The Bengaluru Mobility Challenge, 2024

Call for Participation

Challenge overview

Bengaluru has been ranked the most congested city in India in terms of traffic for several years now. This hackathon is aimed at creating innovative solutions to the traffic management problem in Bengaluru, and is co-sponsored by the Bengaluru Traffic Police, the Centre for Data for Public Good, and the Indian Institute of Science (IISc). The hackathon will be hosted by IEEE DataPort, and will be conducted in two phases, as described below. Teams can participate in either or both phases.

Phase 1: This phase will start in the second week of June and the submissions will be evaluated in the last week of August, 2024. A leader board will be created for Problem 1(a) (see detailed description below), and the top teams will be recognised in a special session at the CyberPhysical Systems Symposium (CyPhySS) 2024, which will be run from July 25-28, 2024 at IISc. The prizes for this phase will be announced in August. However, demos by finalists of this phase, as well as the award ceremony will happen in conjunction with the Symposium of Data for Public Good at IISc in September.

The participants in this phase will be provided with camera feeds from 23 Safe City cameras in northern Bengaluru, around the IISc campus. The task will be to provide short-term (e.g., 30 minutes into the future) predictions of the vehicle counts (by vehicle type) as well as vehicle turning patterns at certain points and junctions of the road network. The predictions may be at different points different from the locations where the camera feeds are available.

Phase 2: In this phase, the participants will be asked to re-identify vehicles seen at some locations of the network at other locations of the network and estimate the origin-destination (O-D) flows for this part of the network for a particular time period. The O-D flow estimates are critical for transportation planning, what-if analysis, etc. This phase will conclude with demos by finalists and announcement of winners on September 20, 2024 in conjunction with the Symposium of Data for Public Good at IISc.

The hackathon will be conducted by an organising committee with members from the Centre of Data for Public Good (CDPG), FSID, IISc, and the Centre for infrastructure, Sustainable Transportation and Urban Planning (CiSTUP), IISc, with help and advice from the Bengaluru Traffic Police, and other industry partners. The participants may interact with the organising committee via IEEE DataPort (by using the messaging service or the comment section). The organising committee may also use an online forum to address any questions or concerns they might have regarding the detailed description of the hackathon and the rules and regulations.

ENTRY IN THIS COMPETITION ASSUMES YOUR ACCEPTANCE OF THE FOLLOWING OFFICIAL RULES.

Eligibility

Each team may have up to five members. An individual can be a member of only one team in each phase. All members of a team need to be residents of India. If the team is selected as a finalist in a phase, at least one of the team members must be physically present at the conference venue to demonstrate their solution.

Rules of Participation

All teams wishing to participate must sign up for the competition on IEEE DataPort by filling the form at the bottom of the web page and requesting access. Each team is allowed to create only one account. Privately sharing code or data outside of teams is not permitted. There are also restrictions on data access and usage as described in detail below. Only one submission per team is allowed. If a team makes multiple submissions before the deadline, the last one will be considered. The same team cannot make submissions from multiple accounts. Team mergers are allowed before the submission deadlines, as long as the total number of members in the merged team does not exceed five.

Data Access and Use

The data provided will be only for use within the competition. The participants give an undertaking that the data will not be used for any purpose other than the hackathon. The participants also must not try and reveal any personally identifiable information present in the videos. In addition, participants must not place this data in the public domain.

Source Code and Publications

The participating teams can make use of any open-source code or software in building their solutions, provided such use is permitted by the owners of the open source code. The participants are free to publish their algorithms, methodologies and other findings in the form of articles. However, they must clearly acknowledge the hackathon in their publications.

All participating teams will be expected to make their source code available to the evaluation committee for scoring purposes. We expect that the algorithms and insights that are generated out of this hackathon will be applicable at a larger scale for the entire city of Bengaluru, and with enhancements and inclusion of other data sets, will help agencies managing traffic build a reliable traffic model for Bengaluru and other cities in India. Therefore, in the interest of the larger public good, we expect the finalists of Phase 1 and Phase 2 to make their code open source at the end of the competition (i.e., after the prizes are announced for Phase 2), under Apache 2.0 licence. However, they must maintain the video data that is shared as private.

