July 2022

Estimating Prevailing Road Conditions by Analyzing Terrain Images from Satellites

Yan Mayster

Brian Shucker

Estimating Prevailing Road Conditions by Analyzing Terrain Images from Satellites

ABSTRACT

The optimal route between destinations at a given time depends on the traffic conditions and speed limits at that time. Detailed and accurate data on traffic conditions and speed limits is sometimes unavailable for a subset of road segments within potential routes. This disclosure describes techniques to determine various parameters about the likely prevailing travel conditions for a given road segment at a given time based on satellite images. The inferred travel conditions can be appropriately factored in when suggesting routing guidance. The satellite images are processed via appropriately trained machine learning models that can detect relevant attributes or physical features. An ensemble model composed of individual trained machine learning models is utilized to analyze the detected attributes and physical features to provide an estimate of road conditions likely to be present at travel time.

KEYWORDS

- Satellite imagery
- Overhead imagery
- Speed limit
- Routing guidance
- Road conditions
- Travel advisory
- Geographic terrain
- Driving directions
- Semantic segmentation
- Instance segmentation

BACKGROUND

Navigation applications utilize maps, geolocation, traffic, and other relevant data to provide optimal routing guidance to users. The optimal route at a given time depends on the traffic conditions and speed limits at that time.

Speed limits can be derived from street imagery or other pertinent geolocation data, especially in parts of the world where speed limit signage is commonplace. However, speed limit information may not be available for streets in many regions of the world. Even in locales where data on speed limits is available, there may be areas in which no data exists for various reasons such as newly constructed road segments, sparse presence of speed limit signs, etc. Moreover, any changes in speed limits result in making the existing data stale until it is updated with the newly set limit.

An alternate approach to obtaining speed limits for a given road segment involves inspecting existing traffic patterns by, with user permission, obtaining movement data from the devices moving through the segment. However, such an approach lacks sufficient accuracy for the purposes of obtaining the legal speed limit since roads are shared by a variety of people, such as vehicle drivers, motorcycle riders, bicyclists, walkers, etc., each traveling at different speeds. As a result, speed limits and travel times obtained via such an approach can be too noisy to be used for navigation and can lead to the derived guidance being inconsistent and confusing.

Further, the manner in which legal speed limits are posted and interpreted can vary depending on the region. For instance, the default speed limits for similar roads can be different in rural, suburban, and urban areas. It is typically assumed that drivers know the types of localities they are driving through and interpret the speed limit guidance accordingly.

Terrain and roads are imaged by the various satellites that orbit the earth. In recent years, the availability, quality, and latency of satellite imagery of the physical terrain has seen significant improvements, making it possible to obtain high resolution images of many parts of the world at frequent intervals, e.g., weekly.

DESCRIPTION

This disclosure describes techniques to determine various parameters about the likely prevailing travel conditions for a given road segment at a given time based on satellite images of the terrain. The inferred travel conditions can be appropriately factored in when providing route guidance, especially for road segments for which the information is unavailable and/or unreliable.

Satellite images of road segments and surrounding terrain can be processed via appropriately trained machine learning models to detect one or more relevant attributes or physical features, such as:

- Width of the road;
- Type of road surface (e.g., paved concrete, gravel, dirt, etc.);
- Major defects in the road surface (e.g., uneven pavement, cracks, potholes, etc.);
- Number and width of lanes marked with lane markers;
- Vehicle traffic;
- Non-automobile objects, such as cattle, bicyclists, pedestrians, etc.;
- Type of built structures (e.g., homes, places of religious worship, schools, shops, factories, office buildings, etc.) present in the vicinity of the road segment;
- Density of built structures in the vicinity of the road segment, etc.

For instance, structures such as schools, places of worship, hospitals, farms, parks, sports fields, factories, etc. can be distinguished because of their distinctive shapes relative to their surroundings as well as various specific distinguishing features such as gates, towers, parking lots, etc.

