NEXT GENERATION INTERNET

Open Call 4

FedIntersect: EU-US federated testbed for cross-atlantic experiments for urban smart intersections

Deliverable 3: Experiment Results and Final Report

Main Authors	Sofia Kleisarchaki (Kentyou), Levent Gürgen (Kentyou), Thomas Calmant (Kentyou), Mahshid Ghasemi (U. Columbia), Jiawei Lu (U. Columbia), Owen Nuckolls (Columbia U.)
Due Date	30 Septembre 2022
Submission Date	9 December 2022
Keywords	Smart intersection, AI, Digital twin, Collision prevention

Deliverable 3: Part I

Analysis, results, and wider impact

1 Abstract

This deliverable provides an overall view of the work conducted during the entire period of the project and focuses on the main results of the project at the second half of the project. It is an update of the Deliverable 2 provided at the mid-term of the project. The main updates include replication of the solution in Santander city, addition of a collision detection mechanism and detection of traffic lights situation. This deliverable presents the realization of our previously proposed architecture, the integration of relevant data sources and the visual demonstration of two smart intersections, in NY and Santander. The digital twin represents the status of the intersection at real-time, provides historical statistics while preserving privacy. It also enables the traffic flow, traffic lights and crowdedness monitoring, while it raises warnings to prevent potential collisions. The replication of this work in two cities indicated the portability of the federating platforms across borders.

2 Project Vision

The objective of this project is to solve an important problem of urban environments; traffic management and road accidents involving pedestrians. In EU and USA, about half of the road accidents occur in intersections, many of them ending up with fatalities. Furthermore, intersections have an important role in the overall traffic flow management in an urban context, thus with economic and environmental consequences.

The project examines the role of intersections in the overall traffic and pedestrians safety. Truly being in line with the NGI Atlantic vision this project federates three important cross-Atlantic platforms, each one allowing US and European researchers respectively to validate their research and innovation; NSF PAWR COSMOS testbed in New York City in the US, Kentyou's IoT platform from Grenoble, France and Smart Santander testbed in Santander, Spain, Europe.

The specific goal of this project is to create a common access to federated COSMOS and Smart Santander smart intersection testbeds under a unified API and data model. A modular and service-oriented approach is being followed, allowing easy access to data (raw and processed) and Machine Learning (ML) and AI services that will be running on the edge. External researchers/innovators will be engaged, who will use the federated testbed and build at least one experiment each.

3 Details on participants (both EU and US)

Kentyou, France, EU

Dr Levent GÜRGEN is the founder and CEO of Kentyou. He has 12+ years of experience as R&D project manager in CEA-LETI. His involvement in European collaborative projects has been awarded in 2016 by the French Ministry of Research and Education as one of the 12 "Stars of Europe". He is also the president of the Urban Technology Alliance, a global smart city initiative. **He is in charge of overseeing the experiment and managing the team.**

Timothy WARD is the CTO of Kentyou. He is also in the board of the directors and chair of the IoT Expert Group in OSGi Alliance, the core technology used by Kentyou. **He is in charge of overseeing the overall experiment and contribute to technical developments**.

Dr Sofia KLEISARCHAKI is a Data Scientist at Kentyou, France. Her research interests include Big Data Analysis, Artificial Intelligence and IoT. **She is the main researcher in charge of the experiment, gathering and analysing data.**

Dr Thomas CALMANT is software architect at Kentyou. He holds a PhD in computer science on autonomous management of secure applications and has more than 5 years of experience as a research engineer at INRIA. **He is in charge of building the connectors with the other frameworks**.

Co-funded by the Horizon 2020
Framework Programme of the European Union

Bertrand COPIGNEAUX oversees operations management at Kentyou. Bertrand has an Engineering Degree (M.Sc.) in Computer Science and has 10+ years of experience in business consultancy and entrepreneurship. **He is in charge of communication, dissemination and exploitation of the experiment results**.

The rest of the team from the other partners are listed below:

University of Cantabria, Spain, EU

- Prof. Luis Muñoz luis@tlmat.unican.es
- Dr. Pablo Soterres psotres@tlmat.unican.es

Columbia University, NY, USA

Professors

- Prof. Xuan Sharon Di sharon.di@columbia.edu
- Prof. Zoran Kostic zk2172@columbia.edu
- Prof. Gil Zussman gil.zussman@columbia.edu
- Prof. Javad Ghaderi jghaderi@columbia.edu

Postdoc, PhD and Master students:

- Mahshid Ghasemi Dehkordi mg4089@columbia.edu
- Owen Nuckolls on2160@columbia.edu
- Jiawei Lu jl5999@columbia.edu
- Igor Kadota ik2496@columbia.edu

Rutgers University, NY, USA

Prof. Ivan Seskar seskar@winlab.rutgers.edu

Co-funded by the Horizon 2020 Framework Programme of the European Union

4 Experiment Results

In this section we briefly summarize the overall results of the entire period of the project duration. Our efforts were focused on:

- Data Integration and Visualization (Section 4.1). A digital twin [1] of the New York's and Santander's intersection was developed, by integrating several data sources (e.g., weather) and analytical results (e.g., traffic) from a street-level view camera. The main purpose of the digital twin is to enable an efficient traffic management and to increase safety.
- Traffic Light Status Detection (Section 4.2). A real-time detection of traffic light status was set in place, in order to enable better traffic management and violations detection.
- Bird's Eye View Video Analysis (Section 4.3). Beyond the street-level video analytic, a dedicated bird's eye view pipeline is deployed in an edge server with the goal to increase the accuracy of traffic/crowd related-information.
- Collision Prevention (Section 4.4). A collision warning between vehicles and pedestrians was developed to monitor risky areas of the intersection. An evaluation of this mechanism was performed (Section 4.5) using synthetic and real datasets.

