



Crowd digital twins: Motivation, architecture, and roadmap

Roberto Casadei *, Giovanni Delnevo , Roberto Girau , Silvia Mirri

Alma Mater Studiorum–Università di Bologna Cesena, Italy

ARTICLE INFO

Keywords:

Human crowds
Crowd digital twin
Human-machine networks
Computational collective intelligence
Crowd management
Crowd safety
Distributed systems

ABSTRACT

Crowd management is the problem of dealing with physical crowds of humans, e.g., for safety reasons or to improve services through crowd-awareness. Information and communication technology can be of great help in supporting the monitoring and management of crowds. Following the “digital twin concept”, a group or crowd of people – as a dynamic physical entity – is amenable to be represented and managed through digital models. Accordingly, in this work we present and develop a general notion of a *Crowd Digital Twin (CDT)*, intended as the digital twin associated to physical gatherings of humans. We motivate it as a fundamental component of crowd management systems, and investigate the functional and non-functional requirements that such a technical solution should satisfy. Then, we consider what the fundamental components of CDT systems are, and accordingly delineate a reference architecture aimed at flexibly supporting the corresponding physical-digital thread. We also discuss methods and tools that might be useful for supporting implementations of such a CDT concept, and delineate a set of challenges that may serve as a roadmap for future research on the topic.

1. Introduction

Background and motivation. Humans often gather in crowds whose size and dynamics pose logistic and safety challenges [1,2]. Crowd management applies information and communication technology (ICT) and engineering methods to support safety and operational efficiency through monitoring, decision support, and control [3,4]. Moreover, emerging visions suggest the opportunity of addressing *cyber-physical collectives and ecosystems* through a global perspective [5]. On a related side, research on *digital twins (DTs)* is suggesting ways to organise *digital models* connected with *physical* entities and processes. For instance, *human DTs* [6–8] have been proposed to model human individuals to support personalised adaptive applications in a wide range of scenarios ranging from e-health to smart mobility. However, to the best of our knowledge, no proposals exist about DTs for *whole groups or crowds of humans*, with the exception of solutions targetting specific aspects of them (e.g., motion prediction [9]).

Contribution. In this paper, we motivate and propose the notion of *crowd digital twin (CDT)*, intended as the digital twin associated to a physical crowd (group) of human individuals. Such a CDT is meant to provide actionable digital models and especially handle the connection between

such models and situated physical crowds in an effective way. Specifically, in this work we elaborate on such a CDT notion, its motivation and goals, and accordingly identify *functional and non-functional requirements* of CDT solutions that are relevant for a variety of scenarios. Then, we also provide *reference architecture views* [10,11] describing a family of concrete CDT architectures where crucial aspects like real-time support and run-time adaptation are promoted. Last but not least, we discuss the suitability of the proposal, and suggest implementation ideas and further research directions. We provide the following main contributions:

- we provide a taxonomy of *crowd-aware services*, namely services that should take into account human gatherings for their functionality or operation;
- based on a *scoping review* [12] of the literature, we motivate, propose, and characterise the notion of *crowd digital twin (CDT)*, with its requirements, architecture, enabling methods and technologies;
- based on the analysis of the state of the art, we identify a *roadmap*, i.e., a set of research challenges, associated to the CDT realisation, suggesting important directions for future research on the topic across fields.

This article builds on two preliminary conference papers: [13] explored the plausibility of the CDT concept and [14] provided a discussion of

Abbreviations: AI, artificial intelligence; DT, digital twin; RL, reinforcement learning; ML, machine learning; BIM, building information model; CIM, city information model; EDT, environment digital twin; HDT, human digital twin; HMI, human machine interface; CCTV, closed-circuit television; CDT, crowd digital twin; CPS, cyber-physical system; GIS, geographic information system; KI, key indicator; ICT, information and communication technology; IoT, internet of things; LIDAR, light detection and ranging; QoS, quality of service; UAV, unmanned aerial vehicle; SLAM, simultaneous localisation and mapping; SIoT, social internet of things.

* Corresponding author.

E-mail addresses: robby.casadei@unibo.it (R. Casadei), giovanni.delnevo@unibo.it (G. Delnevo), roberto.girau@unibo.it (R. Girau), silvia.mirri@unibo.it (S. Mirri).

<https://doi.org/10.1016/j.future.2026.108475>

Received 28 May 2025; Received in revised form 16 February 2026; Accepted 6 March 2026

Available online 11 March 2026

0167-739X/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

requirements and an architectural proposal. This article significantly extends and improves over them, specifically by elaborating on the methodology (by clarifying a set of research questions and how to address them), motivation (by connecting requirements to DT properties), requirements (including meta-requirements), design and implementation, research challenges, and related work on crowd management and DTs.

Paper structure. Section 2 covers the research questions and methodology. Section 3 details the motivation for the CDT. Section 4 analyses the CDT concept and its requirements. Section 5 provides architectural considerations for CDT solutions. Section 6 discusses the proposal. Section 7 present the research challenges. Section 8 covers related work on crowd simulation and crowd management. Section 9 concludes the paper with final considerations.

2. Methodology

By a methodological perspective, this manuscript is based on a *scoping review* [12]. The goal is providing an overview of a large body of literature on topics related to crowd management and digital twins. Such a scoping review is leveraged to motivate and *inform* our technical proposal for the CDT idea. In order to provide a structure for our investigation, we consider a number of research questions and ways to address them.

In this section, we define our research questions (Section 2.1) and then delineate our approach, in terms of methods and materials, to answering them (Section 2.2).

2.1. Research questions

We define the following research questions to guide our investigation:

- (RQ1) *CDT concept and motivation*. What motivates the use of digital twins for crowds? Is the digital twin framework appropriate for crowd management?
 - (RQ1.1) *Crowd applications and gaps*. What are the classes of computer science applications that involve crowds (i.e., possibly large groups of humans)? What are the kinds of state-of-the-art solutions and the main gaps?
 - (RQ1.2) *CDT concept*. What defines the CDT concept?
 - (RQ1.3) *CDT requirements*. What are the main functional and non-functional requirements for CDT implementations?
- (RQ2) *CDT design and implementation*. How could we address the design and implementation of CDT solutions?
 - (RQ2.1) *CDT reference architecture*. What are the main components and relationships among components that provide the intended functionality and non-functional properties of CDTs?
 - (RQ2.2) *CDT enablers*. What are the enabling technologies for CDT implementations? What tools can help to address the CDT requirements and achieve the desired architectural properties?
- (RQ3) *CDT research challenges*. What are the research challenges related to CDTs, namely to the development of designs and implementations able to fulfil the corresponding requirements?

2.2. Methods and materials

We plan to address (RQ1) *CDT concept and motivation* by understanding the landscape of crowd-based applications and their challenges (RQ1.1), then providing the CDT concept (RQ1.2), and identifying significant functional and non-functional requirements (RQ1.3). We aim to substantiate the need for a DT-based solution by mapping the elicited requirements to the peculiar DT features as pointed out by notable DT frameworks.

(RQ1.1) *Crowd applications and gaps* is addressed through a literature survey partially carried out in previous work [13,14], from which we




Fig. 1. Crowd application domains and their relationships.

propose a taxonomy of applications related to crowds. (RQ1.2) *CDT concept* is addressed by considering how the DT concept applies to human crowds, and fixing a reasonable scope based on the identified application types where digital-physical correspondences are mostly useful. (RQ1.3) *CDT requirements* is addressed by extracting requirements from the surveyed papers and applications, elaborating on this basis according to what the DT technology enables, and filling the gaps using systems engineering knowledge [15].

We plan to address (RQ2) *CDT design and implementation* by first devising a high-level, general reference architecture (RQ2.1) and then considering methods and technologies that may help in implementing it while promoting the satisfaction of the aforementioned requirements (RQ2.2).

(RQ2.1) *CDT reference architecture* is addressed by considering best-practices on architectural designs [11], and especially by identifying *architectural views* that are most relevant to convey the CDT characteristics delineated while answering the previous research questions. (RQ2.2) *CDT enablers* is addressed in terms of a literature review tailored to the specific requirements (i.e., by asking, for instance: what are the research threads that address *adaptation* in software?).

The investigation of enabling methods and technologies is also expected to point out *gaps* in the state of the art, suggesting directions for further research that will provide an answer to (RQ3) *CDT research challenges*.

3. Motivation

In this section, we present the “business” (Section 3.1) and technical (Section 3.2) motivation for this research work.

3.1. Social motivation: crowd management

(Physical) *Crowds* are large groups of spatiotemporally co-located people [16]. Crowds can be defined and/or classified based on their *scale* (number of people) and *density* (number of people per unit area) [2]. In this work, we do not set specific thresholds for defining a crowd, which we generally mean as a plurality of people arbitrarily dispersed in some arbitrarily small or large area. The study of crowds is a multi-disciplinary endeavour known as *crowd science* [17].

Due to their scale and implications on services and movement of people, both from safety and performance perspectives, crowds have to be properly taken into consideration and possibly monitored and managed. We define *crowd-aware services* all the services and applications that consider the structure or the dynamics of crowds in their operation. Notable examples include *smart mobility* (cf. distribution of vehicles and passengers) and *pandemic management* [18] (cf. promoting social distancing). Important sub-classes of crowd-aware applications are those where crowds are consumers (subjects) or providers of services.

Crowd management [3,4], a key part of *event management* [17], is the study and application of techniques to monitor and control crowds, leveraging sensing [19], actuation [20], and computational [3,21] solutions. The crowd itself plays a *passive* role in the application. A peculiar aspect of crowd management is *crowd safety* [22], i.e., the application




Fig. 2. A high-level view of the Crowd Digital Twin concept.

of measures to foster the safety and well-being of people in a crowd, including prevention and *emergency/crisis management* [21].

The case where crowds contribute to the provisioning of services is given by *physical crowdsourcing* [23], which exploits the physical distribution of crowds for performing physical tasks. In this case, the crowd members play an *active* role in the application. Physical crowdsourcing examples include *crowd logistics* [24] and *mobile crowdsensing* [18,25,26].

Of course, crowd management and crowdsourcing may also overlap, where the crowd members may take both a passive and active role, or switch roles dynamically. The relationship between these application domains is detailed in Fig. 1. In this work, the focus is on crowd management, i.e., not on crowdsourcing (though parts of the proposal may also generally apply to all crowd-aware services).

Short Answer to (RQ1.1)

Crowd-aware applications are computer systems that explicitly account for the presence and dynamics of human gatherings. Notable sub-classes include *crowdsourcing* (active crowds, e.g., for sensing/logistics) and *crowd management* (passive crowds). Current approaches tend to offer application-specific pipelines with limitations regarding the dynamic integration between physical and digital layers, motivating the adoption of the digital twin paradigm. (cf. Section 3 and 8)

3.2. Technical motivation

The challenges related to crowd management can be addressed by ICT-based solutions (cf. Section 8 for a review) that support the collection, processing, and communication of data. Since crowd management deals with physical crowds situated in physical environments, it is paramount to systematically bridge the physical and the digital worlds: a real-time accurate digital model of a physical crowd (possibly across multiple scales and resolutions) is essential to enable both human and computational elements to perform analysis, decision-making, and operative tasks. This is exactly the problem addressed by the *DT* concept [27]. Also, since the *DT* concept is often aimed at supporting product development across its entire lifecycle, a CDT could also be instrumental for *event planning*—a crucial part of crowd management [17]. In the following, we further develop this notion by eliciting requirements to be addressed and proposing an abstract architecture for guiding implementations.

4. Crowd digital twins: concept and requirements

In this section, we describe the CDT concept and delineate (non-)functional requirements for its implementation.

4.1. Concept

We define a *crowd digital twin (CDT)* as a digital representation of a physical crowd and its surrounding context and environment, enabled by

local devices, environmental devices, and heterogeneous ICT infrastructure, and supporting the analysis, prediction, and possibly the management of their associated spatiotemporal and collective phenomena. Fig. 2 provides a high-level graphical representation of the concept. The definition is meant to comprehensively indicate the main elements of the concept and the general goals. A few elements of the given definition deserve attention:

- For a *physical crowd*, it is meant a group of co-located individuals (generally, humans).
- Since the *spatio-temporal* evolution of the system, rooted in the relationship between the physical crowd and the physical environment, is crucial, we emphasise the role of the latter and accordingly require an *environment digital twin (EDT)* [28] to be conceptually part of the CDT.
- Regarding *collective phenomena*, we consider the ontological characterisation by Wood et al. [29], i.e., (i) *collectives* as pluralities of individuals that are *concrete and continuant particulars* (i.e., not abstract or spatiotemporally discontinuous groups), and (ii) *collective phenomena* as any coherent activity, sharing common purpose or cause, performed by a collective.

