From: Andrew Dienstfrey
To: Liu, Yi-Kai (Fed)
Subject: Re: Getting into more detail
Date: Wednesday, April 12, 2017 5:19:29 PM

Hi Yi-Kai,

I'm so sorry that I did not respond to your note sooner. I received it at the time that I was scrambling to hit two deadlines. In the midst of that chaos the message got lost.

I think your idea is worth pursuing. Somewhere in the training process there must be regularization terms to prevent overfitting. Not sure how they function and what are their pros and cons.

It might be interesting to explore the regularization strategy(ies?) in the context of recent "adversary" results. Not sure if you have run into the later. It turns out that many ANN's are susceptible to making errors in the sense that an extremely corrupted input (in essence, one that is visually indistinguishable from noise as far as human perception is concerned) is "identified" as a known object with high confidence. I haven't looked at the matter carefully so my description of the phenomenon should be taken with a grain of salt. Nevertheless, from a first glance, it seems that the strategy for creating these adversarial inputs is akin to taking a clean/identified input and running backprop in reverse on the input. re: connecting this circle of results to your idea, one thought might be that the regularization filters (that are needed to prevent over-fitting) are also responsible for damping out specificity.

Just thinking out loud. The short answer is that I imagine that there could be some interesting work to be done on the question of stabilizing the training process to prevent overfitting.

Regards,

Andrew

On 03/28/2017 03:44 PM, Liu, Yi-Kai (Fed) wrote:

> Hi Andrew,
>
> I was thinking a bit more about this neuromorphic stuff, and I remembered one other question that we might hope to address: How does overfitting happen (or not happen) when one trains a large neural network?
>
> Maybe, this question could be a good way to motivate the theory part of this IMS project...? First, I think this is one of the major areas where existing theory doesn't do a good job of describing what happens in practice. And, it is one specific question that any good theoretical model of neuromorphic computation should be able to answer. What do you think?
>
> Cheers,
>
> --Yi-Kai
>
> _____
> From: Andrew Dienstfrey <andrew.dienstfrey@nist.gov>
> Sent: Thursday, March 9, 2017 11:04:40 AM
> To: Liu, Yi-Kai (Fed)
> Subject: Re: Getting into more detail
>
> Hi Yi-Kai,
>
> I also found our discussion very helpful. Thank you for summarizing. I will assemble your notes with a few twists and forward to Jeff to give him a heads-up on what we are thinking.
>
> Some articles that I have found (but not yet read) that seem to explore ideas that we discussed are:
>
> * Watkin, "The statistical mechanics of learning a rule" (1993), https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.65.499
> * Mace, "Statistical mechanical analysis of the dynamics of learning in perceptrons" (1998), https://link.springer.com/article/10.1023/A:1008896910704
> * Burger, "Computational Capacity and Energy Consumption of Complex Resistive Switch Networks" (2015), https://arxiv.org/abs/1507.03716
>
> I'll have a look at these and the references you sent. Perhaps we could circle back next week to compare notes.
>
> Regards,
> Andrew
>
> p.s. Unfortunately I will have to miss the SFQ-JJ-Neuromorphic discussion today. Too bad.

On 03/07/2017 12:40 PM, Liu, Yi-Kai (Fed) wrote:
Hi Andrew,

Thanks for the discussion, I thought it was really helpful!

In case it is useful, here are some notes that I wrote down: (though I may have missed some things, particularly some of the things you said)

Theoretical foundations
  What are the fundamental units of a neuromorphic computer?
  "Bits," "operations"?
  High dimensional statistics - sparsity, low dimensional manifolds?
  Complexity theory - seems hard

Standard reference materials
  Data sets, like MNIST, TREC - a lot of work, not as helpful for research?
  Standard examples of neural networks that can run on different hardware?