4.1 Data Integration and Visualization

The first milestone of this project was to create a digital twin of New York's intersection. The objective of the digital twin is to monitor the real-time traffic situation of the intersection and to store these situations for efficient management and further analysis. Providing such a real-time situational awareness of the intersection was achieved by analysing a real-time, street-level view video.

The street-level view video analytic pipeline, deployed in an edge server as part of the COSMOS network, includes multiple modules:

- Object detection and tracking: detects and tracks pedestrians/vehicles. It uses YOLOv4 [2] object detector model and Nvidia DCF-based tracker.
- Calibration: uses multi-area calibration method[3] to convert on-image pixel coordinates to on-ground coordinates.
- Traffic data extractor: analyzes the output of the object detector and tracking module and derives traffic/crowd-related information such as speed and direction.

The obtained information is sent to a central management platform (provided by Kentyou) for further statistical analysis and visualization via Kentyou UI.

Figure 1 shows Kentyou UI which is a visual representation of the digital twin. The top popup window summarizes real-time traffic information, including the current number of vehicles, bicycles, people, average traffic speed, number of accidents this month, as well as weather (temperature, clouds) and traffic violations (speed and red light). The bottom plot shows the number of detected vehicles within the last day, based on stored historical data.

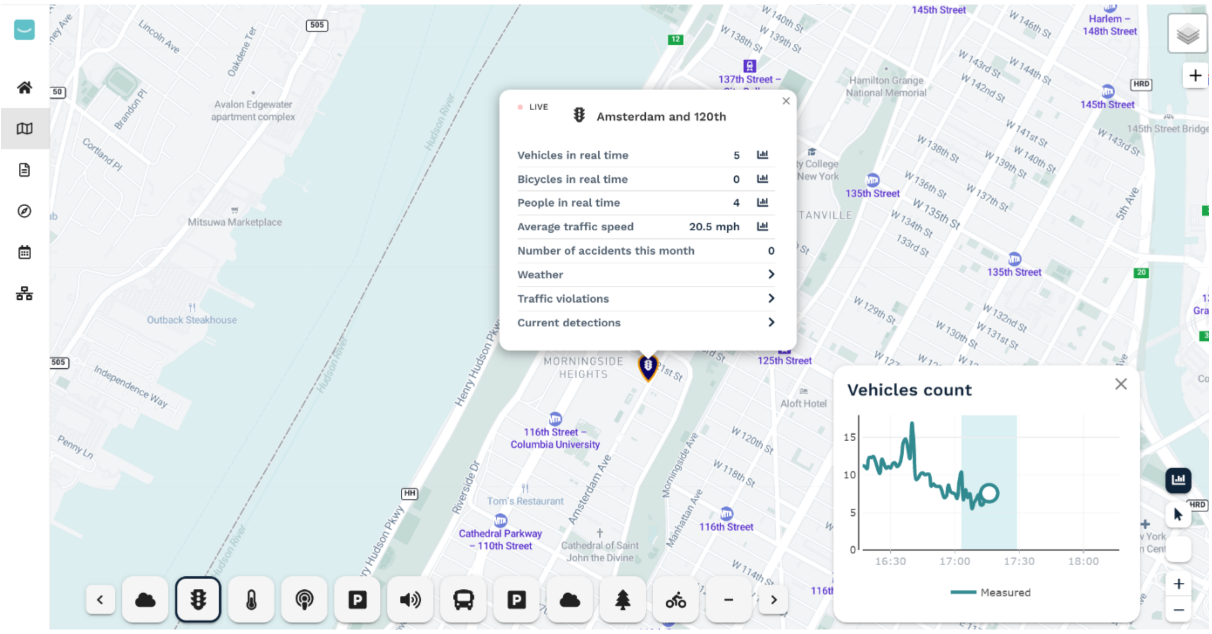

Figure 1 Digital Twin Visualization, New York

Kentyou UI is also providing a real-time representation of the intersection using a dedicated interface (Figure 2) which displays the traffic (detected vehicles, pedestrians, etc) without revealing any private information. The visualization of the traffic includes the detected vehicles (e.g., car, bus, truck), pedestrians, bicycles and the detected status of the traffic lights at real time. More details on the traffic light detection can be found in Section 4.2.