The detected attributes and physical features can then be input to an ensemble model composed of individual trained machine learning models that process the input to provide an estimate for one or more measures of prevailing road conditions, such as an approximate speed limit for the road segment, likelihood of encountering adverse weather, etc. For example, some locales restrict speeds or types of vehicles based on factors such as road width, presence of directional separators, zoning categorization (e.g., residential, commercial, etc.). Similarly, the amount and patterns of vehicular and pedestrian traffic can be predicted based on the presence and distance of various structures in the vicinity.

By taking into account the times at which the satellite images were captured, the time of the day can be used as a factor when generating the estimates. For instance, the imagery can indicate whether traffic is heavier in one direction in the morning and the opposite direction in the evening, corresponding to typical commuting patterns. Further, nighttime images captured by the satellites can be used to determine whether the road segment is well lit in low light natural conditions, such as nights, storms, etc.

Further, the types of structures in the vicinity of the road segment can be used to estimate road conditions at a given time even when satellite imagery does not cover those specific times. For instance, start and end times for schools and religious services can be used to estimate traffic patterns and permissible speeds in the vicinity of schools and places of worship, respectively.

Navigation guidance can be optimized based on the estimates to provide users with optimal routing options for the estimated road conditions likely to be encountered at the time of travel along a given road segment. Such routing can avoid road segments likely to present undesirable or unacceptable travel conditions. The user can also be provided appropriate advisory messages regarding the road conditions that are likely to be encountered during travel along a given road segment within the suggested routing options.