We argue that the proposed definition provides a reasonable trade-off between specificity and broadness, supporting a relevant set of application domains (cf. Section 3.1), while excluding other applications involving humans (e.g., those where humans are not co-located, or where collective phenomena are not considered—e.g., general crowdsourcing).

We use term *situated crowd* to mean the combination of a crowd and its embedding physical environment. Then, a CDT is the system that supports the digital modelling of a situated crowd, by collecting relevant data about human groups and environments, and applying mathematical and computational models to estimate its state with an application-specific degree of fidelity. Applications might leverage various properties of a situated crowd: e.g., its spatial and density distribution, its dynamics at various spatiotemporal resolutions, the probabilities of future states, the states of the tasks assigned to its members, their capabilities, etc.

Short Answer to (RQ1.2) CDT concept

A CDT is a digital representation of a situated physical crowd and its environment, enabled by ICT infrastructure and used to analyse, predict, and manage spatiotemporal and collective crowd phenomena, for safety, organisational, and quality-of-service purposes in crowd-aware applications. (cf. Section 4.1)

4.2. Detailed characterisation of the CDT concept: methodology

To characterise the CDT concept, it is key to clarify its scope and features. Such elements can be formalised as requirements, both functional and non-functional, that the CDT should satisfy. For *requirements elicitation*, we consider related work (cf. Section 8) and the application types introduced in Section 3.1.

To the best of our knowledge, there exist no comprehensive and detailed work on requirements for ICT-based crowd management systems. According to [30], crowd management tasks include (i) crowd counting, (ii) density estimation, (iii) localisation and tracking of individuals/groups within a crowd, and (iv) behaviour monitoring. In [31], (i) situational awareness and (ii) crowd control are mentioned as two key requirements. Martella et al. [32] also mention the need for (i) increased situation awareness and decision-making support, (ii) reliable real-time monitoring and communication, (iii) seamless intervention. For event management, they emphasise the need of (i) considering the different features and phases of the event, (ii) controlling pedestrian flows, bottlenecks, and queues, (iii) predicting and preventing accidents.

4.3. Reference scenarios

4.3.1. Small-scale crowd management in public transport and mobility

A CDT could be implemented to facilitate crowd management within the public transportation network, for city residents and visitors. The CDT can be fed with data from station surveillance cameras, and sensors embedded in vehicles to provide a real-time representation of crowd dynamics and transportation infrastructure. Such a CDT could be exploited by several actors. Commuters may benefit from services to get clues on how to navigate the transit network safely and efficiently (e.g., avoiding overcrowded buses and trains, and considering alternative routes or transportation modes to avoid congestion). Authorities can use the CDT to monitor crowd flow, identify bottlenecks or disruptions, and optimise transit operations. In case of emergency, the CDT can facilitate coordinated responses and communication.

4.3.2. Large-scale crowd management in mass events

Suppose you have a CDT of the crowd attending a mass event, such as a concert, an exhibition, a sport event (e.g., a marathon in a city, or a football match in a stadium), or a political rally. Pre-configured cameras in the environment, as well as wearables and smartphones of the participants, hosting the event application (e.g., installed for accessing the facility), may continuously provide data to populate the CDT representation, constructed by services implemented and deployed across the edge-cloud infrastructure. The organisers of the event as well as public authorities may use it to monitor the situation and possibly respond by quickly alerting emergency and policy services. The participants themselves may use crowd-aware services for navigating across the facility in a safe and efficient manner [33]. In the unfortunate case a stampede occurs, CDT-originating actuations on the participants (e.g., sound alerts, or pre-configured voice instructions) and the environment (e.g., light signals) could support the dispersal of the crowd, e.g. by signalling people far from the origin of the stampede and close to the exits to quickly move out to create spatial basins.

4.4. Stakeholders

We identify the following classes of stakeholders associated to CDTs and their applications:

- *Crowd members (participants)*: the individuals that compose the crowd(s).
- *Affected entities*: the services, people, or vehicles affected by the presence of physical crowds (but are external to them).
- *Workers*: the situated entities involved in the surveillance and management of crowds.
- *Organisers and managers*: the entities responsible for planning, monitoring, controlling, and taking decisions regarding the management of the crowd.
- *Authorities*: the entities accountable for the crowd and environment safety.

4.5. Requirements

In general, the requirements of a digital twin are specific to the application and use case at hand [27]. Aspects like the definition of a bidirectional link, support for real-time, or extreme fidelity may not be always a required (or feasible) goal [27]. By considering potential target use cases (cf. Section 3), in the following, we propose some typical or desirable requirements for CDT implementations. We also consider crowd safety models as proposed in safety science venues, like the *Swiss Cheese Model of Crowd Safety* [22], a multi-layered paradigm based on different levels of protection. We use the MoSCoW (Must have, Should have, Could have, Won't have) method for requirements prioritisation [34] to indicate how much our proposed requirements are essential for the CDT notion.

4.5.1. Functional requirements

We identify the following requirements covering the typical functions of CDTs.

- *(F1) Crowd model synchronisation*. The CDT system must build and maintain digital models of the crowd based on meta-models and sensed data.
- *(F2) Environment and context model synchronisation*. The CDT system must build and maintain digital models of the physical environments (both embedding and surrounding the crowd) and context (e.g., current event phase), based on existing meta-models and sensed data.
- *(F3) Situated crowd model synchronisation*. The crowd and environment models must be combined to obtain a model for the situated crowd, e.g., for monitoring purposes.
- *(F4) Analytic and predictive capabilities*. CDT should be able to analyse the current models, extract relevant metrics, and predict the evolution of a situated crowd, to assess risks and anticipate the occurrence of accidents.
- *(F5) Support for situated crowd control* (digital-to-physical synchronisation). The CDT could provide means for influencing or controlling the crowd, by analysing, planning, and performing actuations on the crowd members or the environment, and fostering participation.
- *(F6) Interactive multi-scale analysis, visualisation, and decision-making support*. The CDT must provide relevant information and effective representations to managers, possibly with means to navigate data and representations to support analysis and decision-making.
- *(F7) Storage of historical data and analytics*. The CDT could store data for a-posteriori, delay-tolerant analysis, e.g., to support retrospective evaluation and model improvement.
- *(F8) Lifecycle management*. The CDT should support all the phases of crowd management, i.e., besides operations, also the crucial phase of planning [17].

Notice how these requirements cover the key digital twin features throughout the entire lifecycle, as well as features for different stakeholders.

4.5.2. Non-functional requirements

We identify the following main non-functional requirements for CDTs.

- *(N1) Dependability*. Since crowd management applications tend to be safety-critical, and that failure to take crowds into consideration may severely impact their quality of service (QoS), it is important that CDT solutions are dependable [35], i.e., available, reliable, safe, and secure.
 - *(N1.1) Availability*. The CDT functionality should remain readily accessible for crowd members and managers.
 - *(N1.2) Reliability*. The CDT functionality should ensure continuity of services (possibly with degraded performance).
 - *(N1.3) Safety*. The CDT must operate with a principle of safety, i.e., with the goal of preventing and mitigating catastrophic consequences.
 - *(N1.4) Security*. Given the potentially serious consequences of attacks, security requirements of confidentiality, integrity, and availability must be supported [36].
- *(N2) Real-time support*. Since accidents may happen abruptly in crowds, soft real-time model synchronisation, simulation, and adaptation are paramount [37] for responsiveness of both prevention and mitigation actions and must be supported.
- *(N3) Scalability*. The CDT should be able to scale with the size of the crowds, environments, and devices involved.
- *(N4) Privacy compliance*. Since crowds are made of humans, the CDT must preserve the privacy of crowd members (e.g., using anonymisation and aggregation of data).
- *(N5) Integrability*. Multiple CDTs could be able to be integrated together and with other systems to consider multiple concurrent crowds for large-scale (e.g., city-scale) applications.

Table 1
Mapping of CDT requirements to core DT characteristics and capabilities.

Requirement	Digital Twin Characteristic/Property
(F1) Crowd model synchronisation	Real-time bi-directional synchronisation [38]
(F2) Environment and context model synchronisation	Context-awareness [39]; Real-time data integration [40]
(F3) Situated crowd model synchronisation	Hybrid model synch.; Spatiotemporal accuracy [41]; Ontologies [42]
(F4) Analytic and predictive capabilities	Analytics & simulation; Machine learning integration [43]
(F5) Support for situated crowd control	Actuation/control loop; Cyber-physical interface [42,44]
(F6) Interactive multi-scale analysis, visualisation, and decision-making support	Multi-resolution modelling; Human-in-the-loop interface [45]
(F7) Storage of historical data and analytics	Historical data archiving [42]; Data lake support [46]
(F8) Lifecycle management	Twin lifecycle support [47]; Versioning & update mechanisms [48]
(N1) Dependability	Robust architecture & design principles [49]
(N1.1) Availability	High availability; Redundancy mechanisms [50]
(N1.2) Reliability	Fault tolerance; Consistent data processing [51]
(N1.3) Safety	Safety-aware decision frameworks [52]
(N1.4) Security	Cybersecurity; Access control; Data encryption [53]
(N2) Real-time support	Low-latency processing & messaging [54]
(N3) Scalability	Scalable architecture (e.g., cloud-native, microservices) [46]
(N4) Privacy compliance	Data anonymisation; GDPR alignment [55]
(N5) Integrability	API support; Modular design (digital twin assets) [56]; Interoperability
(N6) Flexibility, configurability, (self-)adaptability, and maintainability	Modular design (digital twin assets) [56]; DevOps pipelines [56]
(N7) Efficiency	Edge-cloud continuum exploitation [46]; Multi-level simulation [57]

- (N6) Flexibility, configurability, (self-)adaptability, and maintainability. CDTs should be applicable to different contexts and foster and accommodate changes to the environment, infrastructure, software, law, policies, etc.
- (N7) Efficiency. CDTs should adapt and minimise resource usage based on needs and priorities (e.g., taking into account the dependability requirements).

The emphasis is on safety, legal issues, and elements that contribute to them—cf. effectiveness of management and resource usage. Achieving these desired characteristics can also represent relevant challenges (cf. Section 7) and may require investigating specialised techniques (cf. Section 6.2). Functional and non-functional requirements will be discussed in Section 6.1.

Short Answer to (RQ1.3) CDT requirements

A CDT is required to provide a synchronised model of the situated crowd (F1, F2, F3) and to support interactive analysis and decision-making (F4, F6) across the application/event/crowd lifecycle (F8). It could also support analytics of historical data (F7), and, most notably, situated crowd control (F5). Such functionalities must be provided in a dependable (N1), scalable (N3), privacy-compliant (N4), flexible/adaptable way (N6), in real-time (N2), ideally integrating with other systems (N5). (cf. Section 4.3, 4.5 and 5.1)

4.5.3. Mapping CDT requirements to digital twin characteristics and capabilities

In Table 1, we provide a mapping from the identified requirements to the characteristics and capabilities of digital twins in general. This mapping serves as a qualitative evaluation about the suitability of the digital twin approach to the CDT design and implementation. This provides additional useful elements for attempting an answer to (RQ1) CDT concept and motivation.

Short Answer to (RQ1) CDT concept and motivation

CDTs are motivated by the need to monitor, analyse, and manage physical crowds through real-time digital counterparts that enhance safety, efficiency, and service quality in human-centric systems. (cf. Section 3 and 4)

5. Crowd digital twins: reference architecture

In this section, we provide an abstract architectural description for the CDT concept, aimed at satisfying the requirements provided in Section 4.

5.1. Meta-level requirements

We start by providing meta-requirements, dictating general properties about the architecture itself (internal quality requirements) that we would like to promote in a CDT design. We selected the most relevant quality characteristics/attributes from [58] as targets.

- (M1) *Abstraction*. The architecture should not fix too many details to avoid excessive constraints on concrete system designs.
- (M2) *Generality*. The architecture should consider that a CDT may support different kinds of applications involving crowds, with peculiar sets of requirements. For instance, not all of them may require hard real-time support.
- (M3) *Suitability*. This aspect refers to the degree to which an architecture meets the required functional and non-functional features.
- (M4) *Portability*. The architecture should account for highly heterogeneous deployments, and hence be transferred on other hardware, software, and operational environments.

These meta-requirements influence the proposal in the next subsection and will be subject to qualitative evaluation in Section 6.1.

5.2. Reference architecture

Considering the maturity of the CDT concept, we consider two meta-requirements – abstraction and generality of the architecture – and thus propose a *reference architecture*, i.e., a generic architecture that represents a family of concrete architectures for a specific context or application domain [10,59].

Following best-practices on documenting architectural designs [11], it is key to select the *relevant views*. For this contribution, we focus on the *logical* and *deployment* views.

