Research plan — Etienne Grossmann

My research objectives for the coming years address some of the issues raised by the growing amount of digital data and the need to access information in it. This is the situation when data mining through text documents, images, video sequences, computer logs ... Many data mining and machine learning tasks depend on solving a classification problem, e.g. filtering spam, detecting fraudulent transactions, recognizing people from their appearance in images or identifying road signs.

These tasks typically rely on classifiers built from domain-specific heuristics, combined into neural networks, decision trees, boosted ensembles or ad-hoc computational schemes. Present approaches usually address the accuracy of the algorithm, while the complexity is the “loose end” of the problem. However, in practice, a classifier must not only be accurate, but also give results in a timely fashion. This is particularly important not only in data mining, but also in real-time applications such as on-line commerce monitoring, autonomous robot navigation, video surveillance or the video processing tasks on which I have recently worked - face detection and background subtraction. My research plan thus centers on the efficiency o machine learning computations.

I have developed two types of fast classifiers: cascades of classifiers [1] and decision trees [2]. My research on decision trees is inspired by the successful boosting algorithms. Decision trees are known to overfit and my aim is precisely to get a generalization ability at least as good as state-of-the-art ensembles, while achieving better speed. My published work reports on early results that I am keen to push further, together with graduate students and my future colleagues.

In addition to continuing research along these lines, I would like to address the issue of training time. My objective here is to use methods to subsample overabundant training datasets into pieces of manageable size, while attempting to preserve asymptotic properties.

Finally, I should point that I am quick at adopting new research interests: in 2005, I have also formulated and solved, together with my colleagues, new camera calibration problems [4, 3]: we show how to calibrate “shapeless” cameras (i.e. w/ strange lenses or mirrors, insect eyes ...) without a calibration object. In another domain, we also applied machine learning to a long-standing computer vision problem, background subtraction [5, 6]. This problem, like many others in computer vision, still offers fascinating research opportunities, which I am also keen on pursuing.

References

[1] E. Grossmann. Automatic design of cascaded classifiers. In *intl. IAPR workshop on statistical pattern recognition, ICPR*, 2004.

[2] E. Grossmann. AdaTree : boosting a weak classifier into a decision tree. In *Workshop on Learning in Computer Vision and Pattern Recognition, CVPR*, 2004.

[3] D. Nistér, H. Stewenius, and E. Grossmann. Non-parametric self-calibration. In *proc. ICCV*, 2005.

[4] E. Grossmann, E-J Lee, P. Hislop, D. Nistér, and H. Stewénius. Are two rotational flows sufficient to calibrate a smooth non-parametric sensor? In *proc. IEEE CVPR*, 2006.

[5] E. Grossmann, A. Kale, C. Jaynes, and S-C. Cheung. Offline generation of high-quality background subtraction data. In *proc. BMVC, best poster award*, 2005.

[6] E. Grossmann, A. Kale, and C. Jaynes. Towards interactive generation of "ground-truth" in background subtraction from partially labeled examples. In *proc. ICCV VS-PETS workshop*, 2005.