

Technical Disclosure Commons

Defensive Publications Series

April 24, 2017

Machine Learning to Automatically Detect Human Development from Satellite Imagery

Matthew Manolides

Follow this and additional works at: http://www.tdcommons.org/dpubs_series

Recommended Citation

Manolides, Matthew, "Machine Learning to Automatically Detect Human Development from Satellite Imagery", Technical Disclosure Commons, (April 24, 2017)
http://www.tdcommons.org/dpubs_series/481



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

MACHINE LEARNING TO AUTOMATICALLY DETECT HUMAN DEVELOPMENT FROM SATELLITE IMAGERY

Introduction

The present disclosure provides systems and methods that identify geographic areas that are represented with outdated imagery and have experienced human development that is not depicted by such outdated imagery. In particular, mapping applications depend on their ability to provide accurate and reliable imagery of geographic areas to users. In order to maintain such up-to-date information, newer and higher resolution geographic imagery must be acquired on a periodic basis to replace the outdated imagery. However, high resolution geographic imagery can be costly to obtain, and it may not be economically feasible to acquire the latest and best imagery for each geographic area. This can lead to a disproportionate focus on small-scale churn within existing urban geographic areas that have a high population density and user base, at the expense of more remote geographic areas. In turn, development in these more remote geographic areas can go unnoticed for extended periods of time, affecting the overall accuracy and reliability of the mapping applications.

Summary

The present disclosure proposes to solve the challenges described above by using lower resolution geographic imagery to identify large-scale human development (also referred to as “churn”) in certain geographic areas (e.g., remote geographic areas). By using low-resolution imagery to identify those places where large-scale churn has occurred, the costly image acquisition process can be focused on areas where churn has occurred, thereby reducing costs associated with unnecessarily acquiring new imagery of locations that have not changed. In addition, the speed and frequency with which churn is identified can be increased. Geographic areas that are identified as containing churn can be flagged, and newer high-resolution imagery

corresponding to the flagged areas can be acquired (e.g., through interaction with a satellite image acquisition scheduling system). Therefore, remote geographic areas with churn can be identified sooner and updated with newer high-resolution satellite imagery that is more representative of the present state.

In one aspect of the present disclosure, data identifying geographic areas that contain churn can be generated using human operators. More particularly, the human operators can be presented with a low-resolution image corresponding to older existing imagery of a geographic area and a newer low-resolution image of the same geographic area. The human operators can then compare the older and newer low-resolution images to determine whether the newer low-resolution image depicts identifiable churn.

One or more geographic areas (e.g., in the form of cells that correspond to certain areas) can be selected for comparison based on one or more attributes corresponding to the geographic area (e.g., date of existing imagery of a geographic area, census data indicating a size and distribution of population and/or economic output within a geographic area, a number of active users within a geographic area, publicly disclosed infrastructure projects within a geographic area, etc.). For example, geographic areas that have the relatively oldest imagery can be selected for analysis. Once a geographic area is selected, a low-resolution image of the geographic area can be acquired. For example, existing imagery of the selected geographic area can be converted into a desired low-resolution format. As another example, a low-resolution image of the selected geographic area, that corresponds to the date of existing imagery, can be acquired through interaction with a database that stores previously captured imagery from different time periods (e.g., historical satellite imagery). Additionally, a newer low-resolution image of the selected geographic area can be acquired (e.g., from a database that stores recently acquired low

resolution images). For example, the older low-resolution image and/or the newer low-resolution image can be images that were captured by the Landsat program operated by the U.S. Geological Survey.

The older and newer low-resolution images of a selected geographic area can be presented to a human operator side-by-side for comparison. In some implementations, the low-resolution images can be combined to generate an animation (e.g., animated GIF, MPEG, etc.). The animation can transition between the older low-resolution image and the newer low-resolution image of the same geographic area. Generating an animation that transitions between the low-resolution images can allow a human operator to more easily detect specific regions of churn within the selected geographic area. In some implementations, the low-resolution images can be presented one at a time such that a human operator can easily switch between the images (e.g., by pressing the arrow keys).

A human operator can be tasked to determine whether a geographic area contains identifiable churn based on a comparison between older and newer low-resolution images of the geographic area. In some implementations, a human operator can be tasked to select one of several options for each selected geographic area. For example, the human operator can be asked if there has been any visible man-made changes in the imagery, and to select: a) No, b) Yes, a small amount, c) Yes, a medium amount, or d) Yes, a large amount. In some implementations, a human operator can be asked to ignore some types of changes. For example, the human operator can be asked to ignore crop changes, vegetation change, snow/ice, clouds, water levels, etc. In some implementations, more than one human operator can be tasked to review each selected geographic area.

In another aspect of the present disclosure, 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 compare older and newer low-resolution images of a geographic area to determine whether the geographic area contains identifiable churn. In some implementations, the machine-learned models can be trained with an initial set of data identifying a plurality of geographic areas as either containing or not containing identifiable churn based on a comparison between older and newer low-resolution images of each geographic area. In some implementations, the initial set of data can be generated by human operators as described above. Thus, a machine-learned model can be trained based on a training dataset generated by the human operator techniques described above. Once trained, the machine-learned model can receive two or more images from different time periods and output an indication of whether (and/or how much) human development has occurred in the depicted geographic area in the time period between the respective capture dates of the two or more images.

Detailed Description

Figure 1 depicts an example computing system 100 to identify churn 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. For example, the image database 122 can store existing geographic images 124 and newer geographic images 126. The existing geographic images 124 can include high-resolution imagery of one or more geographic areas. The existing geographic images 124 can also include low-resolution imagery of one or more geographic areas. The newer geographic images 126 can include newer low-resolution imagery of selected geographic areas.

