October 05, 2016

ASSESSING BRAND VISIBILITY USING MACHINE LEARNING

Sujoy Banerjee

Michael Greene

ASSESSING BRAND VISIBILITY USING MACHINE LEARNING

Introduction

The present disclosure provides systems and methods that assess brand visibility using imagery and machine learning. In particular, advertisers spend significant amounts each year on public and physical advertising that is aimed at building brand awareness and/or keeping a brand in the consumer's consciousness. Examples of such advertising include "out of home" advertisements such as billboard advertisements, bus station advertisements, mobile advertisements (e.g., taxi toppers or side-of-bus advertisements), packaging (e.g., shopping bags or take-out food containers/cups, etc.), cars, building signs, clothing with logos, etc.

However, it is notoriously difficult to directly assess the effectiveness of advertisement campaigns. In particular, many advertisers only observe the results of their advertisement campaigns through the imperfect lens of aggregate sales figures, which do not provide the granularity to determine the effectiveness of individual advertisements included in the campaign. Therefore, one challenge faced by advertisers is to optimize advertisement campaigns to maximize brand exposure or visibility per dollar spent on advertisements.

One example technique to identify brand visibility includes manual sampling. In particular, an advertiser can hire or recruit people to manually count advertisements that they see during a trial time period. However, this technique has several drawbacks, including cost, scale of coverage, and sampling errors (e.g., the experimenters may not have representative schedules/demographics, or may be primed to see advertising).

Another technique includes the use of surveys which ask a sample of the target audience about their familiarity with brands and/or whether they can recall seeing an advertisement for a specific brand. This technique has the drawback of being quite inaccurate.

Summary

The present disclosure proposes to solve the challenges described above by leveraging one or more machine learning techniques or machine-learned models to perform a large scale statistical assessment of the visibility of an advertising campaign. More particularly, one aspect of the present disclosure provides a processing pipeline that takes input imagery and outputs an indication of a visibility of a brand.

In particular, a series of fixed or mobile cameras (e.g., mobile image capture vehicles) can be operated to collect large amounts of imagery in public settings. Such collection of imagery can capture a corpus of views in locations that are frequented by the advertisement campaign's target audience. The corpus of imagery can be stored in an accessible database of an imagery platform.

A computer vision or machine learning system (such as a system that includes or stores one or more machine-learned models such as neural networks) can be used to review such imagery and detect, count, and/or locate all observations of the brand being assessed. For example, a machine-learned neural network can receive input imagery and output indications of whether a brand or advertisement is depicted in such imagery and, further, the identity of the brand.

Once all instances of the brand being assessed have been detected and located by the machine learning system, a brand visibility score can be computed for the brand based on such

instances. For example, the locations of the brand observations can be correlated with typical population densities to calculate an estimated number of impressions per unit time.

If greater accuracy is desired, traffic flow data (or even demographically-annotated traffic flow data) can be collected from devices such as mobile devices (e.g., by means of a mobile application or device operating system that provides location updates for the device or by means of Wi-Fi or cellular base stations that track attempted connections). The traffic flow data can be overlaid on or otherwise combined with the locations of the detected advertisements to produce more accurate measures of impressions, such as histograms of impression count/probability, or even such histograms broken down by demographic segment. The results can even be analyzed to measure the number of impression contributions of individual ads or ad types, to support cost-benefit analysis.

Detailed Description

Figure 1 depicts an example computing system 100 to assess brand visibility using machine-learned models. The system 100 includes a machine learning computing system 130 and an imagery platform 120 that are communicatively coupled over a network 180.

The imagery platform 120 can include an accessible image database 122 that stores imagery of geographic areas that contain advertisements. For example, the image database 122 can store ground-level images 124. As examples, the ground-level images can be street-level panoramic images, user-submitted photographs, images collected by an image collection vehicle, or other imagery at or near the level of the ground. The image database 122 can also store overhead images 126. As examples, the overhead images can be images captured by aircraft (e.g., planes, drones, etc.), or other imagery taken from an overhead position.

