Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

A FLEXIBLE EDUCATIONAL SIMULATION MODEL TO STUDY AUTONOMOUS LAST-MILE LOGISTICS

Berry Gerrits[1], and Martijn Mes[1]

[1]Dept. of High Tech Business and Entrepreneurship, University of Twente, Enschede, The Netherlands

ABSTRACT

This paper presents an open-source discrete-event simulation model to support simulation education of business and engineering students using an appealing problem setting of automated transport for last-mile logistics. The open-source nature of our approach facilitates customization and further development, fostering collaborative research and education initiatives. Within this context, we develop a flexible approach that utilizes OpenStreetMap to create engaging models for simulation education. The basic functionalities corresponding with last-mile logistics operations are included in the model and allow students to quickly experiment with novel logistics concepts, fleet configurations, and planning algorithms to evaluate operational efficiency, environmental consequences (e.g., carbon emissions), and societal factors (e.g., livability). To illustrate our approach, we focus on conceptualizing and modeling the campus of the University of Twente in Plant Simulation as a final assignment for a graduate-level simulation course. We present several options for the model's use for educational purposes.

1 INTRODUCTION

Discrete-Event Simulation (DES) education is a crucial topic for graduate students in engineering and related fields, as it equips them with the skills needed for effective industry work (Kress 2010). A simulation course provides students with low-cost and risk-free experimentation possibilities to develop knowledge and skills in a simulated environment (Negahban 2024). However, as underlined by Chwif (2001), assessing students' prerequisite knowledge, e.g., concerning programming, before the course is important to ensure their preparedness. Building DES models typically requires a vast amount of programming. Hence, there is a risk of focusing too much on coding skills instead of modeling, experimenting, and analyzing the simulation results. Our goal in simulation education is to familiarize students with the power of simulation, and to teach them how to design a conceptual model and design experiments to support decision-making processes. Therefore, we propose a base model that serves as a starting point for the students in which many basic functionalities are already implemented. In this way, students can focus on simulation-specific elements rather than on boilerplate programming. Moreover, to improve student engagement, we propose a visually pleasing, realistic, and relatable simulation model. The base model allows students to work on an integral case study within a limited time frame. This paper thus focuses on a simulation assignment involving a comprehensive case study such that students may learn by themselves using the Case-Based Learning (CBL) didactic strategy and links to the final step of the SUCCESSFUL framework for designing DES courses (Garcia and Centeno 2009).

The goal of this paper is to present such a base model for CBL in which students need to assess the impact of unmanned parcel delivery systems on their own campus. This strategy involves the students intellectually and emotionally by bringing a real close-to-home situation to the instruction setting (Garcia and Centeno 2009). This realism provides an incentive for students to become more involved in education (Standridge 2000). As such, students link theory to practice, deal with the complexity of a real-world situation and enliven teaching (Jennings 2002). The base model serves as a virtual sandbox to explore, experiment, and optimize autonomous delivery systems. As real-world experimentation with autonomous delivery systems is severely limited due to regulatory issues, this use case provides a textbook example of the added value of simulation to study systems not yet available in the real world. Using geographical information systems and rough estimates for the data requirements, we propose a model that is open-source, quickly to build, easily adjustable, and allows students to get up and running with simulation quickly. We present a detailed, platform-agnostic approach to how to build the base model such that other educators can repeat the process for similar use cases (e.g., a different university campus). We showcase our approach using a case study of the University of Twente and implement our model in the commercial DES software Plant Simulation from Siemens. We select a commercial software tool for reasons of convenience for the student currently unmet in open-source alternatives.

The remainder of this paper is structured as follows. Section 2 presents our approach for a CBL-based simulation course. Section 3 presents a flexible approach to creating a base model for a use case to study autonomous last-mile logistics. In Section 4, we present our use case, and in Section 5, we present our conceptual model. How to use the CBL-based model for educational purposes is discussed in Section 6. The paper closes with conclusions and directions for further research in Section 7.

2 DESIGN OF A CASE-BASED LEARNING SIMULATION COURSE

To provide an example of how a case study can be integrated into a simulation course, we sketch a specific simulation course as used at the University of Twente. In this simulation course, a CBL-based final assignment plays an important role. All concepts and relevant theories about DES should be mastered, and students should also learn with hands-on experience. As such, this course combines theory and application. The theory consists of simulation-specific topics as mentioned in Law (2015), such as fitting distributions, determining the type of simulation, experimental design, warm-up period, and number of replications. These topics are assessed both through an individual exam and by applying the theory in the application part, i.e., the final assignment. As stated in the introduction, we use a base model as a starting point for the CBL-based final assignment. As such, students can work on a comprehensive case study, within a limited time period, while still achieving the intended learning outcomes. These learning outcomes are linked to the general qualifications of the course program. Our intended audience, i.e., business and engineering students, should be able to:

- explain the principles of discrete event simulation;
- describe the steps to be taken when conducting a simulation study focused the on improvement of operations systems (logistics systems and business processes);
- judge whether simulation can be a useful technique to analyze certain operations systems (when to use and when not to use);
- design a simulation model for a given operations system;
- implement the model in an advanced simulation tool (e.g., Plant Simulation);
- verify the model and examine its validity;
- define input and analyze the output of a simulation model;
- design and conduct a structured set of simulation experiments;
- combine simulation with optimization;
- write a well-structured project report on a simulation study.

