Helping remote sensing analysts with active learning

Devis Tuia[1]

[1] Laboratoire des Systèmes d'Information Géographique

Lausanne Institute of Technology (EPFL), Switzerland

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

VNIVERSITAT DE VALÈNCIA

LSSR 2013, Lausanne, Switzerland

The problem: creating learning sets involving human operators

When creating a supervised method, we usually need

- A classifier
- Labeled samples $\{\mathbf{x}_i, y_i\}_{i=1}^n$
  - Labeled samples are costly (time and money)
  - Need to organize field campaigns
  - Involve scrutinizing thousands of unlabeled samples
  - People usually sample randomly in (feature) space or with stratification

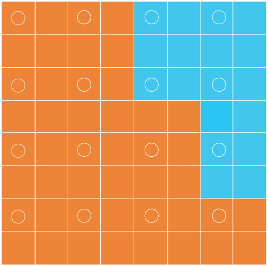

The problem: creating learning sets involving human operators

This is a good strategy, when

- We do not have any labeled examples
- We do not know anything about the phenomenon

  $\rightarrow$ We can rely only on the distribution in the input space

But usually we have some samples

- We can build a suboptimal model
- We can use it to direct the sample scheme

  $\rightarrow$ We use the output space of the suboptimal model

Did you say active learning?

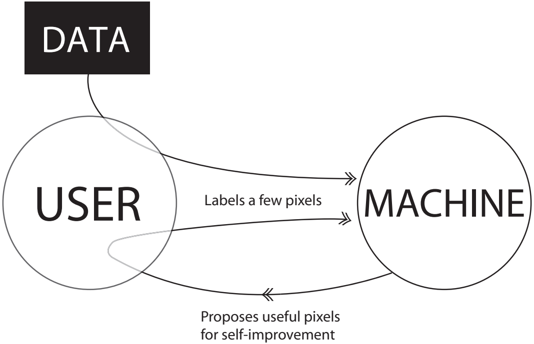

[Tuia, Volpi, Copa, Kanevski, Muñoz-Marí, A survey of active learning algorithms for supervised remote sensing image classification, *IEEE J. Sel. Topics Signal Proc.*, 2011, 5, 606-617]

Be active!

In the context of spatial sampling, this would imply

- Sampling around the class limits,
- Interact with the user to understand the data structure,
- thus minimize the number of pixels to obtain an optimal result

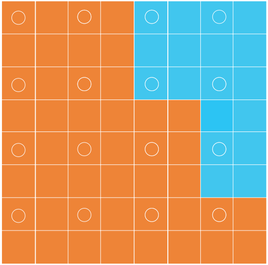

Passive

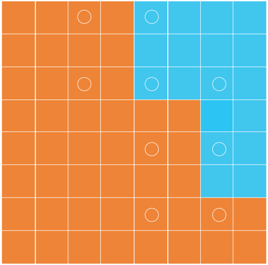

Active

Many effective methods for remote sensing (among others)

- Multiple classifiers [Tuia et al. (TGRS 2009), Di and Crawford (JSTSP, 2011)]
- SVM [Mitra et al. (PRL, 2004), Tuia et al (TGRS, 2009), Patra and Bruzzone (TGRS, 2011), Pasolli et al. (GRSL, 2011) ...]
- Criteria diversity [Demir et al. (TGRS 2011), Volpi et al. (TGRS, 2012), Di and Crawford (TGRS, 2012), Patra and Bruzzone (PRL, 2012)]
- Domain adaptation [Tuia et al. (RSE, 2011), Persello and Bruzzone (TGRS, 2013), Matasci et al. (JSTARS, 2013)]
- Change detection [Demir et al. (TGRS, 2012)]

They optimize the number of labels in controlled scenarios
... however ...

Limitations

- Candidate pixels are always pre-selected on ground truths!

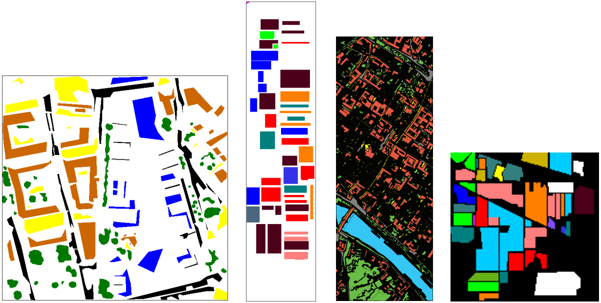

- Users are artificial, thus infallible
  (they always know the correct answer for a query)

When going to the user's side...