The machine learning computing system 130 includes one or more processors 132 and a memory 134. The memory 134 can store data 136 and instructions 138 which are executed by the processor 132 to cause the machine learning computing system 130 to perform operations.

In some instances, the machine learning computing system 130 includes or is otherwise implemented by one or more server computing devices. In instances in which the machine learning computing system 130 includes plural server computing devices, such server computing devices can operate according to sequential computing architectures, parallel computing architectures, or some combination thereof.

The machine learning computing system 130 stores or otherwise includes one or more machine-learned brand recognition models 140. For example, the models 140 can be or can otherwise include various machine-learned models such as neural networks (e.g., deep neural networks) or other multi-layer non-linear models.

According to an aspect of the present disclosure, the machine learning computing system 130 can access images from the imagery platform 120. The machine learning computing system 130 can use one or more models 140 to detect and recognize one or more brands depicted in the images (e.g., one or more advertisements for particular brands). Thus, the machine learning computing system 130 can input the obtained images into one or more models 140 to receive one or more outputs of the one or more models 140 that describe the identity of a brand and one or more locations at which the brand is advertised.

In addition, in some instances, the system 100 further includes a client computing device 102 communicatively coupled over the network 180. The client computing device 102 can be any type of computing device, such as, for example, a personal computing device (e.g., laptop or desktop), a mobile computing device (e.g., smartphone or tablet), a server computing device, or any other type of computing device. The client computing device 102 includes one or more processors 112 and a memory 114. The memory 114 can store data 116 and instructions 118 which are executed by the processor 112 to cause the client computing device 102 to perform operations. For example, the instructions 118 can include instructions associated with a web browser application.

In some instances, the client computing device 102 can receive user input or instructions that identify a particular brand to be assessed and/or a geographic area throughout which the brand is to be assessed. In response to receipt of the user input, the client computing device 102 can request or otherwise cause the machine learning computing system 130 to obtain imagery from the image database 122 and input the retrieved imagery into a model 140 to obtain information about locations at which such brand is advertised. In other instances, the client computing device 102 can obtain the imagery from the imagery platform 120 and then provide it to the machine learning system 130. For example, the obtained imagery can be selected by the user of the client computing device 102 or can otherwise be identified based on user input.

The machine learning computing system 130 can further include a brand visibility scoring engine 142. The brand visibility scoring engine 142 can determine a visibility score for a particular brand based at least in part on the locations at which such brand is advertised, as identified by the one or more machine-learned brand recognition models 140. For example, the brand visibility scoring engine 142 can correlate the locations of the brand observations with typical population densities to calculate an estimated number of impressions per unit time.

In another example, the brand visibility scoring engine 142 can collect or otherwise obtain traffic flow data (or even demographically-annotated traffic flow data) from mobile client computing devices 102 (e.g., by means of a mobile application or device operating system that provides location updates for the device 102 or by means of Wi-Fi or cellular base stations that track attempted connections). The brand visibility scoring engine 142 can overlay or otherwise combine the traffic flow data with the locations of the detected advertisements to produce more accurate measures of impressions, such as histograms of impression count/probability, or even such histograms broken down by demographic segment. The brand visibility scoring engine 142 can also analyze the results to measure the number of impression contributions of individual ads or ad types, to support cost-benefit analysis.

Thus, the client computing device 102 can make use of the machine learning platform 130 as a service for brand advertisement location purposes and/or brand visibility scoring purposes, which may in some instances be referred to as "machine learning as a service."

In some instances, the system 100 further includes a training computing system 150 communicatively coupled over the network 180. The training computing system 150 can be separate from the machine learning computing system 130 or can be a portion of the machine learning computing system 130. The training computing system 150 includes one or more processors 152 and a memory 154. The memory 154 can store data 156 and instructions 158 which are executed by the processor 152 to cause the training computing system 150 to perform operations. In some instances, the training computing system 150 includes or is otherwise implemented by one or more server computing devices.