To achieve these learning outcomes, the course consists of a mixture of theory and application. Besides an individual test on theory, we use a walkthrough tutorial containing intermediate assignments and a final CBL-based assignment. The tutorial, which is freely available online (Mes 2021), aims to (i) support students working with the Plant Simulation software by providing a step-by-step guide on creating various simulation models, and (ii) increase their understanding of the simulation theory by applying the theory in various assignments. To serve a wide audience, the tutorial consists of a basic and an advanced part. The basic part consists of three chapters, with three corresponding assignments, using a general practitioner's office as a running example. After students create a first simple model, the first assignment focuses on implementing a prioritization rule for patients such as to improve overall efficiency. The second assignment focuses on creating more realistic input data by fitting a probability distribution function to a dataset and validating whether the output of the improved model is consistent with real-world observations. In the third assignment, students need to determine a proper warm-up period, run length, and number of replications. In the advanced part of the course, similar assignments are used on a more in-depth level. More details on these assignments can be found in Mes (2021). After students have successfully demonstrated their understanding of the simulation theory and basic modeling skills, the CBL-based final assignment is offered. In the remainder of this paper, we describe the process of creating a flexible base model for the CBL-based final assignment, focused on designing and optimizing autonomous last-mile logistics on a university campus.

3 CREATING A FLEXIBLE BASE MODEL FOR AUTONOMOUS LAST-MILE LOGISTICS

For our simulation course, we create a flexible base model that allows students to study autonomous last-mile logistics (e.g., parcel delivery) in urban areas (e.g., a campus). For the base model, we require four main elements: (i) the physical infrastructure of the area, (ii) data related to deliveries, (iii) logistics concepts utilizing autonomous transport, and (iv) experimentation capabilities to study various scenarios and evaluate the impact on business, environmental, and social factors. The first three elements require a certain level of conceptualization that fits the purposes of our model, namely, allowing students to experiment with automated modes of transport in a realistic yet comprehensible setting. For example, a highly detailed infrastructural representation of a campus is beneficial to studying last-mile logistics at an operational level and allows one to consider all the ins- and outs of a specific area. However, the effort required to create such a representation and obtain the required data defeats our purpose of allowing students to quickly experiment with automated last-mile logistics. Hence, as with every simulation study, we require a balance between the level of detail, the effort required to obtain this level of detail, and the resulting flexibility (i.e., the ability to tailor the model to a specific location and needs). In this paper, we present an approach that captures this balance and allows for flexibility both in terms of different areas as well as experimentation possibilities to educate students on the usefulness of simulation to study last-mile parcel delivery systems. The remainder of this section describes our approach to creating a flexible simulation model in several steps: generating the physical infrastructure (Section 3.1), creating a flexible data model (Section 3.2), and creating delivery scenarios (Section 3.3). For our course, we focus on a specific scenario: e-commerce deliveries (Section 3.4).

3.1 Generating the Physical Infrastructure

The first step to creating a flexible simulation model is the ability to easily generate a (semi-realistic) physical representation. The following three-step approach showcases a generic approach suitable for many DES tools and allows for generating different areas of different shapes and levels of detail. Although we realize that some commercially available DES tools have a geographical tool integrated, not all do, and we aim to present an approach that is generic enough to be used with a wide variety of tools such that different educational institutes can use it. The three-step approach is discussed in the subsequent subsections.

3.1.1 Creating a Base Layer

We use the freely available open-source geographic information system tool QGIS to create a geographical representation for our simulation purposes (QGIS, 2024). We use the QuickMapServices plugin (https://plugins.qgis.org/plugins/quick_map_services/) and use the OpenStreetMap (OSM) Standard map as our base layer. After loading the map, we navigate to our area of interest such that the entire area is in the canvas (see Figure 1a for an example of the campus of the University of Twente).