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 churn 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, convolutional neural networks, etc.) or other classifier models (e.g., support vector machines, random forest classifiers, regression models, etc.) including both linear models and 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 compare an older low-resolution image of a selected geographic area with a newer low-resolution image of the same geographic area, and determine if the geographic area contains identifiable churn. 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 identify which of the corresponding geographic areas contain churn.

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 mapping application.

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 generated by human operators. For example, the training computing system 150 can access images from the imagery platform 120 to present an older low-

resolution image of a selected geographic area and a newer low-resolution image of the same geographic area to the human operator. The human operator can determine if the geographic area contains churn by comparing the images. The determination by the human operator can be stored with the corresponding images in the image database 122. In some instances, the model trainer 160 can train the machine-learned model 140 using imagery from the imagery platform 120 that has been previously identified as depicting churn and/or imagery identified as not depicting churn.

Thus, the imagery platform 120, the training computing system 150, and the machine learning computing system 130 may cooperatively operate to identify geographic areas that contain churn.

Figure 2 illustrates a processing pipeline 200 that takes input imagery 204 and outputs a churn determination 206 for a particular geographic area. More particularly, a low-resolution image corresponding to older existing imagery of a geographic area and a newer low-resolution image of the same geographic area can be input into churn recognition neural network 202 as input imagery 204.

The churn recognition neural network 202 can review such input imagery 204 and identify if the selected geographic area contains churn. For example, the churn recognition neural network 202 can output, for each received pair of low-resolution input images 204, a churn determination 206 of whether the corresponding geographic area contains churn.

Figures

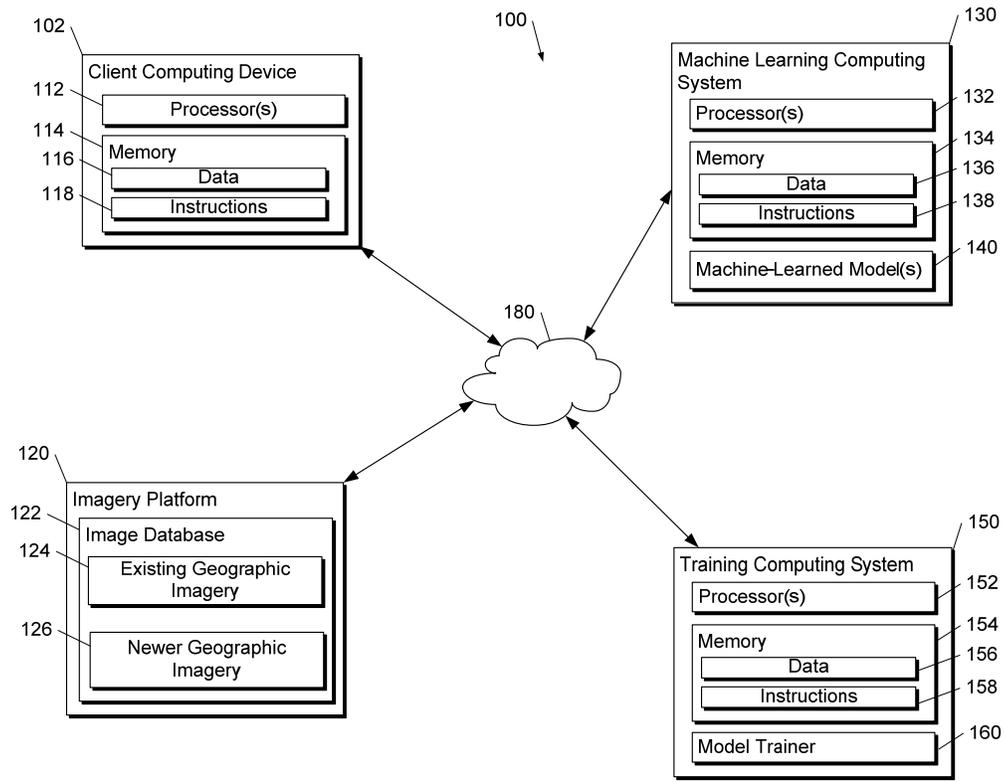


Figure 1

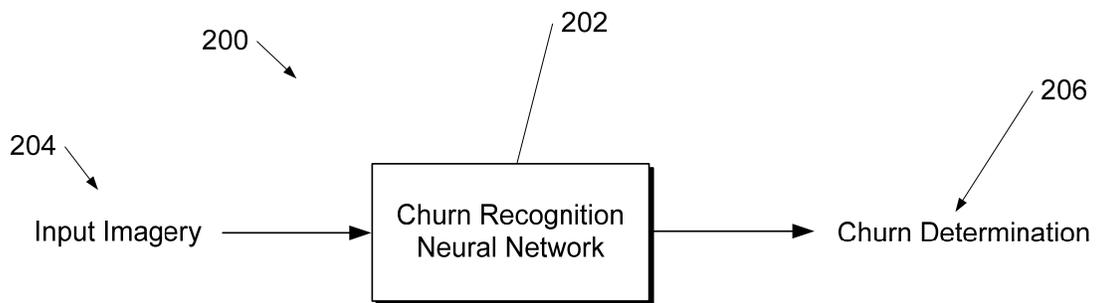


Figure 2

Abstract

The present disclosure describes systems and methods that leverage one or more machine learning techniques or machine-learned models to identify large-scale churn using low-resolution geographic imagery. More particularly, a processing pipeline is provided that takes as input a low-resolution image corresponding to older existing imagery of a geographic area and a newer low-resolution image of the same geographic area. The pipeline can include a churn recognition neural network that compares the input images to identify churn, and output the result of the identification. Keywords associated with the present disclosure include: machine learning; neural network; deep learning; geographic imagery; satellite imagery; aerial imagery; churn; development; infrastructure; roads; automatic detection of change.