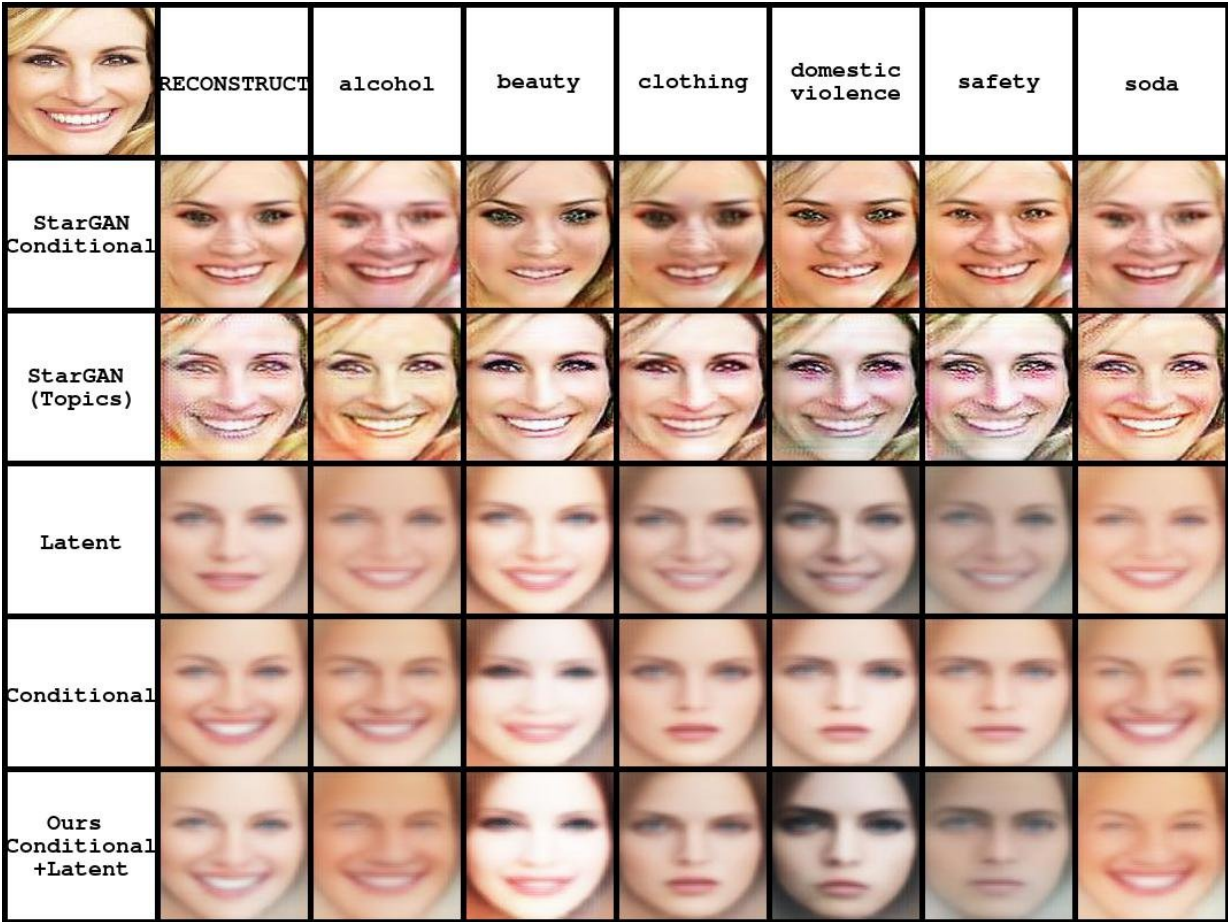


Using machine learning to generate persuasive faces for ads

August 6 2018, by Ingrid Fadelli



Ad faces transformed into 5 different categories. Credit: Thomas & Kovashka

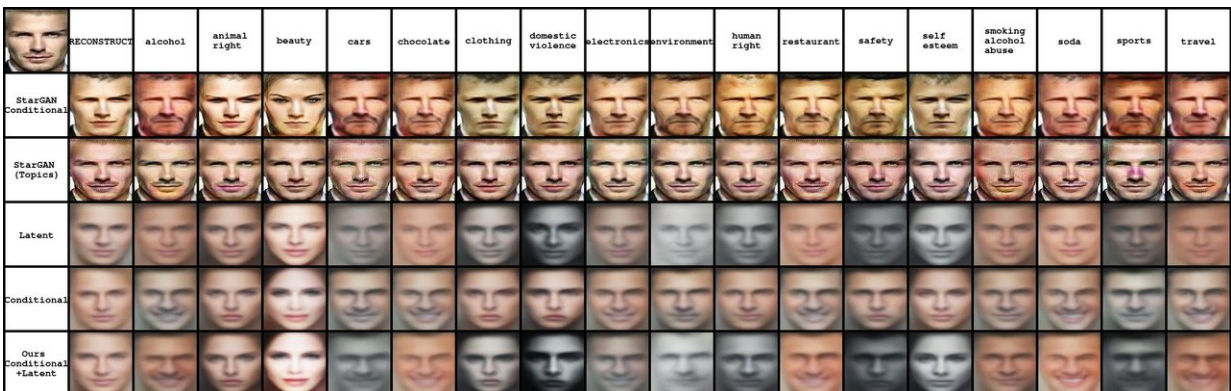
Researchers from the University of Pittsburgh have recently developed a

conditional variational autoencoder that can produce unique faces for advertisements. Their [study](#) is grounded on their previous work, which explored automated methods of better understanding advertisement.

"In our past project, we wanted to see whether machines could decode the complex visual rhetoric found in ads," Christopher Thomas, one of the researchers who carried out the study, told Tech Xplore. "Ads contain puns, metaphors, and other persuasive rhetorical devices that are hard for machines to understand. In this paper, we didn't only want to understand ads, but we wanted to see whether such persuasive content could be automatically generated by computers."

The primary mission of the advertising industry is to promote products or convey ideas using persuasive language and images. Faces, a key aspect of ads, are often portrayed differently depending on the product advertised and message communicated.

In collaboration with his colleague Adriana Kovashka, Thomas used machine learning to generate persuasive faces that would work well for different types of advertisements. They used conditional variational autoencoders, or "generative models," machine learning models which learn to generate synthetic data similar to that which it is trained on.