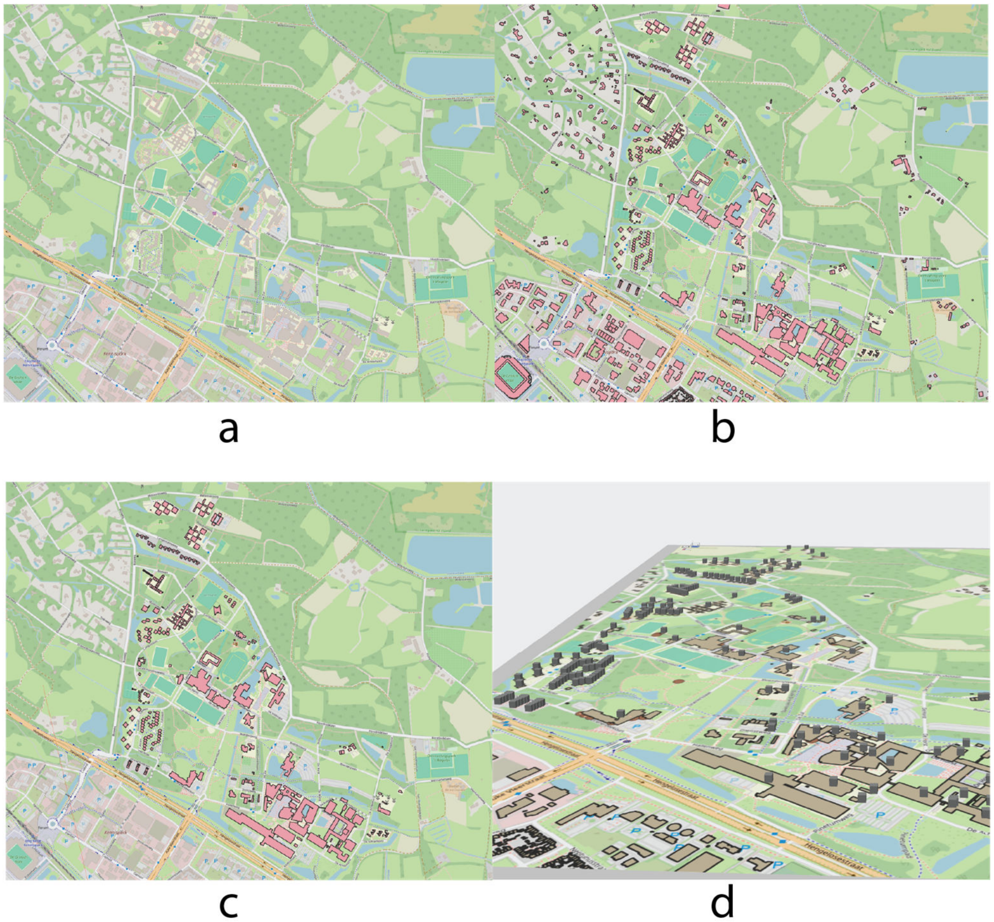

Figure 1: Illustration of creating a physical infrastructure, consisting of (a) the OSM base layer, (b) selected features, (c) filter based on the area of interest, and (d) export to simulation software.

3.1.2 Select Relevant Features and Area of Interest

Using the QuickOSM plugin, see QuickOSM (2022) for details on its usage, we load the features we are interested in (e.g., buildings), see Figure 1b. We then filter the results to include only the features that are within a hand-drawn polygon to narrowly define our area of interest (see Figure 1c). Note that filtering the results allows the user to precisely define which delivery locations should be included. Depending on the output of this step, the user might want to refine the results by excluding certain locations or including additional ones.

3.1.3 Export to Simulation Software

After having obtained the geographical representation of our area of interest and the relevant delivery locations, we require some post-processing such that we are able to create the visualization in our simulation software. The first step is to translate the coordinates (i.e., latitude-longitude) of the selected locations to (x,y)-coordinates for our simulation model. We first obtain the latitude-longitude coordinates of the selected locations by using the *Field Calculator Tool* of QGIS in the *Attribute Table* window. We obtain the latitude using the expression *x($geometry)* and the longitude using *y($geometry)*. This provides a list of real-world coordinates of the delivery locations and can be saved to a spreadsheet. To translate these coordinates into (x,y)-coordinates suitable for our simulation software, we use a Python script (Gerrits, 2024). This script translates the real-world coordinates of the exported locations (i.e., their center coordinates) such that they fit within the world coordinates used in the simulation software. These parameters depend on the simulation tool used. For example, when using Plant Simulation, we found that translating the map of the campus to a world of 2000x4000 pixels provides us with a good view of the area (i.e., not too small, not too big) on a standard laptop screen. Based on these coordinates, we generate simple 3D cubes in the simulation software to resemble the buildings in our area of interest, as shown in Figure 1d. Note that the actual geometry of the buildings can also be exported from QGIS to create a more realistic representation. Still, for the purposes of this study, it suffices to have locations defined by their coordinates with a simple visualization.

3.2 Creating a Flexible Data Model

Opposed to the relatively simple representation of the physical infrastructure, our approach emphasizes the process of obtaining realistic input data. Not only because realistic input data is crucial for any simulation study, the process of obtaining proper input data is often a tedious and difficult yet crucial step. Moreover, it allows students to expand their prior knowledge by solving a realistic problem with its own idiosyncrasies (Negahban 2024). To guide students, we present an approach to creating a flexible data model for the purposes of the course, as discussed in Section 2. Given the limited time-frame of the course, the flexible data model is given beforehand. In the assignment, students still need to critically reflect upon the process illustrated in Figure 2, and verify the resulting input data. When the ability of acquiring input data is an important learning outcome, instructors could decide to let students carry out the steps described below themselves, or to design their own data-retrieval process.

To create instances of home deliveries, we ideally obtain data from logistics service providers active in the area or obtain data from other delivery companies (e.g., food delivery or grocery delivery). We denote this approach as *bottom-up* (i.e., highly detailed data on an individual level). However, as stated at the beginning of the section, this bottom-up approach typically requires much effort and may not be feasible (e.g., due to data protection rights). Hence, for the purposes of our simulation model, we choose to use a *top-down* approach in which general high-level data is translated into less granular data to create rough estimates. This high-level data can typically be easily obtained for most countries. The high-level approach is illustrated in Figure 2, and each step is further discussed in the subsections below.

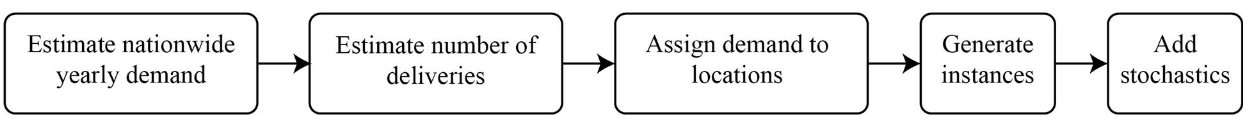

Figure 2: The process of creating a flexible data model.