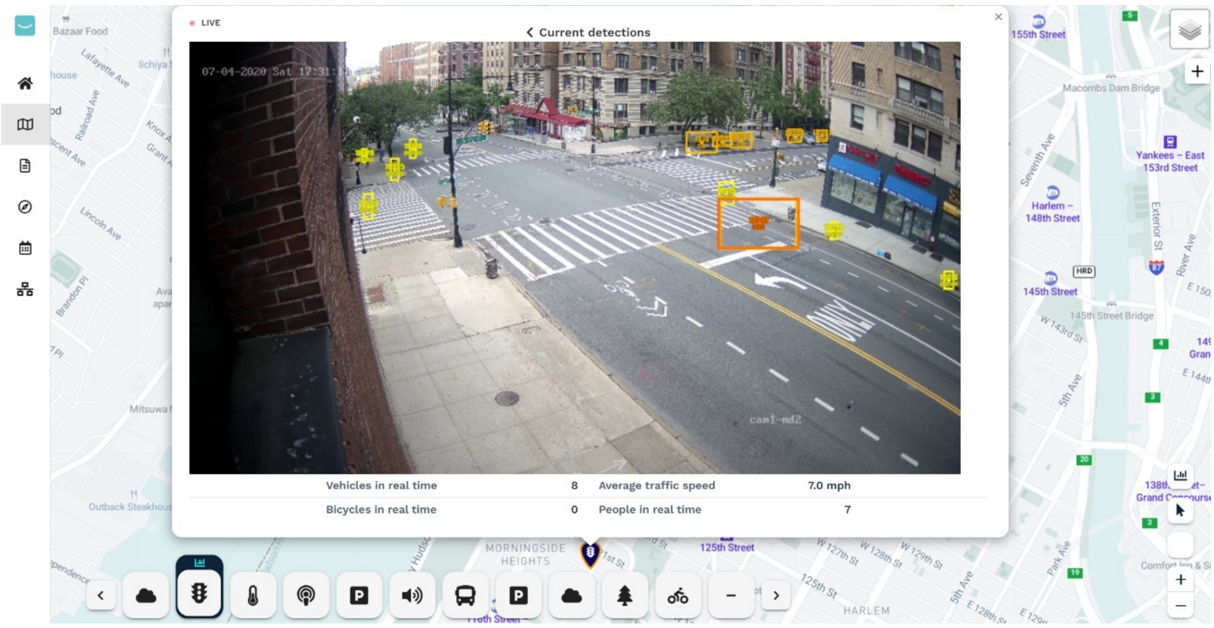

Figure 2 Real-time representation of the intersection, by preserving privacy

Co-funded by the Horizon 2020
Framework Programme of the European Union

A similar visual representation is provided for three intersections at Santander in Spain. Figure 3 shows the locations of these three intersections. The little polygon in front of each mark represents the area which is being monitored by the camera.

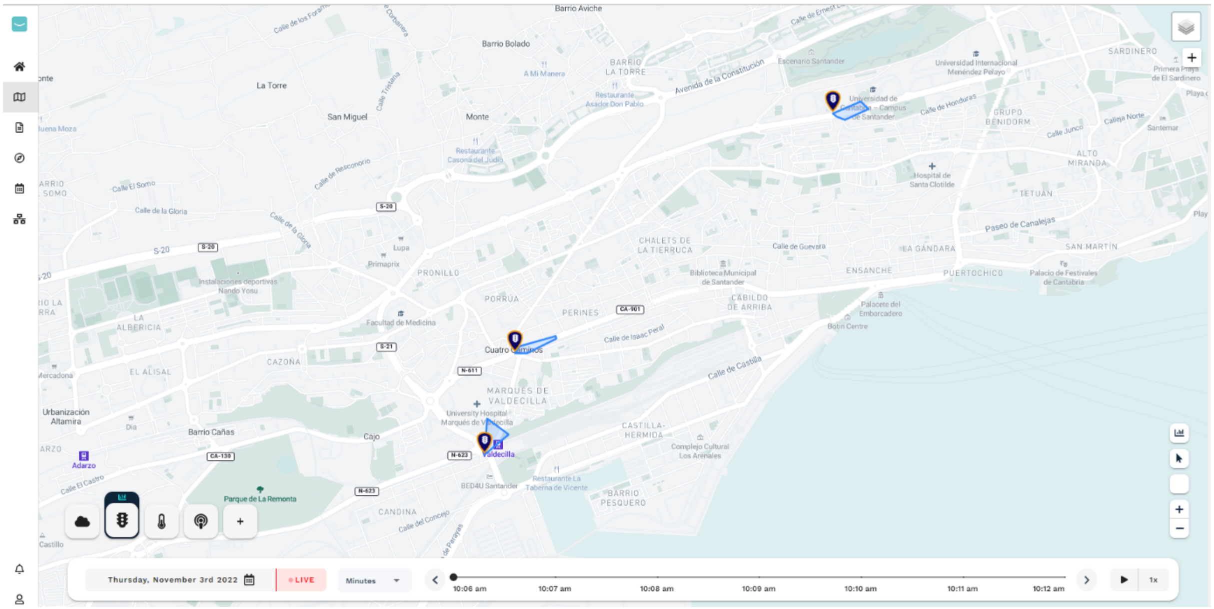

Figure 3 Digital Twin of three Intersections, Santander

Figure 4 illustrates the digital representation of one of these three intersections at Sandander, where the various detected vehicles are being displayed.

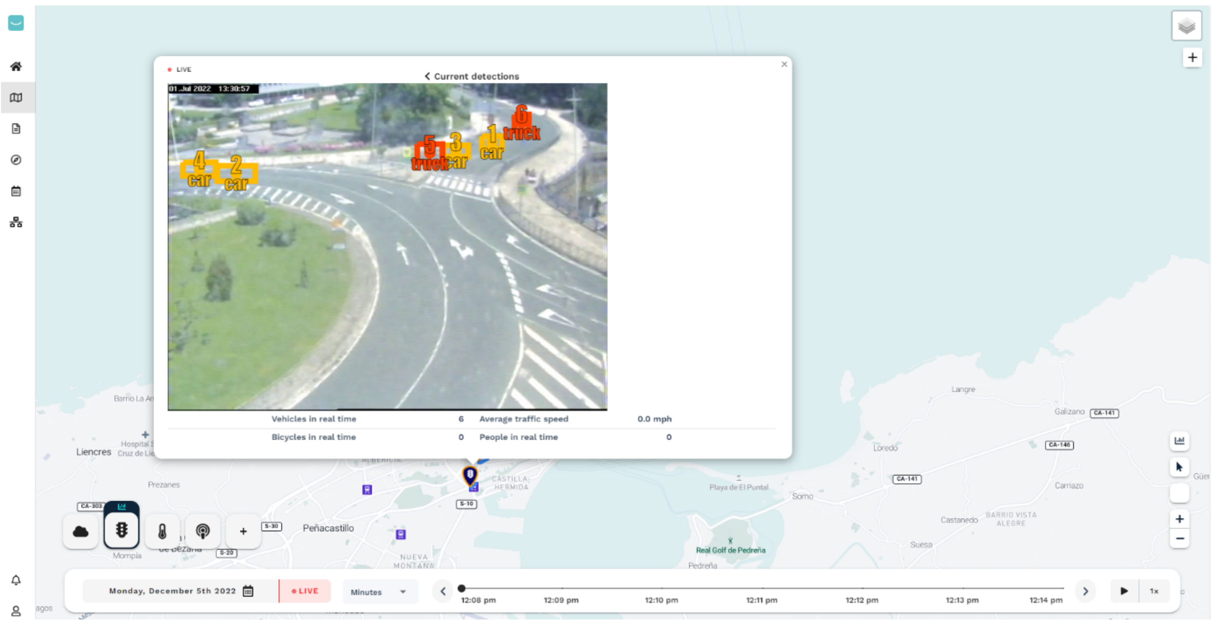

Figure 4 Representation of Valdecilla Sur intersection, Santander

Co-funded by the Horizon 2020
Framework Programme of the European Union

The second milestone of this project was to increase the safety at intersections. The objective of this task is to detect and prevent dangerous situations, such as potential collisions between cars and pedestrians. For this purpose, a collision prevention mechanism (Section 4.3) was developed and evaluated over synthetic and real datasets.

