Environmental Modelling and Software 186 (2025) 106358

Contents lists available at ScienceDirect

Environmental Modelling and Software

journal homepage: www.elsevier.com/locate/envsoft

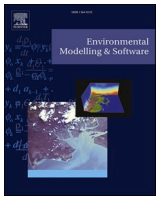

SWMManywhere: A workflow for generation and sensitivity analysis of synthetic urban drainage models, anywhere

Barnaby Dobson [a,*], Tijana Jovanovic [b], Diego Alonso-Álvarez [c], Taher Chegini [d]

[a] Department of Civil Engineering, Imperial College London, UK
[b] British Geological Survey, UK
[c] Department of Computing, Imperial College London, UK
[d] Purdue University, USA

ARTICLE INFO

Keywords:
Algorithmic network generation
Urban drainage
Sensitivity analysis
Synthetic networks
Hydraulic modelling
Network topology

ABSTRACT

Improvements in public geospatial datasets provide opportunities for deriving urban drainage networks and simulation models of these networks (UDMs). We present SWMManywhere, which leverages such datasets for generating synthetic UDMs and creating a Storm Water Management Model for any urban area globally. SWMManywhere's modular and parameterised approach enables customisation to explore hydraulicly feasible network configurations. Key novelties of our workflow are in network topology derivation that accounts for combined effects of impervious area and pipe slope. We assess SWMManywhere by comparing pluvial flooding, drainage network outflows, and design with known networks. The results demonstrate high quality simulations are achievable with a synthetic approach even for large networks. Our sensitivity analysis shows that manholes locations, outfalls, and underlying street network are the most sensitive parameters. We find widespread sensitivity across all parameters without clearly defined values that they should take, thus, recommending an uncertainty driven approach to synthetic drainage network modelling.

Software availability statement

Name of software: SWMManywhere.
Contact information: b.dobson@imperial.ac.uk.
Year first available: 2024.
Licence: BSD-3-Clause.
Programming language: Python (version $\geq$ 3.10).
Hardware required: Windows, MacOS, Linux, 8 GB RAM.
Size: 700 MB (Windows).
Availability:
Code: https://github.com/ImperialCollegeLondon/SWMManywhere (last accessed 2024-10-11, open source).
Documentation: https://imperialcollegelondon.github.io/SWMManywhere/(last accessed 2024-10-11, open source).
Experiments and results: https://github.com/barneydobson/swmmanywhere_paper (last accessed 2024-10-11, open source).

1. Introduction

Urban drainage models (UDMs) are representations of land and pipes that can be simulated with hydraulic models such as Storm Water Management Model (SWMM) (Rossman, 2010). UDMs are essential for managing stormwater, preventing flooding, and ensuring the sustainability of urban water systems (Butler and Davies, 2004). UDM simulations capture the behaviour of drainage networks under various rainfall scenarios, enabling planners to design effective infrastructure and mitigate risks (Bach et al., 2014). However, the development of UDMs typically requires extensive infrastructure records on the connectivity, geometric properties, and elevations of underground pipes (Bach et al., 2020; Chahinian et al., 2019). In cases where these records are unavailable, whether due to ownership issues or simply that they do not exist, the expense of creating a UDM becomes significant because of costly surveying requirements. To forgo this expense, it may be preferable to synthesise a UDM based on the underlying information governing the placement and sizing of drainage pipes, most generally, surface elevation, building locations, and road locations (Chegini and Li, 2022).

The earliest methods to create synthetic UDM exploited the fractal nature of a drainage network and, while simulations were not tested, demonstrated that the broad statistical properties of the network, i.e., distribution of flow path lengths, could be estimated (Ghosh et al.,

* Corresponding author.
E-mail address: barnaby.dobson1@gmail.com (B. Dobson).

https://doi.org/10.1016/j.envsoft.2025.106358
Received 25 October 2024; Received in revised form 10 January 2025; Accepted 30 January 2025
Available online 2 February 2025

2006). When road network and elevation data were incorporated, the accuracy of synthesised UDMs improved and simulations approached those of a real UDM for the same locale (Blumensaat et al., 2012). Generally, UDM synthesis involves three main tasks: delineating surface characteristics, deriving network topology, and hydraulic design, each of which will be reviewed below.

First, surface characteristics, hereafter referred to as sub-catchments, quantify the spatial distribution of stormwater drainage from impervious areas to manholes. Sub-catchment delineation has received the least attention in UDM synthesis literature and most commonly follows a watershed delineation approach (Blumensaat et al., 2012; Warsta et al., 2017). The drained impervious area within a delineated sub-catchment is typically calculated from the area covered by roads and buildings (Mair et al., 2017; Chegini and Li, 2022). However, another simpler method is to assign impervious areas to drain to a nearest manhole (Reyes-Silva et al., 2023). Both methods require identification of manholes, highlighted as critical future work in Blumensaat et al. (2012), but has received little attention since (Bertsch et al., 2017; Chahinian et al., 2019). We place sub-catchment delineation and manhole identification as the first tasks to perform during UDM synthesis, since network topology and hydraulic design should account for the impervious area contributing to a given pipe.

Second, network topology describes the spatial layout of a UDM, connectivity of pipes, and connectivity of the sub-catchments to UDM, i.e., through manholes. Network topology is typically derived by asserting that pipes follow roads, thus dramatically reducing the dimensionality of deriving network topology (Mair et al., 2017; Xu et al., 2021). An efficient network that visits all manholes without redundant pipes can be derived using a shortest path-based algorithm. This algorithm minimizes the total graph cost, with each edge (i.e., a plausible pipe) assigned an individual cost that represents some penalty associated with retaining that edge. Chahinian et al. (2019) provide a detailed exploration minimising costs based on pipe length, pipe adjacent angle, and slope, and highlighting the importance of slope as a cost. Reyes-Silva et al. (2023) derive the UDM by applying the minimum spanning tree (MST) to a full street network to minimise the number of pipes. This approach, however, does not inherently consider slope in the deriving network topology step since MST is only applicable to undirected networks, instead gravitational slope along pipes is enforced as a post-processing correction. A minimum spanning arborescence is an alternative approach to solve such a problem for a directed network (Ray and Sen, 2024; Tarjan, 1977), although has not yet been demonstrated in UDM synthesis. Furthermore, an additional cost to be minimised that has not been tested in the literature is the need to minimise the impervious contributing area to a given pipe and thus more evenly distribute flow throughout the network, much in the manner of the original proposed fractals for network topology (Ghosh et al., 2006).

Third, hydraulic design refers to the selection of pipe diameter, invert levels, and other pipe hydraulic parameters in a synthetic UDM, which typically follows local standards (Chegini and Li, 2022; Duque et al., 2022; Reyes-Silva et al., 2023). Duque et al. (2022), present a methodology for designing sanitary sewer networks by starting from the most upstream pipes, iteratively working downstream, and designing each pipe by minimising the costs subject to the feasibility of design constraints; such an approach could be equally valid for a UDM. If the impervious contributing area of a given pipe is known or has been synthesised, a Rational Method could be applied for determining pipe size. A variety of more extensive global optimisation methods exist for hydraulic design to both minimise costs (e.g., Sun et al. (2011)), or calibrate to observations (e.g., Huang et al. (2022); Sytsma et al. (2022)). However, these calibration studies highlight the inherent equifinality in parameter selection, meaning similar results can be achieved with very different parameters, thus suggesting the uncertainties in such a high dimensionality problem may outweigh any 'optimal' algorithm.

Because UDM synthesis requires a hydraulic design, equifinality must also be inherent to UDM synthesis. We argue that this has been under-recognised in existing UDM synthesis, primarily because of a lack of data and the absence of an automated, end-to-end workflow to assess the impact of parameter selection. A reader may observe results from the supplement of Duque et al. (2022), which demonstrates that small changes in their grid scale for the synthesis algorithm can generate dramatically different sanitary sewer networks.

An equally valid line of inquiry is, therefore, to examine UDM synthesis with sensitivity analysis to quantify the importance of various factors in generating more realistic networks (Pianosi et al., 2016). Sensitivity analysis provides verification by revealing ranges of 'behavioural' parameter values that produce acceptable model outputs, thus informing application of UDM synthesis in data-sparse regions. Additionally, it reveals the dominant controls and key processes governing UDM synthesis by quantifying the relative importance of different parameters. In turn, revealing where uncertainty reduction may be most beneficial. We do not assume that accurate UDM synthesis is possible in every location using solely building, road, and elevation data, particularly given the complexities involved in the gradual expansion of a UDM (Rauch et al., 2017). However, it is impractical to improve UDM synthesis by capturing every possible element involved in network evolution. Instead, sensitivity analysis provides an objective way to guide improvement; prioritising measurements and processes that relate to the most sensitive parameters.

Continual improvements in development and accessibility of remote sensing and open geospatial data at global scale facilitates applicability of UDM synthesis in many locations. Chegini and Li (2022), demonstrate an approach that can perform network topology and the dimensioning part of hydraulic design anywhere in the continental United States. Montalvo et al. (2024) implement the UDM synthesis algorithm presented by Reyes-Silva et al. (2023), in a GIS tool, however this approach does not acquire/process the necessary geospatial data nor is it open-source at time of writing. We believe that a global tool, that is open-source, and tailored to accommodate the uncertainty inherent in the UDM synthesis problem would be of great value to the urban drainage community. Such a tool would enable hydrodynamic method developers to bypass the reproducibility crisis (Stagge et al., 2019; Hutton et al., 2016), by demonstrating their tools on an unlimited suite of UDMs synthesised in cities worldwide, improving on the entirely theoretical test suites that presently exist (Möderl et al., 2009, 2011; Sweetapple et al., 2018). Furthermore, it would contribute towards better representation of urban environments in the regional hydrological cycle, which is a critical and under-represented component of such applications (Coxon et al., 2024).

In this paper we present a workflow, called SWMManywhere, to synthesise UDMs anywhere in the world. We show its versatility by applying it to multiple case studies (i.e. sewer systems of different sizes in two different countries) and to facilitate its global usage we deploy it as an open-source Python tool (Dobson et al., 2024a) with extensive documentation (SWMManywhere documentation, 2024). SWMManywhere provides a parameterised and easy to customise approach for UDM synthesis which allows us to perform sensitivity analysis of synthesised UDMs. We ensure SWMManywhere responds to the need for worldwide application by using open global datasets and including the retrieval and preprocessing of these datasets as part of the algorithm. The methods implemented in SWMManywhere also include a variety of technical novelties in the field of UDM synthesis, including use of a minimum spanning arborescence to derive topology, enabling minimisation of contributing area during this process, and implementing Duque et al. (2022)'s pipe-by-pipe design method for urban drainage networks.

2. Methodology

A high-level overview of our proposed SWMManywhere workflow is shown in Fig. 1.