3.2.1 Estimating Yearly Demand

First, we obtain data on a specific region's total yearly cargo volume (e.g., of a country or state). The reason to start with nationwide data is that this data is typically easily accessible and supplied by logistics service providers on their websites or in annual reports. For example, nationwide data can be obtained from Statista, having data sets on the number of packages delivered in the United Status by LSPs such as Amazon Logistics, FedEx, U.S. Postal Service, and UPS. Similar data can be obtained for other regions. It seems fitting to obtain data from the region that is closest to the area of interest. Research namely indicates that demographics significantly influence the frequency of using home delivery services and includes factors such as health concerns, household members with disabilities (Figliozzi, 2021), age, gender, education (Keeble, 2020), and average age of household members (Liu, 2022). Moreover, if the model requires support for multi-channel deliveries (e.g., e-commerce B2C parcel delivery, food delivery, and grocery delivery), each channel requires an estimate of the yearly demand. Let us denote D_k as the total yearly demand for channel k, where $k = 1, \ldots, K$. Note that data can typically only be obtained from a subset of delivery companies. We may then utilize market share information to extrapolate the yearly volume. That is, if d_{kn} denotes the number of deliveries of company n in channel k and each company n has a market share of s_{kn} (such that $\sum_n s_{kn} = 1$ for each k), then we can estimate D_k using $D_k = \sum_n d_{kn} s_{kn}$ for each channel k.

3.2.2 Estimating Number of Delivery Locations

Next, we estimate the number of delivery locations in the area of interest. Delivery locations may include residential complexes, single-family houses, apartment buildings, and businesses. As discussed in Section 3.1, one may utilize geospatial data sources to estimate both the total number of delivery locations their and coordinates. Depending on the intended focus of the simulation model, different filters can be applied (e.g., residential buildings only) to only obtain relevant locations. Additionally, municipal databases and open data platforms can be used to gather information on the urban infrastructure to estimate the number of locations. We rely on the output of Section 3.1 to get the number of locations. The number of deliveries per location can then be found as follows. Let L denote the total number of delivery locations in a region for which we can retrieve information on the amount of annual deliveries D_k via distribution channel k, and let l denote the number of delivery locations in the area of interest. Then, a simple location-based model to estimate the number of annual deliveries in the area of interest is given by: $D_k\, l/L$. Note that this estimate ignores all demographic aspects and simply assumes that the characteristics of the region as a whole can be applied to the area of interest. Although the location-based model seems a simplistic approach, the data is easy to obtain and provides a rudimentary basis for our educational model. Moreover, when the area of interest resembles the region as a whole quite well, or the data is largely based on the area of interest, this approach seems reasonable to start with. Alternatively, if the population densities are heavily skewed within a region, a population-based model given by $D_k p/P$ may provide a better estimate, where p is the area of interest's total population, and P is the total population on which the yearly data is based.

3.2.3 Assigning Demand to Locations

Having obtained the yearly number of deliveries in our area of interest, either through a location-based model or a population-based model, we can allocate demand to specific locations. Demand allocation is straightforward when using the first approach with the same locations as obtained in Section 3.1.2. If the population-based model is used, the number of people per location must be determined. This may vary greatly depending on the demographics (e.g., apartment buildings and single-family homes). Nevertheless, as a starting point, we can assume that the number of people is uniformly distributed over the number of locations. Note that at this stage, we may want to verify whether the estimates obtained from our *top-down* approach seem reasonable by comparing the yearly number of home deliveries per person (or per location) with *bottom-up* data (e.g., the average number of parcel deliveries per person per year). We may also leave this verification step up to the students as part of the final assignment.

3.2.4 Creating Daily Instances

At this stage, we have a rough estimate of the yearly number of deliveries in our area of interest for every delivery channel. For example, suppose our simulation model simulates a single day of home deliveries. We can then estimate the number of deliveries on a daily basis by uniformly distributing the demand over the year. Allocating demand can either be done on a location level (where we create demand per location based on the number of people) or at a person level (where we create demand per person and each person is allocated to a location). At this stage, we have a simple data model for students to experiment with that is relatively easy to obtain and generic enough to capture different areas and various delivery channels.

3.2.5 Adding Stochasticity

To create more realistic instances, we may want to refine the instances to include additional details to capture the variations in home deliveries (e.g., differences in the day of the week or month of the year) and the stochastic nature (e.g., uncertain arrival times, unknown quantities, and unexpected no-shows). Creating more realistic instances requires additional effort and might be justified depending on the model's educational purposes. For example, historical delivery data can be analyzed to understand demand patterns, including peak delivery times, popular delivery destinations, and seasonal variations. This analysis can be done upfront or can be offered as part of the assignment. Moreover, the recipients-per-location data can be refined if expert knowledge is present. Instead of uniformly distributing the number of recipients over the locations, one may distribute the recipients differently depending on, e.g., the type or size of the locations.