4.2 Traffic Light Status Detection

The research on real-time traffic light status detection with a high accuracy has practical significance and can also potentially have broad prospects for improving road traffic safety. We are using the real-time video stream from the 2nd floor camera for traffic light status detection. Due to the varying levels of luminosity during different times of the day and the sparse feature structure of traffic lights as they occupy less pixels in the fame, the task of status detection becomes challenging. For example, as highlighted in Figure 5, there is a lot of glare around the traffic lights at night, and this can obstruct the view of the entire traffic light.

Figure 5 Video stream frames from day and night highlighting the challenge with traffic light status detection

Since the camera is static, the location of traffic lights can be found offline manually. After locating the traffic lights, we detect their status as follows:

- Apply Top-Hat filter: Top-Hat filter is applied to the gray-scale frame, which helps mitigate the image's uneven illumination.
- Binary threshold: a binary threshold is applied to convert all pixels to either black or white.
- Retrieve the bright pixel percentage rates: the bright pixel percentage rates are obtained for each traffic light color by counting the total number of white pixels in a specific traffic light color and dividing it by the total number of pixels in that traffic light color frame. The light having the highest rate is considered to be the status.

Additionally, the algorithm can also be used to detect the traffic light status from the 1st floor camera, and this can serve as a verification mechanism for the performance of the 2nd floor camera in real-time.

Co-funded by the Horizon 2020 Framework Programme of the European Union

The results of the traffic light status detection are shared with Kentyou at real-time via a dedicated MQTT topic 'ngi/1/traffic_light'.

4.3 Bird's Eye View Video Analysis

This work has been entirely done by the University of Columbia team. The bird's eye view video analytic pipeline, which is deployed in an edge server as part of Columbia University network, also consists of multiple modules:

- Bird's eye view calibration: transforms bird's-eye view videos into calibrated bird's eye videos perpendicular to the ground. It maps a trapezoidal distorted traffic intersection scene into a rectangular one with a uniform scale, which means on-ground distances are obtained from on-image distances by a constant scaling factor.
- Object detection and tracking: a custom-trained YOLOv4 model on a data set recorded from the same camera [2] along with the NVIDIA DCF-based tracker is used for object detection and tracking on the calibrated video frames.
- Traffic data extractor: traffic/crowd-related information such as the position and velocity of pedestrians and vehicles are computed from the data obtained by the object detection and tracking module.

4.3.1 Integration of Street-Level and Bird's Eye View

It is worth mentioning that compared to street-level view, on-ground distance (and hence speed) calculation from a bird's eye view perspective is trivial. Object occlusion rarely happens in a bird's eye view. On the other hand, pedestrians are very small from a bird's eye view, so pedestrian detection from a street-level view is more accurate. To mitigate these issues, we integrate the obtained results from the two cameras (2nd and 12th floor), which results in more accurate traffic/crowd-related information.

- Temporal alignment: find the frames with similar time-stamps from the two cameras for integration.
- Spatial alignment: for accurate on-ground coordinate calculation, we have to use a multi-area calibration method for the 2nd floor camera. Each area is identified by four endpoints. We can find the corresponding endpoints in a bird's-eye view and identify the same area. For each detected pedestrian/vehicle, we determine the area it lies in and its on-ground coordinates with respect to the origin of that area (see Figure 6). Therefore, the objects detected in different views can be aligned spatially.
- Integration: integrate all the detected pedestrians/vehicles detected from both cameras and remove the objects with high occlusion.

Co-funded by the Horizon 2020
Framework Programme of the European Union

Figure 6 Detection areas of 2nd-floor and 12th-floor cameras

4.4 Collision Prevention

The collision prevention mechanism, which was set in place in the previous period of the project and was initially presented in deliverable D2, is improved during the last phase of the project.

The collision prevention mechanism is upgraded in order to simultaneously monitor three different areas of the intersection. These areas, shown in Figure 7 as orange rectangles, are considered the most potentially dangerous ones. Drivers and pedestrians, who are reaching these parts of the intersection, have limited view which leads to a higher risk for an incidence or accident.

Figure 7 Monitored Areas (in orange) for Collision Prevention

Co-funded by the Horizon 2020 Framework Programme of the European Union

Figure 8 illustrates the algorithmic steps of the collision prevention mechanism. Initially, each aforementioned area A is monitored in order to detect any pedestrian P within that area. If a pedestrian is detected, a monitoring of the surrounding area begins and continuous until the pedestrian safely exits the area. During the monitoring of the surrounding area, the algorithm searches for vehicles approaching that area A. If such a vehicle is detected, its distance to the pedestrian is calculated. If the distance is lower than a pre-defined threshold a warning is immediately raised and the rectangle is coloured red.