5.2.1. Logical view

Fig. 3 shows the reference *logical* architecture of CDT systems—the relevant layers/components.

- (L0) *Physical layer*. The physical layer consists of the physical elements of interest, namely the *Humans* (both crowd members and external people) and the *Physical Environment Elements* (e.g., walls, doors, vehicles).




Fig. 3. Logical reference architecture for CDT solutions. The colours are only meant to emphasise related entities.

- (L1) *Cyber-physical layer*. The cyber-physical layer consists of one or more networks of sensors, actuators, and mobile devices (e.g., robots). These represent the means for perceiving and acting on the physical layer. Sensors and actuators may be environmental (e.g., cameras, lights, speakers) or human-attached (cf. wearables).
- (L2) *Autonomic layer*. The bidirectionality requirement of the CDT fosters the adoption of self-* solutions [60], enabling the system to autonomously react to changes to a certain extent. Accordingly, this layer is structured following the classical *MAPE-K* (Monitor-Analyze-Plan-Execute with Knowledge) architecture [60] of the managing system. Notice that the *managed system* consists of the *two underlying layers* (i.e., both the cyber-physical layer and the purely physical layer). The *Knowledge* component includes the relevant representational and operational models for the CDT. The *CDT Core Models* are those that provide the digital shadows of the crowds and the environment; these are built by considering the data provided by the *Monitor* component and by using different kinds of models (e.g., *Surrogate Models* for crowd estimation).
- (L3) *Service layer*. The service layer provides functionalities connecting applications with the managing system. *Visualisation* provides support for inspecting crowd and environment state. *Analytics* deals with storage and offline data analysis, both for external entities and for improving the CDT itself. *Administration* provides an interface for configuring the CDT and setting the *Adaptation Goals* for crowd/environment-directed management actions.
- (L4) *Application layer*. Applications can connect to the CDT through the services provided in the Service layer.

Overall, the layers fill the gap between crowd management and crowd-aware applications and services and the physical world. The centrality of models and adaptation features (which are essential to deal with the dynamic nature of human crowds) is well-represented by the autonomic layer, which mediates between basic cyber-physical support and system-level services.

5.2.2. Deployment view

By a deployment perspective, the fundamental aspect shown in the Fig. 4 is that the CDT spans the entire computing continuum, from the




Fig. 4. Deployment view.

internet of things (IoT) (mist) to the cloud, passing from edge and fog layers. This is important both for real-time and reliability reasons, as further detailed in Section 6.

Short Answer to (RQ2.1) CDT reference architecture

To promote the realisation of the CDT requirements (cf. 1.3), it is arguably important to design CDTs with autonomic capabilities, feedback loops, a rich set of complementary and alternative models, and interfaces organised in layers to harness the digital-physical thread in heterogeneous multi-level architectures (computing continua). (cf. Section 5)

6. Discussion

In this section, we provide a discussion about the proposed requirements and architecture (Section 6.1), and then cover relevant ideas from research for detailed design and implementation (Section 6.2).

6.1. Qualitative validation

6.1.1. Validation of requirements

We consider the question: *are the requirements provided in Section 4.5 relevant and comprehensive w.r.t. the CDT concept and the target applications?*

The model synchronisation requirements – (F1) *Crowd model synchronisation*, (F2) *Environment and context model synchronisation*, (F3) *Situated crowd model synchronisation* – are associated to the DT nature of the proposal. Prediction ((F4) *Analytic and predictive capabilities*) is fundamental for enabling risk assessment, identification of improvement opportunities, and decision-making. The two-way communication requirement of (F5) *Support for situated crowd control*, as per mature DT solutions, is key to enable the actual management of the situated crowd (besides its mere modelling). (F6) *Interactive multi-scale analysis, visualisation, and decision-making support* covers the need of allowing managers to interact with the CDT, hence accounting for human-in-the-loop scenarios (which are essential given the safety-critical nature of crowd-based applications). (F7) *Storage of historical data and analytics* is meant to support historical data management, retrospective analysis, and knowledge extraction. (F8) *Lifecycle management* is meant to make the CDT usable also to support event planning and the continuous improvement of the environment and the management policies.

Non-functional requirements can also be traced to crowd-aware applications and DT maturity models [61]. Since crowd management applications tend to be safety-critical, the CDT should be dependable – (N1) *Dependability* – and avoid people getting harmed (at any level). This has important consequences for the actuation part, its timely assessment and validation. Non-functional requirement (N2) *Real-time support*, at least at the *soft level*¹ [62], may be needed, since quick prediction, model synchronisation, and reaction are paramount when it comes to anticipating or dealing with emergencies (e.g., evacuation or stampedes), as seconds can make a difference. (N3) *Scalability* is also key, since the CDT would be most useful for mass gatherings, which means: large numbers of people, of devices, of physical elements—i.e., stress on the infrastructure. (N4) *Privacy compliance* must also be taken into account, since humans are involved. This may affect the implementation (and even the possibility) of micro-level data storage and analysis. (N5) *Integrability* matches with the highest level of maturity of DT-based solutions, and is deemed essential since (i) within a crowd, multiple human groups exist and may need to be considered; (ii) at large building- or city-scale, multiple separate crowds may concurrently exist and interact; and (iii) it is the interaction with other systems and applications that would create the additional value out of the CDT. The non-functional requirement (N6) *Flexibility, configurability, (self-)adaptability, and maintainability* is key to account for the evolution of models, environments, regulations, and phases of crowd management (cf. (F8) *Lifecycle management*). Finally, (N7) *Efficiency* relates the general need of making effective use of resources and promoting sustainability.

6.1.2. Validation of architecture

We consider the question: *is the architecture provided in Section 5 consistent with the (meta-)requirements identified in Section 4.5?*

The meta-requirement (M1) *Abstraction* is addressed through a high-level *reference architecture* [10,59] not constraining implementations but rather specifying layers of functional components. (M2) *Generality* is supported by considering a comprehensive set of functionalities and services that may serve different applications (e.g., actuation may not be always needed). The meta-requirement (M3) *Suitability* is discussed in the next paragraph. The meta-requirement (M4) *Portability* is addressed by not requiring specific devices and deployments (cf. Fig. 4).

We argue about the suitability of the architecture with respect to the motivations and goals of CDTs (cf. Section 4) as follows. The different layers take into account the main stakeholders considered in Section 4.4: crowd members, affected entities, and workers are part of the environment in (L0) *Physical layer*, perceived and influenced through sensors and actuators in (L1) *Cyber-physical layer*, and captured in models by the *Knowledge* component in (L2) *Autonomic layer*. The modelling requirements (F1) *Crowd model synchronisation*, (F2) *Environment and context model synchronisation*, (F3) *Situated crowd model synchronisation* are addressed through the data provided by (L1) *Cyber-physical layer* and the processing and analysis carried out by the *Monitor* component in (L2) *Autonomic layer*. The bidirectional flow of (F5) *Support for situated crowd control* is supported by the overall activity of the MAPE-K components in (L2) *Autonomic layer*, and the actuation support provided in (L1) *Cyber-physical layer*. (F6) *Interactive multi-scale analysis, visualisation, and decision-making support* and (F7) *Storage of historical data and analytics* are provided in (L3) *Service layer* by dedicated components that interface with lower layers for acquiring information and affecting the system (e.g., its configuration, goals, and policies).

In the future, we plan to evaluate the architecture using other relevant metrics for evaluating software architecture as per Silva et al. [58] and through prototype implementations.

6.2. Detailed design and implementation ideas: methods and technologies

The following methods and technologies suggest means for addressing the development of concrete architectures (or specific functional layers and components) based on the proposal of this paper.

“Mature” techniques and technologies: machine (federated) learning and edge computing. There is a number of “obvious” methods and tools that we briefly mention. *Machine learning* is clearly fundamental for modelling and prediction of crowd state and dynamics; typical uses of machine learning (ML) are for supporting tasks like crowd counting and density estimation [63–65], crowd dynamics analysis [67], and anomaly detection [66]. *Edge computing*, possibly combined with *federated learning*, can be useful for addressing both the real-time (N2) and privacy (N4) requirements in crowd management systems [72]. See Table 2 for a summary of common methods and technologies for typical crowd management tasks.

Sensor fusion and complex event processing. To achieve a fast and accurate estimation of the situated crowd state (cf. the modelling requirements F1, F2, F3), it is essential to integrate the information coming from all the available sensors (e.g., cameras, wearables) [73]. The survey from Li et al. [73] provide a comprehensive overview of sensor fusion in the context of crowd management, discussing how multi-source, multi-modal data could help to improve accuracy and timeliness of context acquisition. Sensor fusion is also an enabler for *complex event processing* and the identification of *macro-events* that aggregate several spatiotemporally correlated micro-events [74].

Multi-level, multi-scale (micro-to-macro), and surrogate modelling and (co-) simulation. Research has shown that crowds can be modelled or simulated at multiple levels and scales [1,57]. This range of possibilities provides flexibility (cf. N6) w.r.t. the desired level of detail, fidelity, real-time, and efficiency (cf. N7). Moreover, given the time-expensiveness of simulations, leveraging *surrogate models* [75] for faster, approximate results may improve real-time support (cf. N2). Regarding the potential need of modelling and simulating different heterogeneous subsystems (at the minimum, the crowd itself and its environment—cf. F1 and F2), the adoption of *co-simulation* techniques [76] is also to be taken into account.

Multi-layer support for high-priority requirements and graceful degradation. The Swiss Cheese Model is a paradigm in safety science that suggests to consider multiple layers of protection to address safety [22]. Similarly, in CDTs, it is key that certain functions (especially those related to the situated crowd monitoring—cf. (F3) *Situated crowd model synchronisation* and (F6) *Interactive multiscale analysis, visualisation, and decision-making support*) are supported by different mechanisms and strategies that may involve different sets of device capabilities. For instance, if cameras are not available or unable to provide images of the crowd (e.g., because smoke or objects are obstructing the vision), a different means for crowd density estimation should be leveraged (e.g., using wearable devices [77] or floor vibration sensors [78]). On one hand, the concurrent adoption of multiple mechanisms to support a given task can be used to improve the accuracy, by combining models and fusing sensor data. On the other hand, different mechanisms could work as *fallbacks* for graceful service degradation. The different layers may be dynamically activated or deactivated by taking into account both the context/situation and the system efficiency (e.g., saving energy in low-risk situations). For this purpose, it may be beneficial to leverage *requirements engineering methods* like RELAX [79] that support uncertainty-aware non-critical requirements relaxation to favour short-term critical requirements.

¹ Soft real-time systems aim to provide best-effort reliability, with some degree of tolerance for missed deadlines.

Table 2
Mapping of crowd-management tasks to methods and enabling technologies.

Task	Methods	Enabling technologies
Crowd counting and density estimation	convolutional neural networks, transformers, generative models, detection-based methods, point-based methods, graph-based models [63–65]	video camera networks, heterogeneous IoT/mobile/crowd sensing (WiFi) [4,19]
Anomaly detection (motion, behaviour)	convolutional neural networks, spatiotemporal models, generative models, graph- and interaction-based models [66,67]	video camera networks, heterogeneous IoT/mobile/crowd sensing (WiFi) [4,19]
Localisation, and tracking	convolutional neural networks, spatiotemporal models, generative models, graph- and interaction-based models, model-based approaches (Bayesian filters) [66,67]	video camera networks, heterogeneous IoT/mobile/crowd sensing (WiFi) [4,19]
Flow prediction / mobility forecasting	agent-based modelling, reinforcement learning, graph neural networks, sequence models, recurrent neural networks (LSTM/GRU) [21,57]	video camera networks, heterogeneous IoT/mobile/crowd sensing (WiFi) [4,19]
Crowd control	direct and indirect steering devices [68–71]	wearable devices (haptics) [69]; environment actuators (lights [70], sounds [71])

Opportunistic sensing and participatory sensing. Opportunistic sensing leverages personal smart devices (e.g., smartphones, wearables) to collect contextual data without explicit user intervention, adapting to users' routines and device availability, while participatory sensing enlists individuals to actively contribute data—such as location traces or environmental measurements—to a shared platform [80]. Both enable fine-grained, large-scale monitoring of crowd dynamics and environmental conditions, complementing static sensors to enhance coverage, resilience, and privacy in CDTs. The social internet of things (SIoT) paradigm [81] further organizes devices into social networks, where trust, shared context, and group affiliations guide privacy-aware data exchange and QoS-aware task orchestration. Combined, these methods support the continuous, distributed sensing needed for real-time model synchronization (F1-F2) and adaptive crowd management.