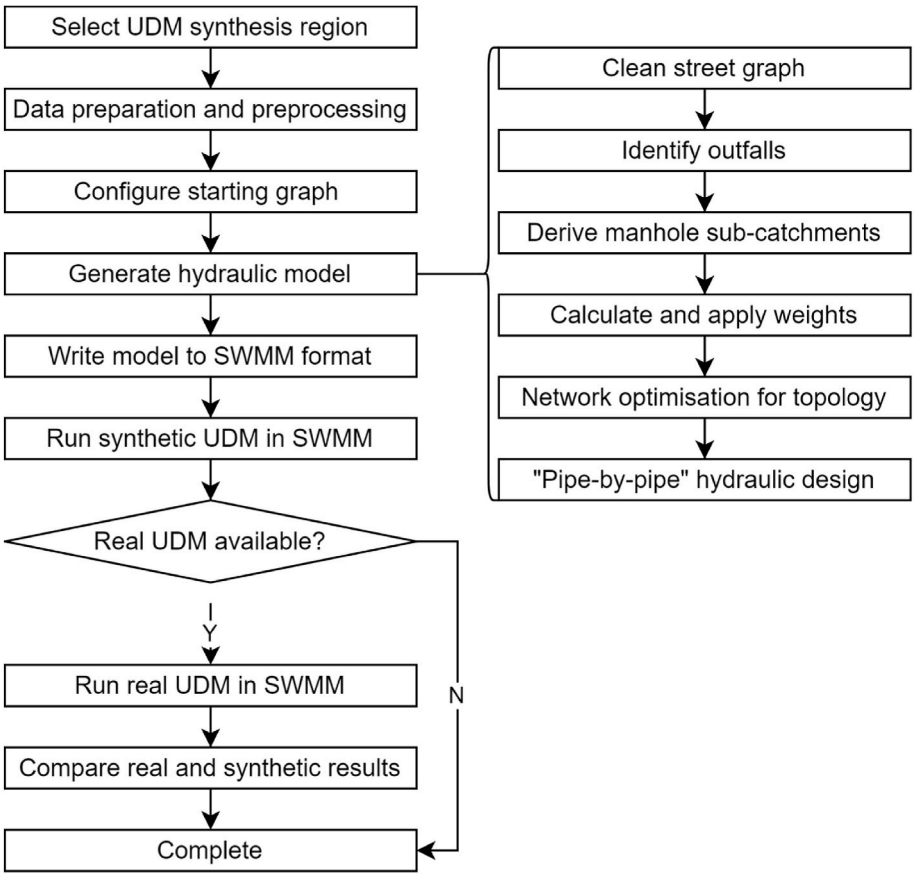

Fig. 1. Overview of the SWMManywhere workflow.

We present an end-to-end workflow for UDM synthesis and comparison against a real UDM, if available, anywhere in the world. The goal of UDM synthesis identified in literature is to set up a hydraulically correct model. The first step in this, for an identified region, is data acquisition and preprocessing, which we specify as part of the workflow. The 'generate hydraulic model' step, is the key step that should ultimately create sub-catchments, a network topology, and hydraulic designs of pipes that together fully describe the synthesised drainage network. Such a drainage network, i.e., the synthetic UDM, will be valid for simulation in a widely used software for design and analysis of drainage networks, such as SWMM, enabling studying different precipitation events and inspecting state variables such as pipe flows and pluvial flooding. In Section 2.1, we describe the theory and processes in our workflow, also explaining how it can be customised with parameter choices and additional processes to vary the nature of the synthesised UDM.

A key hypothesis that must underlie any synthetic UDM approach is that results may be valid in places where a real UDM is not available or not trusted. Because we observed the sensitivity to parameter selection in previous UDM synthesis literature, we intentionally specify SWMManywhere to be highly parameterised and customisable. We use sensitivity analysis on these parameters to demonstrate how they impact synthesised networks. In Section 2.2, we describe how sensitivity analysis can guide the use of SWMManywhere in areas where parameters cannot be estimated a priori based on field data and provide a deeper understanding of UDM synthesis. We apply this analysis to eight UDMs in two different locations.

We discuss UDM synthesis using graph theory terminology:

- A graph represents the UDM, consisting of nodes (manholes) that are connected by edges (pipes). The graph can be either undirected or directed, indicating that connections are between two nodes (undirected) or from one node to another (directed). Pipes are undirected since head may drive uphill flow, however, treating the graph as directed enables better description of preferential flow paths and is also the format required by SWMM, SWMManywhere makes use of both directed and undirected graphs.
- In graph theory, a subgraph refers to a sub-selection of the graph and is called a connected component if every node is in the subgraph reachable from every other node within that subgraph. Connected components are used in SWMManywhere to describe a collection of manholes and pipes that drain to a common outfall.
- The topology of the graph captures the overall arrangement and relationships of nodes and edges and can be thought of as the pipe layout in a UDM. The edges can have costs, for example length, which may be minimised by shortest path algorithms to minimise flow routes. Flow routes should be minimised in any drainage network to drain the catchment as efficiently as possible, preventing the accumulation of water and flooding of manholes. A minimum spanning tree is a subgraph that connects all nodes with the minimum possible total edge cost for an undirected graph, a minimum spanning arborescence is the same for a directed graph, thus these represent the most cost-efficient pipe layouts that may exist for a given area. We will refer to a graph of potential pipe-carrying edges as the street graph, while the optimised layout is the UDM topology.

2.1. SWMManywhere

2.1.1. Data and preparation

Our SWMManywhere approach uses automatically acquired datasets that are of sufficient level of detail to be used in global application: road

locations, river locations, elevation, and building footprints. We stress that these datasets, in particular road locations and building footprints, vary in quality from location to location, thus, the downloaded data for a given case study should always be inspected carefully. Additionally, we enable manually sourced, higher quality data for input, if available, as described in the online documentation (SWMManywhere documentation, 2024). The default datasets are visualised for one of our case studies in Fig. 2, and described in further detail in this section.

2.1.1.1. Streets and rivers. Assuming pipes can only exist in pre-specified plausible locations will dramatically reduce the dimensionality of the shortest-path UDM topology derivation. Thus, we assume that pipes can only exist in certain locations, typically streets due to the common validity of this assumption (Mair et al., 2017). In addition, paved streets constitute one of the key impervious surfaces which must be drained in a UDM. Rivers are also downloaded as these are potential outfall locations of the drainage network, described in further detail in Section 2.1.2. OpenStreetMap (OSM) provides street and river data worldwide and is used in this study.

2.1.1.2. Impervious areas. To calculate the impervious area of a sub-catchment, and thus enable runoff generation, we use road locations, their number of lanes, which are contained within OSM, and building footprints. Recent advances in machine learning have enabled identification of building perimeters from high-resolution satellite data worldwide. Two large datasets, provided by Google and by Microsoft, are now combined and available at (Google-Microsoft Open Buildings, 2024). Both datasets use convolutional neural networks to classify pixels as buildings (Tan and Le, 2019; Sirko et al., 2021) and custom methods to polygonise these pixels into buildings (Sirko et al., 2021); Computer generated building footprints for the (United States, 2024). The authors would note that, although global, this dataset has varying quality in different regions due to the nature of the training data that was used and resolution of satellite imagery. Although other impervious surfaces besides roads and buildings, such as car parks, that must be drained in a UDM are common, we could not identify global datasets describing these, and thus consider them currently outside the scope of this study.

2.1.1.3. Elevation. Elevation data, in the format of a Digital Elevation Model (DEM), is essential to delineate manhole sub-catchments, to calculate the slope along edges, and to identify if there are paths in the street graph that should be ignored, for example, those at that span different hydrological catchments. We use NASADEM, which is a publicly available radar-based global DEM at 30m resolution (Crippen et al., 2016), from Microsoft Planetary Computer (Source et al., 2022). Although higher resolutions around 2m are recommended for UDM (Arrighi and Campo, 2019), such datasets do not yet exist openly at global scales.

2.1.2. Pipe network generation

A key innovation of our approach is to iteratively process a street graph and gradually transform it into a UDM with full hydraulic design. These processes must take a graph and return a graph. The key tasks that they perform are summarised in Fig. 3, discussed further in this section and used for the experiments in this paper.

Table 1 lists all tuneable parameters of our proposed workflow, with reasonable default values for these parameters and their ranges provided in the online documentation (SWMManywhere documentation, 2024). We carry out an extensive sensitivity analysis (see Section 2.2) to identify any behavioural ranges of these parameters and identify their relative importance in UDM synthesis.

2.1.2.1. Data cleaning and manhole identification. As explained in Section 2.1.1, OpenStreetMap (OSM) is the default data source for obtaining street and river graphs. However, the raw OSM data are not directly suitable for UDM generation, so we perform a variety of data cleaning operations to create a more suitable graph for UDM synthesis, as illustrated in Fig. 3a–b. The first significant process in data cleaning is

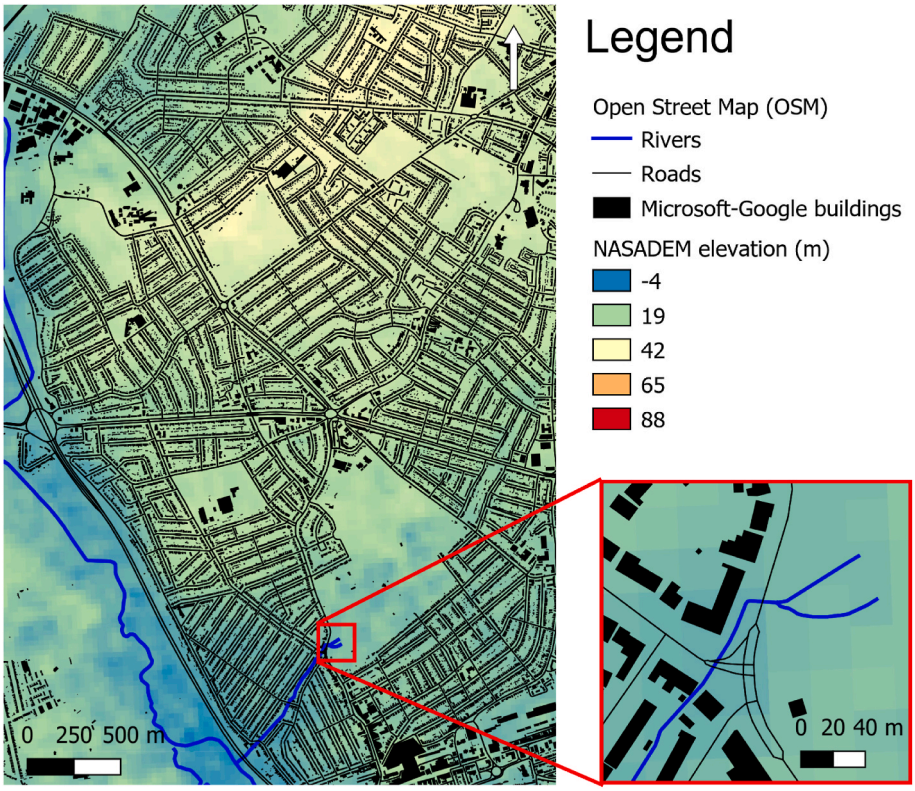

Fig. 2. Visualisation of downloaded data in the Cran Brook, UK. See main text for citations.

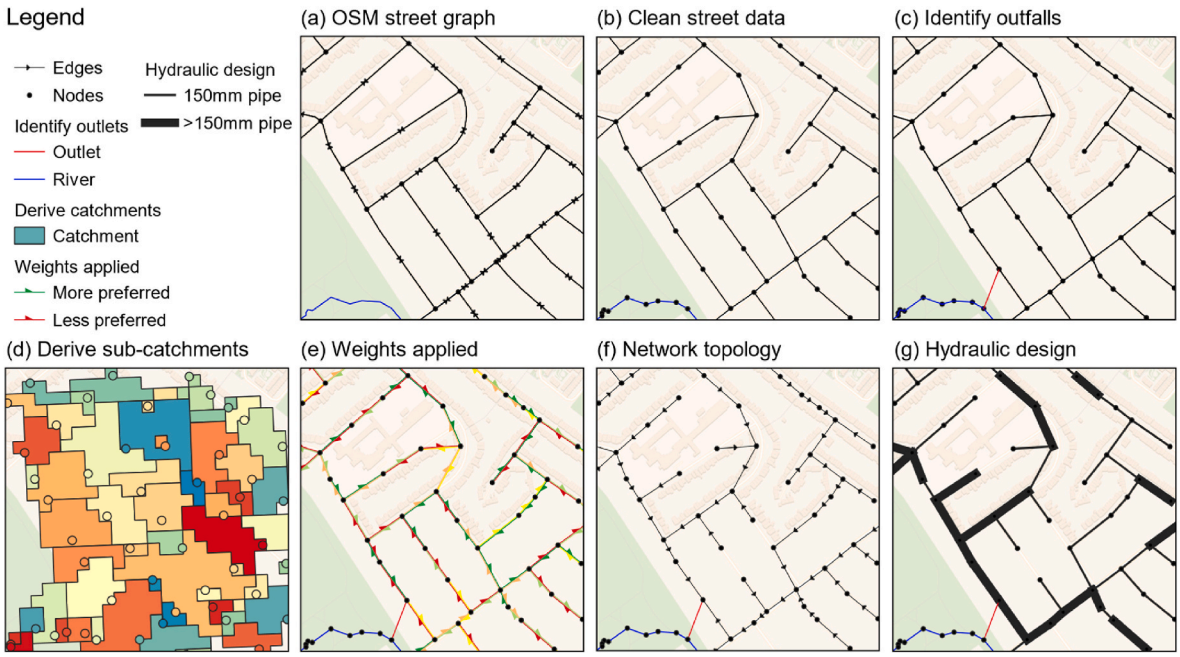

Fig. 3. Visualisation of key iterations to the graph as different processes are applied. Catchments (D) are coloured by the manhole that they drain to.