Theoretical models at all scales (microscopic to macroscopic)
  Models of individual neurons
    How to validate these models, compare with real hardware?
  Models of small networks
    How to measure the state of an entire network? "Tomography"?
    How to simulate an entire network?
  Scaling up to big networks
    Concerns: robustness, fault-tolerance
    Want to integrate out microscopic degrees of freedom
    Want to derive an effective model of the macroscopic system
    Borrow mathematical methods from statistical mechanics, open quantum systems
    What are the key measurements that will ensure that the hardware can scale up?
      E.g., noise spectral density? Time correlation functions?

What do experimentalists need?
  What are questions where they would like guidance from theory?

Finally, here are a couple of references on open quantum systems:
  https://www.amazon.com/Theory-Open-Quantum-Systems/dp/0199213909
  https://arxiv.org/abs/1104.5242
They do some interesting calculations, where they start with a complete microscopic description of a quantum system, and then derive coarse-grained macroscopic descriptions, such as master equations, and convergence to an equilibrium state in the weak coupling limit. Maybe we could try to do something analogous for neuromorphic computing hardware?

Cheers,

--Yi-Kai

From: Andrew Dienstfrey <andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov>
Sent: Monday, March 6, 2017 5:47 PM
To: Liu, Yi-Kai (Fed)
Subject: Re: Getting into more detail

No problem, Yi-Kai. I'll call you tomorrow at 1 your time. Have a good night. --ad

On 03/06/2017 03:46 PM, Liu, Yi-Kai (Fed) wrote:
Hi Andrew,

Cool, tomorrow (Tuesday) morning works for me. Thanks! Can we try later in the morning, say 11am (mountain time) = 1pm (eastern time)? My phone number is 301-314-1850. Also, I'm sorry for being slow to reply!

Cheers,

--Yi-Kai

From: Andrew Dienstfrey
<andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov>
Sent: Monday, March 6, 2017 11:52:50 AM
To: Liu, Yi-Kai (Fed)
Subject: Re: Getting into more detail

Hi Yi-Kai,

Welcome back. What does your schedule look like this week? I am available tomorrow (Tues) am if you would like to talk.

Andrew

On 02/24/2017 03:18 PM, Liu, Yi-Kai (Fed) wrote:
> Hi Jeff,
>
> Hmm, if you'd like to talk sooner, I think I'll have some time on Thursday March 2? Or we could wait until the week of March 6. Sorry my schedule is a bit complicated right now. :(
>
> --Yi-Kai
> _____
> From: Shainline, Jeff (Fed)
> Sent: Friday, February 24, 2017 5:05:34 PM
> To: Liu, Yi-Kai (Fed); Dienstfrey, Andrew M. (Fed)
> Subject: RE: Getting into more detail
>
> Hi Yi-Kai,
>
> Yes, I'd be happy to talk to you about some of the experiments. In addition to that, I wanted to touch base with the theoretical team to make sure we have a good plan for that aspect of the program before developing the six-page proposal for March 17th. We don't need to have everything figured out, but the more we can converge the better.
>
> Let me know what works for you,
> Jeff
>
> -----Original Message-----
> From: Liu, Yi-Kai (Fed)
> Sent: Friday, February 24, 2017 2:33 PM
> To: Dienstfrey, Andrew M. (Fed)
<andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov>
> Cc: Shainline, Jeff (Fed)
<jeffrey.shainline@nist.gov><mailto:jeffrey.shainline@nist.gov><mailto:jeffrey.shainline@nist.gov><mailto:jeffrey.shainline@nist.gov>
> Subject: Re: Getting into more detail
>
> Hi Andrew,
>
> I'm sorry I didn't have any time to talk today... and unfortunately I'll be out of town next week. Could we talk sometime the week of March 6? Tuesday, March 7 is good for me, but other days are ok too. Or do you think that would be too late?
>
> Actually, I wonder if we should talk with some of the experimental people as well? I remember Jeff sent out a big table of the different measurement techniques they are developing, and I'd really like to understand that better...
>
> --Yi-Kai
> _____
> From: Andrew Dienstfrey
<andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov>
> Sent: Friday, February 24, 2017 11:37 AM
> To: Liu, Yi-Kai (Fed)
> Subject: Re: Getting into more detail
>
> Hi Yi-Kai,
>
> Sorry, my week got slammed with other items. I would be happy to discuss today if that works for you.
>
> Andrew
>
>
> On 02/16/2017 08:47 AM, Liu, Yi-Kai (Fed) wrote:
>> Thanks very much for all your work on this, Jeff!
>>
>> Andrew -- yes, I'll be around next week, and would love to chat! Could we try sometime later in the week, say Wed. Feb. 22 or later? My schedule is pretty flexible, let me know when is convenient for you...
>>
>> Cheers,
>>
>> --Yi-Kai
>> _____
>> From: Andrew Dienstfrey
<andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov><mailto:andrew.dienstfrey@nist.gov>
>> Sent: Thursday, February 16, 2017 9:30 AM