3.3 Creating Logistics Scenarios

After completing the steps from Section 3.2, we can focus on the delivery options per channel k. We first start with a benchmark such that we can evaluate the added value of new autonomous logistics concepts. The benchmark is based on the current state-of-the-art of manual delivery. For example, in e-commerce parcel deliveries, we may assume that these parcels have a batched arrival (i.e., many parcels in a delivery van) and that these delivery vans arrive at the area of interest at specific times (i.e., one or more vans per LSP). A similar analogy can be made for grocery deliveries. However, meal deliveries have a different pattern (i.e., one delivery per location) and arrive instantaneously at the area of interest at different times (e.g., during lunch and dinner). We focus on last-mile e-commerce parcel deliveries that are decoupled from the long-haul and are carried out by autonomous vehicles. For our experiments, we assume that delivery vans unload their cargo at the edge of the area at a micro-hub at which automated vehicles are available to transport the parcels to the customers (e.g., via street robots or drones).

3.4 E-commerce Deliveries

To illustrate a relevant logistics scenario, we focus on batched-arrival e-commerce parcels. Let us denote D_d as the daily demand for these parcels in our area of interest. We denote N as the number of daily batches (i.e., the number of delivery vans that arrive). We can then simply uniformly divide D_d over N to allocate parcels to delivery vans. Depending on the specific case study and region at hand, other division strategies can be used if this results in more realistic instances. In the current state-of-the-art, the delivery van delivers the parcels one-by-one using a certain route (e.g., created using commonly available heuristics or exact methods for TSPs). After the trip is completed, the delivery van leaves the area. We denote this procedure as the *benchmark* scenario, where we only consider the operations of the delivery van in the area of interest. Next to the benchmark scenario, students need to experiment with different modes of automated vehicles, e.g., unmanned aerial vehicles (UAVs) and unmanned ground vehicles (UGVs), and how they impact the area of interest. Generally speaking, we can decouple the cargo from the delivery van at the micro-hub, after which (small) automated vehicles take over. Ideally, the delivery van is no longer allowed to enter the region. However, practically speaking, some parcels might still require a human-operated delivery van, e.g., due to odd-shaped parcels, requirement of additional handling, or their heavy size not suitable for automated vehicles. Nevertheless, let us assume that all parcels are decoupled at the micro-hub. At this micro-hub, we have a heterogeneous fleet V of automated vehicles available, such that $V = \sum_m (v_m)$, where v_m denotes the number of vehicles available of type m, and where each vehicle m has certain properties, such as modality (air or road), speed, payload capacity, and battery capacity. In our simulation, we want to experiment with different compositions and scales of V to evaluate the impact on key performance indicators associated with automated vehicles, such as travel time, reliability, energy consumption, greenhouse gas emissions, and risk of casualties (Kroesen et al 2023). Moreover, it allows students to experiment with novel delivery options, such as trigger-based (e.g., deliver when the customer requests it) or location-based (e.g., change the delivery location to the customer's current location).

4 CASE STUDY

To illustrate the presented approach and to create an open-source simulation model for educational purposes, we present a case study focusing on e-commerce parcel delivery on the campus of the University of Twente, see Figure 3 for an overview. This case study serves as an example for a final assignment for our DES course as discussed in Section 2. This case study contains many interesting elements. Next to different businesses, sports facilities, and faculty buildings, the campus households roughly 2500 students spread over 162 residential buildings. Our focus is on the residential and university buildings, resulting in 222 locations requiring parcel delivery services. We use the approach discussed in Section 3.1 to create a geographical representation of the campus, see Figure 4d. Regarding data requirements, we utilize freely available information from a large Dutch parcel delivery company to obtain rough global estimates. For example, the average Dutch household has a yearly demand of 15 parcels. Given that this parcel delivery company has around 50% market share, we estimate the average number of parcels per household at 30. Moreover, the average Dutch person orders 12 parcels annually (McCarthy 2019), and the average Dutch household consists of 2.12 persons (CBS 2023). Based on this information, another estimate is approximately 25 parcels per household per year. Both approaches provide us with an estimate as a base for our case study, under the assumption that the residents of the university campus resemble average Dutch online consumer behavior, which may not be the case since university students are relatively young and may have different online buying behaviors (Sorce 2005). Therefore, for verification purposes, we obtained campus-specific data from the same parcel delivery company. Based on an analysis of this data, students order approximately 11 parcels per year, comparable to the Dutch average. Hence, our rough estimate provides us with a solid starting point for our simulation. Next, we uniformly distribute the yearly demand to obtain daily deliveries and assume that each building has 2.12 people. The latter is obviously not the case with many apartment buildings and dormitories on campus. Nevertheless, it gives us a rough estimate for the 222 locations under consideration and suffices for educational purposes. Furthermore, in a later stage, the input data can be tweaked when more accurate information becomes available and provides an interesting learning opportunity for students to assess the importance of having the right level of detail and input data for the purposes of the simulation study.

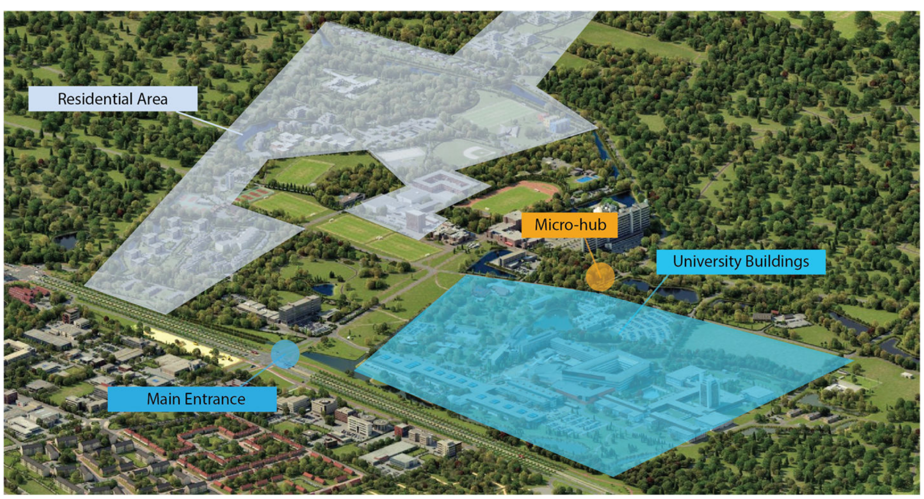

Figure 3: Overview of the campus of the University of Twente.