Table 1
SWMManywhere user adjustable parameters that are tested in sensitivity analysis for this work, a full list of parameters is available in the online documentation (SWMManywhere documentation, 2024). If the variable is described differently from its use in the software to improve clarity, the software term is indicated in brackets.

GROUP	VARIABLE	KEY
SYSTEM DESCRIPTION (MANHOLES AND OUTFALLS)	node merge distance	pNM
	outfall length	pOL
	max street length	pXS
	river buffer distance	pRB
TOPOLOGY DERIVATION	chahinian slope scaling	pSS
	chahinian angle scaling	pAS
	length scaling	pLS
	contributing area scaling	pCS
	chahinian slope exponent	pSE
	chahinian angle exponent	pAE
	length exponent	pLE
	contributing area exponent	pCE
HYDRAULIC DESIGN	max filling ratio (max fr)	pFR
	min v	pMV
	max v	pXV
	min depth	pMD
	max depth	pXD
	design precipitation (precipitation)	pDP

enforcing a maximum edge length, splitting edges that are longer than max street length parameter *(pXS,* Table 1). The second is merging of nodes, which joins nodes together if they are within a specified distance of each other *(node merge distance, pNM*). These two processes jointly control where manholes are located along edges and the frequency with which they occur. The final significant task in data cleaning is buffering street paths in proportion to the number of lanes to create a shapefile of impervious street area. The remainder of processing in this stage is perfunctory, performing tasks such as ensuring consistent geometries, a consistent identification scheme, removing parallel edges, and converting the directed street graph to an undirected graph (as a pipe may flow in a direction opposite to road travel).

2.1.2.2. Sub-catchment outline and surface characteristics. We begin the sub-catchment delineation, by first burning the road network into the DEM of an urban area. This burning process is a common practice in UDM sub-catchment representation (Gironás et al., 2010) and involves lowering the elevation of grid cells in the DEM that contain roads. Then, we hydrologically condition the DEM by breaching depressions (Lindsay, 2016a). Upon conditioning the DEM, we compute the flow direction, using the D8 method (O'Callaghan and Mark, 1984), and slope. Because our workflow includes generation of a SWMM model, a sub-catchment width parameter is required. We follow the approach proposed by InfoWorks (Subcatchment Data Fields (InfoWorks), 2024) to compute the width of a sub-catchment based on the radius of a circle with area equal to the area of the sub-catchment.

2.1.2.3. Outfall identification. Another key feature of a UDM is outfall locations. The first step in identifying outfalls requires assessing the topology of the graph to ensure its hydrologic feasibility (Seo and Schmidt, 2013; Li and Willems, 2020). To achieve this, we remove edges that cross the boundaries of hydrological catchments (defined as the largest non-overlapping drainage basins in the study region), because these are unlikely to carry a pipe. Then, we identify potential outfall locations by assuming that outfalls may only exist within a specified distance of a river (*river buffer distance, pRB*). Although other factors such as environmental considerations affect selection of the outfall location, in this study, we only account for the vicinity of water bodies. We incorporate the relative construction cost of outfalls in our workflow, by assigning weights to the identified outfall locations based the length of the pipe that connects the network to the river (*outfall length, pOL*). If no potential outfalls are identified the node with the lowest elevation is used as the outfall. On the other hand, in cases where multiple plausible outfalls are identified, we retain them all at this step and determine the outfall during the network topology derivation step.

2.1.2.4. Calculating weights and network topology. The network topology can be derived as a minimisation problem of overall graph cost. This minimisation should start with a graph of potential edges (i.e., the graph up to this point) and return a directed graph of edges (pipes) that visit all nodes (manholes), minimising overall graph cost, without retaining redundant edges, which is also referred to as a minimum spanning arborescence (MSA).

The first step to take in network topology derivation is to identify

how each edge contributes to the overall graph cost. As identified by Chahinian et al. (2019), it is plausible that each of pipe length, pipe slope and pipe adjacent angle (the angle at which two joining pipes meet) are important to consider for minimisation. We further propose that the total contributing area carried by pipes in the derived network should also be minimised. As with *ibid.*, these factors do not necessarily contribute to the overall graph cost symmetrically or proportionally. For example, while both negative slopes and overly steep positive slopes are penalised, the penalisation on negative slopes increases with slope more sharply than for positive slopes, because the former becomes hydraulically impractical more quickly than the latter. It is not apparent which of these factors (slope, angle, area, and length) are more important to minimise than others and so we combine each factor to be varied, as in *ibid.* We deviate from *ibid.* by assigning both a linear and exponential scaling parameter to each factor (rather than solely linear), enabling high customisation of how overall graph cost is calculated (i.e., parameters in the *topology derivation* group). We calculate each individual factor, apply scaling parameters, and sum these into an overall cost for each edge in the graph, producing a graph such as that visualised in Fig. 3e.

The network topology is then derived using an implementation based on the shortest-path algorithm proposed by Tarjan (1977), to find the MSA of the graph. The algorithm starts from a designated "waste" node that all potential outfall locations are connected to, either directly or through river paths, thus enabling all connected component subgraphs to be handled in a single pass. The algorithm initializes a priority queue with the waste node's incoming edges, sorted by their costs. At each step, the minimum cost edge is extracted from the priority queue. If the node it leads to is not already included in the arborescence being constructed, that node and edge are added to the arborescence. The node is marked with its parent, and any edges incoming to that node are added to the priority queue. This process continues until all nodes are included in the arborescence. The final arborescence represents the UDM network topology, where the selected edges correspond to the pipes that must be hydraulically designed.

2.1.2.5. Hydraulic design. Duque et al. (2022), propose a "pipe-by-pipe" method to design sanitary sewer network pipes, the method starts at the most upstream pipes, designing each pipe in terms of diameter and depth under a set of design constraints (see *hydraulic design* group), and continues iterating downstream. They demonstrate that this method is comparable to an optimal dynamic programming-based approach, although is significantly more efficient. We adapt the pipe-by-pipe approach to make it suitable for a SWMManywhere approach:

- Rather than deriving the design flow from household waste generation, we use a Rational method that calculates the design flow as the entire impervious area in sub-catchments upstream of the pipe being designed multiplied by a *design precipitation, pDP,* amount.
- Inspection of any large real UDM will commonly reveal pipes travelling in an uphill direction, as measured by surface elevation. Wherever possible the pipe's elevation will be such that they flow downhill despite the surface elevation, however there is no guarantee that any hydraulically feasible design will exist. Because Duque et al. (2022) derive network topology using hydrological flow paths a feasible design will always exist, however this is not the case for SWMManywhere, which uses streets for pipe locations and accounts for factors besides slope during network topology derivation. To accommodate this, we include a surcharge feasibility constraint, which allows a pipe to be designed for flow under surcharge, provided this is the only way to reach a feasible hydraulic design.
- To provide better performance, we assess all designs for a pipe rather than selecting the first feasible design. The selected design first aims to satisfy feasibility constraints, and if no feasible design exists, picking the most feasible design. It then minimizes depth, diameter, and excavation cost, as calculated in Duque et al. (2022).

The final product is a fully described UDM, complete with sub-catchments and hydraulic designs, thus sufficient to be simulated in software such as SWMM.

2.1.3. Measuring effectiveness of UDM synthesis

The question of 'how realistic is a synthesised UDM' is most sensibly assessed by comparing synthesised results against a real UDM. UDM synthesis in a sensitivity analysis context requires understanding why we see the results that we see. Thus, an extensive suite of allowable performance metrics is provided covering a variety of different measures and variables, see Table 2 for a list of the metrics used in this study. We define metrics that measure performance for different elements of UDM synthesis. System description metrics assess the synthesised UDM in terms of properties that describe infrastructure, topology metrics investigate the layout of the graph, and design metrics assess the derived diameter of pipes. Furthermore, the UDM is simulated in SWMM and thus simulated flow, and flooding can also be compared.

Comparing flow and/or flooding simulations is typical in the UDM synthesis literature (Reyes-Silva et al., 2023; Blumensaat et al., 2012). We create a timeseries of total flooded volume to assess flooding simulation performance across the entire network. Meanwhile, flows are assessed at the system outfall, as with (Blumensaat et al., 2012). However, unlike existing literature, which assumes that the outfall locations of the network being synthesised are known, we do not make this assumption, as this information is not globally available. Instead, we identify where synthetic manholes fall inside sub-catchments of the real network. From these classified manholes we identify the most commonly represented outfall, and sub-select only that connected component for comparison purposes.

Because the reasons for performing sensitivity analysis are to understand how parameters change behaviours in UDM synthesis, the most common measure of performance we use is the relative error (*relerror* measure in Table 2, equation (1)), which is simple to understand and provides directionality in terms of over/under estimation,

$$relerror = \frac{mean(synthetic) - mean(real)}{mean(real)} \tag{1}$$

where *synthetic* is the synthetic UDM data to be compared against the *real* data. We omit a conventional *time* component of the metric because the same equation can equally be used for timeseries or design properties (such as average diameter) alike. In cases of comparing flow or flooding timeseries, we also include the Nash-Sutcliffe Efficiency and Kling-Gupta Efficiency, because these are commonly used and so will provide users who are familiar with them a more nuanced grasp of the

Table 2
List of metrics implemented in SWMManywhere.

CATEGORY	MEASURE	VARIABLE	KEY
SYSTEM DESCRIPTION	relerror	length	mRL
	relerror	npipes	mRP
	relerror	nmanholes	mRM
TOPOLOGY	deltacon0	–	mD0
	laplacian distance	–	mLD
	vertex edge distance	–	mVD
	kstest	edge betweenness	mKE
	kstest	node betweenness	mKN
DESIGN	relerror	diameter	mRD
	kstest	diameter	mKD
SIMULATION	nse	flow	mNQ
	kge	flow	mKQ
	relerror	flow	mRQ
	nse	flooding	mNF
	kge	flooding	mKF
	relerror	flooding	mRF

synthetic UDM's performance.

A further set of measures that can be used for synthetic networks are those that test the topological similarity of the derived vs real network, we implement a variety of those presented in existing literature (Wills and Meyer, 2020; Chegini and Li, 2022), see *topology* category in Table 2.

2.1.4. Implementation

SWMManywhere is a highly modular workflow and in this study, we implement it in Python and publish it as an open-source tool (Dobson et al., 2024a). We note that the workflow is general and can be implemented in any other programming language. In our implementation of SWMManywhere the minimal required user input is the bounding box of the target urban area, and all remaining steps Fig. 1 are automated. The bounding box should be provided in terms of latitudes and longitudes in WGS 84 geographic coordinate system (EPSG:4326). SWMManywhere will reproject all downloaded data into the Universal Transverse Mercator (UTM) coordinate system. UTM uses a coordinate system with metre as its unit, and thus can provide accurate distance and area calculations in contrast to WGS 84. The UTM is split into zones and the zone ultimately used in a SWMManywhere run is calculated based on the UTM zone of the bounding box.

As we described in Section 2.1.2 the most complex step in the workflow is the pipe network generation, i.e., iteratively applying various graph operations to the initial graph street to generate the final UDM. These operations, referred to in our implementation as graph functions, have a variety of parameters that need to be specified, as listed in Table 1. Considering the importance of graph functions and to accommodate flexibility in applying them, our implementation allows adding, removing, or changing their order without modifying the code. Structuring code into graph functions enable easy reuse of code, customisation, and introduction of new processing steps. Graph functions are wrapped in a class for validation, enabling SWMManywhere to identify if a set of graph functions to be applied is valid a priori. Graph functions are stored in a registry object to enable easy access. We provide a description of all graph functions in the documentation online (SWMManywhere documentation, 2024). However, users may also customise the selection and order of graph functions, or create new ones, as described in the online documentation, to fit their requirements.