The training computing system 150 can include a model trainer 160 that trains the machine-learned models 140 stored at the machine learning computing system 130 using various training or learning techniques, such as, for example, backwards propagation. In particular, the model trainer 160 can train a model 140 based on a set of training examples. In some instances, the training examples can be provided or otherwise selected by the client computing device 102 (e.g., from the imagery platform 120). In some instances, the model trainer 160 can train the machine-learned model 140 using imagery that is labelled with known brand identities.

Thus, the imagery platform 120, the training computing system 150, and the machine learning computing system 130 may cooperatively operate to enable the user to select portions of imagery, instruct training of a model 140 based on the selected training imagery, select additional imagery, and then instruct use of the trained model 140 to receive indications of locations at which a particular brand is advertised. The brand visibility scoring engine 142 can use the detected locations to better assess a visibility of brand, which can be useful, for example, in assessing advertising campaign effectiveness and/or strategy.

Figure 2 illustrates a processing pipeline 200 that takes input imagery 204 and outputs a brand visibility score 210 for a particular brand. More particularly, a series of fixed or mobile cameras (e.g., mobile image capture vehicles) can be operated to collect large amounts of imagery in public settings. Such collection of imagery can capture a corpus of views in locations that are frequented by the ad campaigns target audience. The corpus of imagery can be stored in an accessible database of an imagery platform. Some or all of the imagery can be input into a brand recognition neural network 202 as input imagery 204.

The brand recognition neural network 202 can review such input imagery 204 and detect and count all observations of the brand being assessed. For example, the brand recognition neural network 202 can output, for each received input image 204, indications 206 of whether a brand or advertisement is depicted in such image and, further, the identity of the brand. A respective location associated with each image can then be assigned to each detected brand instance.

Once all instances of the brand being assessed have been detected and located through outputs 206, a brand visibility scoring engine 208 can compute a brand visibility score 210 for the brand based on such instances. For example, the brand visibility scoring engine 208 can correlate the locations of the brand observations with typical population densities to calculate an estimated number of impressions per unit time.

In another example, the brand visibility scoring engine 208 can collect or otherwise obtain traffic flow data (or even demographically-annotated traffic flow data) from mobile client computing devices (e.g., by means of a mobile application or device operating system that provides location updates for the device or by means of Wi-Fi or cellular base stations that track attempted connections). The brand visibility scoring engine 208 can overlay or otherwise combine the traffic flow data with the locations of the detected advertisements to produce more accurate measures of impressions, such as histograms of impression count/probability, or even such histograms broken down by demographic segment. The brand visibility scoring engine 208 can also analyze the results to measure the number of impression contributions of individual ads or ad types, to support cost-benefit analysis.

