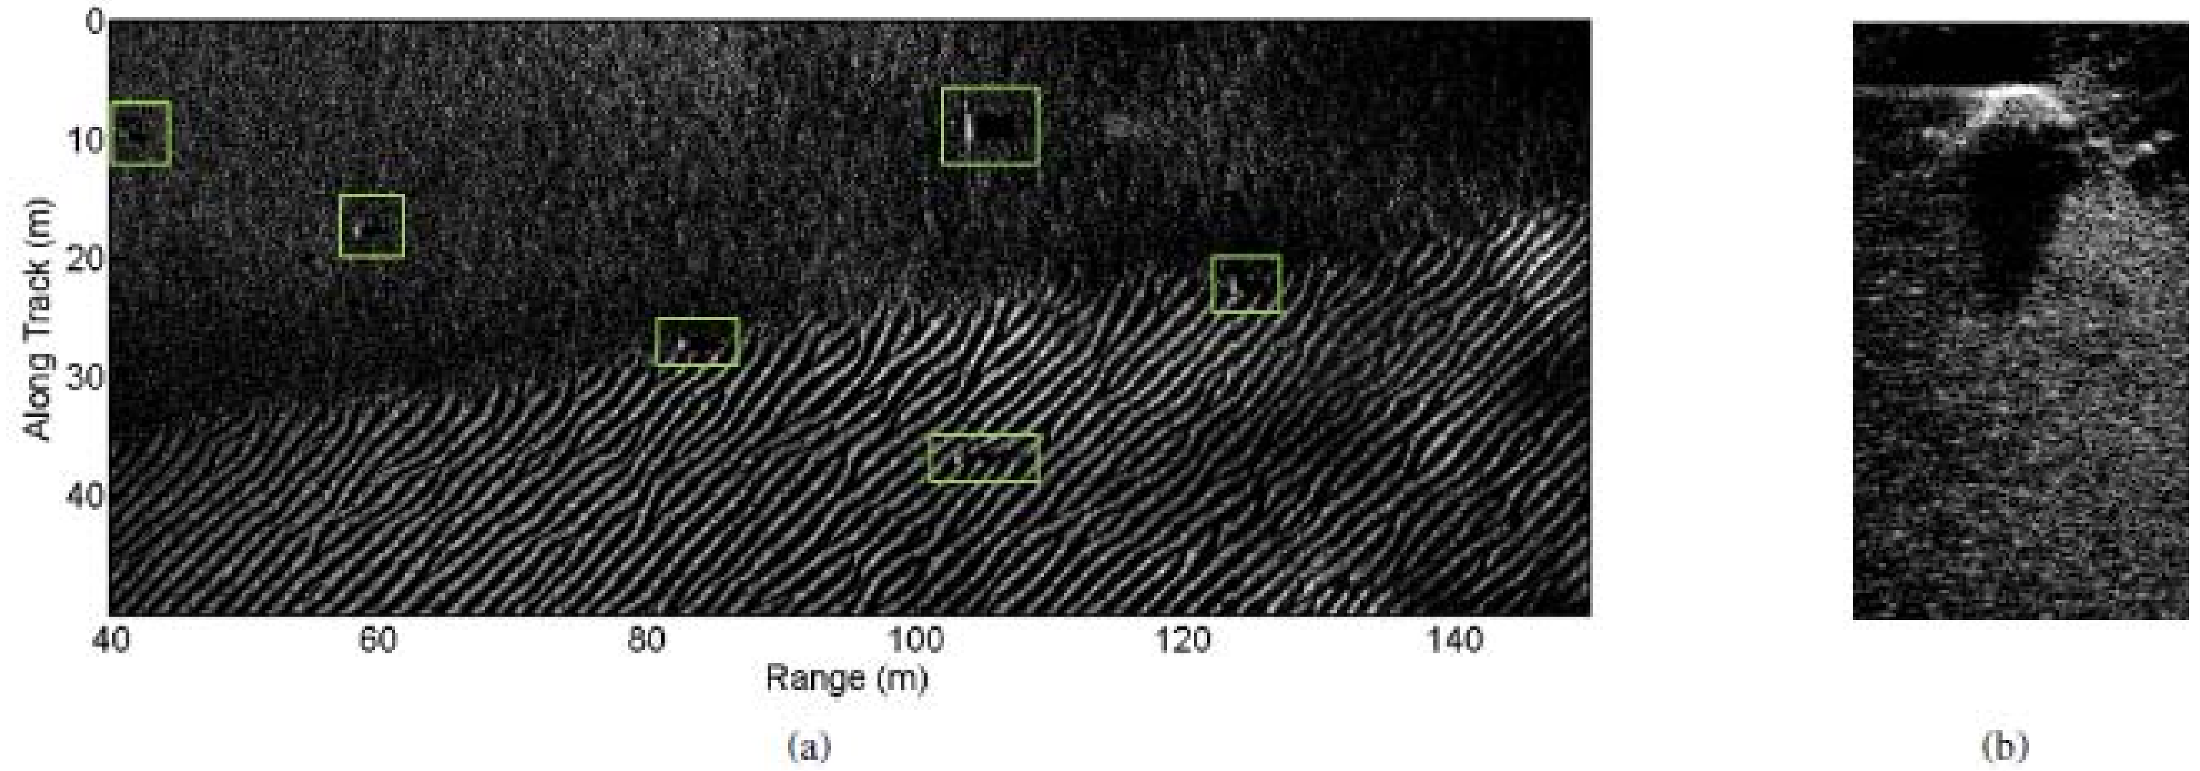

Figure 1. (a) A typical SAS image with mines indicated in green boxes. (b) A SAS image chip of a typical alarm.

Issues with labeling:
- Other unknown objects may already be present.