Processes or operations described in Sections 2.1.1 and 2.1.2 that are used by but not implemented natively within the tool are described in Table 3.

To accommodate ease of use for a wide range of users with different levels of programming experience, we provide a command-line interface (CLI) for the software. More experienced users can take advantage of the modularity of SWMManywhere for more advanced customisation of the workflow. The CLI works with a configuration file that enables a user to change parameter values and functionality. The minimal essential requirements that this file must contain are a project name, a base directory, and a bounding box. The configuration file provides a centralised location to perform customisations including changing the selection/ordering of graph functions, changing parameter values (see Table 1), file locations of a real network to compare against (see Section 2.1.3) and which metrics to calculate (see Table 2), a starting graph if not using downloaded street data, and any running settings for the SWMM simulation. A variety of online tutorials explain the procedure to make such customisations.

SWMManywhere provides a capability to write simple SWMM model files (typically with a .inp file extension) that have been synthesised via the various graph functions used. It also provides a wrapper of the PySWMM software (McDonnell et al., 2020), which enables calling the SWMM software and interacting with its simulations from Python. Thus, in addition to the UDM synthesis, writing, running, and calculating metrics if real network information exists are carried out during the command line call.

Precipitation data is frequently identified as a critical factor in UDM simulations (Ochoa-Rodriguez et al., 2015), however, there are currently no open global datasets that provide the high frequency monitoring needed to drive these models and so precipitation must be user-provided if deviating from the default storm provided as part of the tool.

Table 3
List of tools for specific tasks in our SWMManywhere implementation.

Task	Software	Reference
DEM conditioning	Whitebox	Lindsay (2016b)
Flow direction calculations	Whitebox	Lindsay (2016b)
Sub-catchment delineation from flow directions	PyFlwDir	Eilander (2022)
Sub-catchment slope calculation	PyFlwDir	Eilander (2022)
OSM data retrieval	OSMnx	Boeing (2017)
Graph operations	Networkx	Hagberg et al. (2008)

2.2. Sensitivity analysis

Sensitivity analysis is a critical tool for assessing how variations in model outputs can be attributed to variations in input parameters, often quantified through sensitivity indices that describe the relative importance of each parameter (Pianosi et al., 2016). In this study we employ the Sobol method (Sobol, 1993), a widely adopted global sensitivity analysis known for its robustness, particularly when a sufficiently large number of samples is used to ensure reliable results (Pianosi et al., 2016). The Sobol method is a model-independent technique that decomposes the total variance of the output into contributions from individual input parameters and their interactions (Saltelli et al., 2008, 2019; Pianosi et al., 2016). In this study, we use the Sobol method to calculate first-order indices, which quantify the direct impact of individual parameters on the output; second-order indices, which represent the joint effects of parameter pairs; and total-order indices, which encompass the overall contribution of a parameter, including its individual impact and all interactions with other parameters. This decomposition enables a comprehensive understanding of the parameter contributions to model variability, making the Sobol method particularly suitable for exploring complex, high-dimensional systems. As defined in Table 1, there are a wide variety of parameters that must be selected, many of which do not have values that could be easily measured or derived, and thus are useful candidates for sensitivity analysis. We note that SWMManywhere is a workflow rather than a model, however, sensitivity analysis is equally applicable to a parameterised workflow as to a model.

In general, it is recognised that, to robustly conduct sensitivity analysis, a global method should be used, and the variability of the calculated indices should be checked to ensure that the number of samples is sufficient (Saltelli et al., 2019). Because of the presumed high level of dependency in the SWMManywhere workflow (for example, hydraulic design depends entirely on network topology, which in turn depends on outfall and manhole identification), we also calculate second order indices to better capture interactions between parameters (Herman and Usher, 2017). To implement sensitivity analysis for SWMManywhere, we use the SALib software (Herman and Usher, 2017), which provides a variety of global methods for sampling parameter ranges and calculating sensitivity indices natively in Python. 18 parameters are sampled, indicated in Tables 1 and 16 metrics, Table 2, are evaluated. Thus, the overall approach is to perform parameter sampling, run SWMManywhere with the parameters of each sample, calculate the performance metrics between the synthesised and real UDMs, and calculate sensitivity indices.

Sobol sensitivity analysis that includes second order interactions with SALib requires taking samples equal to,

$$N = n^*(2m + 2) \tag{2}$$

where N is the total number of samples (or SWMManywhere calls), m is the number of parameters, and n is the number of Sobol sequence samples to generate (preferably a power of 2). In this experiment we set n to 2^{10} (1024), m is 18, resulting in N of 38912. As demonstrated in Pianosi et al. (2016), this many evaluations (m * 1000) are towards the upper limit of what is found in the literature. We take this opportunity to note a further benefit of sensitivity analysis in the context of SWMManywhere as a global tool, which is that testing it under such a large and diverse range of parameters further guarantees robustness of the software implementation in locations not tested.

3. Case studies

In this study, we evaluate our proposed workflow by comparing the SWMM simulation results obtain from using the synthetic UDM with those of the real UDM for the Cran Brook, London, UK (Babovic and Mijic, 2019). We then perform sensitivity analysis in other locations to examine the transferability of results and parameters. We use seven UDMs around the town of Bellinge, Denmark, as delineated by Farina et al. (2023). These data were selected because they are openly available and demonstrate results over a wide range of scales (Pedersen et al., 2021). The properties of the case study networks are presented in Table 4.

A key feature of many UDM is the presence of hydraulic structures such as weirs, orifices, storages, or pumps. There is extensive evidence from the sewer network simplification literature that capturing and parameterising these structures is critical towards reproducing the behaviour of the real network (Thrysøe et al., 2019; Dobson et al., 2022). However, SWMManywhere currently does not attempt to estimate the locations or hydraulic properties of any such structures. We acknowledge that this could be a significant limitation and hope to add this behaviour in future work. In this paper, we replace the hydraulic structures and storage nodes in the real models with simple and uniform nodes to better assess the SWMManywhere workflow as designed. Furthermore, the hydraulic properties of sub-catchments (specifically, the Manning's roughness coefficient and depression storage of both impervious and pervious areas) are a significant source of uncertainty and typically calibrated or set arbitrarily (Deletic et al., 2012), thus we set these at the same values for all networks to ensure comparability.

The precipitation event we use to demonstrate SWMManywhere is the largest storm in the openly available Bellinge data (Pedersen et al., 2021).

While SWMManywhere is designed to run on consumer-grade computer systems, the computational demands of sensitivity analysis, which require many model evaluations, necessitated the use of High-Performance Computing environments. These systems provided the processing power and scalability needed to efficiently execute the extensive simulations required for robust sensitivity analysis. Workflow evaluations were performed on the Imperial College High Performance Computing facilities (see Acknowledgements). Hardware used was typically AMD EPYC 7742 (128 cores, 1 TB RAM per node), although this varied based on availability. On a single core of one of these machines, for a single workflow evaluation in the Cran Brook case study, downloading and data preprocessing takes 10 s (except for buildings, which are national datasets and so download speeds will vary significantly depending on the country), evaluating graph functions and writing the UDM to SWMM format takes around 2 min (deriving sub-catchments and deriving network topology are the slowest individual steps at 20 s each), simulation in SWMM takes 8 min, and evaluating metrics takes 2 min (dominated by the *K-S test for node betweenness*, *mKN* which took over 1 min). Run times for all Bellinge case studies were dramatically quicker owing to the far smaller UDM sizes.

The code used to perform this experiment and results required to reproduce the figure results in Section 4 are openly shared in a separate repository (Dobson et al., 2024b).

Table 4
Summary of networks tested and their properties.

Network	Number of nodes	Number of edges	Impervious percentage
Cran Brook	6931	6965	27
Bellinge 1	142	150	36
Bellinge 2	118	117	33
Bellinge 3	52	51	25
Bellinge 4	46	45	32
Bellinge 5	45	46	40
Bellinge 6	36	35	32
Bellinge 7	15	14	34

4. Results

4.1. Proof-of-concept examination

We focus this section on a synthesised model selected from our sensitivity analysis sampling (Section 4.2) that performed well across a range of metrics to assess SWMManywhere's ability to generate a high-quality UDM and raise methodological points of interest. Fig. 4 plots a flow and flooding timeseries and diameter, elevation, slope, and travel time distributions for the Cran Brook network to provide a detailed comparison of the real and synthetic UDM.

Fig. 4a shows flow at the network outfall, while Fig. 4b shows the total flooded volume across the network. We see that the maximum values of both are captured with accuracy, however, the falling limb for both recedes more quickly in the synthetic network than in the real UDM simulations. The synthetic network has consistently larger diameters (Fig. 4f) than the real, while chamber floor elevations (Fig. 4d) are well matched. Synthetic pipe slopes (Fig. 4e) are lower, although we observe that these are primarily within the grey dashed lines which show the target design range (Chahinian et al., 2019). The average travel time from each node to the outfall (Fig. 4c) shows distinctively different patterns across the distribution, with a good match for the quickest third of nodes, the synthetic UDM quicker for the middle third (because of the larger diameters), and the real UDM quicker for the slowest third. Although not shown, the total runoff from manhole sub-catchments in both the real and synthetic models is within 2% of each other, which is true for all synthesised UDMs.

4.2. Sensitivity analysis, Cran Brook

SWMManywhere parameters, see Table 1, were sampled using a Sobol sampling scheme, see Section 2.2, to enable a global sensitivity analysis using the Sobol method. The findings of this analysis for Cran Brook, the sensitivity indices, are presented in Fig. 5a–b. Two other locations (Fig. 5c–f) are discussed in Section 4.2, with other locations in full presented in Supplemental Fig. S1.

In Fig. 5 we show the first order variance (S1) of each metric (as listed in Table 2) attributable to each parameter (Fig. 5a), and the total variance (ST) attributable (Fig. 5b). We see that sensitivity is widespread, with every parameter exhibiting at a total variance attributable of >1% for at least one metric. We also see that sensitivity is overwhelmingly occurring through interactions (i.e., sensitivities present in Fig. 5b but not in Fig. 5a), with some notable exceptions: *node merge distance (pNM), minimum velocity (pMV), minimum* (pMD) and *maximum chamber depth (pXD)*. First order sensitivity indicates that the parameter is sensitive regardless of other parameter values, *pMV, pXD,* and *pMD* are evidently dominant in their influence on pipe design while *pNM* is discussed below. Second order variance indices were also calculated, but not included because their confidence intervals were prohibitively large.