5 CONCEPTUAL MODEL

Before presenting the final assignment, we first describe the conceptual model for the base model. In this section, we describe the configurability of the model related to inputs (Section 5.1), outputs (Section 5.2), and model logic (Section 5.3).

5.1 Inputs

To get students quickly up and running with the model, we provide several input parameters and their corresponding (estimated) values. Possibly, students can refine these parameters based on a specific use case or when (more detailed) data is available. Moreover, based on the specifics of the final assignment (see Section 6), students can add additional input parameters if necessary. The starting input parameters are listed in Table 1.

Table 1: Input parameters of the base model and their values.

Input	Description	Default Value
DecouplingPointLocation	Coordinates of decoupling point / micro-hub	(x,y)-value
AvgCustomersPerBuilding	Initialize number of customer for each building	2.12
AvgYearlyDeliveriesPerCustomer	Yearly demand per customer	25
StdDevYearlyDeliveriesPerCustomer	Standard deviation of yearly demand per customer	1.5
NumBatches	Number of delivery vans that arrive during a day	3
DeliveryBatches	Time-windows used to create a random arrival time per delivery van	Table
BenchmarkScenario	Boolean to run Scenario 1 (True) or Scenario 2 (False)	False
BenchmarkElectricVan	Boolean to select electric van (True) or diesel van (False) for Scenario 1	True
StartTimeWindow	Earliest start of delivery time window for customer	Equal to arrival time of delivery van
EndTimeWindow	Latest closing time of delivery window of customer	Should be set such that it is feasible to deliver all parcels
AdhereToTimeWindows	Enforces delivery of each customer within time-window	True
VehicleCharacteristics	Sets the (technical) properties per delivery mode, e.g., speed, parcel capacity, (un)load time per parcel, operating costs per hour, battery consumption, emissions	Table

5.2 Outputs

Students can use different outputs of the model for further data analysis to gain insights to be used in their written report to support their claims on the (optimal) design and configuration of the automated last-mile logistics system. To provide a starting point for students, we create an output table in the simulation software such that students can evaluate various KPIs. Students are free to use any combination of these outputs or add their own to support their analysis in the written report. The first set of outputs is chosen such that students can develop insights into the People, Planet, and Profit (PPP) aspects of the automated last-mile logistics system they configured and implemented. The initial output table has the following values (per delivery): id of parcel, id of delivery van, customer location, start of delivery time-window, end of delivery time-window, time of arrival at customer, time of delivery at customer, total time in system, assigned to vehicle (Boolean), *successfullDelivery* (Boolean), and *vehicleUsed* (e.g., street robot or drone).

Using these outputs (in combination with relevant inputs), students can derive KPIs regarding, e.g., total emissions, energy consumption, and delivery costs. Although we provide this initial table, the actual data gathering in the model and calculation required to find correct values for these KPIs is part of the final assignment and thus not provided beforehand.

5.3 Model Logic

To get students up and running with the model, we provide basic functionality such that the students start with a working model to minimize boilerplate coding. The starting model logic is such that the model runs without any modification, but results in poor performance. Students need to first derive a conceptual model of the base model to showcase their conceptualization skills. Second, they need extend the code in the model to achieve good performance, e.g., by implementing an intelligent scheduling algorithm. Improving the performance also teaches the student's some conceptualization, e.g. how vehicles are dispatched.

For the starting model, we distinguish between two scenarios. In Scenario 1, a delivery van goes to all customers in a one-by-one fashion, to reflect how the current delivery system works on campus. In Scenario 2, the delivery van goes to the micro-hub, unloads all parcels and leaves the campus. In Scenario 1, if the parameter *EnforceTimeWindows* is set to *True*, and if the delivery van arrives at a location *before* the time-window is open, it simply waits until the time-window opens. If a time-window is already closed, we assume that the delivery failed. Alternatively, if the parameter *EnforceTimeWindows* is set to *False*, we assume that all parcels can be delivered, yielding good performance, but this setting most probably does not reflect a realistic scenario. Students should keep this in mind and build upon these starting scenarios to create a realistic simulation model. In Scenario 2, all parcels arrive at the micro-hub, but nothing is delivered yet. Students should include their own logic in deciding which modality should deliver which parcel at what time. This logic could for example be triggered when the delivery van arrives at the micro-hub or whenever an autonomous vehicle returns to the micro-hub. At these triggers, students should program their own logic to dispatch the available fleet.