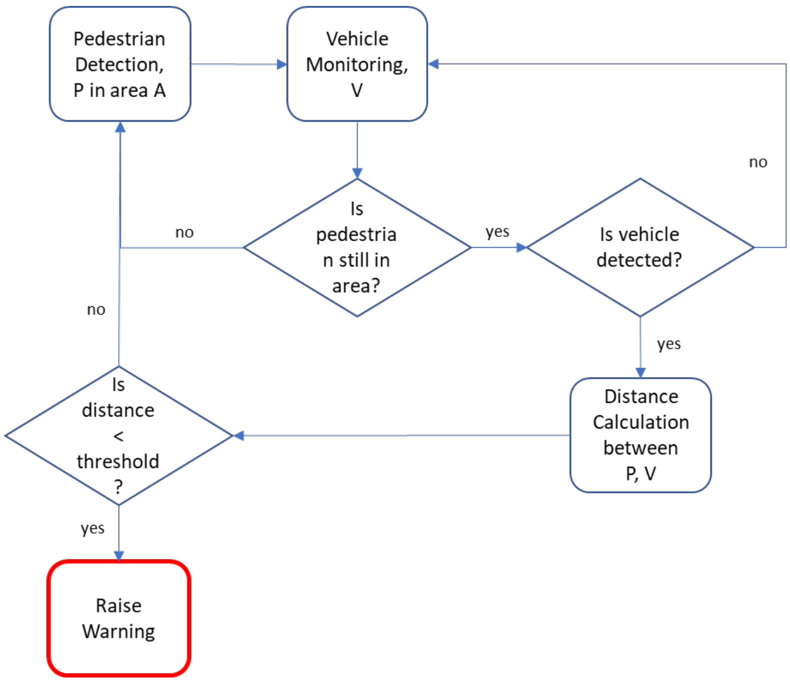

Figure 8 Algorithmic Steps of Collision Prevention Mechanism

The results of the prevention mechanism are thoroughly presented and discussed in the next Section 4.5.

4.5 Discussion and Analysis on Results

In this section we will mainly focus on analysing the results of the collision prevention mechanism. These results support our initial hypothesis that we can increase the safety at the intersection and prevent potential accidents.

4.5.1 Collision Prevention

The collision prevention mechanism was evaluated in a qualitative and quantitative manner using synthetic and real datasets. The next sections present the datasets, as well as the results of this analysis.

Co-funded by the Horizon 2020 Framework Programme of the European Union

4.5.1.1 Datasets

Real dataset

Several datasets were collected, corresponding to different periods of time (2021 and April-June 2022), from the intersection in New York. Given that the project was evolving, the collected datasets have different format and contain slightly different type of information. For instance, the latest datasets track, not only the time of an object's detection, but also the time of when the message was received by Kentyou. The different formats are understandable by the prevention mechanism, as well as by the visual interface, since appropriate parsers have been developed.

It is worth mentioning that the real datasets contain no accidents. To this end, it was difficult to test the sensitivity of our collision prevention algorithm; its ability not to miss any accident. For the purpose of evaluating the misses or false negatives of the algorithm, a synthetic dataset was created.

Synthetic dataset

A synthetic dataset was created containing accidents, occurring at different parts of the intersection. Figure 9 illustrates some of these accidents covering the three predefined monitoring areas.

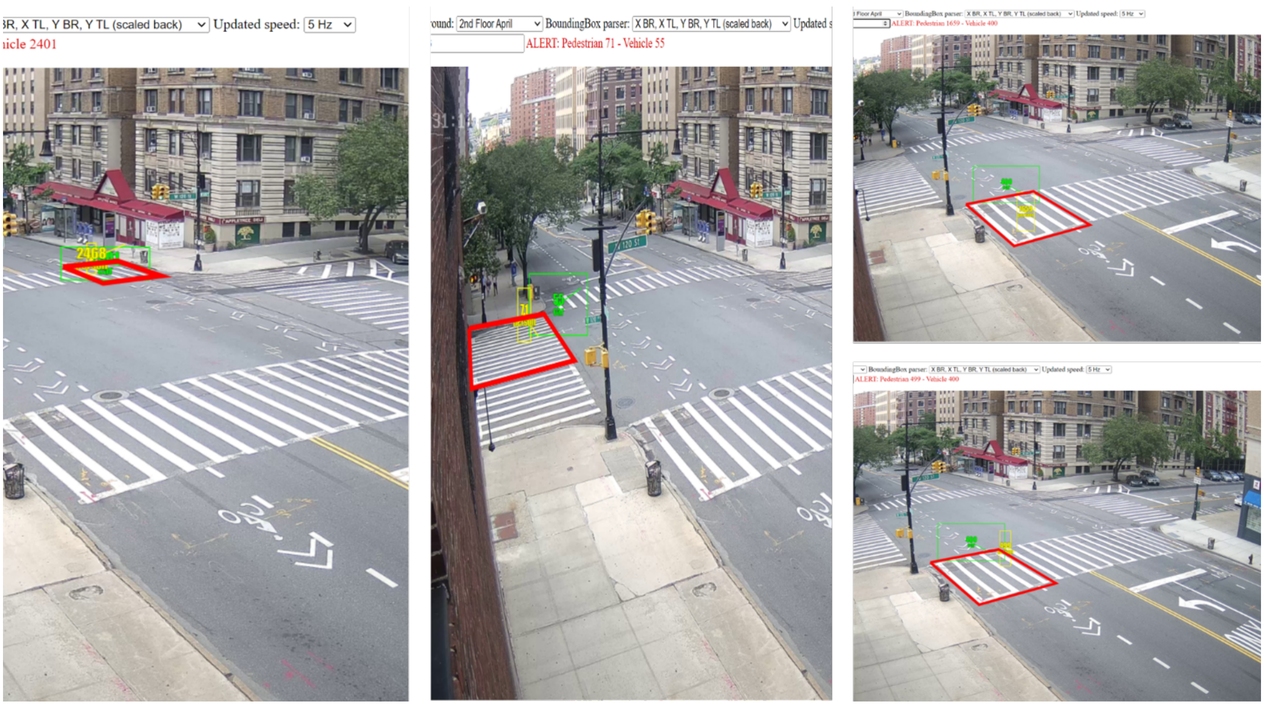

Figure 9 Synthetic Dataset with Collisions

In order to create the synthetic datasets, a two-steps approach was followed. First, trajectories of vehicles and pedestrians, crossing the monitoring areas were identified from

Co-funded by the Horizon 2020
Framework Programme of the European Union

the real datasets. Second, these trajectories were synchronized in order to end up in a collision within the monitoring area. **Figure 9** shows the moment of some of those collisions.