Collective computing [82] and aggregate computing [83]. The aggregate computing paradigm, as a method for artificial collective intelligence engineering [84], enables to express self-organisation through decentralised collective computation at the level of IoT devices. Thus, it is useful for providing, e.g., crowd density estimation and control from the bottom-up, without necessarily involving fog and cloud layers. This could also be instrumental for edge-cloud incremental, approximate, spatial computing, possibly leveraging amorphous computers [85] of resource-constrained devices [65].

Context-awareness, adaptation, and self- features.* There exist several contributions in the autonomic and self-* systems community that support the proposed requirements and architecture. Indeed, the proposal of this paper is centred around an architecture-based adaptation approach [86]. MAPE-K patterns [87] suggest ways to organise and decentralise the autonomic functionality to address requirements of different infrastructures—cf. (N6) *Flexibility, configurability, (self-)adaptability, and maintainability*. The thread on *self-improving system integration* [88] provides means to address (N5) *Integrability*.

Crowd steering devices. Though not extensive, there exist some proposals of mechanisms for crowd-targeting actuation, via so-called *crowd-steering devices* [68]. One approach for affecting crowds is by broadcasting signals to a large portion of their individuals: this can be done through environmental stimuli [20], e.g., based on arrow-like signals [68], lights [70], or sounds [71]. Notice that experiments have typically carried out in specific environments (e.g., corridors) and it may not be clear how to influence crowds in less constrained or more complex environments. Another way to influence a crowd is based on tactile feedback from wearables [69], but we are not aware of studies aimed at large crowds.

Uncertainty-aware approaches. Situated crowds as systems are characterised by large *uncertainty* regarding the events that may happen and the overall environment-crowd dynamics. It is therefore crucial to

take uncertainty into account, and hence adopt *uncertainty-aware methods* [89]. It is key to identify the different sources of uncertainty (i.e., managed system, environment, goals, resources, model, and adaptation function uncertainties [90]) and to adopt methods to deal with them (e.g., data quality management, probabilistic analysis, verification, learning, and adaptation [89]).

Short Answer to (RQ2.2) CDT enablers

Enabling methods and technologies include: *sensor fusion, surrogate and multi-level modelling, collective computing, autonomic, self-*, and uncertainty-aware methods*, and *layered solutions* to ensure safety (through risk management), flexibility, scalability, adaptability, and efficient implementations. (cf. Sections 5 and 6.2)

Short Answer to (RQ2) CDT design and implementation

The design and implementation of CDTs should be addressed by layered, modular, and autonomic architectures where the different CDT requirements are addressed (i) across infrastructural layers (cf. edge-cloud continuum), (ii) by multiple layers of heterogeneous solutions that may exploit the available resources and organise the distributed computation in multiple complementary or alternative ways. (cf. Sections 5 and 6.2)

6.3. Case study: mass event planning, execution, and post-analysis

This subsection details a case study of CDT realisation and application based on the large-scale crowd management reference scenario introduced in Section 4.3.2.

6.3.1. Motivation and objectives

Mass events like concerts, festivals, marathons, and large exhibitions gather hundreds or thousands of people in a constrained environment, posing strict safety challenges and normative requirements (cf. N1, N4). The conceptual and architectural framework of CDT can be instantiated in such scenarios to support the entire lifecycle of a situated crowd, covering both *planning, execution, and post-execution analysis* of the event (cf. F1 to F8).

6.3.2. Pre-event planning

Before the event, organisers (cf. Section 4.4) must consider (as constraints or as alternatives to be evaluated) at least (i) a *target location, preliminary schedule, and context information* for the event and the organisation (cf. (L0) *Physical layer, F2, N5*), (ii) the *available infrastructure* for the system (cf. (L1) *Cyber-physical layer, F2, F5*), (iii) the *preliminary information about the expected crowd* like an estimated number of participants (cf. L0, F1, F3), and (iv) the relevant *regulations* providing the normative frame (cf. F2). In the following, we consider some arbitrary

event which is large-scale (cf. *N3*), in terms of expected participants and spatial extension, and that includes both outdoor and indoor locations (cf. *N6*). A planning application and dashboard in (*L4*) *Application layer* can support the organisers, by interfacing with (*L3*) *Service layer*, in the CDT construction and configuration during the pre-event analysis and planning.

Data ingestion and environment modelling. The first step of the CDT construction begins by constructing the EDT (cf. *F2*). The EDT can be built using a wide range of data and models, overall contributing to annotated maps of the environment, e.g., with geographic information system (GIS) layers (cf. *N5*). Data sources and concrete model formats may include *architectural blueprints*, *building information models (BIMs)*, *city information models (CIMs)*, and/or *simultaneous localisation and mapping (SLAM)* systems. These models populate the Knowledge component of (*L2*) *Autonomic layer* and serve as a basis for the situated crowd modelling (cf. *F1*, *F3*) and analytical (cf. *F4*) activities.

Scenario generation and simulation. The next step involves using data about the expected participants and ingress/egress mobility flows to generate models and simulations of virtual crowds over the EDT (cf. *F3*, *F4*). This generation activity will be based on existing crowd models (assumed to be available at the Knowledge component), possibly at different scales, enabling a wide set of investigations depending on the required accuracy (cf. *N1*, *N6*). For instance, macro-level flow models can be used to evaluate large-scale ingress/egress patterns and detect bottlenecks and attention points. On such hot spots (cf. *N7*), micro-level models can be used to better understand the dynamics and identify possible intervention points (e.g., adjustments to the environment, monitoring and/or management and/or operational recommendations). From this analysis, *key indicators (KIs)* should be determined, covering, e.g., expected queue lengths at gates, overall crowd density limits, congestion risks at junctures or corridors, estimated evacuation time under different hazard scenarios. All the resulting information and models will converge on the Knowledge component and support further inspection and reasoning (cf. *F4*).

Operational policy design. The constructed model, inferred information, as well as early decisions and constraints can be used to define the *operational policy*, to be submitted to authorities for approval, to guide recruitment of workers, and to inform the workers themselves regarding all surveillance and control activities at the live event. The policy will report on the KIs and may define environmental amendments (if possible), monitoring objectives, early warning indicators (cf. *F4*), triggers for manual (human-based), automatic (CDT-driven), and hybrid interventions (cf. *F5*), responsibilities, and the schedule. This contributes to lifecycle support (*F8*) and to preparing adaptation-related constraints and goals stored in the Knowledge component (*L2*).

6.3.3. Event execution

Right before and during the event, the CDT provides a continuously synchronised representation of the evolving situated crowd, supporting real-time monitoring (cf. *F1*, *F2*, *F3*), forecasting (cf. *F4*), and decision-making (cf. *F5*, *F6*), while satisfying real-time (cf. *N2*) and privacy (cf. *N4*) constraints. Integration of any contextual information (e.g., weather forecasts, city logistics—cf. *N5*) is also essential for effective operations.

Sensing, data fusion, analysis, and real-time crowd state estimation. The (*L1*) *Cyber-physical layer* includes devices and sensors (e.g., cameras with anonymised analytics, light detection and ranging (LIDAR) and thermal sensors, volunteering wearables) that support perception of the physical layer state, which includes the environment and the people located therein (cf. *F1*, *F2*, *F3*). The sensed data converges into the Monitor component, which implements a sensor fusion pipeline to determine aggregated measures estimating density, flow, group formations at the desired scale and accuracy. The Monitor then updates the Knowledge

component, where an up-to-date situated crowd model becomes available (cf. *F3*). The Monitor also feeds the Analyze component, which can trigger simulations and predictive analyses (cf. *F4*), in order to (i) provide forecasts and further information to the managers (cf. *F6*), and (ii) determine whether to activate autonomous reactions (e.g., early warnings) or consider the need for adaptation through the Plan component (cf. *F5*).

The application layer (*L4*) may provide managers and authorities with visual analytics tools (cf. *F6*) enabling inspection of the current situation. A dashboard can provide 2D/3D representations of the crowd at different spatial resolutions, heatmaps of crowd density, and outputs of predictive analyses such as forecasts and risk estimates (cf. *F4*). Concurrently, all relevant data (events, system states, decisions) is stored in the Knowledge component for history tracking and retrospective analytics (cf. *F7*, *F8*).

Decision-making and actuations. Decision-making and actuation can be of one or more of three different modalities: (i) human-based, (ii) human-in-the-loop-based, and (iii) autonomous. Managers can use information from the dashboard of a “monitoring and control panel” to coordinate human workers or to activate specific CDT functions (cf. *F6*). When the Analyze component identifies the need for adaptation, it triggers the Plan component, which determines adaptation plans (cf. *F5*), consistent with safety and dependability requirements (cf. *N1*, *N1.3*). Such adaptation plans can be of two types: (i) those requiring human approval or refinement, or (ii) those which can be autonomously executed. Plans of the latter kind are forwarded to the Execute component, which orchestrates the cyber-physical layer to run interventions targeting: (i) the environment (e.g., dynamic reconfiguration of routing signage, light/sound cues), (ii) the crowd (e.g., personalised guidance via mobile apps, push notifications, wearable haptics), (iii) the staff (e.g., notifications to authorities or workers). The decision-support dashboard should display real-time and predictive views and the status of ongoing adaptations (cf. *F6*), including the rationale behind model outputs, thereby enabling human-in-the-loop oversight (cf. *N1.3*).

Handling disruptions and emergencies. The CDT can support preventive preparation and alerting by detecting precursors of incidents and running short-horizon predictive models (cf. *F4*). If an incident occurs, it can quickly notify managers, workers, and authorities, and may activate emergency plans (e.g., activating lights and sounds to support effective crowd dispersal strategies). In all cases, CDT behaviour adheres to the operational policy and contingency plans configured in the Knowledge component (cf. *F8*).

6.3.4. Post-event analysis

Once the event has concluded, the historical data tracked in the Knowledge component supports retrospective analysis for improving CDT models and organisational policies (cf. *F7*, *F8*, *N6*). Examples include the re-assessment of high-risk episodes, the evaluation of compliance to regulations and organisational policies, the benchmarking between planned vs. observed crowd behaviour, and the re-training of models (e.g., density estimation, crowd flow prediction, anomaly detection). Such insights inform future events, realising a continuous learning cycle that improves both the CDT and organisational processes.

6.4. Limitations and threats to validity

Internal validity relates to the extent to which the design of the study may undermine the conclusions [91]. Common threats include bias, missing factors, and model simplification. We mitigated (i) bias, by involving multiple researchers with different expertise (self-* systems, machine learning, networking, human-computer interaction), (ii) missing factors, by considering related work on crowd management and DTs (cf. Sections 4.2 and 8), (iii) model simplification, relatively to the state and goals of the reference architectural proposal, by considering

multiple views and assessing detailed design/implementation ideas (cf. [Section 6.2](#)).

In future work, the former two threats can be further reduced by systematic literature reviews, investigating the areas identified by the scoping review in this article. Instead, model simplification threats are meant to be addressed in future work by more comprehensive studies of the mentioned enabling technologies and methods, more detailed architectural descriptions, and, crucially, prototypes and simulations for assessing the feasibility and effectiveness of the proposal.

External validity refers to the extent to which the results can be generalised [91]. One threat is given by the inherent dependency of DT-based system requirements, level of maturity model, and architecture to the context and application at hand. We have mitigated this by considering the most comprehensive scenarios (e.g., deployments across the entire edge-cloud continuum) and considering different levels of requirements according to the MoSCoW approach (cf. [Section 4.5](#)).

7. Research roadmap

On the basis of the CDT concept, requirements, architecture ([Section 4](#)), and enabling technologies ([Section 6.2](#)), we extract the most interesting *research challenges*. We believe that addressing these challenges is necessary in order to fulfil the CDT vision. Moreover, these are non-trivial challenges that are expected to stimulate researchers for years to come. Notice that challenges may also be related with one another (e.g., modelling is crucial for early detection).

7.1. Seamless, adaptive multi-level/multi-scale modelling

Challenge overview. Multi-level modelling is fundamental for the reuse of models and for achieving the desired efficiency-accuracy tradeoffs that are essential for crowd management [57]. The challenge here is how to effectively *integrate* different models together and how to *adapt* the models (e.g., their spatiotemporal resolution or composition) according to changes in the environment. Indeed, for instance, changes in the available infrastructure (e.g., connectivity or device failures) may require the adoption of different models for improved performance or gentle degradation. In the general respect of the adoption of model-driven engineering for DTs, a recent systematic mapping study by Lehner et al. [92] points out, as challenges, the need for (i) more mature, industry-ready approaches, (ii) increased automation, (iii) reuse of models and, especially, (iv) the adoption of such techniques for new kinds of DTs, such as those representing systems-of-systems—which is the case of our CDT proposal.

