
From: jeffrey E. <jeevacation@gmail.com>
Sent: Wednesday, October 11, 2017 6:02 PM
To: Misha Gromov
Subject: Fwd:

----- Forwarded message -----

From: Joscha Bach <[REDACTED]>
Date: Wed, Oct 11, 2017 at 7:55 PM
Subject: Re:
=o: Jeffrey Epstein <jeevacatio=@gmail.com <mailto:jeevacation@gmail.com>>

After skimming their paper, the idea seemed =nexciting to me at first: basically, if we have enough feature dimensions =e can almost always find a linear separation. This is also related to how =support Vector Machines work: they project the data into an extremely high-=imensional space, find a separating hyperplane with linear regression, and=then project that plane back into the original space as the separator. A s=milar idea is behind Echo State networks, which use a randomly wired recur=ent neural network and then only train the output layer with a single line=r regression. The authors take an existing trained neural network, and whenever it makes = mistake, they train a linear classifier on the network state and data, i.=. they try to find out when the network goes wrong. Instead of improving t=e network (which is also likely to make it worse in other cases), they add=an additional layer to it. For engineering, this makes a lot of sense, bec=use large neural networks are cheap to use and deploy but expensive to tra=n.

On a more philosophical level, it is tempting to ask if that might be a gen=ral learning principle for brains: when you don't perform well, add mo=e control structure on top. It probably makes sense whenever you are confi=ent that training the existing structure won't improve it that much, b=t unless training the weights in an existing network, it also adds quite a=few milliseconds to the processing time. There is probably an optimal trad=off for this. The other thing is that the new layer is a linear classifier=only (at least in this paper), and it is creating a local override on the =system's results, instead of integrating with it, somewhat similar to h=w reasoning might override our subconscious behavior. One of the drawbacks=is that this won't allow us to use the new layer for simulating/unders=anding the structure of the domain modeled by the rest of the network.

– Joscha

> On Oct 10, 2017, at 09:43, jeffrey E. <jeevacation@gmail.com <mailto:jeevacatio=@gmail.com>> wrote:

>
> <https://www.sciencedaily.com/htm>
> <<https://www.sciencedaily.com/releases/2017/08/170821102725.=tm>>
>
> --

> please note
> 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
> JEE 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
> <<mailto:jeevacation@gmail.com>> , and destroy this communication and
> all copies thereof, including all attachments. copyright -all rights
> reserved

--

=A0 please note

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 JEE 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. copyright -all rights reserved

--94eb2c0c7c8693da21055b493771-- conversation-id 28259 date-last-viewed 0 date-received 1507744906 flags 8590195713 gmail-label-ids 7 6 remote-id 757805