For the purpose of synchronizing the trajectories of a vehicle and a pedestrian in order to collide, we implemented a simple but efficient algorithm. Initially, the algorithm finds the location and moment of collision of two objects (i.e., vehicle, pedestrian). For finding the location, the algorithm calculates the Intersection Over the Union (IoU) of the bounding boxes of the two objects. When the value of IoU is maximized, the two objects have been collided revealing their location and moment of collision. Then, the trajectory before the collision for both objects is synchronized in order the two objects to end up to the same place at the same time. Figure 10 shows the trajectory of two objects, being synchronized in order to collide.

Figure 10 Trajectory Synchronization of Synthetic Data

4.5.1.2 Qualitative and quantitative results

The collision prevention mechanism was evaluated for its sensitivity and specificity. Sensitivity refers to the ability of the algorithm to detect the collisions with few false negative results, and thus fewer cases of collisions being missed. Specificity refers to the ability of the algorithm to detect actual collisions by generating few false positive results and thus fewer cases where a collision is detected without actually having one.

To test the sensitivity of our algorithm, the synthetic datasets are used, since the real dataset contains no incidences. The algorithm was calibrated in such a way in order not to miss any collisions. The calibration includes the size of the monitoring area as well as the threshold which defines the minimum distance of two objects before raising an alert. After an appropriate tuning, the algorithm achieved a 100% sensitivity. Avoiding false negatives is critical for our system, since it results to undetected collisions and inability to prevent accidents and to ensure security at intersections.

To test the specificity of our algorithm, the real datasets are used. Since the real datasets contain no collisions, they give a great opportunity to test how many false positives the algorithm will generate in the presence of no incidences. It is worth mentioning that the algorithm is tuned to achieve high sensitivity. The high sensitivity causes a low specificity, as any two objects relatively close between each other in the monitoring area will trigger an alert. To this end, we expect that the system will raise alerts when no collision has actually

Co-funded by the Horizon 2020
Framework Programme of the European Union

occurred. However, we are willing to receive some false alerts in order to ensure that we will not miss any real incidence. Our system is targeting not only to detect collisions but also to prevent potential ones.

Table 1 summarizes the number of false alerts (third column) generated by our algorithm for each one of the datasets (first column) collected from the New York's intersection. Real data of approximately 30 hours were collected in a period of 5 days.

Table 1 False Alerts per Real Dataset

Day	Time duration of dataset	Number of False Alerts
27/04/2022	30 minutes	1
10/05/2022	~8 hours	217
16/06/2022	~7:30 hours	3
18/06/2022	~20 minutes	0
21/06/2022	~14 hours	23

It is worth noticing that on average we receive 1.2 false alerts per hour (when excluding May 10th). This number is relatively low if we consider the high sensitivity configuration of the system. Thus, the trade-off between not missing real alerts while not receiving too many false ones is satisfactory and reliable.

Below, we visualize some false alerts. Figure 11 shows an alert which was raised by the system on April 27th due to the close proximity of a car with a pedestrian. Looking the frames before and after the alert, there is no significant danger for any of the participants and this alert could definitely categorized as a false alert.

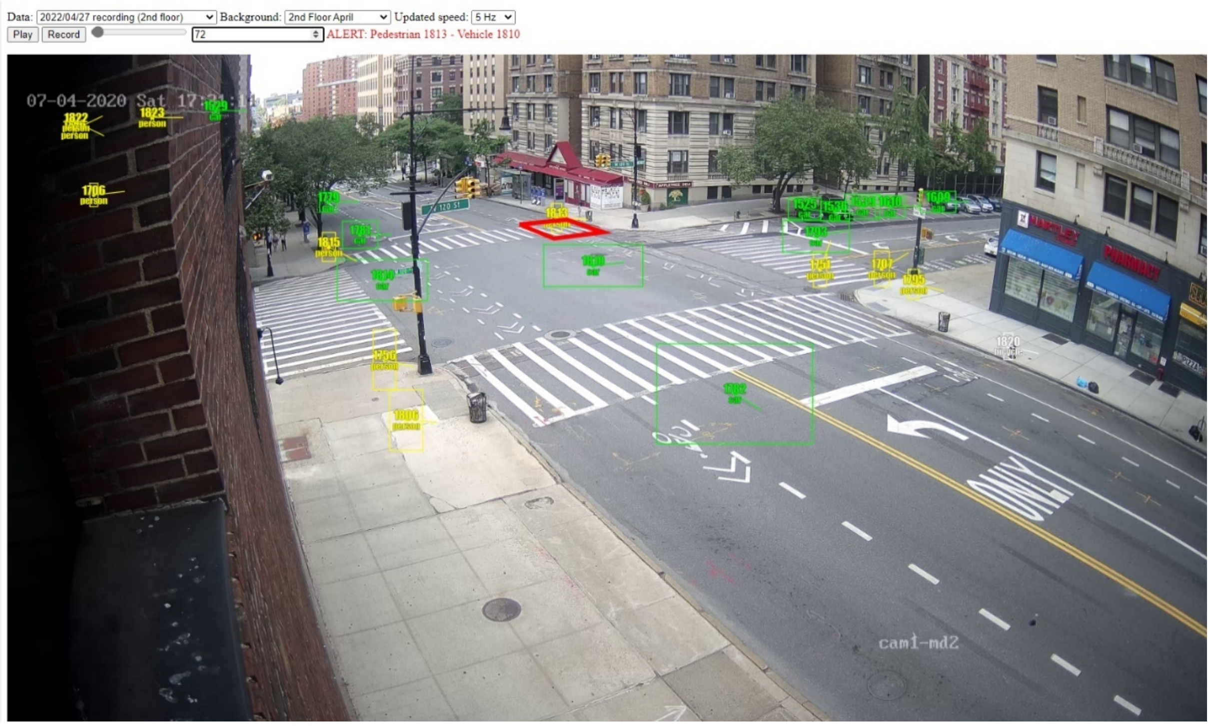

Co-funded by the Horizon 2020 Framework Programme of the European Union

Figure 11 False alert on April 27th: A car is approaching a pedestrian.