Potential solution paths. Research on design patterns for multi-level modelling and simulation could prove especially useful. Serena et al. [93] provide an early catalogue on multi-level modelling patterns, classifying them into orchestration patterns, execution policy patterns, information exchange patterns, and multiscale patterns. The multi-scale pattern class further includes (i) *spatial aggregation/disaggregation patterns* and (ii) *adaptive resolution patterns*. However, only limited work (e.g. [94]) has been carried out on the latter aspect, which we deem important for CDTs. One suggestion is to investigate how self-adaptive software research contributions [86] could be integrated into *hybrid modelling frameworks* [95]. In this sense, conceptual and operational frameworks for reasoning about models, like *Models@Runtime* [96,97], can support dynamic model switching and model refinement based on causal linkage between models and reality/goals.

Models@Runtime also more generally supports adaptation and uncertainty management across the CDT functionality. Indeed, such methodology combines model-driven engineering with computational reflection, hence promoting the definition of causally connected self-representations and reasoning for self-adaptation. By a language-based software engineering perspective, an interesting concept to be investi-

gated is *domain-specific integration languages* [98], whose goal is to reuse and compose existing models into an integrated solution.

7.2. Early event detection and warning

Challenge overview. One major goal in crowd management systems is the reduction of the risks of crowd disasters and mitigating consequences. Early event detection and warning is essential for both prevention and timely situation management. This activity builds on real-time data, prediction models, risk mitigation models, and experience of the managerial personnel [99]. Though there exists preliminary work on the topic [100], solutions are typically limited in scope, and hence this is still an open challenge, which would be key to address using an interdisciplinary approach. For instance, Ramesh et al. [100] propose the use of mobile ad-hoc networks and participatory sensing for estimating the risk of stampedes by distributed fusion of data from accelerometers, GPS, acoustic sensors, and video sensors. The use of opportunistic networks and collective computing paradigms have also been proposed by Casadei et al. [33], but one challenge in such distributed computations (where information flows and consensus are typically involved) is related to reactivity.

Potential solution paths. For addressing the issues related to the reactivity of collective estimation and consensus, one path is suggested by *real-time aggregate computing* [101], where the idea is to automatically derive proofs from collective control programs providing real-time guarantees. More generally, efforts should be concentrated on the detection of incident precursors [102] and the implementation of efficient short-horizon predictive models and simulations. For the former aspect, it will be key to explore combined frameworks based on machine learning techniques applied to *historical incident data* [67] (cf. [Section 6.3](#)), *context learning/calibration/transfer* [103], *knowledge distillation* into efficient models [104], and attention mechanisms [105]. For the latter, ideas include the application of *context-aware* and *probabilistic* motion prediction models [106,107] as well as approaches to model adaptation (cf. [Section 7.1](#)) and to efficient stacking of multiple models [108].

7.3. Proactive and anticipative crowd management (crowd control planning and/or execution)

Challenge overview. The difficulty of handling emergency situations like evacuations and stampedes suggests to focus on prevention through risk-reducing control actions.

Anticipative crowd management [109] builds on early “preliminary event” detection (cf. [Section 7.2](#)) and is crucial for mitigating risks. One safety measure for preventive stampede management include the crowd steering and flow control (e.g., via acoustic and optical signals [109]) aimed at the creation of room and the avoidance of occlusions in escape points [109]. Another strategy is the detection of crowd forces to predict the occurrence of stampedes and then apply preventive strategies for pressure relief and crowd pressure shockwave hindrance [100].

More generally, *proactive crowd control* (i.e., *autonomous crowd control*), which may be limited to planning (e.g., when human confirmation is required) or even extended to execution (for fully autonomous implementations), is an action that may happen (if enabled) during the different phases of an event or situation (i.e., before, during, and after). It is fundamental to configure the corresponding policies taking into account the broader context of operation, including national and supranational regulations. Indeed, major issues with crowd control are related to law, accountability, and ethical concerns: these aspects are discussed in detail in [Section 7.6](#).

The implementation of effective autonomous control planning and execution is challenging in many respects:

- *reliability of inputs of crowd control*: to make accurate predictions, accurate models reflecting the actual physical system state should be in place, and their *level of confidence* should be taken into account;

- *on the construction of reliable inputs for crowd control*: accurate models of crowd states might need to be inferred from noisy, partial, or delayed sensor inputs, ideally in an uncertainty-aware fashion (Section 6.2);
- *coordinated multi-agent actuation*: CDTs are typically *heterogeneous* systems where many kinds of devices and agents are in place and, thus, the space of possible orchestration plans might be huge to explore, also requiring probabilistic reasoning;
- *dynamic planning*: possible and unpredictable events may happen at any time, possibly requiring a revisitation of existing control plans, or a complete re-planning;
- *multi-objective policy synthesis*: control policies should find a (typically complex) balance among safety, efficiency, normative constraints, and ethical considerations (cf. Section 7.6);
- *reliability and generalisability*: it is key to understand how much the control policies are robust with respect to variations in context, namely across different environments, crowd types, and event dynamics (cf. bias-variance tradeoff);
- *explainability*: it is key the control plans are interpretable by human operators and stakeholders, for supporting timely confirmation, amendments, logging, and post-audit analysis;
- *adversarial resilience*: when safety is to be maximised, it is important to take into account the possibility of malicious behaviours or manipulations that might compromise the control goals.

Given the difficulty of predicting the effects of crowd control activities, it is hard to see at the present time something but human-designed, well-defined management policies taking into account the possibility of limited, soft, autonomous control actions in safe conditions, to ensure zero harm by the system. However, further research might help to devise reliable strategies to be incorporated more widely.

Potential solution paths. The potential solution paths include those indicated for early event detection and warning (see Section 7.2). Furthermore, for crowd management planning and execution, it is fundamental to investigate methods able to integrate confidence levels and risk levels into operational models and to *quantify trust under uncertainty* [110], techniques for explainable models [45], and effective visual analytics [111] to maximise safety while promoting human-in-the-loop oversight (see also Section 7.5).

7.4. CDT security and privacy

7.4.1. Security

Challenge overview. Security is critical for supporting the crucial safety requirements of CDTs. This also entails protecting infrastructure, data, and software. It is required that CDT implementations include countermeasures to preserve *confidentiality, integrity, and availability (CIA triad)*. One challenge is due to the distributed, heterogeneous, and dynamic nature of data sources and actors – including IoT devices, wearables, humans, and deployed services – which imply a *very large attack surface*. Examples of threats include data injection (e.g., false crowd signals), model manipulation (e.g., malicious inputs to learning algorithms), and actuation hijacking (e.g., malicious activation of emergency protocols).

Potential solution paths. Countermeasures might include the adoption *security-by-design* and *zero-trust models* throughout the CDT layers, formal verification of adaptation and control logic, and multi-level access control. In particular, adversarial robustness of machine learning components (e.g., those used for predictive modelling and event detection) must be ensured, as even subtle perturbations may cause dangerous misclassifications or false positives in crowd event recognition. Run-time monitoring and anomaly detection systems should be integrated across the CDT layers to enable real-time identification of policy violations, data drift, or signal tampering. Moreover, explainability and transparency in the CDT decision pipeline (e.g., being able to trace why

a certain intervention was triggered) can support post-incident audits, improve operator trust, and discourage exploitation by contributing to *accountability* (see next for more about it). The previous considerations on multi-level support of functionality and graceful degradation are also relevant when facing attacks to availability (e.g., occlusion to cameras' fields of view). Future research should explore formal threat models specific to CDTs, that incorporate both the cyber-physical and the socio-technical dimensions, also through quantitative security metrics.

7.4.2. Privacy

Challenge overview. There exists a trade-off between privacy and safety/security measures. Indeed, tracking individuals or leveraging granular data (e.g., trajectories, biometric cues, or behavioural fingerprints) may support effective crowd monitoring, surveillance, and early detection of trigger events. However, the extraction of such data adds privacy risks as they can potentially enable re-identification even if anonymisation techniques are applied.

Potential solution paths. Research must explore privacy-preserving computation paradigms (e.g., aggregation techniques, differential privacy) that support large sets of CDT functionality without exposing raw data. Furthermore, privacy should be viewed as a dynamic, context-dependent issue rather than a static policy. CDTs should allow for adjustable privacy levels according to risks, event criticality, and crowd features (e.g. density). Methods to acquire and consider explicit privacy preferences by individuals and groups could also affect decision-support in terms of recomputed privacy-functionality tradeoffs.

7.5. Human-“crowd machine” interaction

Challenge overview. Human-“crowd machine” interaction represents a complex frontier in the development of CDTs. Unlike traditional human-machine collaboration, CDTs must account for the collective behaviours, emotions, and dynamics of large groups, making real-time modelling and prediction significantly more challenging [112]. Integrating human-centric design principles into CDTs is essential yet challenging. Traditional digital twins have often prioritised system efficiency over human factors, leading to interfaces and models that may not adequately reflect human needs or facilitate intuitive interaction [113].

Potential solution paths. Promising directions for future human-CDT interaction emphasise designs centred on human abilities and needs, rather than solely on technological capabilities.

- *Intuitive human machine interfaces (HMIs)*. Developing multimodal and user-adaptive HMIs, incorporating voice commands, gestures, gaze tracking, and natural language processing, can significantly improve interaction fluidity. These modalities allow operators to interact with complex crowd systems in more natural ways, reducing cognitive load and enhancing situational awareness [114].
- *Human-centred visualisations*. Interfaces should prioritise dynamic and interpretable visual representations of crowd behaviour, such as real-time heat maps, crowd density overlays, and projected flow trajectories. Such visual tools enable operators to understand evolving patterns at a glance, thereby supporting timely and informed decisions [115].
- *Human-in-the-loop feedback mechanisms*. Rather than fully automating decisions, CDTs should incorporate human-in-the-loop capabilities. These allow experts to validate simulations, adjust control parameters, and intervene when automated suggestions diverge from context-specific knowledge or social considerations [112].
- *Situational and role-based context awareness*. Interaction models should adapt to the user's context—whether they are a first responder, planner, or event manager—by offering tailored control dashboards and filters. Context-aware interaction ensures that users are not overwhelmed with irrelevant data and can access the most critical information quickly [105].

- *Participatory and co-design approaches.* Involving end-users in the early stages of interface and interaction design leads to systems that are more usable, trusted, and aligned with real-world operational demands. Participatory design also helps uncover domain-specific interaction patterns and workflows that generic solutions might miss [116].

Essentially, interactive visual analytics is key to support operations, and its effective design is crucial given the safety-critical nature and the human-in-the-loop time requirements of typical CDT applications.

7.6. Philosophical and social aspects of crowd management: autonomy, objective evaluation, accountability, ethics, etc.

The deployment of CDT introduces a complex interplay of safety, autonomy, and ethical considerations. As these systems increasingly influence real-world decisions, particularly in high-stakes environments like public safety and urban planning, it is imperative to address the multifaceted challenges they present.

Reliability and guarantees. Ensuring the reliability and safety of CDT requires formal verification methods to guarantee that system behaviours align with intended outcomes [117]. Given the complexity of modelling human behaviours and interactions, establishing proofs for system correctness becomes challenging. The absence of such guarantees can lead to *unintended consequences*, especially in critical applications like emergency evacuations or public event management. Approaches for quantitative risk assessment should be investigated (see Section 7.3).

Autonomy. In CDTs, it is crucial to properly design the levels of system autonomy, and to take into account safety, ethical, and normative constraints (see next). A potential increase in control autonomy poses several challenges, like maintaining human oversight [118]. While automation can enhance efficiency, it may also lead to scenarios where human operators are sidelined, reducing their ability to intervene in critical situations. Ensuring that CDTs function as decision-support tools rather than decision-makers is crucial to preserve human agency. Depending on regulations, different boundaries may be considered, and hence exhibiting forms of *adjustable autonomy* [119] could be important for flexible CDTs implementations.

Normative and ethical aspects. Since CDTs could work in private and public spaces, and deal with human safety, it is crucial to consider normative and ethical aspects, both in design and in operation (cf. context awareness). Veltmeijer and Gerritsen [120] investigated the legal and ethical implications of artificial intelligence (AI) in automated crowd analysis, from a European perspective. They focused on three main aspects: i) privacy, ii) bias and accuracy, and iii) accountability. The complex landscape of privacy involves a crucial dual aspect. On the one hand, individuals should be made aware if and how their data will be used; on the other hand, extra protection is required for sensitive data or personal data of children and other vulnerable groups [53]. Then, there is a risk that these systems may inadvertently perpetuate existing *biases* present in the data they are trained on [121]. Such biases can lead to discriminatory outcomes, particularly against marginalised groups. Developers must implement strategies to identify and mitigate biases. Finally, the opacity of DT implementations and algorithms may obscure the decision-making process, making it difficult to ascertain responsibility when outcomes are unfavourable [122]. Even if the CDT considers a *human in the loop*, who is ultimately responsible for decisions to act, mistakes and adverse consequences may still happen, and may not be easy to evaluate a posteriori the appropriateness of the course of action. Veltmeijer and Gerritsen recommend avoiding automated decision-making and adhering to regulations like the AI Act [123] as much as possible. The formal framework of *SLEEC (Social, Legal, Ethical, Empathetic, Cultural) rules* [124] can help to operationalise ethical reasoning, whereas

research on *normative multi-agent systems* [125] can offer solutions for ruling autonomous behaviour.