- Users are real operators

Figure 48B. The sequence of steps in producing National Wetlands Inventory maps. *(Photograph by Donald W. Woodard, U.S. Fish and Wildlife Service.)*

Source: water.usgs.gov

Source: english-heritage.org.uk

- Sometimes they cannot provide a good answer

>> Question: are AL methods effective for real users?

When going to the user's side...

These are typical pixels retrieved by an AL algorithm, when processing an entire VHR image:

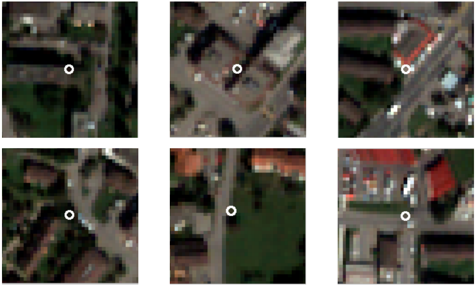

>> Contradiction between AL goal (find uncertain pixels) and users's (label them)

These pixels are difficult to label

Consequences: this can lead to

- A greater number of pixels to be screened (even vs random sampling)
- Erroneous labels
- Frustration, fatigue

Solutions:

- Constrain AL in predefined (no-border/shadow/...) areas
- Learning the skills of the user and maximize them! >> Today!

Active learning with user's confidence (AL-UC)

What we want is to avoid bad states [Judah et al., ICML 2011].

A bad state is a query that <u>for the user</u> is

- Ambiguous
- Unfamiliar

$>>$ Query that the user does not want to re-encouter.

Intuition: let's learn to recognize bad states!

Active learning with user's confidence (AL-UC)

In parallel to the active learning process, we train a confidence classifier.

- Binary: $y^C = \{$-1: can label, 1: can't label$\}$
- Probabilistic outputs $p(\hat{y}^C = 1|x^C)$
- Bad state when exceeding a probability threshold θ.

Active learning with user's confidence (AL-UC)

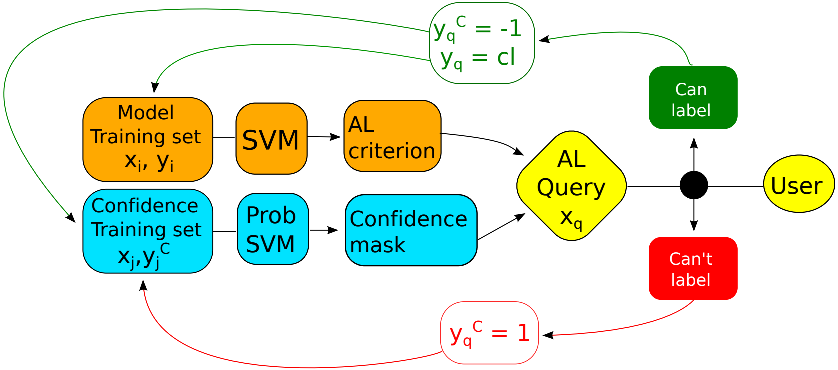

Case study: VHR image of Zurich (Switzerland)

- QuickBird, 2.4m
- 9 classes (see →)
- User starts identifying 5 pixels/class ($y^C = -1$)
- 20 pixels added per iteration
- SVM classifier.
- 5 runs per strategy
  - Random sampling (RS)
  - Classical AL
  - AL-UC
- AL using the ALtoolbox `http://code.google.com/p/altoolbox/`

