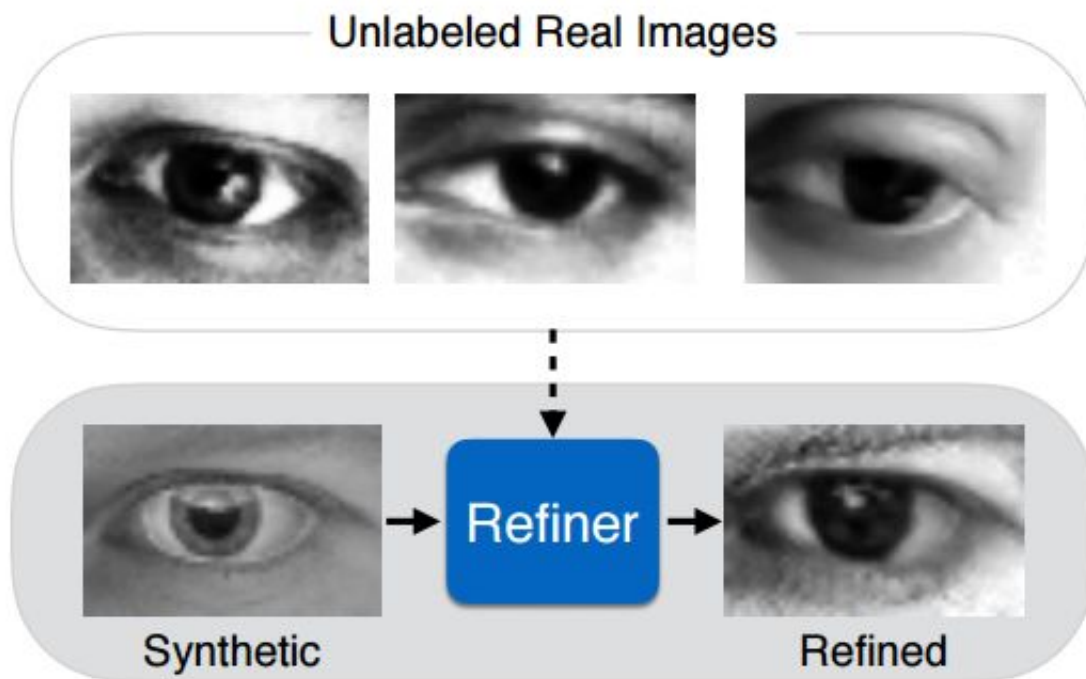


Apple AI research paper is from vision expert and team

December 28 2016, by Nancy Owano



Simulated+Unsupervised (S+U) learning. The task is to learn a model that improves the realism of synthetic images from a simulator using unlabeled real data, while preserving the annotation information. Credit: Ashish Shrivastava, et al. arXiv:1612.07828 [cs.CV]

(Tech Xplore)—Apple is usually in the news over some product launch, future iPhone speculations, or patent filing. Not this week. Apple made the news over, wait for this, a research paper.

Turns out this is a research [paper](#) that describes a technique for improving artificial intelligence. The focus is on computer vision and pattern recognition.

How well do machines see [images](#)? How well do they interpret them? The researchers are in that area of enquiry.

The [research paper](#) is a big deal, said Paul Lilly, senior editor, *Hot Hardware*. Why? Apple, he said, "[joins](#) the fray after having published its first AI paper this month." It was submitted in November.

By fray, Lilly was referring to the world's biggest technology companies, including Microsoft, IBM, Facebook, Google, paying attention to the growing fields of machine learning and [artificial intelligence](#).

There is another reason that this move has drawn the attention of tech watchers. "Apple has kept its research tight lipped and out of the public eye. Publishing this paper can be seen as an indication that Apple wants a more visible presence in the field of AI," said Lilly.

Don Reisinger in *Fortune* similarly noted that "Scientists around the world have long criticized Apple for not publishing research about artificial [intelligence](#)." (Nonetheless, Apple's competitors may generally publish their own papers on a number of topics, but they, too, he added, keep some advancements secret.)

Hot Hardware pointed out that Apple's researchers looked at a method that involves "a simulator generating synthetic images that are put through a refiner. The result is then sent to a discriminator that must figure out which are real and which are synthetic."

AppleInsider talked about the use of synthetic, or computer generated, images.

"Compared to training models based solely on real-world images, those leveraging synthetic data are often more efficient because computer generated images are usually labelled. For example, a synthetic image of an eye or hand is annotated as such, while real-world images depicting similar material are unknown to the algorithm and thus need to be described by a human operator."

Thing is, relying on simulated images may not prove successful. *AppleInsider* said, "[computer](#) generated content is sometimes not realistic enough to provide an accurate learning set. To help bridge the gap, Apple proposes a system of refining a simulator's output through "Simulated+Unsupervised learning."

The authors said in their paper that they proposed S+U learning to refine a simulator's output with unlabeled real data. Their method involved an adversarial network.

Avaneesh Pandey, *International Business Times*, explained what was going on. He said, "the researchers used a technique known as adversarial learning, wherein two competing neural networks basically try to outsmart each other. In this case, the two [neural networks](#) are the generator, which, as the name suggests, generates realistic images, and the discriminator, whose function is to distinguish between generated and real [images](#)."

The [paper](#) was made public on the *arXiv*. "Learning from Simulated and Unsupervised Images through Adversarial Training" is the title of the [research](#) paper by vision expert Ashish Shrivastava and team including Tomas Pfister, Oncel Tuzel, Wenda Wang, Russ Webb and Josh Susskind, said *AppleInsider*.

More information: Learning from Simulated and Unsupervised Images through Adversarial Training, arXiv:1612.07828 [cs.CV]

arxiv.org/abs/1612.07828v1

Abstract

With recent progress in graphics, it has become more tractable to train models on synthetic images, potentially avoiding the need for expensive annotations. However, learning from synthetic images may not achieve the desired performance due to a gap between synthetic and real image distributions. To reduce this gap, we propose Simulated+Unsupervised (S+U) learning, where the task is to learn a model to improve the realism of a simulator's output using unlabeled real data, while preserving the annotation information from the simulator. We develop a method for S+U learning that uses an adversarial network similar to Generative Adversarial Networks (GANs), but with synthetic images as inputs instead of random vectors. We make several key modifications to the standard GAN algorithm to preserve annotations, avoid artifacts and stabilize training: (i) a 'self-regularization' term, (ii) a local adversarial loss, and (iii) updating the discriminator using a history of refined images. We show that this enables generation of highly realistic images, which we demonstrate both qualitatively and with a user study. We quantitatively evaluate the generated images by training models for gaze estimation and hand pose estimation. We show a significant improvement over using synthetic images, and achieve state-of-the-art results on the MPIIGaze dataset without any labeled real data.

© 2016 Tech Xplore

Citation: Apple AI research paper is from vision expert and team (2016, December 28) retrieved 4 May 2024 from <https://techxplore.com/news/2016-12-apple-ai-paper-vision-expert.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