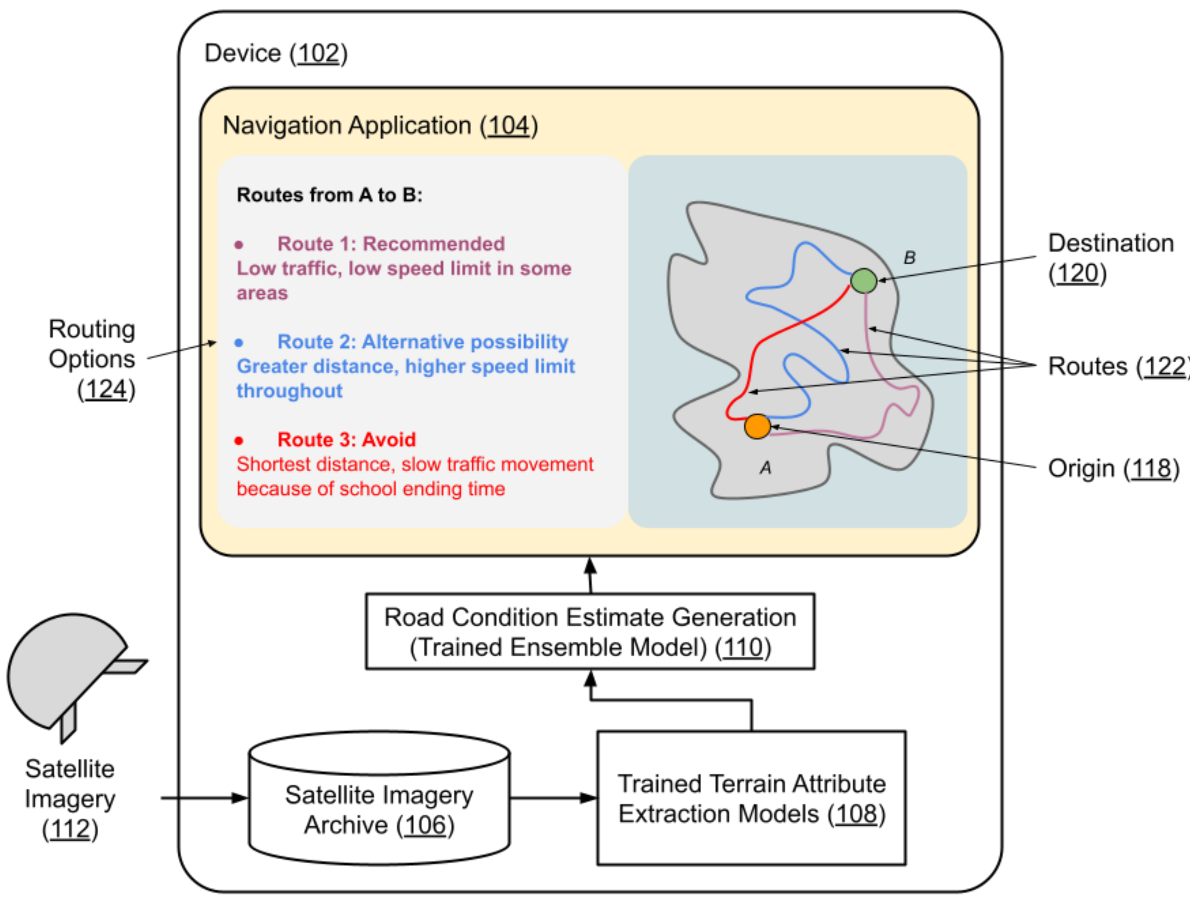

Fig. 1: Generating routing guidance by estimating road conditions from satellite imagery

Fig. 1 shows an example operational implementation of the techniques described in this disclosure. A user seeks routing guidance from a navigation application (104) on a device (102) to go from an origin A (118) to a destination B (120). Routing options (124) with details and route visualization (122) presented to the user are derived from processing images from an archive (106) of satellite images (112). The processing occurs in two steps: in the first step, one or more trained machine learning models are employed to extract attributes of the terrain in the images (108). In the second step, an ensemble model (110) is utilized to estimate prevailing road conditions at the likely time of travel.

As Fig. 1 shows, route 1 is recommended as the optimal option because of shorter distance combined with low estimated traffic, despite having low estimated speed limits for some segments within the route. In contrast, route 3 is estimated to present slow moving traffic since it passes through school zones at a time when the school day ends, thus making it the slowest option despite being the shortest in terms of distance. The practical recommended speed limit in certain cases may be lower than the inferred legal speed limit, e.g., if it is detected that a particular road segment is damaged (e.g., has potholes, cracks, waterlogging, etc.) or is otherwise modified, e.g., due to construction activity, etc. Detection of bad road conditions, e.g., caused by recent weather effects or construction activity, can be taken into account to provide additional driving advisories.

Operational implementations can additionally include mechanisms for the users to provide feedback about the accuracy and usefulness of the estimated road conditions. With user permission, collective user feedback gathered using such mechanisms along with deviations from recommended routes can be utilized to flag erroneous estimates. The errors can be utilized during model training to improve model operation and refine future estimates.

Analyzing changes over time in the archive of captured satellite images can help detect changes in road conditions that occur gradually over long time periods, such as weeks, months, or years. In contrast, sudden changes in road conditions can be detected with the use of Synthetic

Aperture Radar (SAR) imagery of an appropriate resolution. While SAR imagery is not visual like RGB colors, it has the advantage of being useful even when the conditions are cloudy.

A machine learning model employing SAR imagery needs to be carefully designed and trained to ensure that it can distinguish accurately between changes in road conditions and presence and movement of permissible objects (e.g., uncommon vehicles such as a double decker bus). Accurate distinctions between the two types of short-term changes can be achieved by averaging SAR results over multiple images since moving vehicles or other objects are unlikely to result in persistent changes over time in the same pixel. Suspected short-term changes in the road conditions as inferred by analyzing SAR imagery can be flagged as potentially dangerous. Accordingly, suitable driver advisories can be generated and displayed depending on the extent and confidence in the prediction.

The techniques described in this disclosure can be incorporated within any application that includes navigation and mapping capabilities. The satellite image data can be stored and/or processed external to the user device, such as in the cloud.

Implementation of the techniques can improve routing guidance even when high-quality non-satellite data on road parameters is unavailable or inaccurate, thus enhancing the user experience (UX) of map and navigation applications. Moreover, use of the techniques can potentially result in better traffic routing and more optimally distributed utilization of the existing road network.

CONCLUSION

This disclosure describes techniques to determine various parameters about the likely prevailing travel conditions for a given road segment at a given time based on satellite images. The inferred travel conditions can be appropriately factored in when suggesting routing guidance. The satellite images are processed via appropriately trained machine learning models that can detect relevant attributes or physical features. An ensemble model composed of individual trained machine learning models is utilized to analyze the detected attributes and physical features to provide an estimate of road conditions likely to be present at travel time.