Short Answer to (RQ3) CDT research challenges

The challenges involved in realising the CDT vision and their implementation are related to reliably and efficiently achieving key functions (e.g., safety) through a proper integration of models, methods, and actors (i.e., humans and machines). Importantly, some challenges are not just technological but also philosophical, cultural, and inter-disciplinary (cf. law, autonomy, guarantees, accountability, and ethics). (cf. Section 6 and 7)

8. Related work

To the best of our knowledge, there exists no proposal for *DT-based* solutions for crowd modelling and management.

There exist, however, work on DT for cyber-physical systems (CPSs). For instance, Josifovska et al. [126] propose a framework for DTs of CPSs around four main building blocks: (i) the *physical entity platform*, consisting of physical objects, physical nodes, and humans; (ii) the *virtual entity platform*, maintaining semantic models regarding the geometry, physical features, behaviour, business rules, and processes of the interacting entities; (iii) the *data management platform*, consisting of data models and data management methods; and (iv) the *service platform*, consisting of service models and service management layers. In comparison, our CDT proposal is specialised for crowd management and emphasises the adaptive component of the solution through the autonomic layer.

So, as related work, we cover contributions related to *architectures and simulation platforms for crowds*, which may work as inspiration or components for the proposed CDT design and implementation. The authors observe that *crowd simulators* tend to be tailored to specific ranges of numbers of simulated pedestrians, and that *crisis simulators* tend to focus on a limited set of target environments. Different simulators may support different feature sets, such as navigation methods, agent behaviour, collision avoidance, emotional aspects, and realistic group dynamics. Also, the authors note that: multi-agent-based modelling is raising, with increasing crowd sizes supported; game engines are often leveraged; ML and especially reinforcement learning (RL) can be effective in reproducing human-like behaviour.

8.1. Architectures for crowd monitoring and management

Wijermans et al. [3] propose *InCrowd*, an integrated framework for organising crowd management decision-support systems. Their proposal is based on four subsystems: (i) the crowd-interaction subsystem consists of sensors and actuators, and provides the interface between the physical crowd and the decision-support system; (ii) the mining subsystem analyses the measured crowd state to obtain a refined crowd state; (iii) the predicting subsystem aims at computing the future crowd state; and (iv) the decision-making subsystem is responsible for determining and enacting management decisions. InCrowd and CDT are similar, but in CDT we emphasise the DT and autonomic aspects. In particular, the DT component is relevant for (i) supporting analysis and design throughout the entire lifecycle of applications and events involving crowds, and (ii) supporting bidirectional synchronisation (that we address on the digital → physical direction through the autonomic layer).

Darsena et al. [19] consider crowd management in the context of public transportation systems. Their abstract architecture is organised around two main components: an analysis/monitoring component and a control component. The analysis component performs data pre-processing, sensor fusion, feature-based data fusion, and decision fusion, and provides real-time data to the control component. The control

component has a prediction module that, leveraging real-time monitoring data, Internet data, and historical data, provides information to the decision-making and control module. Upon this basis, the authors provide a reference architecture for intelligent transportation systems organised in three layers: sensing and actuation subsystem, communication subsystem, and monitoring, prediction, and control subsystem. Our proposed architecture (Section 5) is more general, not focussing on transportation systems only, is less ad-hoc through the reuse of the MAPE-K architectural pattern, and details on the kinds of models adopted.

Perez and Zeadally [4] propose a *Communication Architecture for CRowd Management (CACROM)* to support emergency scenarios. The goal is to help event planners in the monitoring of crowds, while also considering the possibility of infrastructure disruption (both at the core network and at the edge). The proposal is structured in three segments: the power segment, the communication segment, and the edge device segment. As challenges, the authors indicate power resilience, infrastructure availability, security, privacy, and participation. CACROM focusses on the communication and networking aspects of crowd management, whereas our CDT reference architecture provides a more abstract, end-to-end solution. CACROM could be integrated at the CDT levels *LO* and *LI*.

Jiang et al. [2], propose a ultra large-scale crowd monitoring system architecture, combining air-sky-ground control. Their focus is on the networking architecture, based on five infrastructural layers: private cloud, edge cloud, gas and ultrasonic network, unmanned aerial vehicle (UAV) network, and closed-circuit television (CCTV) and infrared network. By a functional point of view, the focus is on crowd detection using ML models analysing camera images and providing people counting results. The work somewhat neglects the need of modelling crowd dynamics, relating crowds to the surrounding environment, and crowd management. However, the proposed architecture may support part of the digital thread contributing to a CDT.

The 2017 study by Martella et al. [32] discusses limitations of state-of-the-art crowd management solutions, found mostly associated to (i) limitations due to bias, simplification, inaccuracy, invalidity of technical solutions of what-if analysis, simulation, estimation models, decision-making models; (ii) limitations due to insufficient data quality and automation, and (iii) limitations due to scarcity of actuators and means for timely intervention. We observe that these limitations are mostly technical and addressable: the CDT proposal can provide a path for addressing them through the improving feedback loops enabled by the DT and autonomic frameworks.

8.2. Technologies for crowd management

Sharma et al. [127] provide a review of technologies for crowd management. They organise contributions following a task and methodological/technological area classification: (i) for crowd modelling: flow-based, entity-based, and agent-based approaches; (ii) for crowd information acquisition: vision technologies, wireless/radio-frequency technologies, and social media/web data mining. The techniques in Section 6.2 integrate and extend those proposals (i.e., existing modelling approaches and data acquisition technologies) to rather focus on complementary and integrative scientific proposals and paradigms—fundamental to address the CDT challenges beyond “plain” crowd management.

Darsena et al. [19] survey sensing technologies for crowd management specifically to the public transportation domain. They consider infrared sensors, pressure/load sensors, optical cameras, thermal cameras, light detection and ranging (LiDAR), acoustic/ultrasound sensors, device-free radio-frequency sensors, and device-aided radio-frequency sensors (e.g., RFID/NFC, Bluetooth, WiFi, cellular). The mentioned technologies can be integrated at CDT level *LI* to provide further data sources for the Monitor component.

8.3. Crowd simulation

A large amount of research has been conducted on crowd simulation, as also witnessed by recent surveys on the topic.

Sidiropoulos et al. [21] review contributions on crowd simulation, with special focus on crisis management. Their analysis framework is based on four elements: key supported features (e.g., wayfinding, environment realism, kind of approach), supported scenarios, simulation size (number of agents), support for parallel execution, and software platform (e.g., MATLAB, Unity, Menge). These kinds of contributions fit the CDT architecture and can be integrated therein as sub-components (cf. simulation models and tools in the Analyze and Knowledge components, and lifecycle management platforms for organisers and managers).

Dang et al. [128] provide a review on *dense* crowd simulation. They classify research contributions based on Hoogendoorn’s classification into the strategic, tactical, operational levels [129]. Regarding future directions, they mention the need of collecting high-density crowd datasets, the use of deep learning techniques to predict human behaviour, the integration of different modelling levels to dynamically deal with varying densities, and the support for multi-scale simulation by hybrid modelling and consistent dynamic dis/aggregation of crowd representation.

Serena et al. [57] provide a review of multi-level modelling and simulation of human mobility and behaviour. They detail on the types of simulation paradigms: discrete-event simulation (DES), agent-based models (ABM), cellular automata (CA), continuous simulation models, Monte-Carlo methods. Then, they review focus on three aspects: epidemic modelling, traffic modelling, and crowd mobility modelling. On crowd mobility, they review social force models, centrifugal models, agent-based models, and group dynamics models. They conclude that different kinds of (micro, macro, or multi-level) models may fit different kinds of applications. We agree on that, and specifically to CDTs, we foresee that *context awareness* would be key to dynamically choose the appropriate model composition to achieve the desired trade-offs.

8.4. Human digital twins

A *human digital twin (HDT)* [6–8] applies the DT concept to humans, hence creating a digital shadow of (a part of) a human, and possibly supporting ways to influence it. Motivated by recent trends in human-centric engineering, HDTs help to provide insights on human behaviour and tackle challenges in human-machine networks [7].

The HDT involves replicating the individual human body, in whole or in part, within a virtual space, reflecting both its physical state, psychologically and physiologically, in real-time [130]. A HDT should be uniquely identified, evolvable, synchronously connected with its physical twin, in relationship with the virtual representation of its surrounding environment, actionable (supporting analysis and prediction), authorised, privacy-preserving, and secure [131].

A HDT can be employed to monitor the health and well-being of a person, provide personalised support, and improve the safety of individuals, e.g., by preventing hazardous situations such as a driving accident, or in emergencies, where the HDT can automatically send an emergency call to the designated emergency services [6].

In the literature, various applications of the HDT are found, particularly many applications are related to the healthcare and industrial fields, while fewer are associated with the sports and transport sectors. The research work in [132] is related to a project sponsored by the US Air Force, which aims to create a “twin” of an airplane pilot using a variety of data from wearable sensors processed by artificial intelligence. These data are applied to the musculoskeletal model of the pilot, the anthropometric model, and the physiological model to produce results such as personalised training to improve the pilot’s performance, injury prevention, physiological predictions, and optimisation of the cabin ergonomics and clothing. Such micro-level models *might* be relevant as

well for modelling (specific parts of) crowds in specific contexts (e.g., when extreme accuracy is required).

The HDT notion is related but quite distinct from CDT by means and goals. The micro level of detail aimed by HDT is typically excessive for CDTs, which generally requires real-time macroscopic models. Can the CDT be designed in terms of a collection of HDTs? Yes, in principle, but it may lead to inefficient solutions. It should also be noticed that CDT models typically leverage aggregated data (e.g., from cameras) that enable an immediate estimation of macroscopic/global features of a crowd, without recurring to computational processing of micro-level data.

8.5. DTs of groups/swarms of physical entities

Limited research has been carried out on the use of DTs for modelling *whole crowds of humans*. One proposal is the *Artax* framework [9], which aims to support crowd motion prediction by storing and analysing mobility traces. However, we see this as a particular functionality that we consider part of a more comprehensive CDT framework.

Very related to the use of a DT for modelling groups of entities is the use of *multiple interconnected DTs*, aimed at exhibiting forms of (artificial) collective intelligence [84,133]. Indeed, research has proposed various concepts along the lines of combining multiple DTs (cf., groups, hierarchies, or social networks of DTs [134]) or having DTs capture pluralities of physical assets. Term *Social Digital Twin* has been proposed to refer to a “social network of DTs” [134], i.e., a construct based on the connection of different DTs around a social notion like friendship. In [135], the authors propose the notion of a *collective DT* capturing the bidirectional link with an entire group of self-organising physical devices. Various other works propose the notion of a *virtual sensor* as a digital shadow associated with an entire *group* of physical sensors [136–139]. A virtual sensor can provide the benefit of supporting reliable and “predictable” sensing over a larger ad-hoc network of low-cost “unpredictable” sensors [137], or load balancing [136]. Another use is for load balancing and data sharing [136].

DT-based solutions have also been proposed for monitoring and control of swarms of robots and UAVs [140–143]. Most of the time, however, there is not an explicit notion of a DT associated to an entire group of devices: much common is the use of the DTs of the individual devices, or an integration of data from multiple devices to build digital representations of an overall context. For instance, in [140], a federated DT framework is proposed, where local DT models are aggregated to refine the DTs of UAVs and target objects. In few works, such as [141], there is indeed a reified digital model of a swarm “as a whole”, but generally its characteristics and implementation are not detailed. The CDT framework could in principle be adapted or specialised to *cyber-physical crowds* in scenarios like swarm robotics or vehicular networks. Future work may investigate on this opportunity.

9. Conclusion

In this paper, we provide the motivation, concept, requirements, reference architecture, and a discussion of challenges related to crowd digital twins (CDTs). The goal of CDTs is to support crowd-aware applications and, most specifically, crowd management – which includes crowd monitoring/analysis and control. It aims to achieve such a goal using the *digital twin* paradigm, and specifically a solution that integrates available infrastructure and different kinds of paradigms. The proposals have been qualitatively evaluated and discussed, also w.r.t. detailed design ideas based on enabling technologies.