- AUV navigation errors complicate process.
- ➢ *Manually assigned labels may be incorrect.*

Proposed solution:
- Assign a labeling error to each training data point:

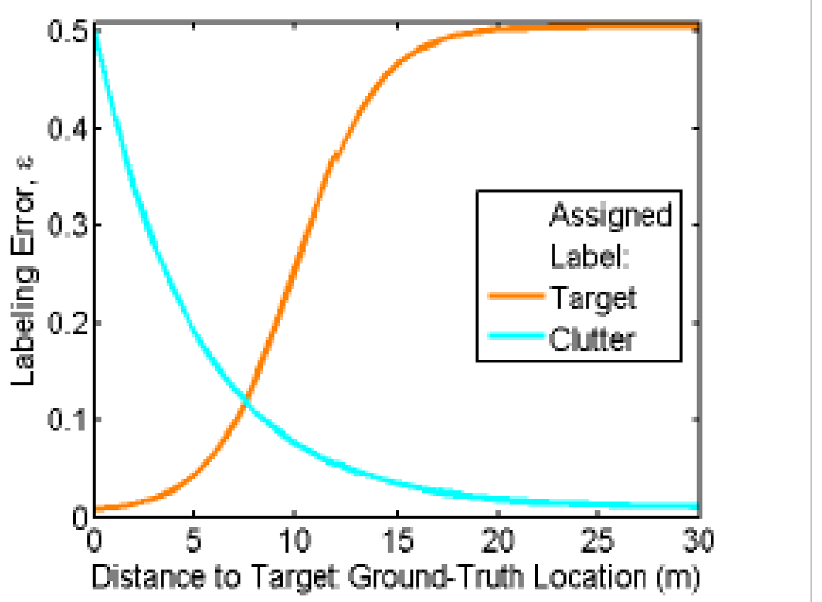

Figure 2. Definition for the labeling error.