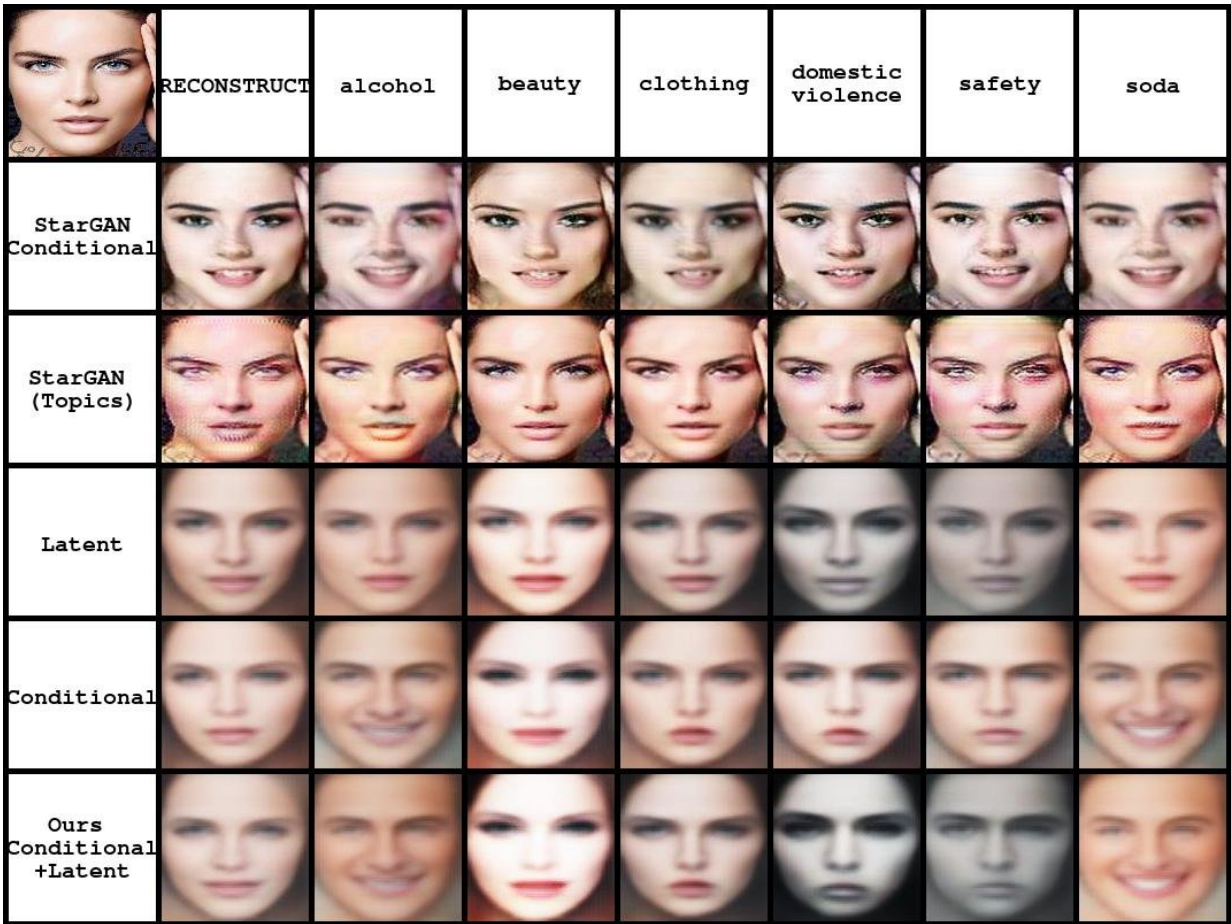
Ad faces transformed into 17 different categories. Credit: Thomas & Kovashka

"In computer vision, autoencoders work by taking an image and learning to represent that image as a few numbers," Thomas said. "Then, a second piece of the model, the decoder, learns to take those numbers and reproduce the original image from it. You can almost think of it as a form of compression, in which a large image is represented by a few numbers."

When this type of machine learning model is trained on a large enough dataset, it starts to represent semantic aspects within the numbers. For instance, in the model developed by Thomas and Kovashka, one number would control the shape of a face, another the shade of the skin, and so on for other semantic features.

However, if the researchers wanted the model to capture whether a person is wearing glasses, but the training dataset did not contain enough pictures of people with glasses, this property would be lost when the image is reconstructed. Thus, they developed a conditional autoencoder, meaning that they could add other numbers to the model that it had not acquired alone, representing semantic features that might be relevant to particular advertisements.

"The cool part of this is that once we trained the model to represent faces in 100 numbers, if we then change some of those numbers and "decode" them, we can change the face," Thomas said. "We can thus transform existing faces so that they look the same but have different attributes, such as eyeglasses, smiling or not, etc., just by changing some of the numbers that our model uses to represent them."