Across metrics the *node merge distance* is the most sensitive parameter, impacting both system description metrics (top rows), topology metrics (middle rows), and simulation metrics (bottom rows). It is a sensitive parameter because it interacts with three key elements in SWMManywhere:

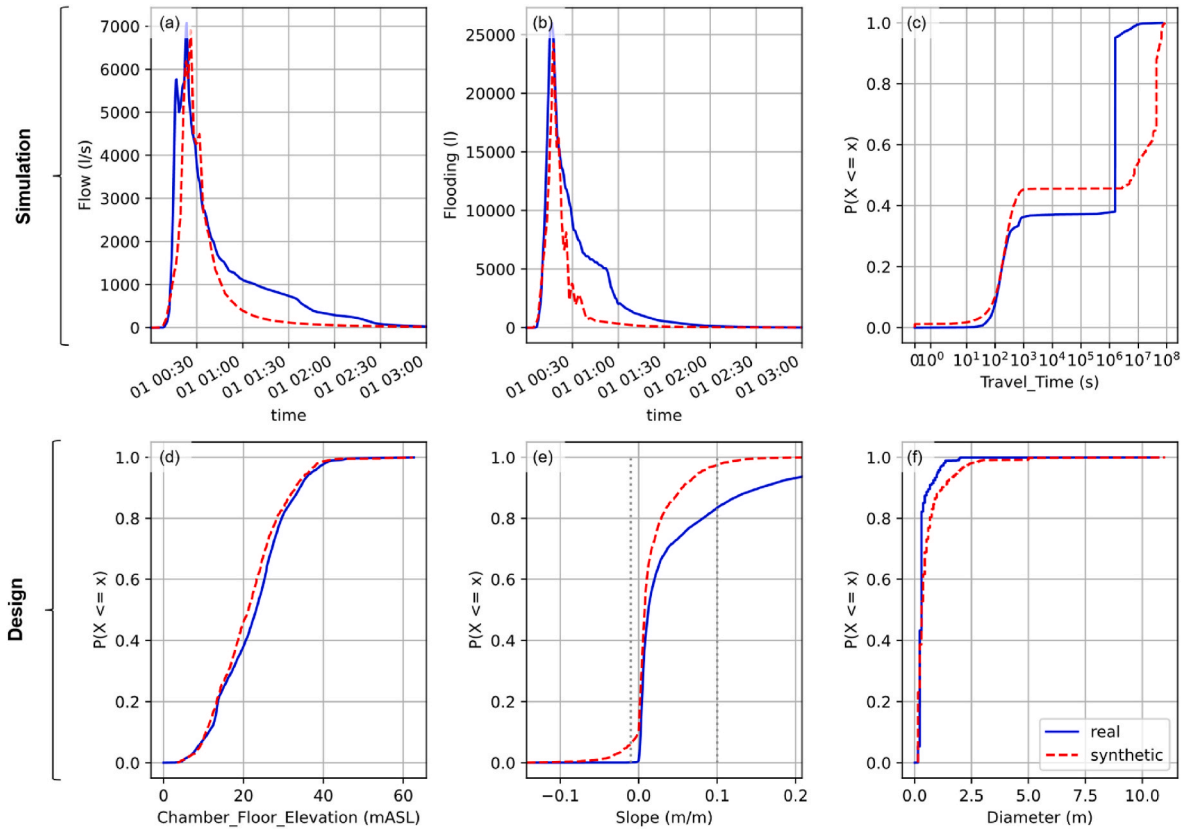

Fig. 4. Demonstration plot of a high performance synthetic UDM in the Cran Brook case study. Red dashed lines represent synthetic data, while solid blue represents the real UDM simulations. Grey lines on the slope plot (e) show the target design range. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

- It influences manhole placement, which is also impacted by *max street length* (*pXS*, another sensitive parameter).
- It impacts which nodes can drain to where, which is also impacted by *river buffer distance* (*pRB*, another sensitive parameter).
- It can significantly alter how the road layout is translated into potential pipes, which no other parameter does. For example, two adjacent roads running parallel may not have a driving connection between them, and thus no potential pipe may span them, but if these nodes are merged then a pipe could span them.

In addition to the dominance of these three parameters (*node merge distance*, *max street length* and *river buffer distance*), we see many other intuitively appropriate sensitivities. System metrics are sensitive to system parameters, topological metrics to topological parameters, and design metrics to design parameters. We also find that flow simulation metrics are sensitive to parameters relating to network topology, indicating that topology is primarily influencing the global behaviour of the UDM. In contrast, we see that flood simulation metrics are more sensitive to design parameters, indicating that design is primarily influencing the local behaviour of the UDM.

We provide a specific examination of *node merge distance*, selected because it is the most sensitive parameter, to illustrate how SWMManywhere can be used to provide a detailed parametric exploration. We plot the parameter value of *node merge distance* against each metric value for all sampled points in Fig. 6. We see clear evidence of first order sensitivity in line with those reported in Fig. 5a–b.

Fig. 6 indicates some evidence of identifiable parameter values, e.g., we observe that the *node merge distance* should fall around 25m for the simulation flow metrics. However, this value can vary depending on which metric is used, for example, System metrics (*mRL, mRP,* and *mRM*) and some topology metrics (*mVD, mKE, mKN)* indicate better performance when *node merge distance* is around 0–15m.

It can be more informative to investigate whether parameters are identifiable using a Gaussian kernel density estimate (KDE), which can be used to create a cumulative density function for each parameter and can be weighted by different metrics. For example, in Fig. 7, top left panel (node merge distance, *pNM)*, weighting the KDE by flow simulations (red lines), we can see that 75% of the distribution indicates that *node merge distance* should be greater than 20m, agreeing with Fig. 6. The KDE plots further highlight identifiability for a range of other parameters, however, as with *node merge distance*, in most cases there can be disagreement depending on which metric is used for weighting.

4.3. Sensitivity analysis, other locations

In Fig. 5c–f we show the sensitivity analysis results for the Bellinge 1 and 7 networks, which are the largest and smallest networks in the Bellinge dataset, respectively. Results for the other Bellinge UDMs are shown in Supplemental Fig. S1, although do not meaningfully differ from the observations below.

In general, we see many similarities between the two analyses, including high sensitivity of *node merge distance*, more sensitivity occurring through interactions (Fig. 5 c and e) than in first order (Fig. 5 d and f). In contrast, we see less significance of network topology parameters across all metrics, which can likely be attributed to smaller sizes of both Bellinge networks compared to the Cran Brook network resulting in availability of fewer topological configurations, thus less significance in terms of impacting metrics. We note that no indices for flooding metrics could be calculated because the networks did not experience flooding under the precipitation timeseries used.

In Fig. 8 we show the KDE estimates of parameter distributions for each network, weighted by the NSE flow metric (*mNQ*). In contrast to

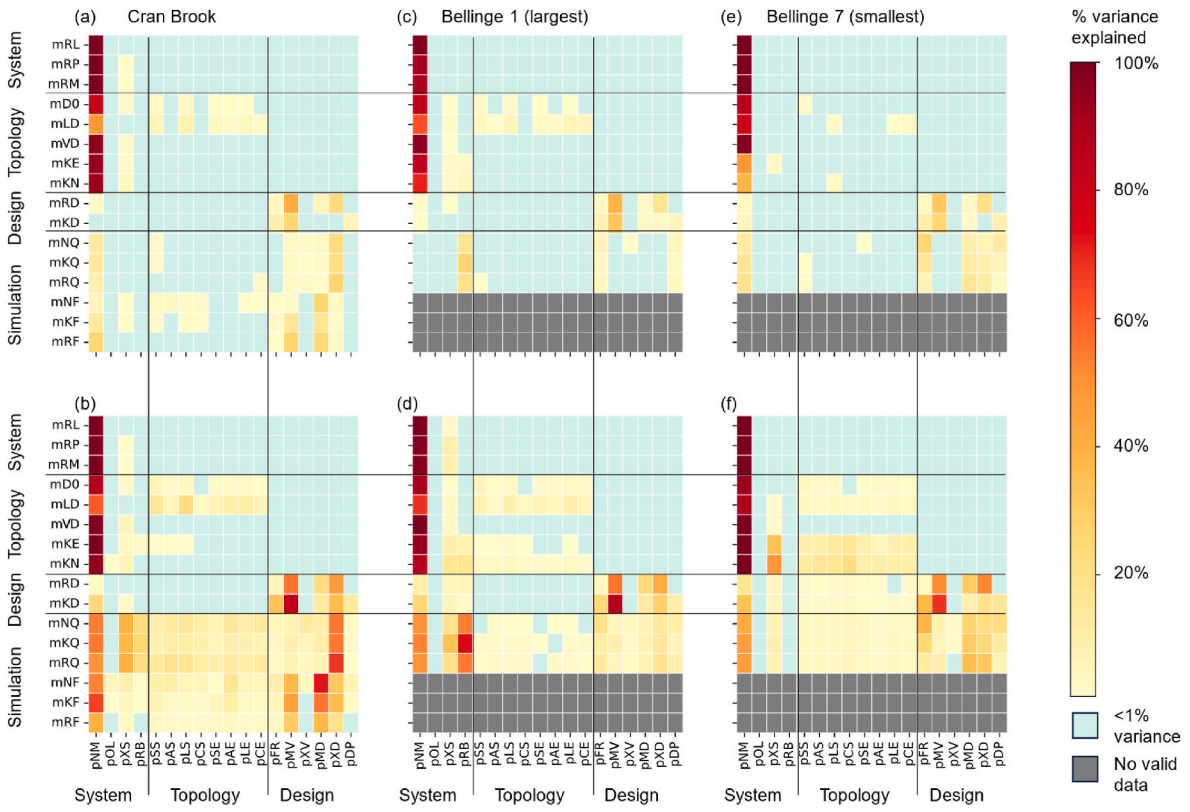

Fig. 5. Heatmap demonstrating the sensitivity indices (a) of first order variance (S1) of a metric (y-axis) attributable to a parameter (x-axis) based on simulations in the Cran Brook network, (b) of total variance (ST) attributable to a parameter based on simulations in the Cran Brook network. Red indicates more sensitive and yellow less sensitive. Blue indicates less than 1% variance explained. (c, d, e, f), shows equivalent of (a) and (b) respectively, but for the Bellinge 1 and Bellinge 7 networks. Grey indicates no sensitivity indices could be calculated because no flooding occurred. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

Fig. 7, we see stronger agreement across networks than across metrics for weighting. For example, *max fr (pFR), precipitation (pDP), max depth (pXD),* tend towards aligned parameter distributions across most networks. Despite this, we also see some sensitive parameters with less agreement across networks. For example, min *v (pMV), max street length (pXS),* and *node merge distance (pNM)* have unaligned but identifiable parameter values. We also see that many parameters do not have clearly defined distributions for most networks but do for one/few. For example, *river buffer distance* appears to have a clearly defined distribution for Cran Brook and Bellinge 2, while no other networks do. We do not see clear parameter distributions for network topology parameters, despite their sensitivity demonstrated in Fig. 5.

5. Discussion

5.1. Development of SWMManywhere

In this study, we provide new insight into the intricacies of UDM synthesis through application of our proposed SWMManywhere workflow and performing extensive sensitivity analysis.

We demonstrate in Fig. 4 that reasonable simulations are achievable, with NSE values of >0.7 for both outfall flow and flooding, Fig. 6. The implementation based on Duque et al. (2022)'s pipe-by-pipe method results in high efficiency UDMs, although we observe that the synthesised UDMs are perhaps too efficient as they drain the network too quickly, see flow and flood simulations in Fig. 4. This finding suggests that inefficiencies seen in the real network reflect the additional constraints not captured by data sources that we employ in our workflow, for example, incremental construction of the network (Rauch et al., 2017). Despite such shortcomings, we recommend that sensitivity analysis is a necessary precursor to introducing any further complexity.

The sensitivity analysis results, Fig. 5, show widespread and intuitively sensible sensitivity of all parameters, including those introduced as technical innovations of this paper: enabling slope inclusion within the MSA and contributing area as a factor in the network topology derivation (*pSS, pCS, pSE, pCE*). We also highlight the dominant sensitivity of *node merge distance* (*pNM*), which influences manhole locations, outfall locations, and underlying street graph preprocessing. Manhole locations are well established to be important factors (Chahinian et al., 2019; Blumensaat et al., 2012), while to our knowledge this is the first UDM synthesis study that treats outfall locations as an explicit unknown. The importance of the underlying street graph that *node merge distance* controls is intuitively sensible, as a UDM synthesis can only ever be as good as the potential pipe-carrying locations that it begins with. These findings indicate a promising outlook for UDM synthesis, as manhole locations, outfall locations, and the street graph are all surface elements whose estimation will only improve with improved satellite imagery and machine learning. Furthermore, in cases where SWMManywhere is to be applied to a local area, surveying these surface elements is likely to be far less costly than a below-ground network survey.

