From:	Joscha Bach ███████████████████
Sent:	Wednesday, October 11, 2017 5:56 PM
To:	Jeffrey Epstein
Subject:	Re:

After skimming their paper, the idea seemed unexciting to me at first: =asically, if we have enough feature dimensions we can almost always =ind a linear separation. This is also related to how Support Vector =achines work: they project the data into an extremely high-dimensional =pace, find a separating hyperplane with linear regression, and then =roject that plane back into the original space as the separator. A =imilar idea is behind Echo State networks, which use a randomly wired =ecurrent neural network and then only train the output layer with a =ingle linear regression.
The authors take an existing trained neural network, and whenever it =akes a mistake, they train a linear classifier on the network state and =ata, i.e. they try to find out when the network goes wrong. Instead of =mproving the network (which is also likely to make it worse in other =ases), they add an additional layer to it. For engineering, this makes = lot of sense, because large neural networks are cheap to use and =eploy but expensive to train.

On a more philosophical level, it is tempting to ask if that might be a =eneral learning principle for brains: when you don't perform well, add =ore control structure on top. It probably makes sense whenever you are =onfident that training the existing structure won't improve it that =uch, but unless training the weights in an existing network, it also =dds quite a few milliseconds to the processing time. There is probably =n optimal tradeoff for this. The other thing is that the new layer is a =inear classifier only (at least in this paper), and it is creating a =ocal override on the system's results, instead of integrating with it, =omewhat similar to how reasoning might override our subconscious =ehavior. One of the drawbacks is that this won't allow us to use the =ew layer for simulating/understanding the structure of the domain =odeled by the rest of the network.

– Joscha

> On Oct 10, 2017, at 09:43, jeffrey E. <jeevacation@gmail.com> wrote:
>
> https://www.sciencedaily.com/releases/2017/08/170821102725.htm
>
> --
>   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

<?xml version=.0" encoding=TF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN" "http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version=.0">
<dict>
        <key>conversation-id</key>

```
        <integer>28259</integer>
        <key>date-last-viewed</key>
        <integer>0</integer>
        <key>date-received</key>
        <integer>1507744549</integer>
        <key>flags</key>
        <integer>8590195717</integer>
        <key>gmail-label-ids</key>
        <array>
                <integer>6</integer>
                <integer>2</integer>
        </array>
        <key>remote-id</key>
        <string>757802</string>
</dict>
</plist>
```