Prospective participants (teams) can either register for both phases or just one phase. At the end of each phase, five teams will be selected as finalists. All finalists will receive a nominal award (see below), The teams participating in Phase 1 will have the option to continue to participate in the second phase of the hackathon, or stop after Phase 1. If the finalists of Phase 1 decide not to participate in Phase 2, they would need to agree to make their code available under Apache 2.0 licence to participants of Phase 2. In this way, the participants of Phase 2, if they so choose, can build on the ideas from Phase 1. If the participants make any further enhancements to their code beyond the competition, they need not make the new versions available in open source.

Submission Guidelines

Submissions should be made as a zip file containing the code, final trained models and charts containing the sample results. Some of the **mandatory** components of your submission are indicated below. More details of each of these components can be found later in this section.

1. All code and notebooks
2. All created/stored/used models
3. A README.md file containing instructions for execution and other information.
4. A team_name_report.pdf document containing details such as the members of the team, the approach used, etc.

Your submission must **not** contain -

1. Input videos in the zip.
2. Modified names or folder structure of the input videos.
3. Processed video inputs (for e.g., compressed videos, etc.). If preprocessing was done on the videos, then a notebook containing the exact procedure must be provided.

Code and notebooks must mandatorily follow the following guidelines:

1. Break your submission into numerous notebooks (or scripts) serving one large functionality. For example, have a notebook just to perform preprocessing, one notebook just for training, one notebook for evaluation, etc.
2. Make the code modular (for e.g. by using Kedro), have neat and understandable function names and utilise object-oriented abstractions to make the code readable and modular.
3. Leave sufficient comments clearly explaining the functionality of these modular blocks.
4. Clearly indicate the inputs and outputs and demarcate an area of code for handling inputs and outputs
5. Use a fixed seed wherever necessary to make the results reproducible.

The created/used/stored models

1. All components which are using neural networks or machine learning must have relevant models (for e.g., yolo, svm, torch, etc.) and code and explanation for the same must be provided.
2. Any model used must comply with the supported model formats for the networks used, e.g., TensorFlow uses pickled or .h5 format, Torch uses .pt format, Yolo uses .pt or .onnx, etc.

The README.md file should contain the following information:

1. Names and descriptions of the various notebooks and/or scripts used.
2. Clear instructions on how to run your code, what inputs to provide, etc.
3. Clear references to open-source models used, for e.g, if a huggingface model is used.
4. System requirements to run your code such as required GPU cores, RAM, CPU, etc. This needs to be consistent with the specs of the workstation that will be used for evaluation. Please see the section entitled “Evaluation.”

The pdf document (named team_name_report.pdf) should contain the following information -

1. Details of the members of the team
2. Introduction to the problem statement (summarising your understanding of the problem).
3. Methodology - showing a clear solution architecture block diagram, describing the solution in detail.

4. Results - clearly showing the performance of the solution (e.g., charts of showing the learning metrics over time, etc.) as well as prediction results on validation and test datasets.
5. Conclusions, including any shortcomings of the proposed solution
6. References - to relevant literature, software used, etc.

Evaluation

Submissions will be evaluated by a panel of judges consisting of experts from the fields of transportation and data science, coming from academia, industry, and the Bengaluru Traffic Police. The scoring criteria will include

1. Accuracy of the predictions or estimates in comparison with ground truth or other independently obtained statistics that are available with the organisers,
2. Novelty of the approach as determined by the panel based on the submitted code, documentation, and other descriptions
3. The hardware and computational requirements (i.e., cost and ease of deployment in practice), determined by running the solution on a standard platform. For a fair comparison, evaluations will be conducted on a workstation with the following specifications.
    i. CPU - Core i9
    ii. GPU - RTX4090
    iii. RAM - 64GB
    iv. SATA - 500GB

    It is imperative that all submissions are able to run within the above specs.
4. The final presentation and demo.