5.2. Transferability of SWMManywhere

A key motivation for performing sensitivity analysis is in identifying behavioural parameter ranges. In an ideal case, behavioural parameter ranges align across metrics and locations, thus implying that parameter choices are good under any condition. Figs. 7 and 8 demonstrate that we do not see clear evidence of this for either metrics or locations respectively. Fig. 5 demonstrates that parameters are generally sensitive through interactions, rather than through first order effects, which limits

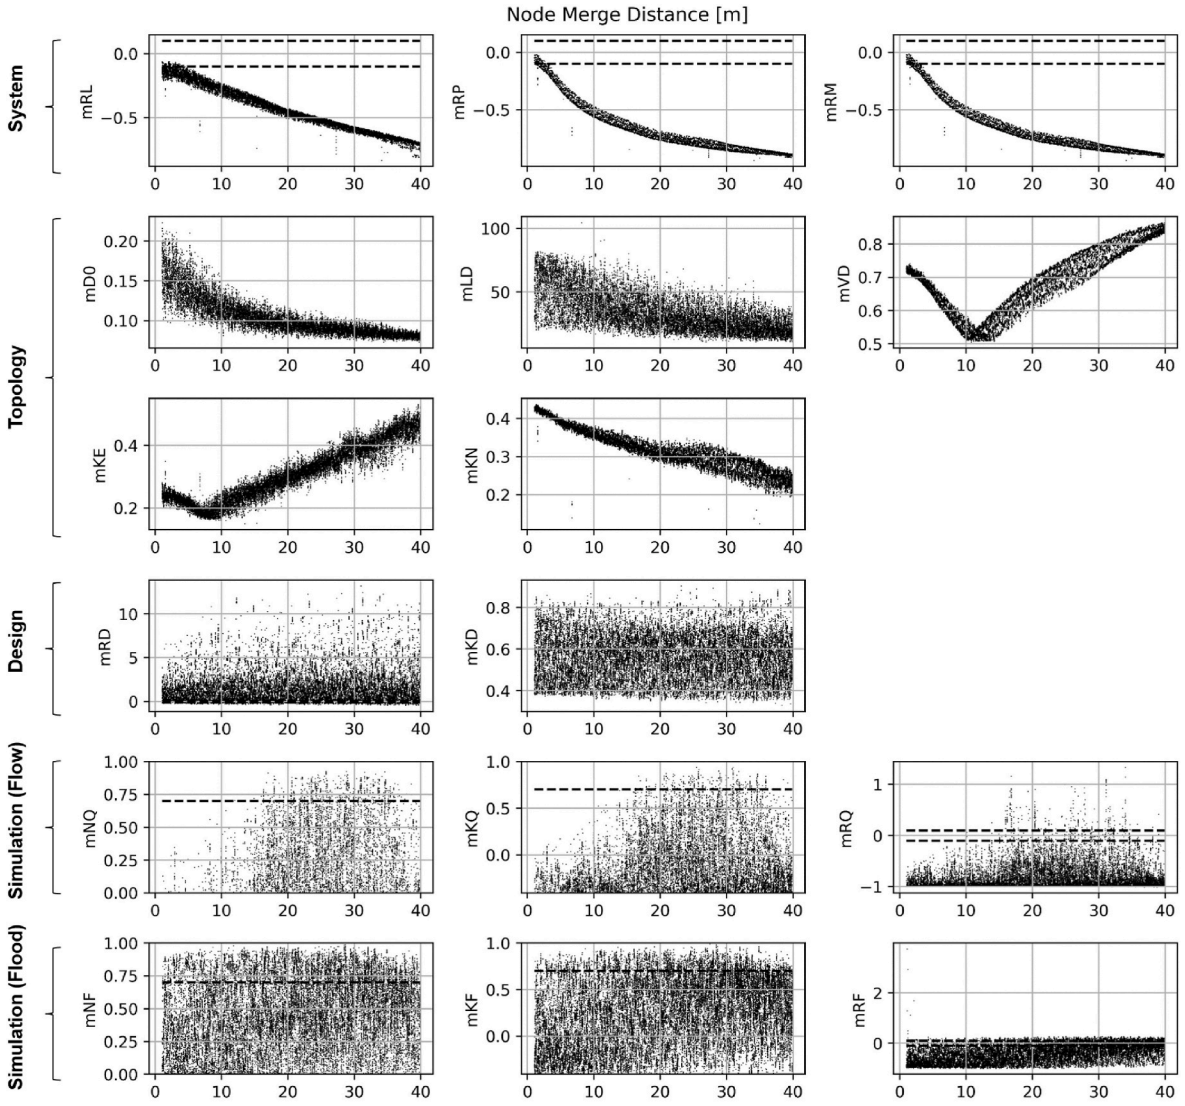

Fig. 6. Node merge distance parameter value plotted against all evaluated performance metrics. Based on simulations in the Cran Brook network. Panels showing NSE or KGE are clipped at 0.0 and −0.41 respectively, which indicates the performance of taking the mean value of observations as the simulation. Dashed lines for KGE, NSE, relative error indicate a region of 'behavioural' performance, that is, values > 0.7 (KGE, NSE) or within ± 0.1 (relative error).

the ability to provide clear advice on behavioural ranges. Although second order effects were calculated, the confidence intervals were such that no conclusions could be drawn, and we estimate that multiple magnitudes more samples would be required to provide definitive results, which is outside the scope of this paper, but future work may investigate. Parameters that have high first order sensitivity provide some clear advice, for example, *node merge distance* values should be between 20m and 35m to provide good flow simulation metrics in a variety of locations. However, the dominant finding on parameter transferability is that the panacea of arriving to a single 'correct' UDM through a synthesis approach is false. Indeed, we believe that widespread findings of equifinality in hydraulic calibration of real UDMs (Huang et al., 2022; Sytsma et al., 2022), supports that the idea of having a single 'correct' UDM that is based on survey information (rather than synthesis) is also false. We propose that data uncertainty will remain a fundamental element of UDM, synthetic or otherwise, for the foreseeable future.

We therefore advocate for an approach to UDM that is uncertainty driven. Rather than narrowly focussing on aligning synthetic with real UDMs, a synthetic UDM may most appropriately be considered one hypothesis for the underlying system. Further developments to SWMManywhere should thus seek to synthesise UDMs that are plausible hypotheses, possibly through a more iterative approach that refines the UDM based on simulation data only, for example, by assessing SWMM continuity errors. In a no field data setting, the focus could shift to, for example, various climate scenarios or future urban development, to study an ensemble of plausible UDMs and explore their likely outcomes. If such studies would prohibitively increase simulation time, a range of network complexity reduction techniques have been demonstrated for UDMs (Farina et al., 2023; Palmitessa et al., 2022), including as part of an uncertainty ensemble approach (Dobson et al., 2022; Thrysøe et al., 2019).

To ensure that our proposed workflow is applicable anywhere, we have deliberately avoided some data, for example design regulations or a higher resolution DEM, that are commonly available at national scales. In local applications, however, we do recommend exploring utilising the best quality data available and adapting the underlying datasets or parameter assumptions, which our SWMManywhere implementation supports. Nevertheless, we caution that this may not provide the certainty that one might expect. For example, as our KDE parameter estimates demonstrate in Bellinge, Denmark (Fig. 8), design regulations, which are parameters that are country and region specific, do not show

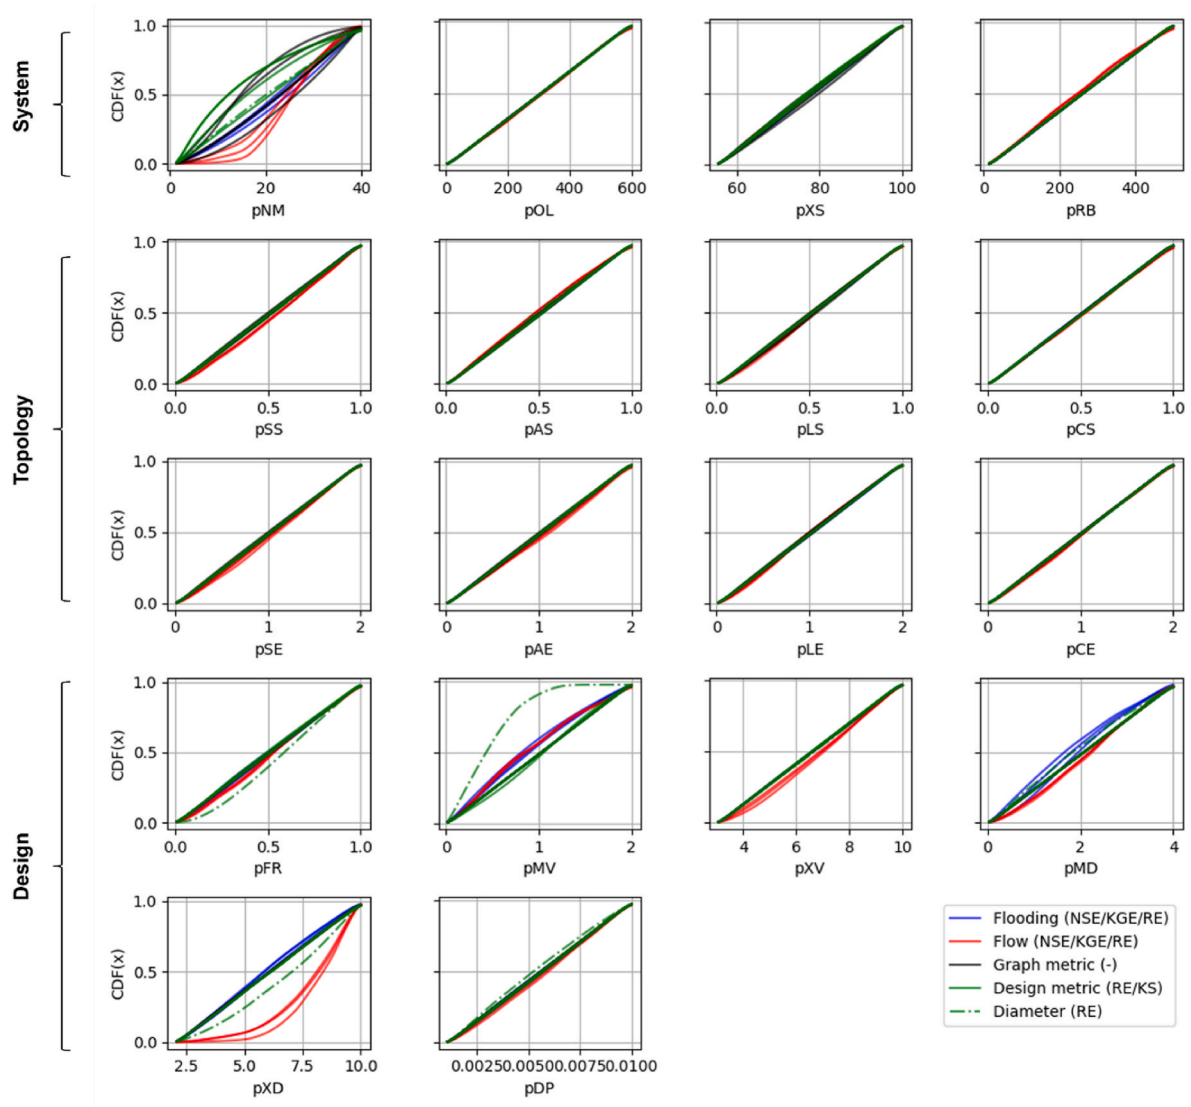

Fig. 7. Gaussian KDE cumulative density functions for parameters, weighted by different metrics evaluated on the Cran Brook network. Metrics are grouped by colour to highlight different behaviours. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

agreement across UDMs in the same locale.

5.3. Outlook and limitations

The modular graph function-based architecture of SWMManywhere makes it easy to extend or customise, thus we hope that it may become a centralised location for the synthetic UDM community. We see a variety of potential improvements that may be introduced to increase the realism of synthesised UDMs, although we stress the importance of performing sensitivity analysis before investing time in creating complicated customisations. As highlighted by the importance of pipe-carrying locations above, a key starting point for reducing uncertainty would be to improve street location surveying, particularly in regions poorly served by OSM. In addition, other surface factors such as hydraulic structures are well established to play a dominant role in UDM behaviour (Thrysøe et al., 2019; Dobson et al., 2022). SWMManywhere does not attempt to synthesise these structures, and they are omitted from the analysis performed in this paper. However, structures such as weirs may be identifiable from satellite imagery, while including pumps may form an additional element in the network topology derivation, as has been demonstrated for sanitary sewer networks (Khurelbaatar et al., 2021).