>> To: Shainline, Jeff (Fed); Liu, Yi-Kai (Fed)
>> Subject: Re: Getting into more detail
>>
>> Nicely done, Jeff. Looking forward to developing this proposal further.
>>
>> Yi-Kai, perhaps you and I can think on this some and then discuss over a phone call. Are you around next week?
>>
>> Andrew
>>
>> On 02/15/2017 05:39 PM, Shainline, Jeff (Fed) wrote:
>> Hi Andrew, Hi Yi-Kai,
>>
>> I'd like to get more concrete about what the theoretical component of this IMS will look like. During the elevator pitch, I briefly mentioned that we'd like to work toward a new set of benchmarks and metrics by which neuromorphic computers can be gauged. On first pass, I wrote down that these should include network metrics (clustering, path length, cycle length), computational metrics (speed and energy consumption per operation and for standard tasks), and information metrics (entropy, mutual information applied at the device level and system level). Please let me know if you have suggestions or can offer anything more specific.
>>
>> Also, one way I've been thinking about this problem is in terms of one computer watching another computer. The general concept is described in the attached Word document. It is shamefully superficial, but hopefully we can refine our thinking before the proposal in mid-April. Any comments or feedback would be greatly appreciated.
>>
>> Thanks,
>> Jeff
>>
>>
>>
>> --
>> Andrew Dienstfrey
>> Manager, Virtual Measurement Systems Program Information Technology
>> Laboratory National Institute of Standards & Technology
>> 325 Broadway, Mail Stop 891.01
>> Boulder, CO 80305-3328
>> Tel: 303-497-7485
>> Fax: 303-497-3012
>>
>>
>>
>>
> --
> Andrew Dienstfrey
> Manager, Virtual Measurement Systems Program Information Technology Laboratory National Institute of Standards & Technology
> 325 Broadway, Mail Stop 891.01
> Boulder, CO 80305-3328
> Tel: 303-497-7485
> Fax: 303-497-3012
>
>
>

--
Andrew Dienstfrey
Manager, Virtual Measurement Systems Program
Information Technology Laboratory
National Institute of Standards & Technology
325 Broadway, Mail Stop 891.01
Boulder, CO 80305-3328
Tel: 303-497-7485
Fax: 303-497-3012

--
Andrew Dienstfrey
Manager, Virtual Measurement Systems Program
Information Technology Laboratory
National Institute of Standards & Technology
325 Broadway, Mail Stop 891.01
Boulder, CO 80305-3328
Tel: 303-497-7485
Fax: 303-497-3012

--
Andrew Dienstfrey
Manager, Virtual Measurement Systems Program
Information Technology Laboratory
National Institute of Standards & Technology
325 Broadway, Mail Stop 891.01
Boulder, CO 80305-3328
Tel: 303-497-7485
Fax: 303-497-3012

--
Andrew Dienstfrey
Manager, Virtual Measurement Systems Program
Information Technology Laboratory
National Institute of Standards & Technology
325 Broadway, Mail Stop 891.01
Boulder, CO 80305-3328
Tel: 303-497-7485
Fax: 303-497-3012