The weightage for each will be announced ahead of time. In all cases, the decision of the panel of judges will be considered final.

Prizes

Phase 1:

First prize: Rs. 2.5 lakhs in cash.

Second prize: Rs. 1.5 lakhs in cash.

Special recognition prize Rs. 25,000 to remaining finalists.

(These prizes are subject to any applicable mandatory tax deductions.)

Phase 2:

First prize: Rs. 4.0 lakhs in cash.

Second prize: Rs. 2.5 lakhs in cash.

Special recognition prize Rs. 25,000 to remaining finalists.

(These prizes are subject to any applicable mandatory tax deductions.)

Important Dates

Registration deadline for teams to participate in Phase 1	**June 30, 2024**
Final submission of entries for Phase 1	**August 19, 2024**
Award announcement of Phase 1	**August 26, 2024**
Registration of teams to participate in Phase 2	**August 19, 2024**
Final submission of entries for Phase 2	**September 9, 2024**
Award announcement of Phase 2	**September 20, 2024**

Detailed Problem Description

The participating teams will be provided video clips from surveillance cameras mounted at various locations in a road network around IISc campus. The details of the data provided to the participating teams are described in the section entitled Detailed "Description of the Training Data." The participating teams are expected to use these video clips for designing and implementing their solutions for the problems listed below. For scoring and ranking purposes, the submitted solutions will be tested on video clips that are previously unseen by the participating teams. The scoring criteria are described in the section entitled "Evaluation." The detailed problem descriptions for Phase 1 and Phase 2 are given below.

Phase I:

For Phase I, the participating teams are required to solve two problems.

Problem 1: The goal is to count the vehicles of different classes that pass through the camera view and predict the counts for the future. Seven vehicle classes are considered, namely, **'Cars'**, **'Bus', 'Truck', 'Three-Wheeler'**, **'Two-Wheeler', 'LCV',** and **'Bicycle'**. (See section entitled "Detailed Description of Training Data" for more details).

At mid-block sections only counts in two directions are to be estimated and predicted. However, at junctions, the turning counts between all approaches visible in the video must be estimated and predicted. For a given video clip from a camera with a known location, the outputs of the solutions submitted by the participants should include:

(a) The **cumulative observed number of vehicles by vehicle class** from the start of the clip to the end of the clip.
(b) The **predicted cumulative number of vehicles by vehicle class** from the end time of the clip up to 30 minutes into the future.

For mid-block sections, the output should include the number of observed vehicles in two directions. However, at junctions, the output should include the number of vehicles that turn between each pair of approaches within the view of the camera. The submitted solutions will be tested by the panel of judges on 30-minute clips from the same set of cameras from different dates than the ones in the training data and the locations can be either at mid-block sections or at junctions/intersections.

Problem 2: The goal in this problem is to predict the counts (turning counts in the case of junctions) by vehicle class at "unseen junctions" and "unseen mid-block locations" (i.e., junctions and mid-block locations where no camera feed is provided in the training data) for a period equal to the 30 min window of the given clips + 30 minutes into the future. The locations of the "unseen junctions" and "unseen mid-block locations" are given in the section entitled "Detailed description of the training data." The output of the solution should include:

(a) The **predicted number of vehicles by vehicle class** at the previously unseen location from the starting time to the ending time of the 30-minute clips provided.
(b) The **predicted cumulative number of vehicles by vehicle class** at the previously unseen location from the end time of the 30-minute clips provided up to 30 minutes into the future.

For mid-block sections only counts in two directions are to be predicted. However, at junctions, the turning counts between all approaches must be predicted. A subset of unseen locations (see the section entitled "Detailed Description of Training Data" for more details) will be selected for evaluation of the submitted solutions. Therefore, it is important that the participating teams incorporate all of these locations in their traffic model. A sample image frame from each of the "unseen locations" will be provided to the participating teams before the submission deadline in case any view specific parameters need to be specified.