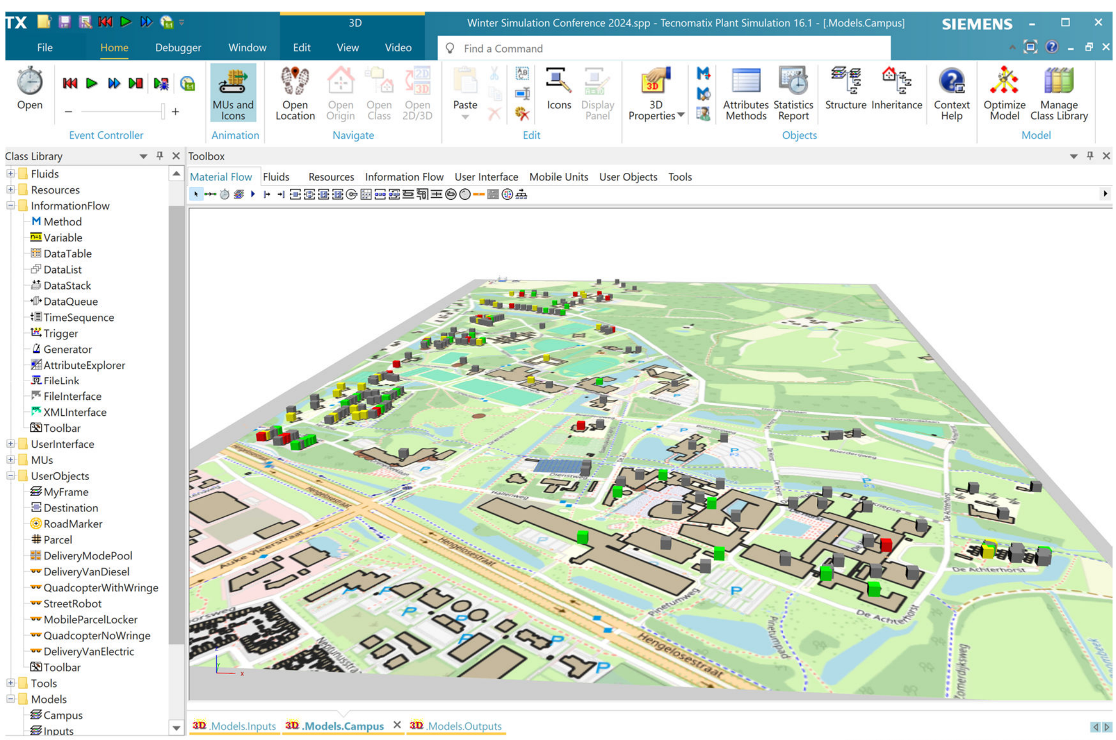

Figure 4: Screenshot of implemented base model in Plant Simulation.

6 USES OF THE CBL-BASED MODEL

This section presents several uses of the presented simulation model for educational purposes. In our course, the final assignment consists of two connected parts: (i) the actual assignment (what students need to do) and (ii) the case study from a problem owner's perspective. The former describes what is expected from the students, i.e., what they have to hand-in (deliverables like simulation model, report, and supplementary material), required elements within these deliverables, and how the different elements will be assessed. The latter provides a case study description from the perspective of the problem owners. For example, the campus manager might be unsatisfied with the current parcel delivery system, resulting in too much congestion, noise complaints, and unsuccessful deliveries. The problem owners challenge students to show how (different forms of) autonomous transport in combination with a micro-hub can help overcome these challenges and, therefore, requests a simulation study to be performed to aid in the decision-making process. Next to achieving the learning outcomes of executing a simulation study and writing a well-structured report about the study and its findings, students should also demonstrate they master simulation theory, e.g., by selecting proper warm-up periods, simulation length, and number of replications.

With the proposed base model, several specific assignments can be offered, depending on the background of the students and the specific learning outcomes. Typically, students should demonstrate that they are able to (i) familiarize themselves with an existing model, (ii) extend the software through programming their own code, (iii) experiment with the model, and (iv) collect and extract model output for further analysis. A first assignment would be to focus on Scenario 2. Students need to schedule the decoupled parcels at the micro-hub given a certain fleet of autonomous delivery modes. Students need to demonstrate that they are able to program, and test effective dispatching logic, possibly demonstrating trade-offs such as balancing customer satisfaction and operational costs. Moreover, they should demonstrate the PPP benefits of Scenario 2 compared to benchmark Scenario 1. Additionally, students could experiment with optimizing the required fleet size and composition, and assess the impact on the delivery performance. This latter aspect focuses on tactical decision-making. Students could also focus on strategic decision-making by optimizing the location(s) of the micro-hub(s) and assessing how many micro-hubs are required and at which locations to attain a certain logistics performance in terms of PPP. For more technically-oriented students, the final assignment could also focus on experimenting with the technical properties of the autonomous fleet. The input parameters regarding speed, capacity, and battery consumption can be tweaked to experiment with existing or new-to-be-designed robotic solutions. Moreover, students could experiment with an 'optimal' design of the autonomous fleet and assess whether it is feasible to design such a fleet. Lastly, students could also focus on implementing a cooperative approach with the customers, e.g., my means of a (negotiation)-mechanism that tries to influence the selected delivery windows of customers to create more efficient delivery plans. Testing the student's experience with the proposed simulation course and CBL-assignment is scheduled to be done in Fall 2024.