- Treat labels as samples from a Bernoulli distribution parameterized by labeling error:

$$y_i' \sim \mathcal{B}(1-\epsilon_i) = \begin{cases} y_i, & \text{with probability } 1-\epsilon_i; \\ 1-y_i, & \text{with probability } \epsilon_i. \end{cases}$$

- Draw multiple sets of labels to exploit the "Wisdom of Crowds": The average prediction of a group of sufficiently informed people is often more accurate than one supposed expert.

$$p(y_* = 1 | \boldsymbol{x}_*, \{\boldsymbol{w}_{(m)}\}_{m=1}^{M}) = \frac{1}{M} \sum_{m=1}^{M} p(y_* = 1 | \boldsymbol{x}_*, \boldsymbol{w}_{(m)})$$

Experimental results:

Data collected at sea off coast of Latvia at two sites (Riga and Liepaja).

Detection: highlight-shadow patterns characteristic of mines.
- About 300 data points per mission.

Features:
- 5 for size and shape.
- 6 for contextual information.

Classifier: Logistic regression.

Table 1. AUC (mean ± one standard deviation from the six trials) for each of the two test sites.

Method	Riga	Liepāja
$\epsilon = 0$	0.9389 ± 0.0237	0.9775 ± 0.0165
$\epsilon \neq 0$	0.9653 ± 0.0108	0.9823 ± 0.0263
Proposed	**0.9726 ± 0.0052**	**0.9932 ± 0.0033**