Figures

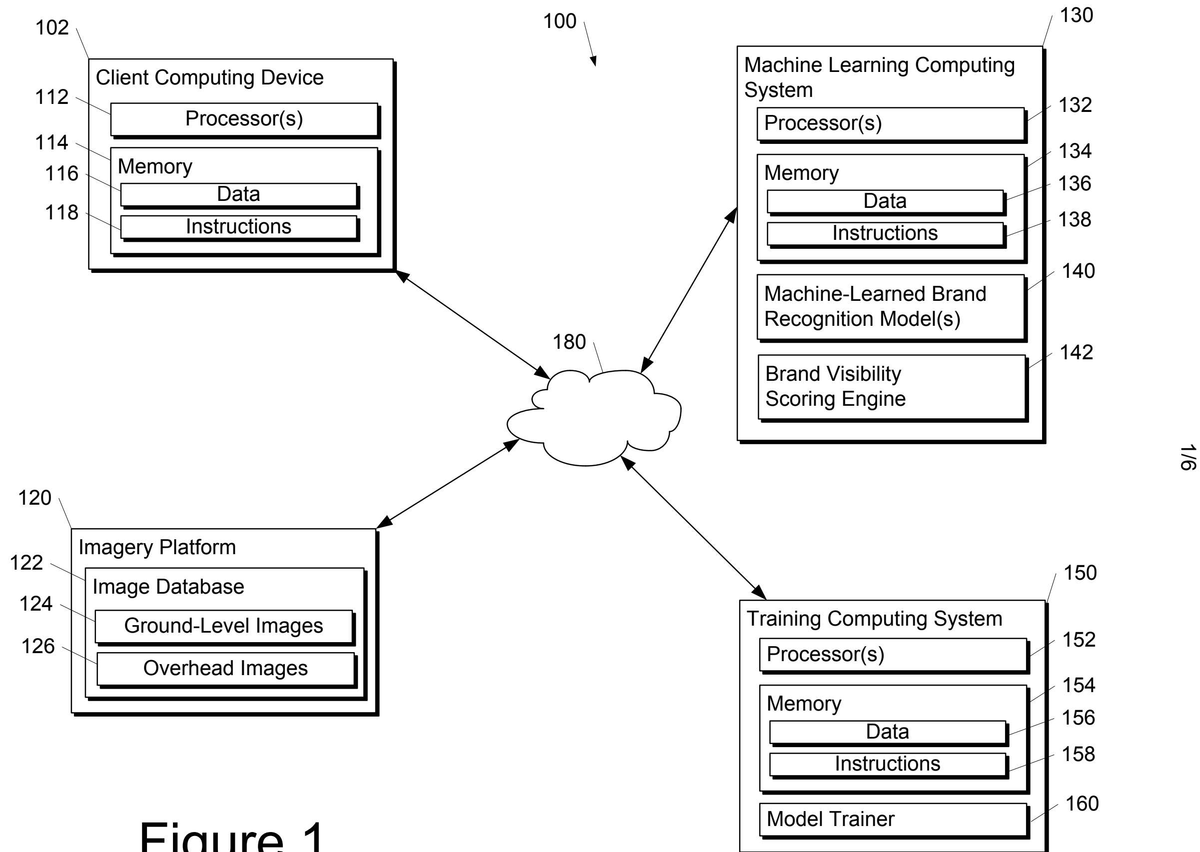

Figure 1

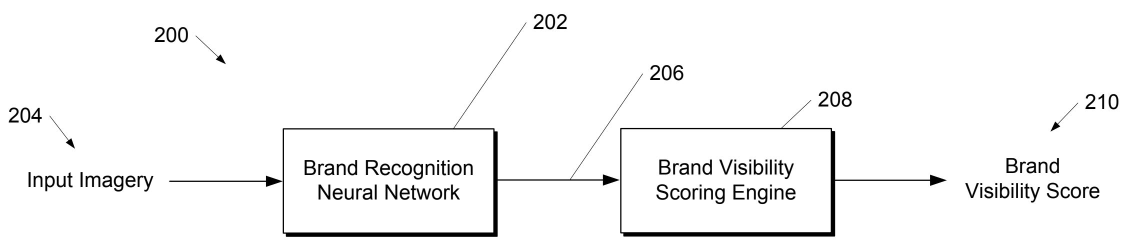

Figure 2

Abstract

The present disclosure describes systems and methods that leverage one or more machine learning techniques or machine-learned models to perform a large scale statistical assessment of the visibility of an advertising campaign. More particularly, a processing pipeline is provided that takes input imagery and outputs an indication of a visibility of a brand. The pipeline can include a brand recognition neural network and a brand visibility scoring engine. Keywords associated with the present disclosure include: machine learning; neural network; deep learning; model; training; image; imagery; advertisement; advertising; ads; campaign, brand; logo; visibility; effectiveness; score; location; traffic; emblem; trademark; label.