Our CDT proposal is not just a technical proposal but an integrative vision of research that should be carried out to enable the management of crowds while achieving the desired trade-offs of functional and non-functional requirements. We foresee that multi-scale modelling and composition will be crucial for both offline and online analysis and

simulation able to provide responses efficiently and at the desired fidelity. On a related side, collective computing techniques could support self-organisation processes as well as functional alternatives when available infrastructure is scarce. Dependability, security and privacy, and social/ethical challenges must also be considered beyond functionality. We emphasise the role of autonomic features for resiliently addressing context change and uncertainty. Future work will be dedicated to provide proofs-of-concept, quantitative evaluations of the architecture, and investigations on the mentioned challenges and promising technical solutions.

CRedit authorship contribution statement

Roberto Casadei: Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Conceptualization; **Giovanni Delnevo:** Writing – review & editing, Writing – original draft, Investigation, Conceptualization; **Roberto Girau:** Writing – review & editing, Writing – original draft, Investigation, Conceptualization; **Silvia Mirri:** Writing – review & editing, Project administration, Investigation, Funding acquisition, 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.

Data availability

No data was used for the research described in the article.

Acknowledgements

This work was supported by project ECOSISTER (ECS_00000033) under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU.

References

- [1] A. Corbetta, F. Toschi, Physics of human crowds, *Annu. Rev. Condens. Matter Phys.* 14 (1) (2023) 311–333. <https://doi.org/10.1146/annurev-conmatphys-031620-100450>
- [2] Y. Jiang, Y. Miao, B.A. Alzahrani, A. Barnawi, R. Alotaibi, L. Hu, Ultra large-scale crowd monitoring system architecture and design issues, *IEEE Internet Things J.* 8 (13) (2021) 10356–10366. <https://doi.org/10.1109/JIOT.2021.3076257>
- [3] N. Wijermans, C. Conrado, M. van Steen, C. Martella, J. Li, A landscape of crowd-management support: an integrative approach, *Saf. Sci.* 86 (2016) 142–164. <https://doi.org/10.1016/j.ssci.2016.02.027>
- [4] A.J. Perez, S. Zeadally, A communication architecture for crowd management in emergency and disruptive scenarios, *IEEE Commun. Mag.* 57 (4) (2019) 54–60. <https://doi.org/10.1109/mcom.2019.1800626>
- [5] R. Casadei, G. Aguzzi, G. Audrito, F. Damiani, D. Pianini, G. Scarso, G. Torta, M. Viroli, Software engineering for collective cyber-physical ecosystems, *ACM Trans. Softw. Eng. Methodol.* 34 (5) (2025) 153:1–153:40. <https://doi.org/10.1145/3712004>
- [6] H. Pascual, X.M. Bruin, A. Alonso, J. Cerdà, A systematic review on human modeling: digging into human digital twin implementations (2023). <https://doi.org/10.48550/ARXIV.2302.03593>
- [7] B. Wang, H. Zhou, X. Li, G. Yang, P. Zheng, C. Song, Y. Yuan, T. Wuest, H. Yang, L. Wang, Human digital twin in the context of industry 5.0, *Robotics Comput. Integr. Manuf.* 85 (2024) 102626. <https://doi.org/10.1016/j.rcim.2023.102626>
- [8] M.E. Miller, E. Spatz, A unified view of a human digital twin, *Human-Intell. Syst. Integr.* 4 (1–2) (2022) 23–33. <https://doi.org/10.1007/s42454-022-00041-x>
- [9] I. Alba, J. Troya, C. Canal, Towards a Digital Twin System for Human Crowd Motion Prediction, *Springer Nature*, 2025, p. 359–372. https://doi.org/10.1007/978-3-031-79059-1_22
- [10] L. Garcés, S. Martínez-Fernández, L. Oliveira, P.H.D. Valle, C.P. Ayala, X. Franch, E.Y. Nakagawa, Three decades of software reference architectures: a systematic mapping study, *J. Syst. Softw.* 179 (2021) 111004. <https://doi.org/10.1016/j.jss.2021.111004>
- [11] P.C. Clements, D. Garlan, R. Little, R.L. Nord, J.A. Stafford, Documenting software architectures: views and beyond, in: L.A. Clarke, L. Dillon, W.F. Tichy (Eds.), 25th Int. Conf. on Software Engineering, IEEE, 2003, pp. 740–741. <https://doi.org/10.1109/ICSE.2003.1201264>

- [12] M.T. Pham, A. Rajić, J.D. Greig, J.M. Sargeant, A. Papadopoulos, S.A. McEwen, A scoping review of scoping reviews: advancing the approach and enhancing the consistency, *Res. Synth. Methods* 5 (4) (2014) 371–385. <https://doi.org/10.1002/jrsm.1123>
- [13] R. Casadei, G. Delnevo, R. Girau, S. Mirri, Modelling groups of humans: towards crowd digital twins, in: *IEEE Symposium on Computers and Communications, IEEE*, 2024, pp. 1–6. <https://doi.org/10.1109/ISCC61673.2024.10733562>
- [14] R. Casadei, G. Delnevo, R. Girau, S. Mirri, Crowd digital twins: high-level requirements and architecture, in: *22nd IEEE Consumer Communications & Networking Conference, IEEE*, 2025, pp. 1–4. <https://doi.org/10.1109/CCNC54725.2025.10975899>
- [15] R.D. Adcock, A. Squires, R. Adcock, J.M. Wing, et al. (Eds.), *Guide to the Systems Engineering Body of Knowledge (SEBoK)*, The Trustees of the Stevens Institute of Technology, 2.9 ed., 2023. <https://www.sebokwiki.org>
- [16] D. Helbing, A. Johansson, *Pedestrian, Crowd and Evacuation Dynamics*, Springer New York, 2009, p. 6476–6495. https://doi.org/10.1007/978-0-387-30440-3_382
- [17] G.K. Still, *Introduction to Crowd Science*, CRC Press, 2014. <https://books.google.it/books?id=QgcQEQAQBAJ>
- [18] D. De, S. Ghosh, A. Mukherjee, Socialsense: mobile crowd sensing-based physical distance monitoring model leveraging federated learning for pandemic, *Internet Things* 23 (2023) 100872. <https://doi.org/10.1016/j.iot.2023.100872>. <https://doi.org/10.1016/J.IOT.2023.100872>
- [19] D. Darsena, G. Gelli, I. Iudice, F. Verde, Sensing technologies for crowd management, adaptation, and information dissemination in public transportation systems: a review, *IEEE Sens. J.* 23 (1) (2023) 68–87. <https://doi.org/10.1109/jsen.2022.3223297>
- [20] C. Feliciani, S. Tanida, X. Jia, K. Nishinari, Influencing pedestrian route choice through environmental stimuli: a long-term ecological experiment, *J. Disaster Res.* 19 (2) (2024) 325–335. <https://doi.org/10.20965/jdr.2024.p0325>
- [21] G. Sidropoulos, C. Kiourt, L. Moussiades, Crowd simulation for crisis management: the outcomes of the last decade, *Mach. Learn. Appl.* 2 (2020) 100009. <https://doi.org/10.1016/j.mlwa.2020.100009>
- [22] M. Haghighi, M. Coughlan, B. Crabb, A. Dierickx, C. Feliciani, et al., A roadmap for the future of crowd safety research and practice, *Saf. Sci.* 168 (2023) 106292. <https://doi.org/10.1016/j.ssci.2023.106292>
- [23] A. Sadilek, J. Krumm, E. Horvitz, Crowdphysics: planned and opportunistic crowdsourcing for physical tasks, in: E. Kiciman, N.B. Ellison, B. Hogan, P. Resnick, I. Soboroff (Eds.), *7th Int. Conf. on Weblogs & Social Media*, AAAI Press, 2013. <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/view/6121>. <https://doi.org/10.1609/icwsm.v7i1.14430>
- [24] A. Mladenow, C. Bauer, C. Strauss, “Crowd logistics”: the contribution of social crowds in logistics activities, *Int. J. Web Inf. Syst.* 12 (3) (2016) 379–396. <https://doi.org/10.1108/ijwis-04-2016-0020>
- [25] L. Wang, D. Zhang, Y. Wang, C. Chen, X. Han, A. M'hamed, Sparse mobile crowd-sensing: challenges and opportunities, *IEEE Commun. Mag.* 54 (7) (2016) 161–167. <https://doi.org/10.1109/MCOM.2016.7509395>
- [26] Z. Yu, H. Ma, B. Guo, Z. Yang, Crowdsensing 2.0, *Commun. ACM* 64 (11) (2021) 76–80. <https://doi.org/10.1145/3481621>
- [27] S. Mihai, M. Yaqoob, H.V. Dang, et al., Digital twins: a survey on enabling technologies, challenges, trends and future prospects, *IEEE Commun. Surv. Tutor.* 24 (4) (2022) 2255–2291. <https://doi.org/10.1109/COMST.2022.3208773>
- [28] B.S. Barn, The sociotechnical digital twin: on the gap between social and technical feasibility, in: *IEEE 24th Conf. on Business Informatics, IEEE*, 2022. <https://doi.org/10.1109/cbi54897.2022.00009>
- [29] Z. Wood, A. Galton, A taxonomy of collective phenomena, *Appl. Ontol.* 4 (3–4) (2009) 267–292. <https://doi.org/10.3233/ao-2009-0071>
- [30] W. Halboob, H. Altafari, A. Derhab, J. Almuhtadi, Crowd management intelligence framework: umrah use case, *IEEE Access* 12 (2024) 6752–6767. <https://doi.org/10.1109/access.2024.3350188>
- [31] T. Franke, P. Lukowicz, U. Blanke, Smart crowds in smart cities: real life, city scale deployments of a smartphone based participatory crowd management platform, *J. Internet Serv. Appl.* 6 (1) (2015). <https://doi.org/10.1186/s13174-015-0040-6>
- [32] C. Martella, J. Li, C. Conrado, A. Vermeeren, On current crowd management practices and the need for increased situation awareness, prediction, and intervention, *Saf. Sci.* 91 (2017) 381–393. <https://doi.org/10.1016/j.ssci.2016.09.006>
- [33] R. Casadei, G. Fortino, D. Pianini, W. Russo, C. Savaglio, M. Viroli, A development approach for collective opportunistic edge-of-things services, *Inf. Sci.* 498 (2019) 154–169. <https://doi.org/10.1016/J.IINS.2019.05.058>
- [34] T. Kravchenko, T. Bogdanova, T. Shevgunov, *Ranking Requirements Using MoSCoW Methodology in Practice*, Springer International Publishing, 2022, p. 188–199. https://doi.org/10.1007/978-3-031-09073-8_18
- [35] J.C. Laprie, *Dependability: Basic Concepts and Terminology*, Springer Vienna, 1992, p. 3–245. https://doi.org/10.1007/978-3-7091-9170-5_1
- [36] K.L.M. Ang, S.J.K. Phooi, E. Ngarhamike, Towards crowdsourcing internet of things (crowd-IoT): architectures, security and applications, *Fut. Internet* 14 (2) (2022) 49. <https://doi.org/10.3390/fi14020049>
- [37] C.A.S. Pouw, F. Toschi, F. van Schadewijk, A. Corbetta, Monitoring physical distancing for crowd management: real-time trajectory and group analysis, *PLoS ONE* 15 (10) (2020) e0240963. <https://doi.org/10.1371/journal.pone.0240963>
- [38] J. Yu, A. Alhilal, P. Hui, D.H.K. Tsang, Bi-directional digital twin and edge computing in the metaverse, *IEEE Internet Things Mag.* 7 (3) (2024) 106–112. <https://doi.org/10.1109/iotm.001.2300173>
- [39] N. Sahlab, D. Braun, C. Köhler, N. Jazdi, M. Weyrich, Extending the intelligent digital twin with a context modeling service: a decision support use case, *Procedia CIRP* 107 (2022) 463–468. <https://doi.org/10.1016/j.procir.2022.05.009>
- [40] M. Djonov, M. Galabov, Real-time data integration AWS infrastructure for digital twin, in: *Proceedings of the 21st International Conference on Computer Systems and Technologies, CompSysTech '20*, ACM, 2020, p. 223–228. <https://doi.org/10.1145/3407982.3407994>
- [41] Y. Li, S. Chen, K. Hwang, X. Ji, Z. Lei, Y. Zhu, F. Ye, M. Liu, Spatio-temporal data fusion techniques for modeling digital twin city, *Geo spatial Inf. Sci.* (2024) 1–24. <https://doi.org/10.1080/10095020.2024.2350175>
- [42] B.R. Baricelli, E. Casiraghi, D. Fogli, A survey on digital twin: definitions, characteristics, applications, and design implications, *IEEE Access* 7 (2019) 167653–167671. <https://doi.org/10.1109/access.2019.2953499>
- [43] T. Goodwin, J. Xu, N. Celik, C.-H. Chen, Real-time digital twin-based optimization with predictive simulation learning, *J. Simul.* 18 (1) (2022) 47–64. <https://doi.org/10.1080/17477778.2022.2046520>
- [44] L. Schweiger, L. Barth, Properties and characteristics of digital twins: review of industrial definitions, *SN Comput. Sci.* 4 (5) (2023). <https://doi.org/10.1007/s42979-023-01937-4>
- [45] N. Zhang, R. Bahsoon, N. Tziritis, G. Theodoropoulos, Explainable Human-in-the-Loop Dynamic Data-Driven Digital Twins, *Springer Nature*, 2024, p. 233–243. https://doi.org/10.1007/978-3-031-52670-1_23
- [46] M.S. Dihan, A.I. Akash, Z. Tasneem, P. Das, S.K. Das, M.R. Islam, M.M. Islam, F.R. Badal, M.F. Ali, M.H. Ahmed, S.H. Abhi, S.K. Sarker, M.M. Hasan, Digital twin: data exploration, architecture, implementation and future, *Heliyon* 10 (5) (2024) e26503. <https://doi.org/10.1016/j.heliyon.2024.e26503>
- [47] G.N. Schroeder, C. Steinmetz, C.E. Pereira, D.B. Espindola, Digital twin data modeling with automationML and a communication methodology for data exchange, *IFAC-PapersOnLine* 49 (30) (2016) 12–17. <https://doi.org/10.1016/j.ifacol.2016.11.115>
- [48] E. Ferko, A. Bucaioni, P. Pelliccione, M. Behnam, Standardisation in digital twin architectures in manufacturing, in: *IEEE 20th International Conference on Software Architecture, IEEE*, 2023, p. 70–81. <https://doi.org/10.1109/icsa56044.2023.00015>
- [49] D. Preuveneers, W. Joosen, E. Ilie-Zudor, Robust digital twin compositions for industry 4.0 smart manufacturing systems, in: *IEEE 22nd International Enterprise Distributed Object Computing Workshop, IEEE*, 2018, p. 69–78. <https://doi.org/10.1109/edocw.2018.00021>
- [50] E. Brucherseifer, H. Winter, A. Mentges, M. Mühlhäuser, M. Hellmann, Digital twin conceptual framework for improving critical infrastructure resilience, at - Automatisierungstechnik 69 (12) (2021) 1062–1080. <https://doi.org/10.1515/auto-2021-0104>
- [51] D. Saxena, A.K. Singh, A self-healing and fault-tolerant cloud-based digital twin processing management model, *IEEE Trans. Industr. Inform.* 21 (5) (2025) 4233–4242. <https://doi.org/10.1109/tii.2025.3540498>
- [52] K. Long, C. Ma, H. Li, Z. Li, H. Huang, H. Shi, Z. Huang, Z. Sheng, L. Shi, P. Li, S. Chen, X. Li, AI-enabled digital twin framework for safe and sustainable intelligent transportation, *Sustainability* 17 (10) (2025) 4391. <https://doi.org/10.3390/su17104391>
- [53] G. Sirigu, B. Carminati, E. Ferrari, Privacy and security issues for human digital twins, in: *IEEE 4th Int. Conf. on Trust, Privacy and Security in Intelligent Systems, and Applications, IEEE*, 2022, p. 1–9. <https://doi.org/10.1109/tps-isa56441.2022.00011>
- [54] T.Q. Duong, D. Van Huynh, S.R. Khosravirad, V. Sharma, O.A. Dobre, H. Shin, From digital twin to metaverse: the role of 6G ultra-reliable and low-latency communications with multi-Tier computing, *IEEE Wireless Commun.* 30 (3) (2023) 140–146. <https://doi.org/10.1109/mwc.014.2200371>
- [55] C.J. Vera-Arenas, ‘Human Digital Twins’ and Blockchain: Some Challenges and Solutions for Digital Identity and Privacy, *Springer Nature*, 2025, p. 473–488. https://doi.org/10.1007/978-3-031-74889-9_21
- [56] R. Miao, S. Liu, Y. Sun, M. Du, J. Bao, Streamlining digital twin development and operation with DTops, *J. Intell. Manuf.* (2025). <https://doi.org/10.1007/s10845-025-02608-2>
- [57] L. Serena, M. Marzolla, G. D’Angelo, S. Ferretti, A review of multilevel modeling and simulation for human mobility and behavior, *Simul. Model. Pract. Theory* 127 (2023) 102780. <https://doi.org/10.1016/j.simpat.2023.102780>
- [58] S. Silva, A. Tuyishime, T. Santilli, P. Pelliccione, L. Iovino, Quality metrics in software architecture, in: *20th IEEE Int. Conf. on Software Architecture, IEEE*, 2023, pp. 58–69. <https://doi.org/10.1109/ICSA56044.2023.00014>
- [59] A. Bucaioni, A.D. Salle, L. Iovino, I. Malavolta, P. Pelliccione, Reference architectures modelling and compliance checking, *Softw. Syst. Model.* 22 (3) (2023) 891–917. <https://doi.org/10.1007/S10270-022-01022-Z>
- [60] J.O. Kephart, D.M. Chess, The vision of autonomic computing, *Computer* 36 (1) (2003) 41–50. <https://doi.org/10.1109/MC.2003.1160055>
- [61] Y. Liu, J. Feng, J. Lu, S. Zhou, A review of digital twin capabilities, technologies, and applications based on the maturity model, *Adv. Eng. Inform.* 62 (2024) 102592. <https://doi.org/10.1016/j.aei.2024.102592>
- [62] Kluwer Academic Publishers, 2002. <https://doi.org/10.1007/b116085>
- [63] M.A. Khan, H. Menouar, R. Hamila, Revisiting crowd counting: state-of-the-art, trends, and future perspectives, *Image Vis. Comput.* 129 (2023) 104597. <https://doi.org/10.1016/j.imavis.2022.104597>
- [64] B. Li, H. Huang, A. Zhang, P. Liu, C. Liu, Approaches on crowd counting and density estimation: a review, *Pattern Anal. Appl.* 24 (3) (2021) 853–874. <https://doi.org/10.1007/s10044-021-00959-z>
- [65] K.P. Chong, C.H. Lai, W. Ling, Z. Lu, R. Wang, R. Wang, Y. Yu, A. Testa, C. Lam, S. Tang, G. Delnevo, R. Casadei, R. Girau, S. Mirri, Experiments of crowd detection for crowd digital twins, in: *22nd IEEE Consumer Communications & Networking Conference, IEEE*, 2025, pp. 1–4. <https://doi.org/10.1109/CCNC54725.2025.10976164>