Ad faces transformed into 5 different categories. Credit: Thomas & Kovashka

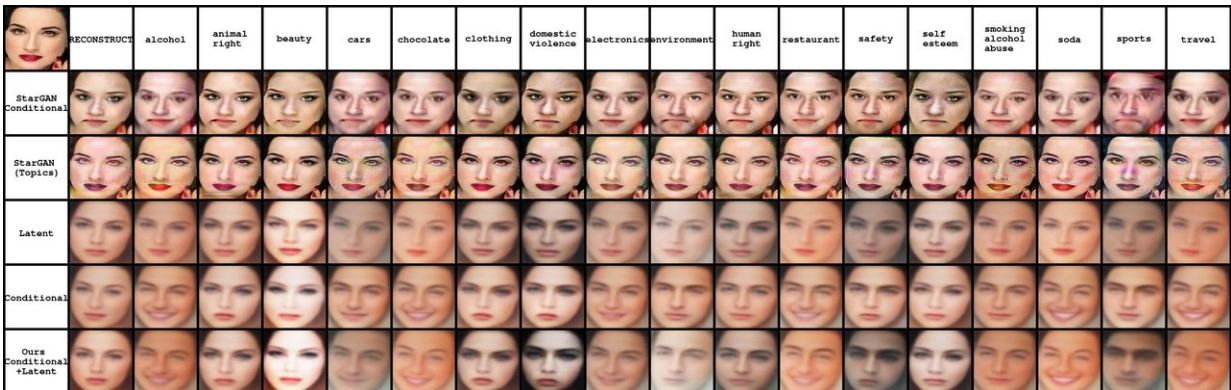
Training generative models for computer vision can be a challenging task, requiring large image datasets and often failing when trained on very diverse data, such as ads. Thomas and Kovashka overcame these limitations by using an autoencoder that required less data and could cope with the sizeable variance found in advertisement.

"Even so, because there wasn't enough data, it didn't always capture the concepts we wanted it to in its representations," says Thomas. "Thus, we

deliberately injected semantics into its representation, which improved results significantly."

Their findings suggest that in the future, advertisers will be able to create customized and targeted ads that are tailored to individual customers. For instance, they could generate faces with facial features that match those of the viewer, so that they identify more with the subject.

"This kind of automatic, fine-grained ad customization could have huge implications for online advertisers," says Thomas. "In addition, an advertiser who doesn't want to hire an extra [model](#) for their ad or to do manual editing may be able to transform an existing face from another ad into a face appropriate for their type of ad."



Ad faces transformed into 17 different categories. Credit: Thomas & Kovashka

The researchers are now exploring ways in which they could improve their generated images so that they match the quality of those produced using larger amounts of data. To do this, they will need to design other generative models that are more robust when trained on highly varied

and limited data.

"Another possible line of research is generating other objects besides faces, or even generating entire ads that are meaningful and interesting," says Thomas. "This would require the development of new techniques for modeling rhetorical structure in a generative framework, combined with text understanding and generation."

More information: Persuasive Faces: Generating Faces in Advertisements, arXiv: 1807.09882v1 [cs, CV].
arxiv.org/abs/1807.09882

Abstract

In this paper, we examine the visual variability of objects across different ad categories, i.e. what causes an advertisement to be visually persuasive. We focus on modeling and generating faces which appear to come from different types of ads. For example, if faces in beauty ads tend to be women wearing lipstick, a generative model should portray this distinct visual appearance. Training generative models which capture such category-specific differences is challenging because of the highly diverse appearance of faces in ads and the relatively limited amount of available training data. To address these problems, we propose a conditional variational autoencoder which makes use of predicted semantic attributes and facial expressions as a supervisory signal when training. We show how our model can be used to produce visually distinct faces which appear to be from a fixed ad topic category. Our human studies and quantitative and qualitative experiments confirm that our method greatly outperforms a variety of baselines, including two variations of a state-of-the-art generative adversarial network, for transforming faces to be more ad-category appropriate. Finally, we show preliminary generation results for other types of objects, conditioned on an ad topic.

© 2018 Tech Xplore

Citation: Using machine learning to generate persuasive faces for ads (2018, August 6) retrieved 23 April 2024 from <https://techxplore.com/news/2018-08-machine-persuasive-ads.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.