Where is my Advisor?

Anonymous CVPR submission
Address Line2
Paper ID 9

Abstract

*We present a **novel** method to a **unique** problem solved in a **new** way: we use vision technology to aid in the process of vision research. Specifically, we develop an **innovative** algorithm to determine the presence of our advisor. By creating a system that will identify when our advisor's office door is open, we can attempt to corner him only when necessary, remaining at our desk for longer periods of time to do **actual work** [8]. We present a synopsis of the vision literature to date along with a working algorithm tested in two environments.*

***Keywords**: David Forsyth*

1. Introduction

We study the problem of computer vision because a careful observer will note that the problem of vision is obviously useful. [Some more vacuous statements]

As careful students in the field of computer vision, we adopt the following techniques originally conceived by other authors [1, 2, 3, 4, 5]:

1. **Shop for Datasets**. In many cases, the general problem is too hard – maybe even unsolvable. Current methods don't work well, and, even worse, your particular algorthm does not work well with the traditional datasets. Solution: find a better dataset. Surf the web. Or, better yet, make your own. Comparisons become difficult if nobody else has seen your data. (This is an important step of the grad student flow chart and resumé building, outlined in Figure 3)

2. **Incorporate Non-Silicon Intelligence**. We affectionately refer to this line of work as "Grad Student in the Loop". Most algorithms work better if a grad student can analyze the errors on the test set and make modifications of the original code. We outline this technique further in Figure 2.

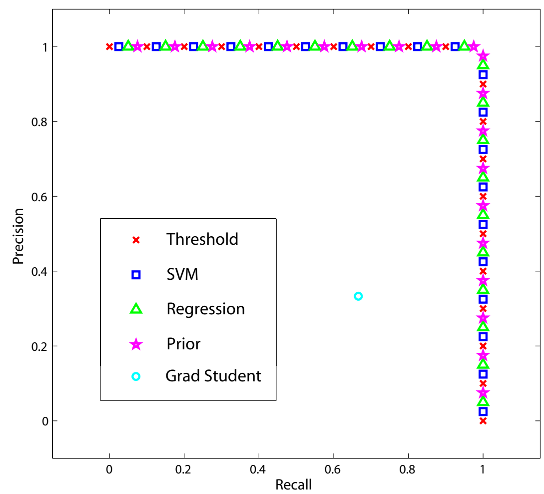

Figure 7. *The precision recall curve for the classifier at the alternate office. Since the door is never open in the dataset, all methods have comparable results.*

3. **Ignore Theory**. Elegant solutions to research problems provide long-term employment for theoriticians. However, we focus on immediate tasks, and thus ignore theory. Often, a threshold or a hard-coded constant can eliminate many hours, days, months or even years of mathematical contemplation.

4. **Abuse Theory**. Often, reviewers dislike hard coded constants. To please them, we seek to make an overly complicated theoretical model, then agressively simplify to justify our kludge. [REF Nicolas' figure]

2. Results

[Insert Results Here]

Figure 1. *Advisor apparently involved in some sort of ritual social interaction with another subject.*

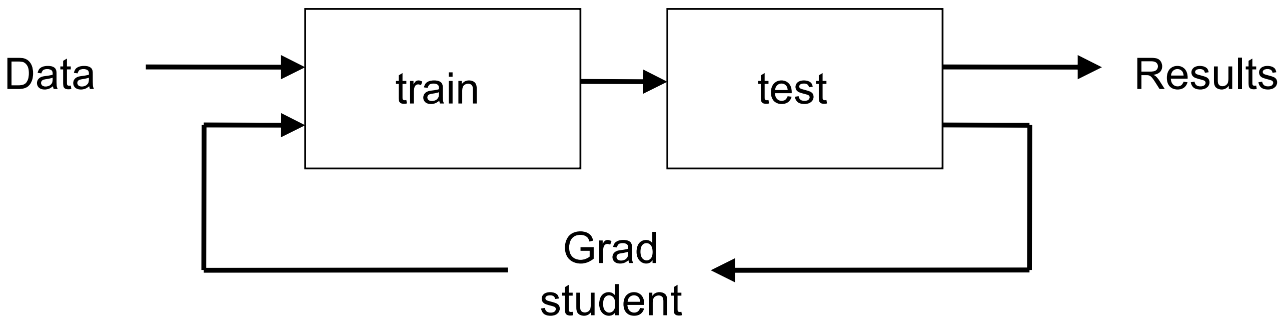

Figure 2. *"Grad student in the loop"*

References

[1] IEEE Computer Society International Conference on Computer Vision and Pattern Recognition, 1983 – 2005. 1

[2] European Conference on Computer Vision, 1990 – 2004. 1

[3] IEEE Conference on Computer Vision, 1987 – 2005. 1

[4] IEEE Transactions on Pattern Analysis and Machine Intelligence, 1979 – 2005. 1

[5] International Journal of Computer Vision, 1987 – 2005. 1

[6] Paul S. Heckbert. Ray tracing Jell-O brand gelatin. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 73–74, July 1987.