An immediate extension to the behaviour of SWMManywhere that we are exploring is around water quality. Recent literature demonstrates the feasibility of quantifying urban pollution deposition on roads (Revitt et al., 2022), which could be included as an optional extension, thus providing the much-needed transport element to link deposition with in-river pollution. We anticipate that a key difficulty of such an approach would be in identifying validation data, since sampling water quality at urban drainage outfalls during a storm is dangerous to do in person. For example, while the English Environment Agency's harmonised water quality sampling database contains over 60 million pollution sample records since 2000 (Open water quality archive datasets (WIMS), 2024), just 0.4% are of urban drainage outfalls, and only 25% of these have occurred since 2010, reflecting the diminishing focus on non-compliance monitoring observed across England (Dobson et al., 2021). Further improvements towards capturing water quality may also focus on representing combined or misconnected systems, linking with the synthetic sanitary sewer network literature (Duque et al., 2022).

A clear limitation of the case studies demonstrated in this paper is that they are based in temperate and wealthy European countries, this is particularly problematic when considering that street graph and building footprint data uncertainty will be more significant in nearly any other type of region. We reached out to a variety of urban drainage

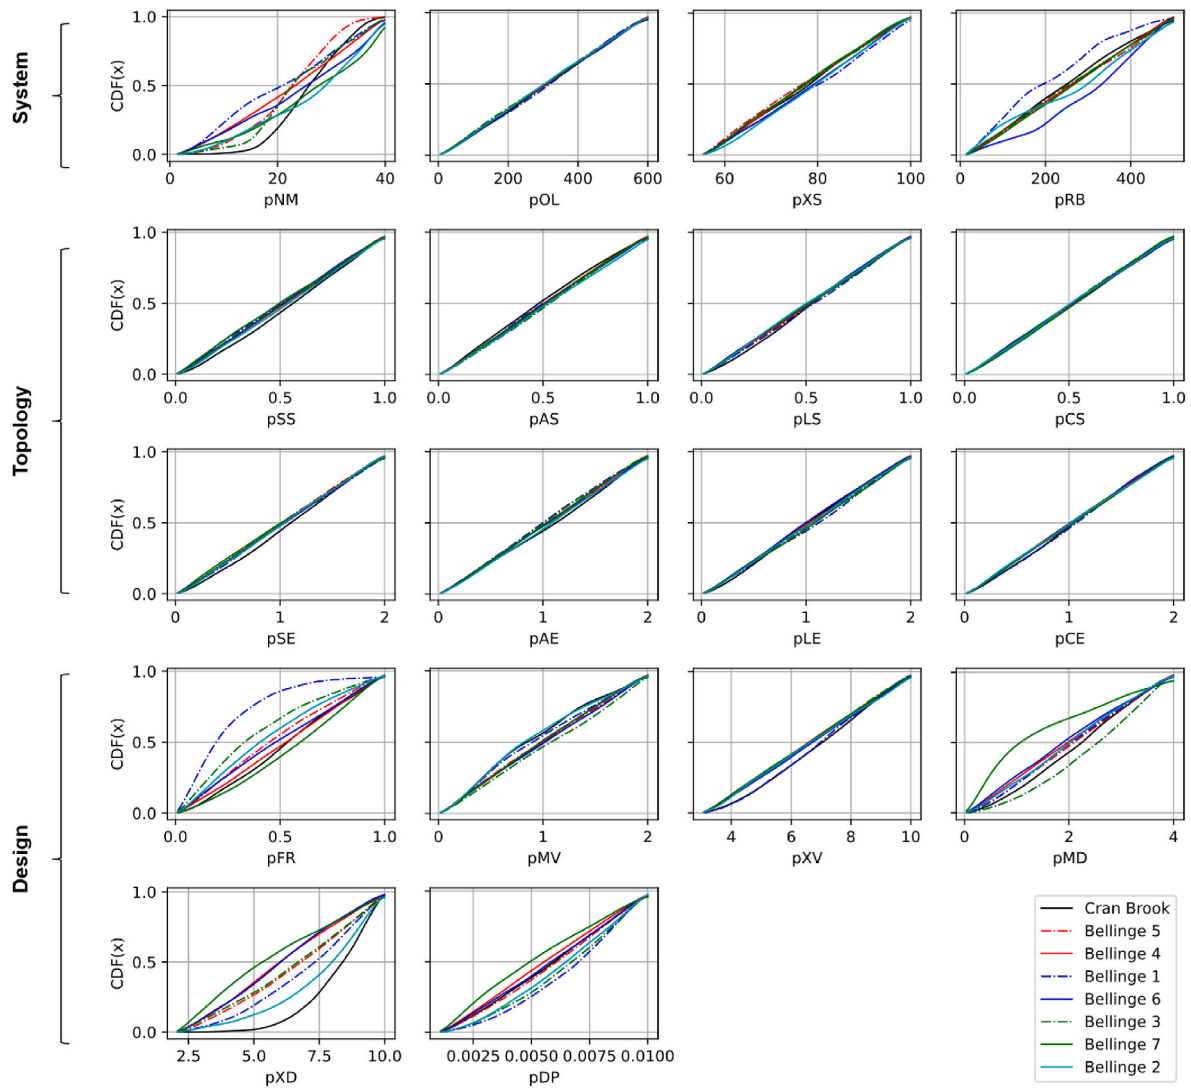

Fig. 8. Gaussian KDE cumulative density function for parameters for each location. Weighted by outfall flow NSE.

modellers in both industry and research but were not able to extend our case study selection further and instead identified a significant paucity of publicly available reliable SWMM models. While the SWMM website hosts a variety of useful example models, these aim to build understanding about representing different elements of drainage networks, rather than providing a suite of real test cases. Ultimately, SWMManywhere and other synthetic UDM tools will not be trusted at global scales until they have been demonstrated on a wide variety of case studies. Thus, we call for collaborators who can either share their SWMM models openly or are willing to demonstrate SWMManywhere for their models to reach out that we may create a more robust demonstration of UDM synthesis and verify that understandings created are sustainable.

6. Conclusion

In this paper we present a workflow, SWMManywhere, that can synthesise an urban drainage network model (UDM) and simulate it in SWMM, anywhere globally. We test the parameters of SWMManywhere using sensitivity analysis to understand the dominant processes involved in synthesising UDMs. Our results revealed three key findings:

1. The SWMManywhere workflow can synthesise high quality UDM at a range of spatial scales. The parameters that can be used to tune SWMManywhere behave in intuitively sensible ways, verifying its implementation.
2. We find that parameters controlling surface elements such as manhole locations, street layout, and network outfalls are the most sensitive, and thus should be the key focus of uncertainty reduction. Encouragingly, the identification of these elements is also the most likely to improve in the foreseeable future.
3. UDM synthesis is sensitive to all parameters and these parameters primarily influence outputs through second order or higher interactions, revealing UDM synthesis to be a more complex process than previously recognised. Additionally, we recommend that an ensemble approach would be appropriate for practical applications to better reflect the inherent uncertainty of the underlying system.

We hope that the urban drainage community will use SWMManywhere to further explore the complexity of UDM synthesis, develop robust interventions in areas without existing UDM data, and do so in a more open and reproducible scientific environment.

CRediT authorship contribution statement

Barnaby Dobson: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding

acquisition, Formal analysis, Data curation, Conceptualization. **Tijana Jovanovic:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Data curation, Conceptualization. **Diego Alonso-Álvarez:** Writing – review & editing, Writing – original draft, Software, Methodology. **Taher Chegini:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

BD, TJ, DA, TC were involved in theoretical formulation of the SWMManywhere methodology. BD, DA, TC were involved in software development, with further support provided by Imperial College's Research Software Engineering service. BD, TJ, TC were involved in the experimental formulation and sensitivity analysis. BD, TJ, TC, DA were involved in drafting and editing the manuscript. We are also grateful to Ana Mijic for their insightful comments on the manuscript that have improved the paper.

BD is funded through the Imperial College Research Fellowship scheme, which also funded the software development. TJ time was supported by the UK Natural Environment Research Council-funded CAMELLIA project (grant no. NE/S003495/1). TJ publishes with the permission of the Executive Director of the British Geological Survey. We acknowledge computational resources and support provided by the Imperial College Research Computing Service (http://doi.org/10.14469/hpc/2232).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.envsoft.2025.106358.

Data availability

Code and data used is open-source and indicated in the paper where it can be accessed.

References

Arrighi, C., Campo, L., 2019. Effects of digital terrain model uncertainties on high-resolution urban flood damage assessment. J. Flood Risk Manag. 12. https://doi.org/10.1111/jfr3.12530.

Babovic, F., Mijic, A., 2019. The development of adaptation pathways for the long-term planning of urban drainage systems. J. Flood Risk Manag. 12. https://doi.org/10.1111/jfr3.12538.

Bach, P.M., Rauch, W., Mikkelsen, P.S., McCarthy, D.T., Deletic, A., 2014. A critical review of integrated urban water modelling - urban drainage and beyond. Environ. Model. Software 54, 88–107. https://doi.org/10.1016/j.envsoft.2013.12.018.

Bach, P.M., Kuller, M., McCarthy, D.T., Deletic, A., 2020. A spatial planning-support system for generating decentralised urban stormwater management schemes. Sci. Total Environ. 726, 138282. https://doi.org/10.1016/j.scitotenv.2020.138282.

Bertsch, R., Glenis, V., Kilsby, C., 2017. Urban flood simulation using synthetic storm drain networks. Water (Switzerland) 9. https://doi.org/10.3390/w9120925.

Blumensaat, F., Wolfram, M., Krebs, P., 2012. Sewer model development under minimum data requirements. Environ. Earth Sci. 65, 1427–1437. https://doi.org/10.1007/s12665-011-1146-1.

Boeing, G., 2017. OSMnx: new methods for acquiring, constructing, analyzing, and visualizing complex street networks. Comput. Environ. Urban Syst. 65, 126–139. https://doi.org/10.1016/j.compenvurbsys.2017.05.004.

Butler, D., Davies, J.w., 2004. Urban Drainage, second ed., p. 566

Chahinian, N., Delenne, C., Commandré, B., Derras, M., Deruelle, L., Bailly, J.S., 2019. Automatic mapping of urban wastewater networks based on manhole cover locations. Comput. Environ. Urban Syst. 78, 101370. https://doi.org/10.1016/j.compenvurbsys.2019.101370.

Chegini, T., Li, H.Y., 2022. An algorithm for deriving the topology of belowground urban stormwater networks. Hydrol. Earth Syst. Sci. 26, 4279–4300. https://doi.org/10.5194/hess-26-4279-2022.

Computer generated building footprints for the United States. https://github.com/Microsoft/USBuildingFootprints?tab=readme-ov-file. (Accessed 30 July 2024) last access.

Coxon, G., McMillan, H., Bloomfield, J.P., Bolotin, L., Dean, J.F., Kelleher, C., Slater, L., Zheng, Y., 2024. Wastewater discharges and urban land cover dominate urban hydrology signals across England and Wales. Environ. Res. Lett. 19, 084016. https://doi.org/10.1088/1748-9326/ad5bf2.

Crippen, R., Buckley, S., Agram, P., Belz, E., Gurrola, E., Hensley, S., Kobrick, M., Lavalle, M., Martin, J., Neumann, M., Nguyen, Q., Rosen, P., Shimada, J., Simard, M., Tung, W., 2016. Nasadem global elevation model: methods and progress. Int. Arch. Photogram. Rem. Sens. Spatial Inf. Sci. XLI-B4, 125–128. https://doi.org/10.5194/isprs-archives-XLI-B4-125-2016.

Deletic, A., Dotto, C.B.S., McCarthy, D.T., Kleidorfer, M., Freni, G., Mannina, G., Uhl, M., Henrichs, M., Fletcher, T.D., Rauch, W., Bertrand-Krajewski, J.L., Tait, S., 2012. Assessing uncertainties in urban drainage models. Phys. Chem. Earth 42–44, 3–10. https://doi.org/10.1016/j.pce.2011.04.007.