The submitted solutions will be tested by the panel of judges on 30-minute clips from cameras unseen in the training data. The clips will be from different dates than the ones in the training data and the locations of the unseen cameras can be either at mid-block sections or at junctions/intersections.

Phase II:

In Phase II, teams will work on the problem of re-identifying a vehicle across multiple cameras. Given approximately synchronised 30-minute video clips from a subset of the 23 cameras used the training data, your objective will be to count the number of vehicles that reappear in another camera after initially appearing in one of them. For example, consider the second row and third column of the following matrix, where the entry is highlighted in red. You will have to find the number of vehicles that appear in the feed from Camera 3 after they have appeared

in the feed from Camera 2. Multiple appearances of the same vehicle can be counted as a unique re-identification.

		Reappeared in			
Originally appeared in		Cam 1	Cam 2	Cam 3	…
	Cam 1				
	Cam 2			4	
	Cam 3				
	…				

Specifically, given a subset of the cameras, the output of the solution for Phase II should include:

(a) The cumulative numbers of vehicles (by vehicle class) that are re-identified as explained above, in the format of a matrix, from the start of the clip to the end of the clip.
(b) For each reidentified vehicle counted in the matrix, the corresponding image frames of the two cameras where the vehicle appears (shown with the same unique ID in both image frames), along with the frame number from the start of the clip.

The submitted solutions will be tested by the panel of judges on 30-minute video clips from a subset of the cameras. However, the clips will be from different dates than the ones in the training data and some of the cameras may not be from the set of cameras used in the training dataset.

Detailed Description of the Training Data

The data for the hackathon comprises video segments from different surveillance cameras that are part of the Bengaluru Safe City Project of Bangalore Police. There are over 7500 such cameras in Bengaluru and we have selected 23 cameras from the road networks around the IISc campus. The selected loop (highlighted in beige colour) and the camera locations (indicated by indigo markers) are shown in the figure below. The map is also available here.