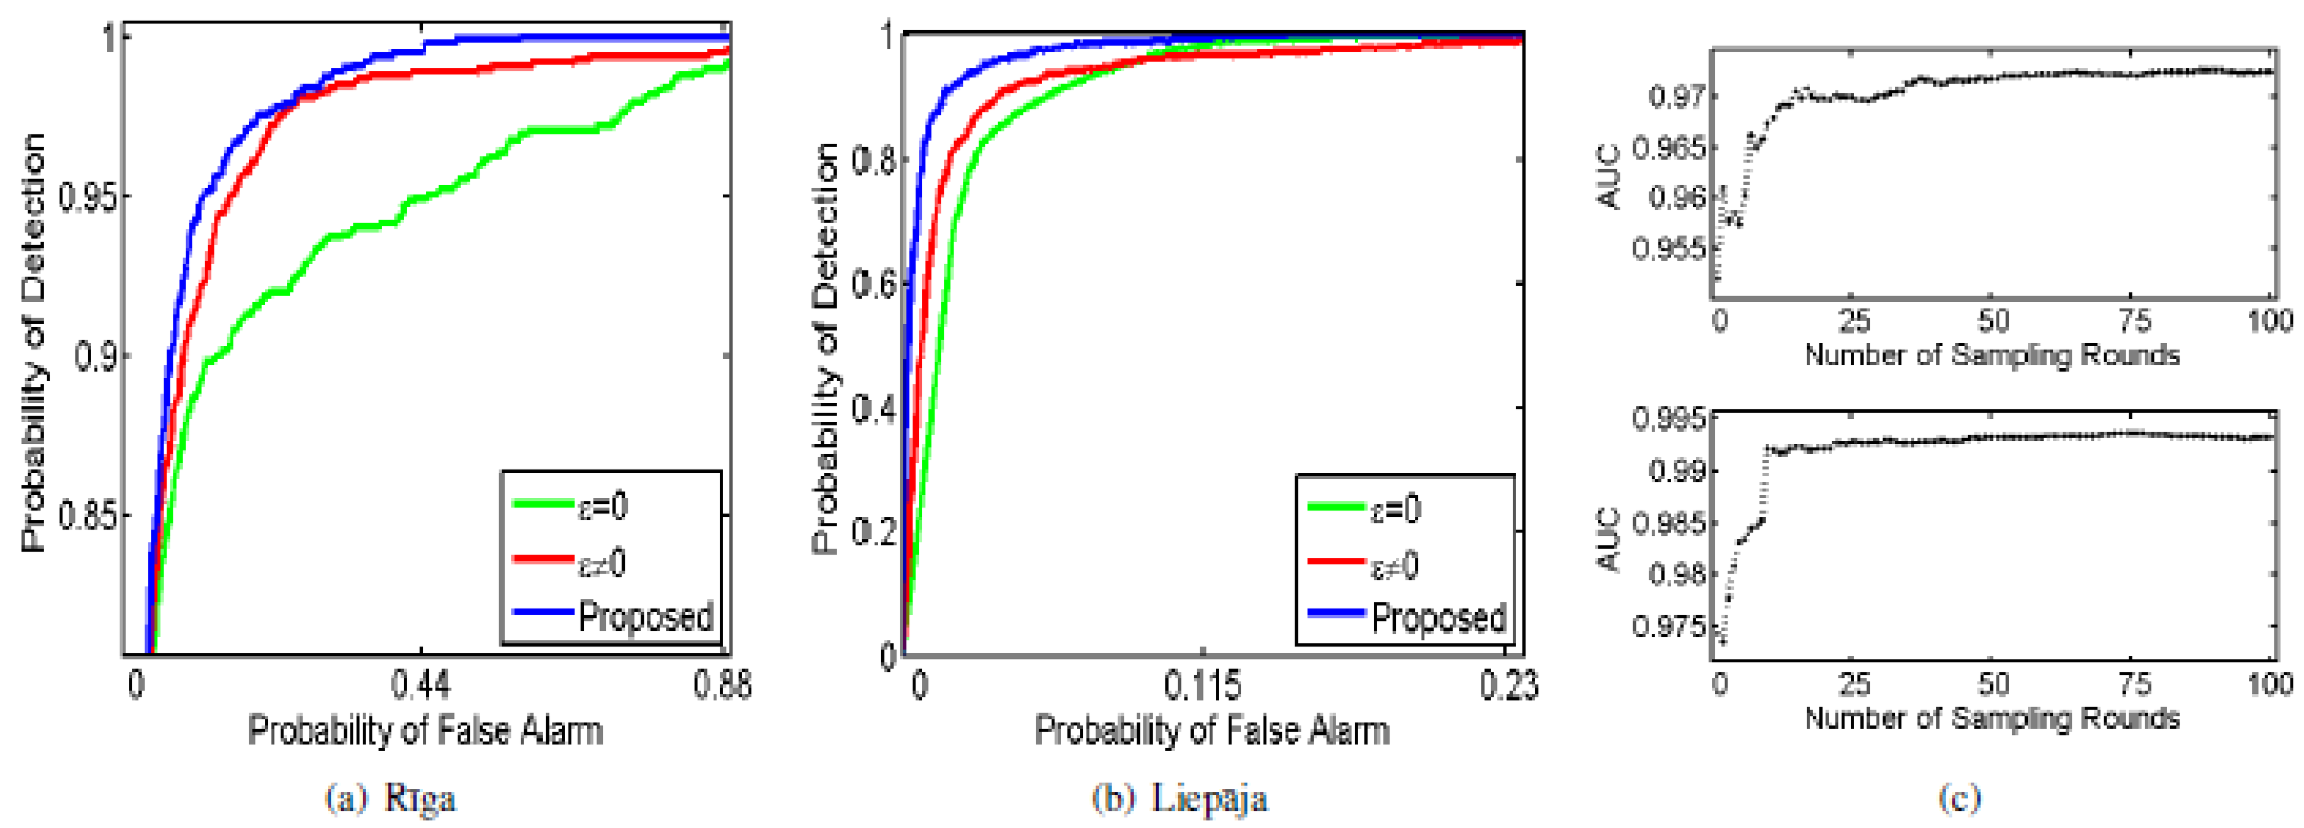

Figure 3. (a-b) Average ROC curves and (c) evolution of the AUC for the proposed method (top: Rīga; bottom: Liepāja).