7 CONCLUSIONS AND FURTHER RESEARCH

This paper presented an open-source discrete-event simulation (DES) model for use in simulation education, considering last-mile logistics using automated modes of transport. Because building DES models typically requires a vast amount of programming, we proposed a base model that serves as a starting point for Case-Based Learning (CBL) in which many basic functionalities are already implemented so that students can focus on simulation-specific elements rather than on boilerplate programming. We focused on the process of designing a flexible and engaging case-based learning assignment for students to show they master discrete-event simulation. We illustrated our approach using a case study of parcel deliveries on the campus of the University of Twente using autonomous transport modes. We presented several options for the model's use as a final assignment in a CBL-based simulation course for business and engineering students. We invite researchers and educators to use and extend the model, for example by: (i) assessing the applicability of the presented approach in other simulation software, (ii) assessing the feasibility of the presented approach for different campuses and urban areas, and (iii) extending the model to include different delivery channels, e.g., food- and groceries delivery.

REFERENCES

CBS. 2023. Huishoudens Nu. https://www.cbs.nl/nl-nl/visualisaties/dashboard-bevolking/woonsituatie/huishoudens-nu, accessed 11th April 2024.

Chwif, L., M. Pereira-Barretto, and R.J. Paul. 2001. "Assessment of student preparation for discrete event simulation courses". *In Proceedings of the Winter Simulation Conference*, December 9th-12th, Arlington, VA, USA, 1624–1631.

Figliozzi, M., and A. Unnikrishnan. 2021. "Exploring The Impact Of Socio-Demographic Characteristics, Health Concerns, And Product Type On Home Delivery Rates And Expenditures During A Strict COVID-19 Lockdown Period: A Case Study From Portland, OR". *Transportation Research Part A: Policy and Practice, 153*, 1-19.

Garcia, H., and M.A. Centeno. 2009. "S.U.C.C.E.S.S.F.U.L.: A Framework For Designing Discrete Event Simulation Courses". *In Proceedings of the Winter Simulation Conference*, December 13th-16th, Austin, TX, USA, 289–298.

Gerrits, B,. 2024. Plant Simulation Base Model to Study Autonomous Logistics on University Campus. https://github.com/DistributeCompany/Plant-Simulation-Base-Model-to-Study-Autonomous-Logistics-on-University-Campus, accessed 11th April.

Jennings, D. 2002. "Strategic Management: An Evaluation Of The Use Of Three Learning Methods". *Journal of Management Development*, 21(9/10): 655-665.

Keeble, M., J. Adams, G. Sacks, L. Vanderlee, C.M. White, D. Hammond, and T. Burgoine. 2020. "Use Of Online Food Delivery Services To Order Food Prepared Away-From-Home And Associated Sociodemographic Characteristics: A Cross-Sectional, Multi-Country Analysis". I*nternational Journal of Environmental Research and Public Health, 17(14),* 1–17.

Kress, R., A. Cemerlic, J. Kress, and J. Varghese. 2010. "Discrete Event Simulation Class for Engineering Graduate Students". *In Proceedings of the Winter Simulation Conference*, December 5th-8th, Baltimore, MD, USA, 344-352.

Kroesen, M., D. Milakis, and B. van Wee. 2023. "Automated Vehicles: Changes in expert opinions over time". *Transport Policy. 136*, 1–10.

Law, A. M. (2015). *Simulation Modeling and Analysis.* 5th ed. New York: McGraw-Hill, Inc.

Liu, Y., K., Wang, P. Loa, and K.M. Habib. 2022. "Modelling the Frequency of Home Deliveries: An Induced Travel Demand Contribution of Aggrandized E-shopping in Toronto during COVID-19 Pandemics". 10.48550/arXiv.2209.10664.

McCarthy, N. 2019. Online Shopping: Where The Parcels Are Piling Up. https://www.statista.com/chart/18396/average-number-of-parcels-received-per-person/, accessed 11th April 2024.

Mes, M. 2021. Simulation Modelling Using Practical Examples: A Plant Simulation Tutorial. https://www.utwente.nl/en/bms/iebis/staff/mes/plantsimulation/, accessed 11th April 2024.

Negahban, A. 2024. "Simulation in Engineering Education: the Transition From Physical Experimentation to Digital Immersive Simulated Environments. *Simulation*. https://doi.org/10.1177/00375497241229757

QGIS. 2024. A Free and Open Source Geographic Information System. https://www.qgis.org/en/site/, accessed 11th April.

QuickOSM. 2022. QuickOSM User Guide and Documentation. https://docs.3liz.org/QuickOSM/, accessed 11th April 2024.

Sorce, P., V. Perotti, and S. Widrick. 2005. "Attitude and age differences in online buying". *International Journal of Retail and Distribution Management*, 33(2), 122-132.

AUTHOR BIOGRAPHIES

BERRY GERRITS is a researcher within the Industrial Engineering and Business Information Systems section at the High Tech Business and Entrepreneurship department at the University of Twente, the Netherlands. He holds a PhD in Industrial Engineering (2023) and his research interests are self-organizing logistics, automated transport systems, (agent-based) simulation, and bio-inspired AI. His email address is b.gerrits@utwente.nl.

MARTIJN R.K. MES is a full professor within the Industrial Engineering and Business Information Systems section at the High Tech Business and Entrepreneurship department at the University of Twente, the Netherlands. He holds a MSc in Applied Mathematics (2002) and a PhD in Industrial Engineering and Management at the University of Twente (2008). After finishing his PhD, Martijn did his postdoc at Princeton University. His research interests are transportation, multi-agent systems, stochastic optimization, discrete event simulation, and simulation optimization. His email address is m.r.k.mes@utwente.nl.