Figure 12 illustrates a false alert on June 16th, which although there is no real incidence, it is quite reasonable and could definitely prevent a potential accident. The system detects a pedestrian in the middle of the crossing while a truck is passing behind her/him.

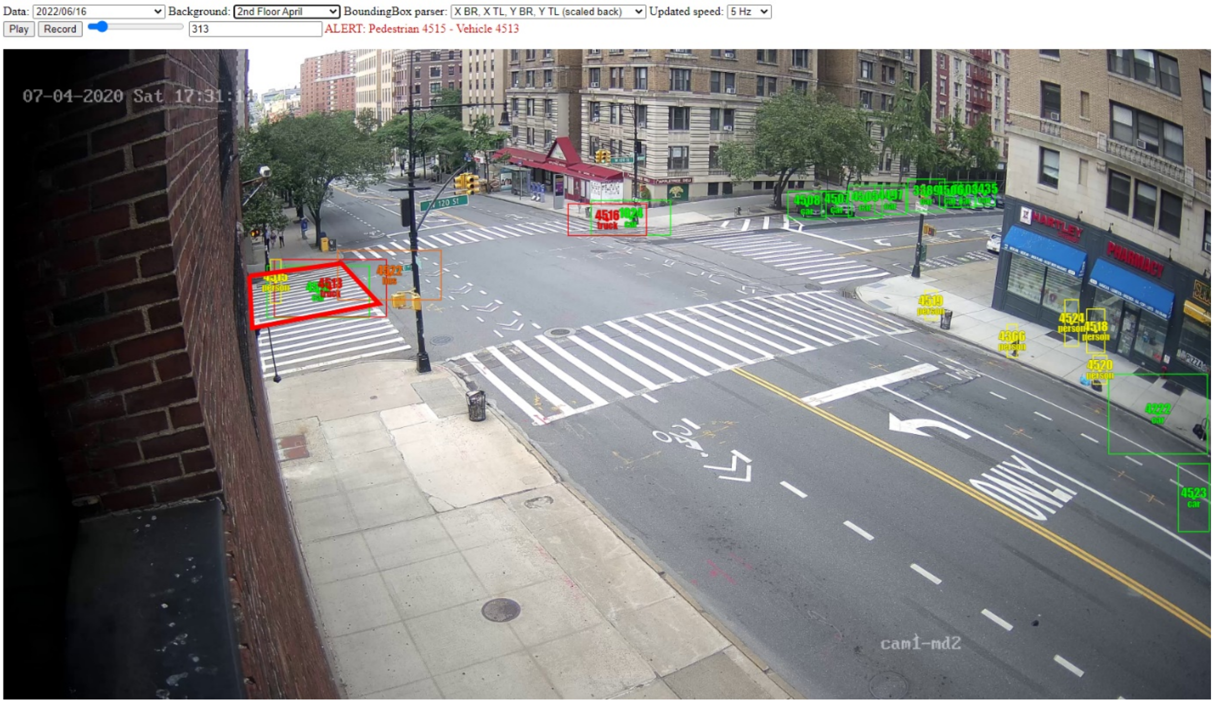

Figure 12 Alert on June 16th : A pedestrian is in the middle of the crossing while a truck is passing behind her/him.

An interesting observation is done for May 10th during which an unusual number of false alerts (217 in total) was raised. During that day, a road construction was taking place with a truck positioned in the middle of the street, causing a traffic disruption (see Figure 13). This disruption forced pedestrians and bicycles to use the same lanes, as well as the vehicles to approach pedestrians in their effort to avoid the construction truck. As a result, the warning mechanism reported an increased number of alerts.

Co-funded by the Horizon 2020 Framework Programme of the European Union

Figure 13 Construction at the Intersection

In the context of our quantitative analysis, we measured the computational demands of our system in order to detect a collision. The performance analysis was done on the historical dataset. All experiments were conducted on a 2 GHz IntelCore i7 processor with 16 GB memory, which runs Windows operating system.

The average processing time to detect a collision found to be approximately 2.5 ms.

Co-funded by the Horizon 2020
Framework Programme of the European Union

5 Present and Foreseen TRL

The smart intersection feature of the platform started with the TRL level of 4 and reached to the TRL level of 6 thanks to the validation of the feature in the relevant intersection environment. It will soon reach to the level 7 once we finalise all the with the demonstration of the prototype with the operational environment.

6 Exploitation, Dissemination and Communication Status

Since the project start, we have been communicating the project's progress and results via different channels, such as:

Kentyou's web page:

- Project start: https://kentyou.com/2022/03/23/kentyou-new-project-smarter-and-safer-road-intersections-in-new-york-and-santander/
- Presentation to NYC (2nd half of the project : https://kentyou.com/2022/10/06/kentyou-presents-fedintersect-to-ny-city-cto/

Urban Technology Alliance:

- News: http://www.urbantechnologyalliance.org/2022/05/16/city-challenge-safer-road-intersections-from-ny-to-santander/
- Newsletter: https://mailchi.mp/c5f497a24a09/uta-newsletter-13871350

Conferences:

- M. Ghasemi, S. Kleisarchaki, T. Calmant, L. Gürgen, J. Ghaderi, Z. Kostic, G. Zussman. "Demo: Real-time Camera Analytics for Enhancing Traffic Intersection Safety", MobiSys'2022.
- M. Ghasemi, S. Kleisarchaki, T. Calmant, J. Lu, S. Ojha, Z. Kostic, L. Gürgen, J. Ghaderi, G. Zussman. Demo: Real-time Multi-Camera Analytics for Traffic Information Extraction and Visualization. Percom 2023 (under submission).

