Twitter Thread by Vladimir Haltakov



Vladimir Haltakov
@haltakov

Machine Learning Paper Reviews ■■

Check out this thread for short reviews of some interesting Machine Learning and Computer Vision papers. I explain the basic ideas and main takeaways of each paper in a Twitter thread.

■ I'm adding new reviews all the time! ■

AlexNet - the paper that started the deep learning revolution in Computer Vision! ■■■■

https://t.co/aE8259ca1N

> It's finally time for some paper review! \U0001f4dc\U0001f50d\U0001f9d0
>
> I promised the other day to start posting threads with summaries of papers that had a big impact on the field of ML and CV.
>
> Here is the first one - the AlexNet paper! pic.twitter.com/QNLPIMZSIa
>
> — Vladimir Haltakov (@haltakov) September 28, 2020

DenseNet - reducing the size and complexity of CNNs by adding dense connections between layers. ■■

https://t.co/ZaGeHPzoTs

> ML paper review time - DenseNet! \U0001f578\ufe0f
>
> This paper won the Best Paper Award at the 2017 Conference on Computer Vision and Pattern Recognition (CVPR) - the best conference for computer vision problems.
>
> It introduces a new CNN architecture where the layers are densely connected. pic.twitter.com/DuHytaoXia
>
> — Vladimir Haltakov (@haltakov) October 15, 2020

Playing for data - generating synthetic GT from a video game (GTA V) and using it to improving semantic segmentation models. ■■

https://t.co/tgNvLnZgT1

> Time for another ML paper review - generating synthetic ground truth data from video games! \U0001f3ae
>
> I love this paper, because it pushes the boundaries of creating realistic synthetic ground truth data and shows that you can use it for training and improve your model.
>
> Details \U0001f447 pic.twitter.com/fBgORYG8Lz
>
> — Vladimir Haltakov (@haltakov) October 5, 2020

Transformers for image recognition - a new paper with the potential to replace convolutions with a transformer. ■

https://t.co/K3uN0zE2DA

> Another paper review, but a little different this time... \U0001f937\u200d\u2642\ufe0f
>
> The paper is not published yet, but is submitted for review at ICLR 2021. It is getting a lot of attention from the CV/ML community, though, and many speculate that it is the end of CNNs... \U0001f447https://t.co/bh6wUxYfxu pic.twitter.com/dZGBYB8A5U
>
> — Vladimir Haltakov (@haltakov) October 5, 2020

ResNet allowed to train very deep and simpler networks for the first time. ■■

https://t.co/BCWxz7gfPe

> Here comes a review of one of the fundamental papers in CV and ML - ResNet \u21a9\ufe0f
>
> Formulating the classical CNN network as a residual network allowed training much deeper (but simpler) networks. ResNet won the first place in the ILSVRC and COCO 2015 challenges.
>
> Thread \U0001f447 pic.twitter.com/bGEEz08HpE
>
> — Vladimir Haltakov (@haltakov) October 22, 2020