

From: "[REDACTED]" on behalf of Ben Goertzel <[REDACTED]>
To: Jeffrey Epstein <jeevacation@gmail.com>
Cc: Itamar Arel [REDACTED]
Subject: Re: Distinguishing Cats from Dogs
Date: Sat, 10 Oct 2009 13:59:09 +0000

Hi Jeffrey,

First: Regarding MIT ... IMO the best vision work being done there now is Tomaso Poggio's work, which I suppose you're familiar with. He has some biologically realistic neural nets that classify images and videos, and that appear to effectively emulate the way humans classify images when they see them very briefly. (But not how humans classify images when they see them at length ... because this requires feedback connections which Poggio's networks don't model.) Based on discussions with some people who are involved in trying to commercialize Poggio's work, my strong impression is that Poggio's networks provide inferior classification results to Itamar's system

Next, about the "cats vs. dogs" task: of course it *isn't* that interesting....

People talk about "distinguishing cats vs. dogs" because Jeff Hawkins likes to give speeches about how hard the problem of distinguishing cats vs. dogs is (because Hawkins' Numenta vision system can't solve it). And of course if a famous guy like Hawkins speechifies about it, it must be important ;-D

So when Itamar first built his vision system he loaded in some cats vs. dogs pictures to see if it would distinguish them OK, and it seemed to ... and then he didn't pursue that anymore, since it was of no practical use and not terribly interesting [except to prove a point to people].... But it's not much work for him to download some more pictures of cats vs. dogs and run the algorithm on them again though.

However, the machines his algorithm is running on are currently doing some practical video classification work as part of a contract with ITT, so this will wait a week or so till the machines are available...

(the code now is running on ordinary PCs, though there is an almost-complete port to the Nvidia GPU supercomputer).

About MIT's and other peoples' vision processing results -- as Itamar pointed out at your house, the quality of video or image classification systems has to be considered carefully. For problems like face recognition or car recognition [or recognition of any specific class of objects], extremely high success rates can be achieved by hand-tuning a set of feature extractors, and then building a system with a pipeline like

video or image ==> hand-tuned feature extractors ==> machine learning system ==> image classification

Cassio and his Brazilian team built a face recognizer like this a couple years ago, for a Brazilian government customer to use with security cameras.

What's a harder problem is making a vision system that can solve ANY video classification problem WITHOUT creation of hand-tuned feature extractors.

Itamar's system does not outperform the competition if one assumes use of hand-tuned feature extractors. I think it does outperform the competition if one assumes hand-tuned feature extractors are not allowed.

But that's not even the main point ("outperforming" the competition on classification metrics). The main point -- from my perspective -- is that his vision system allows easy integration with a cognition system (like mine).

For instance, the face recognizer we built for the Brazilian government uses Support Vector Machines (SVMs) for the machine learning component (together with some hand-tuned feature extractors).

That works OK for face recognition under appropriate lighting conditions... but there's no way for SVMs to interact with a cognition system except by feeding the cognition system its outputs.

So if you use SVMs for vision processing, there is no way to use cognition to help vision handle the really difficult perception problem (say, where you have to use both structure *and* apparent function to identify any object). But in Itamar's architecture, it is straightforward to link his vision system in with my cognition system... which is important.

As I said to you before, I think much of the "trick" to how intelligence works is: It's not that each component of the intelligent system has to be SO awesome, but rather that the different components have to be interconnected in a "synergetic" way that helps them all work together.

If this is correct, then work on fine-tuning particular components (vision systems, language systems, etc.) is largely misdirected, and effort should be better spent on building holistic integrated intelligent systems. But of course this is more expensive and takes longer and it's harder to measure progress, so people tend to focus on particular components instead...

While we are not trying to model the brain in any detail, it makes sense for us to pay attention to high-level systemic properties of the brain -- and massive interconnectedness between regions carrying out different functions, is one of the brain's very notable properties.

Sorry for the long answer; brevity is not one of my virtues ;-)

ben

On Sat, Oct 10, 2009 at 9:37 AM, Jeffrey Epstein <jeevacation@gmail.com> wrote:
> videos are fine. // however I have seen the robot at the mit media lab.
> doing far more complex recognition many years ago.. i 'm not sure why you
> consider this interesting at all
>

> On Fri, Oct 9, 2009 at 3:51 PM, Itamar Arel [REDACTED] wrote:

>>
>> Jeffrey,
>> Thanks, again, for inviting us to talk to you - I enjoyed our discussion
>> very much.
>> Ben mentioned that you would like to see a demonstration of our system
>> distinguishing between cats and dogs. I have a deadline for delivering a
>> surveillance demo based on the same system to a defense contractor exactly a
>> week from today. I plan to spend a few days after that putting together a
>> video that clearly demonstrates our system's capabilities in distinguishing
>> cats and dogs. I hope that time frame is acceptable to you.
>> Meanwhile, I'm attaching a short video demonstrating our emotion
>> recognition system, which was trained to identify when a person is smiling,
>> pouting and either awake or asleep. Sending videos is the easy way to show
>> you what the system does. If you want to run our system directly rather than
>> just watching videos, you'll need to have MATLAB installed; our system is
>> currently implemented as a set of MATLAB scripts, rather than as a
>> standalone executable program. We also have a GPU version of the code in
>> development, which is bound to be integrated with Ben's system after our
>> collaborative project gets off the ground.

>> - Itamar

>>

>> --

>> =====

>> Itamar Arel, Ph.D., M.B.A.
>> Associate Professor
>> Director, Machine Intelligence Lab (<http://mil.engr.utk.edu>)
>> Office: [REDACTED]
>> Mobile: [REDACTED]
>>
>> Min H. Kao Department of Electrical Engineering & Computer Science
>> The University of Tennessee

>> =====
>
>
>
> --

> *****
> The information contained in this communication is
> confidential, may be attorney-client privileged, may
> constitute inside information, and is intended only for
> the use of the addressee. It is the property of
> Jeffrey Epstein
> Unauthorized use, disclosure or copying of this
> communication or any part thereof is strictly prohibited
> and may be unlawful. If you have received this
> communication in error, please notify us immediately by
> return e-mail or by e-mail to jeevacation@gmail.com, and
> destroy this communication and all copies thereof,
> including all attachments.
>

--
Ben Goertzel, PhD
CEO, Novamente LLC and Biomind LLC
Director of Research, SIAI
[REDACTED]

"Truth is a pathless land" -- Jiddu Krishnamurti