Dobson, B., Jovanovic, T., Chen, Y., Paschalis, A., Butler, A., Mijic, A., 2021. Integrated modelling to support analysis of COVID-19 impacts on London's water system and in-river water quality. Front. Water 3, 26. https://doi.org/10.3389/frwa.2021.641462.

Dobson, B., Watson-Hill, H., Muhandes, S., Borup, M., Mijic, A., 2022. A reduced complexity model with graph partitioning for rapid hydraulic assessment of sewer networks. Water Resour. Res. 58. https://doi.org/10.1029/2021WR030778.

Dobson, B., Alonso-Álvarez, D., Chegini, T., 2024a. SWMManywhere. https://doi.org/10.5281/zenodo.13837741. September.

Dobson, B., Alonso-Álvarez, D., Chegini, T., 2024b. SWMManywhere sensitivity analysis. https://doi.org/10.5281/zenodo.13918627. September.

Duque, N., Bach, P.M., Scholten, L., Fappiano, F., Maurer, M., 2022. A simplified sanitary sewer system generator for exploratory modelling at city-scale. Water Res. 209, 117903. https://doi.org/10.1016/j.watres.2021.117903.

Eilander, D., 2022. pyFlwDir: Fast Methods to Work with Hydro-And Topography Data in Pure python, vol. 10. Zenodo [code].

Farina, A., Di Nardo, A., Gargano, R., van der Werf, J.A., Greco, R., 2023. A simplified approach for the hydrological simulation of urban drainage systems with SWMM. J. Hydrol. 623, 129757. https://doi.org/10.1016/j.jhydrol.2023.129757.

Ghosh, I., Hellweger, F.L., Fritch, T.G., 2006. Fractal generation of artificial sewer networks for hydrologic simulations. Proc. 2006 ESRI Int. User Conf. 1–12.

Gironás, J., Niemann, J.D., Roesner, L.A., Rodriguez, F., Andrieu, H., 2010. Evaluation of methods for representing urban terrain in storm-water modeling. J. Hydrol. Eng. 15, 1–14. https://doi.org/10.1061/(ASCE)HE.1943-5584.0000142.

Google-MS Open Build. https://beta.source.coop/repositories/vida/google-microsoft-open-buildings, last access: 30 July 2024.

Hagberg, A.A., Schult, D.A., Swart, P.J., 2008. Exploring network structure, dynamics, and function using NetworkX. In: Proceedings of the 7th Python in Science Conference, pp. 11–15.

Herman, J., Usher, W., 2017. SALib: an open-source Python library for Sensitivity Analysis. J. Open Source Softw. 2, 97. https://doi.org/10.21105/joss.00097.

Huang, Y., Zhang, J., Zheng, F., Jia, Y., Kapelan, Z., Savic, D., 2022. Exploring the performance of ensemble smoothers to calibrate urban drainage models. Water Resour. Res. 58, 1–22. https://doi.org/10.1029/2022WR032440.

Hutton, C., Wagener, T., Freer, J., Han, D., Duffy, C., Arheimer, B., 2016. Most computational hydrology is not reproducible, so is it really science? Water Resour. Res. 52, 7548–7555. https://doi.org/10.1002/2016WR019285.

Khurelbaatar, G., Al Marzuqi, B., Van Afferden, M., Müller, R.A., Friesen, J., 2021. Data reduced method for cost comparison of wastewater management scenarios–case study for two settlements in Jordan and Oman. Front. Environ. Sci. 9. https://doi.org/10.3389/fenvs.2021.626634.

Li, X., Willems, P., 2020. A hybrid model for fast and probabilistic urban pluvial flood prediction. Water Resour. Res. 1–26. https://doi.org/10.1029/2019wr025128.

Lindsay, J.B., 2016a. Efficient hybrid breaching-filling sink removal methods for flow path enforcement in digital elevation models. Hydrol. Process. 30, 846–857. https://doi.org/10.1002/hyp.10648.

Lindsay, J.B., 2016b. Whitebox GAT: a case study in geomorphometric analysis. Comput. Geosci. 95, 75–84.

Mair, M., Zischg, J., Rauch, W., Sitzenfrei, R., 2017. Where to find water pipes and sewers?-On the correlation of infrastructure networks in the urban environment. Water (Switzerland) 9. https://doi.org/10.3390/w9020146.

McDonnell, B., Ratliff, K., Tryby, M., Wu, J., Mullapudi, A., 2020. PySWMM: the Python interface to stormwater management model (SWMM). J. Open Source Softw. 5, 2292. https://doi.org/10.21105/joss.02292.

Möderl, M., Butler, D., Rauch, W., 2009. A stochastic approach for automatic generation of urban drainage systems. Water Sci. Technol. 59, 1137–1143. https://doi.org/10.2166/wst.2009.097.

Möderl, M., Sitzenfrei, R., Fetz, T., Fleischhacker, E., Rauch, W., 2011. Systematic generation of virtual networks for water supply. Water Resour. Res. 47, 1–10. https://doi.org/10.1029/2009WR008951.

Montalvo, C., Reyes-Silva, J.D., Sañudo, E., Cea, L., Puertas, J., 2024. Urban pluvial flood modelling in the absence of sewer drainage network data: a physics-based approach. J. Hydrol. 634, 131043. https://doi.org/10.1016/j.jhydrol.2024.131043.

Ochoa-Rodriguez, S., Wang, L.P., Gires, A., Pina, R.D., Reinoso-Rondinel, R., Bruni, G., Ichiba, A., Gaitan, S., Cristiano, E., Van Assel, J., Kroll, S., Murlà-Tuyls, D., Tisserand, B., Schertzer, D., Tchiguirinskaia, I., Onof, C., Willems, P., Ten Veldhuis, M.C., 2015. Impact of spatial and temporal resolution of rainfall inputs on

urban hydrodynamic modelling outputs: a multi-catchment investigation. J. Hydrol. 531, 389–407. https://doi.org/10.1016/j.jhydrol.2015.05.035.
Open water quality archive datasets (WIMS). https://environment.data.gov.uk/water-quality/view/download. (Accessed 21 August 2024) last access.
O'Callaghan, J.F., Mark, D.M., 1984. The extraction of drainage networks from digital elevation data. Comput. Vis. Graph Image Process 28, 323–344. https://doi.org/10.1016/S0734-189X(84)80011-0.
Palmitessa, R., Grum, M., Engsig-Karup, A.P., Löwe, R., 2022. Accelerating hydrodynamic simulations of urban drainage systems with physics-guided machine learning. Water Res. 223, 118972.
Pedersen, A., Pedersen, J., Vigueras-Rodriguez, A., Brink-Kjær, A., Borup, M., Mikkelsen, P., 2021. The Bellinge data set: open data and models for community-wide urban drainage systems research. Earth Syst. Sci. Data Discuss. 13, 4779–4798. https://doi.org/10.5194/essd-2021-8.
Pianosi, F., Beven, K., Freer, J., Hall, J.W., Rougier, J., Stephenson, D.B., Wagener, T., 2016. Sensitivity analysis of environmental models: a systematic review with practical workflow. Environ. Model. Software 79, 214–232. https://doi.org/10.1016/j.envsoft.2016.02.008.
Rauch, W., Urich, C., Bach, P.M., Rogers, B.C., de Haan, F.J., Brown, R.R., Mair, M., McCarthy, D.T., Kleidorfer, M., Sitzenfrei, R., Deletic, A., 2017. Modelling transitions in urban water systems. Water Res. 126, 501–514. https://doi.org/10.1016/j.watres.2017.09.039.
Ray, G., Sen, A., 2024. Minimal spanning arborescence. arXiv Prepr. arXiv2401.13238.
Revitt, D.M., Ellis, J.B., Gilbert, N., Bryden, J., Lundy, L., 2022. Development and application of an innovative approach to predicting pollutant concentrations in highway runoff. Sci. Total Environ. 825, 153815.
Reyes-Silva, J.D., Novoa, D., Helm, B., Krebs, P., 2023. An evaluation framework for urban pluvial flooding based on open-access data. Water (Switzerland) 15. https://doi.org/10.3390/w15010046.
Rossman, L.A., 2010. Storm water management model user's manual, version 5.0. Cincinnati 72–73.
Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S., 2008. Global sensitivity analysis. The Primer. Wiley, John & Sons, pp. 1–292. https://doi.org/10.1002/9780470725184.
Saltelli, A., Aleksankina, K., Becker, W., Fennell, P., Ferretti, F., Holst, N., Li, S., Wu, Q., 2019. Why so many published sensitivity analyses are false: a systematic review of sensitivity analysis practices. Environ. Model. Software 114, 29–39. https://doi.org/10.1016/j.envsoft.2019.01.012.
Seo, Y., Schmidt, A.R., 2013. Network configuration and hydrograph sensitivity to storm kinematics. Water Resour. Res. 49, 1812–1827. https://doi.org/10.1002/wrcr.20115.
Sirko, W., Kashubin, S., Ritter, M., Annkah, A., Bouchareb, Y.S.E., Dauphin, Y.N., Keysers, D., Neumann, M., Cissé, M., Quinn, J., 2021. Continental-scale building detection from high resolution satellite imagery. CoRR 1 abs/2107.
Sobol, I.M., 1993. Sensitivity estimates for nonlinear mathematical models. Math. Model. Comput. Exp. 1, 407.
Source, M.O., McFarland, M., Emanuele, R., Morris, D., Augspurger, T., 2022. microsoft/PlanetaryComputer. October. https://doi.org/10.5281/zenodo.7261897. October 2022.
Stagge, J.H., Rosenberg, D.E., Abdallah, A.M., Akbar, H., Attallah, N.A., James, R., 2019. Assessing data availability and research reproducibility in hydrology and water resources. Sci. Data 6, 190030. https://doi.org/10.1038/sdata.2019.30.
Subcatchment Data Fields (InfoWorks). https://help2.innovyze.com/infoworksicm/Content/HTML/ICM_IL/Subcatchment_Data_Fields.htm. (Accessed 3 June 2024) last access.
Sun, S., Djordjevic, S., Khu, S.T., 2011. A general framework for flood risk-based storm sewer network design. Urban Water J. 8, 13–27. https://doi.org/10.1080/1573062X.2010.542819.
Sweetapple, C., Fu, G., Farmani, R., Meng, F., Ward, S., Butler, D., 2018. Attribute-based intervention development for increasing resilience of urban drainage systems. Water Sci. Technol. 77, 1757–1764. https://doi.org/10.2166/wst.2018.070.
SWMManywhere documentation. https://imperialcollegelondon.github.io/SWMManywhere/. (Accessed 4 October 2024).
Sytsma, A., Crompton, O., Panos, C., Thompson, S., Mathias Kondolf, G., 2022. Quantifying the uncertainty created by non-transferable model calibrations across climate and land cover scenarios: a case study with SWMM. Water Resour. Res. 58, 1–21. https://doi.org/10.1029/2021WR031603.
Tan, M., Le, Q.V., 2019. EfficientNet: rethinking model scaling for convolutional neural networks. CoRR 1 abs/1905.
Tarjan, R.E., 1977. Finding optimum branchings. Networks 7, 25–35. https://doi.org/10.1002/net.3230070103.
Thrysøe, C., Arnbjerg-Nielsen, K., Borup, M., 2019. Identifying fit-for-purpose lumped surrogate models for large urban drainage systems using GLUE. J. Hydrol. 568, 517–533. https://doi.org/10.1016/j.jhydrol.2018.11.005.
Warsta, L., Niemi, T.J., Taka, M., Krebs, G., Haahti, K., Koivusalo, H., Kokkonen, T., 2017. Development and application of an automated subcatchment generator for SWMM using open data. Urban Water J. 14, 954–963. https://doi.org/10.1080/1573062X.2017.1325496.
Wills, P., Meyer, F.G., 2020. Metrics for graph comparison: a practitioner's guide. PLoS One 15, e0228728. https://doi.org/10.1371/journal.pone.0228728.
Xu, Z., Dong, X., Zhao, Y., Du, P., 2021. Enhancing resilience of urban stormwater systems: cost-effectiveness analysis of structural characteristics. Urban Water J. 18, 850–859. https://doi.org/10.1080/1573062X.2021.1941139.