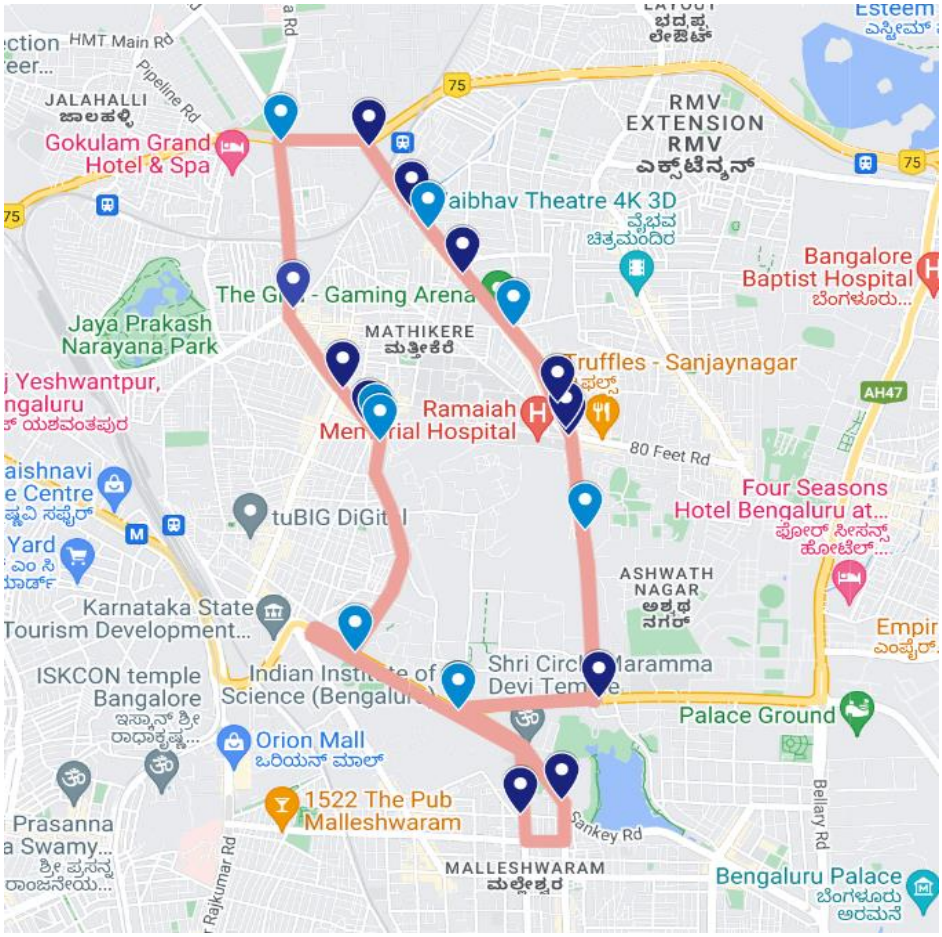

Camera and junction locations: Each indigo marker on the map in the figure represents a location where the camera is fixed. Using these cameras, much of the traffic flow around it can be monitored. In addition to the indigo markers, there are blue markers which represent other important junctions (from a traffic perspective). Camera feeds from these junctions are not available. However, the participating teams need to incorporate these junctions in their model so that vehicle counts and turning patterns can be predicted at these locations. The details of the cameras and the “unseen locations” are available in csv format in a folder called “Locations” in the “competition data files” section.

Video clips: Video streams from each of the cameras have been captured from 7:30 AM to 10:30 AM during the period May 18, 2024 - May 31, 2024. The videos are in the form of 15-minute segments from 7:30 to 10:30 AM. The videos are labelled according to their camera name which can be found in the table below, along with the timestamp (YYYY-MM-DDTHH:MM:SS_X), where X denotes the index of the 15-minute segments, e.g., Mattikere_JN_HD_1_time_2024-05-20T07:30:02_002.mp4 denotes the video segment from the camera called Mattikere_JN_HD_1 on May 20, 2024, from 7:45-8:00 AM. The videos are of resolution 1280x720 and can be accessed here. Some cameras were not functioning during this time period, and hence some video segments will be missing. This reflects a practical real-world situation. The participating teams are free to use all the data or a subset of the data for their work. The videos are available in a folder called “Videos” in the “competition dataset files” section. For easier download, zipped files of the same videos can be found under the folder called “Videos_Zipped.” Each zipped file contains video clips from all cameras on a particular date, and the date is reflected in the name of the file.

Camera orientations: PDF files depicting the orientations of the cameras can be found in a folder called “Camera Views” in the “competition data files” section. The orientation files can be related to the table below using the SITE ID column. For example, “478” in the filename (CENTRAL-478) STATION JUNCTION (PS-SADASHIVANAGARA).pdf is the Site ID for Stn_HD_1. The participating teams may also make use of Google Street View to gain a better understanding of the location where the cameras are installed.

Vehicle classes: There are 7 classes of vehicles to be considered for prediction. These classes are: **‘Car’** (including ‘Hatchback’, ‘Sedan’, ‘SUV or Sport Utility Vehicle’, and ‘MUV or Multi Utility Vehicle’), **‘Bus’, ‘Truck’, ‘Three-Wheeler’**, **‘Two-Wheeler’, ‘LCV’** (including 'Mini-Bus’, ‘Mini-Truck’, ‘Tempo-Traveller’) and **‘Bicycle’**. Example images of different classes of vehicles can be found in a folder called “Vehicle Classes” in the “competition data files” section.

The camera details (indicated by indigo markers in the map above) can be found in csv format in a folder called “Locations” in the “competition data files” section. The “unseen locations” (indicated by blue markers in the map above), which need to be considered for prediction, can also be found in the same folder.