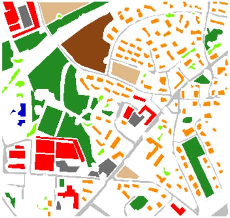

Interactive testing environment

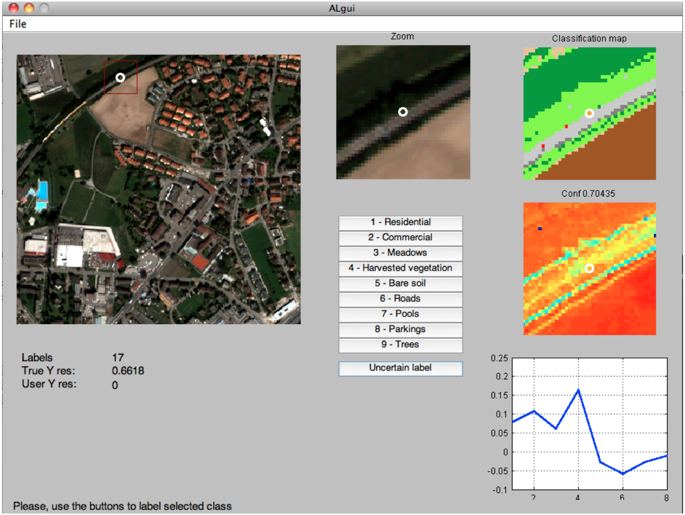

Iteration # 2

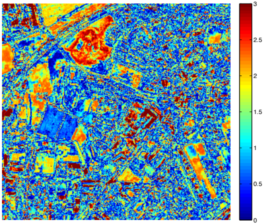

AL criterion

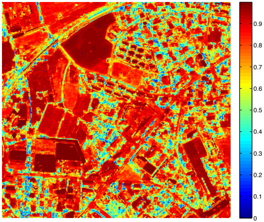

Confidence map

Confidence mask at $\theta = 0.6$

Iteration # 5

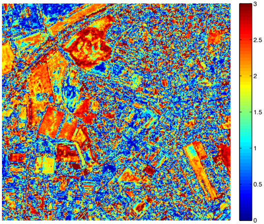

AL criterion

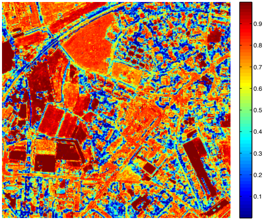

Confidence map

Confidence mask
at $\theta = 0.6$

Numerical comparison: each 20 valid labels

At a first glance, no numerical difference!

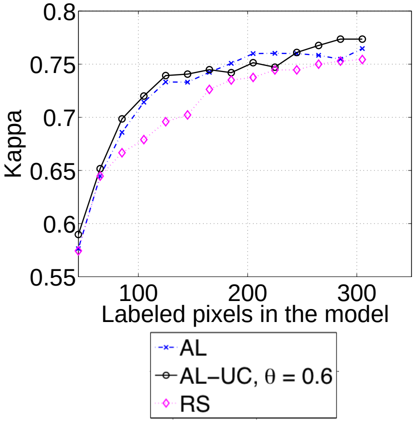

But this is what would be achieved by an infallible user.

Numerical comparison

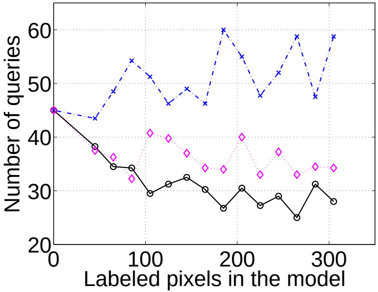

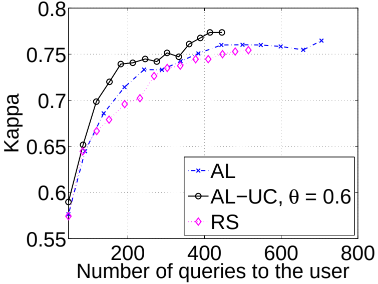

↑ Average queries per iteration

↑ Real effort (# of clicks)

Is it only because I know the task?

- Five users performed the task
- 3 users are familiar with photointerpretation
- 2 users are not
- Single run per model
- Same image and setting.

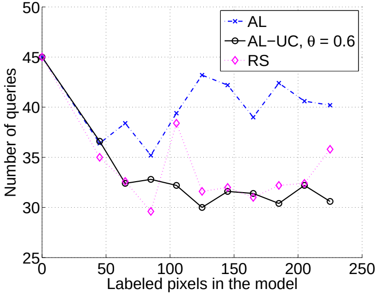

Summary

Today's talk:

- Users are the target for active learning
- Classical AL often return difficult pixels
- We integrated user's skills in the AL process
- User's effort is strongly reduced, AL keeps its advantages
- Want to know more?

  [Tuia, D. and Muñoz-Marí, J., 'Learning user's confidence with active learning', TGRS, 2013]

New opportunities:

- Discovery of new classes
- Crowdsourcing

Thank you!

http://wiki.epfl.ch/eo-adapt/
http://devis.tuia.googlepages.com
devis.tuia@epfl.ch

 

Funded by:

Swiss National Science Foundation