Smart Cities Connect, Washington, September

An annual event bringing together important smart city actors in USA.

- https://fall.smartcitiesconnect.org/program/index_2022.html#detail (Panel on Safe Reliable Daily Travel Using Emerging Technology)

Smart City Expo, Barcelona

- The project results have been presented in a demo session in the Smart City Expo event in Barcelona on November 14 – 17, https://www.smartcityexpo.com. This is a reference smart city event worldwide.

Co-funded by the Horizon 2020
Framework Programme of the European Union

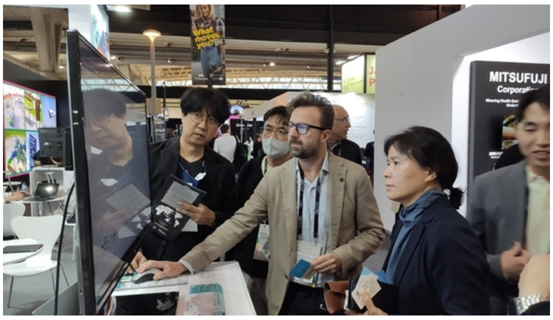

Demo at the Smart City Expo, Barcelona

Presentation at Smart Cities Connect, Washington

7 Impacts

Impact 1: Enhanced EU – US cooperation in Next Generation Internet, including policy cooperation.

The project allowed building a joint EU-US demonstration which is presented in Mobisys conference to the scientific committee and in Smart Cities Connect event to the smart city community in US. The results will also be presented to the New York City's CTO office to build a sustainable relationship between Kentyou, Columbia University and New York City.

Kentyou will exploit the results of this project within EU funded research and innovation projects.

Kentyou is invited to present the outcomes of this project and lessons learnt at the next EU-US workshop on the continuum computing that will take place virtually on November 14th and November 16th.

Impact 2: Reinforced collaboration and increased synergies between the Next Generation Internet and the US Internet programmes.

The results of our project is strongly contributing to US NSF funded project of the Columbia University, NSF CPS award 20-38984, "CPS: Medium: Hybrid Twins for Urban Transportation: From Intersections to Citywide Management."

It is a $1.2M NSF Grant to Improve Traffic Management in Real Time. The objective of the project is to a virtual replica, or digital twin, of New York City that will continuously learn and dynamically update itself as the city traffic environment changes in real time. The twin will help traffic managers to monitor traffic patterns as they happen and quickly come up with adaptive management strategies. The researchers will use Columbia's COSMOS, the only beyond-5G testbed in New York City, to get real-time traffic data, leveraging Cosmos's rich sensor data and deep computational capabilities, as well as the Kentyou Digital Twin of intersections resulting from this FedIntersect project.

Moreover, we are also exploring how the project results can contribute to the COSM-IC project (NSF award number 2029295), which is about interconnecting the COSMOS testbed with other international testbeds

We are doing regular bi-weekly meetings between the participants for enforced collaboration, which will sustain after the project end.

Impact 3: Developing interoperable solutions and joint demonstrators, contributions to standards.

The outcomes of the project, including connectors to public brokers and other frameworks, will be provided as open source to the Eclipse sensiNact project. We will leverage the existing networks of the partners, the open-source community of Eclipse Foundation, as well as the NGI community to attract developers and experimenters to use the federated platform. Joint demonstrations will be planned and performed to additional EU and USA municipalities, further enriching the collaboration of the partners. Last but not least, the work will influence the standardization work on OSGi Working Group of the Eclipse Foundation to transfer the results of the project into its next specification (Kentyou chair of the IoT Expert Group).

Impact 4: An EU – US ecosystem of top researchers, hi-tech start-ups / SMEs and Internet-related communities collaborating on the evolution of the Internet

The relations between the partners will be sustained after the project and in particular via the Urban Technology Alliance (UTA), which is an alliance bringing together city scale testbeds around the world, to exchange best practices and lessons learnt. Kentyou, University of Cantabria and Santander city are founding members of the UTA. The project will disseminate the results of this experimentation with the international UTA community composed of many urban areas, research institutions and industrial partners from all around the world. We will also explore opportunities of replicating the experimentation with some of the UTA members to increase the impact of this project.

We have already started to disseminate the project results in the community via blog posts, newsletters and social networks.

Under the UTA or other frameworks such as Horizon Europe, we expect to make the federated experimentation testbed available for EU and US researchers for further experimentations beyond the project lifetime.

8 Conclusion and Future Work

The deliverable presented the implementation of all use cases defined in D1, based on the architecture described in D2. All necessary developments of the back-end and front-end are being finalized by Kentyou and Columbia university teams. Moreover, the deliverable demonstrated the virtual representation (digital twin) of two intersections, at New York and Santander. The digital twin demonstrated its ability to monitor the real-time situation of traffic, crowdedness, weather and its capability of raising alerts in case of a potential collision. Historical statistical data were also stored and became available through the visual interface. Both teams are keen on continuing the collaboration beyond the NGI. Atlantic project funding. Planning for the future work is going. Potential follow-up topics have been already identified.

Co-funded by the Horizon 2020
Framework Programme of the European Union

9 References

[1] M. Ghasemi, S. Kleisarchaki, T. Calmant, L. Gürgen (Kentyou), J. Ghaderi, Z. Kostic, G. Zussman. "Demo: Real-time Camera Analytics for Enhancing Traffic Intersection Safety", MobiSys'2022.

[2] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," arXiv preprint arXiv:2004.10934, 2020.

[3] M. Ghasemi, Z. Kostic, J. Ghaderi, and G. Zussman, "Auto-SDA: Automated video-based social distancing analyzer," in Proc. ACM HotEdgeVideo'21, 2021

Co-funded by the Horizon 2020
Framework Programme of the European Union