[7] Daniel Aguayo Jeremy Stribling and Maxwell Krohn. Rooter: A methodology for the typical unification of access points and redundancy. Also known as SciGen. http://pdos.csail.mit.edu/scigen/ Accepted then rejected, WMSCI, 2005.

[8] Ryan White, Okan Arikan, Himanshu Arora, Alexander Berg, Tamara Berg, Jim Chan, Jaety Edwards, Ali Farhadi, Leslie Ikemoto, Nazli Ikizler, Adam Kirk, Anthony Lobay, Nicolas Loeff, Michael Maire, Hossein Mobahi, Deva Ramanan, Alexander Sorokin, Jai Vasanth, Linus Wong, and Celal Zifti. Where is my advisor? Submitted, Computer Vision and Pattern Recognition, 2006. 1

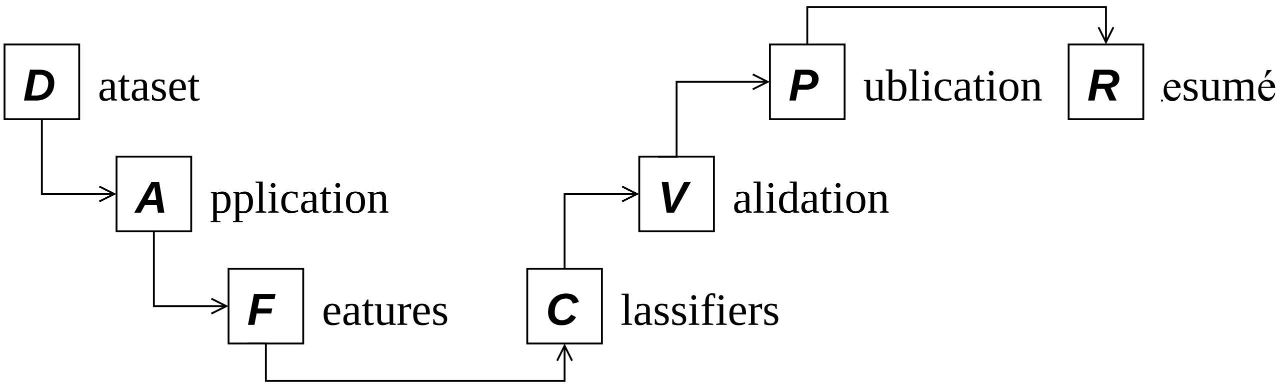

Figure 3. *Ah, the life of a grad student. This life shows the typical progression of a research project. First, we find a dataset. Then, depending on the specific characteristics of the dataset, we pick an application, select features, make a classifier, validate and publish. This entire process is driven by the final step: a bullet on a resumé.*

Figure 4. *Shopping for datasets: our advisor happens to have multiple offices. In this case, our detection task becomes significantly easier — in the course of data collection, we never observed an open door. See figure 7 for a precision recall curve. This figure clearly shows that our method is scale and aspect invariant.*

Figure 5. ***Our advisor is the Baby-eating Bishop of Bath and Wells:*** *On the **left** is the top-ranked image from feeding "baby eating bishop bath wells" to Google's popular image search engine. On the **right** is a randomly chosen picture of our advisor. Using Lowe's SIFT features to detect keypoints (shown with black stars), we find that corresponding keypoints match perfectly. In fact, our results are independent of the choice of keypoints and similarity metric, suggesting that (1) our face detection system is robust, and (2) our advisor is the Baby-eating Bishop of Bath and Wells.*

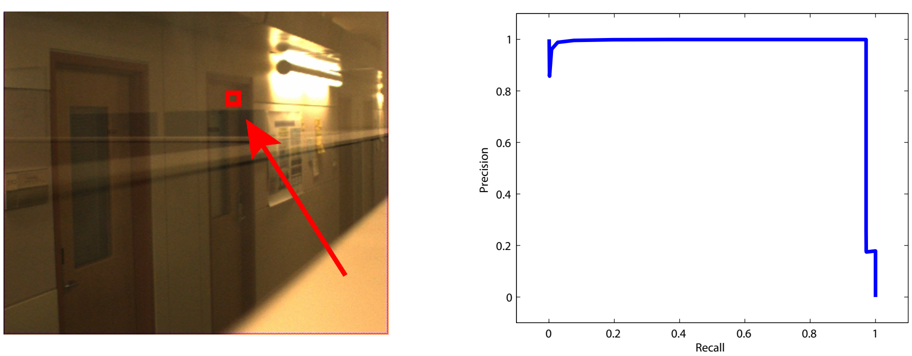

Figure 6. *We build a classifier for the state of our advisor's primary office door by thresholding the intensity of a single pixel. Typically, when the door is open, the intensity is brighter than when the door is closed. On the **left** we identify the pixel and on the **right**, we show the precision recall curve. Our method is not perfect – our method makes errors when the door is open and the lights are out. We believe that this method is general: most advisors have doors that look different from the inside of their office.*

**THIS PAGE
FOR SALE
$99.95**

**THIS PAGE
FOR SALE
$99.95**