- [66] S. Xie, X. Zhang, J. Cai, Video crowd detection and abnormal behavior model detection based on machine learning method, *Neural Comput. Appl.* 31 (S1) (2018) 175–184. <https://doi.org/10.1007/s00521-018-3692-x>
- [67] A. Bahamid, A. Mohd Ibrahim, A review on crowd analysis of evacuation and abnormality detection based on machine learning systems, *Neural Comput. Appl.* 34 (24) (2022) 21641–21655. <https://doi.org/10.1007/s00521-022-07758-5>
- [68] A. Corbetta, W. Kroneman, M. Donners, A. Haans, et al., A large-scale real-life crowd steering experiment via arrow-like stimuli, *Collect. Dyn.* 5 (2020). <https://doi.org/10.17815/cd.2020.34>
- [69] A. Koiliias, C. Mousas, C. Anagnostopoulos, I feel a moving crowd surrounds me: exploring tactile feedback during immersive walking in a virtual crowd, *Comput. Animat. Virtual Worlds* 31 (4–5) (2020). <https://doi.org/10.1002/cav.1963>
- [70] B. Ryan, M. Hallewell, N. Hughes, N. Coad, J. Eliseu, et al., Field testing two innovative lighting interventions to influence waiting behaviours and movements on stairways in train stations, *Light. Res. Technol.* 54 (3) (2021) 237–259. <https://doi.org/10.1177/14771535211016811>
- [71] T. Senan, A. Corbetta, B. Hengeveld, Towards sound-based crowd management: investigating sonification for pedestrian steering, in: *AudioMostly 2022*, AM '22, ACM, 2022. <https://doi.org/10.1145/3561212.3561233>
- [72] A. Siddiqi, W.Z. Khan, M.H. Alkinani, E.A. Aldhahri, M.K. Khan, Edge-assisted federated learning framework for smart crowd management, *Internet Things* 27 (2024) 101253. <https://doi.org/10.1016/j.iot.2024.101253>
- [73] X. Li, Q. Yu, B. Alzahrani, A. Barnawi, A. Alhindi, D. Alghazzawi, Y. Miao, Data fusion for intelligent crowd monitoring and management systems: a survey, *IEEE Access* 9 (2021) 47069–47083. <https://doi.org/10.1109/access.2021.3060631>
- [74] M.L. Itria, M. Kocsis-Magyar, A. Ceccarelli, P. Lollini, G. Giunta, A. Bon-davalli, Identification of critical situations via event processing and event trust analysis, *Knowl. Inf. Syst.* 52 (1) (2016) 147–178. <https://doi.org/10.1007/s10115-016-1009-x>
- [75] R.B. Gramacy, *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*, CRC Press, 2020. <https://doi.org/10.1201/9780367815493>
- [76] C. Gomes, C. Thule, D. Broman, P.G. Larsen, H. Vangheluwe, Co-simulation: a survey, *ACM Comput. Surv.* 51 (3) (2018) 1–33. <https://doi.org/10.1145/3179993>
- [77] S. Mustapha, A. Kassir, K. Hassoun, B.A.A. Modad, H. Abi-Rached, Z. Dawy, Joint crowd management and structural health monitoring using fiber optic and wearable sensing, *IEEE Commun. Mag.* 57 (4) (2019) 62–67. <https://doi.org/10.1109/MCOM.2019.1800631>
- [78] Y. Dong, Y. Wu, Y.-C. Chang, J. Aggarwal, J.R. Codling, H. Latapie, P. Zhang, H.Y. Noh, Context-aware crowd monitoring for sports games using crowd-induced floor vibrations, *Data-Centric Eng.* 5 (2024). <https://doi.org/10.1017/dce.2024.28>
- [79] J. Whittle, P. Sawyer, N. Bencomo, B.H.C. Cheng, J.-M. Bruel, RELAX: a language to address uncertainty in self-adaptive systems requirement, *Requir. Eng.* 15 (2) (2010) 177–196. <https://doi.org/10.1007/s00766-010-0101-0>
- [80] N.D. Lane, S.B. Eisenman, M. Musolesi, E. Miluzzo, A.T. Campbell, Urban sensing systems: opportunistic or participatory?, in: *Proceedings of the 9th Workshop on Mobile Computing Systems and Applications*, 2008, pp. 11–16.
- [81] R. Girau, M. Anedda, R. Presta, S. Corpino, P. Ruiu, M. Fadda, C.-T. Lam, D. Giusto, Definition and implementation of the cloud infrastructure for the integration of the human digital twin in the social internet of things, *Comput. Netw.* 251 (2024) 110632.
- [82] G.D. Abowd, Beyond weiser: from ubiquitous to collective computing, *Computer* 49 (1) (2016) 17–23. <https://doi.org/10.1109/MC.2016.22>
- [83] G. Audrito, R. Casadei, G. Torta, A general framework and decentralised algorithms for collective computational processes, *Fut. Gener. Comput. Syst.* 158 (2024) 11–27. <https://doi.org/10.1016/j.future.2024.04.020>
- [84] R. Casadei, Artificial collective intelligence engineering: a survey of concepts and perspectives, *Artif. Life* 29 (4) (2023) 433–467. <https://doi.org/10.1162/ARTLA.00408>
- [85] H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsy, T.F. Knight, Jr., R. Nagpal, E. Rauch, G.J. Sussman, R. Weiss, Amorphous computing, *Commun. ACM* 43 (5) (2000) 74–82. <https://doi.org/10.1145/332833.332842>
- [86] D. Weyns, *An Introduction to Self-Adaptive Systems: A Contemporary Software Engineering Perspective*, John Wiley & Sons, 2020.
- [87] D. Weyns, B.R. Schmerl, V. Grassi, S. Malek, et al., On patterns for decentralized control in self-adaptive systems, in: R. de Lemos, H. Giese, H.A. Müller, M. Shaw (Eds.), *International Seminar, Revised Selected and Invited Papers*, 7475 of LNCS, Springer, 2010, pp. 76–107. https://doi.org/10.1007/978-3-642-35813-5_4
- [88] K.L. Bellman, J. Botev, A. Diaconescu, L. Esterle, C. Gruhl, C. Landauer, P.R. Lewis, P.R. Nelson, E. Pournaras, A. Stein, S. Tomforde, Self-improving system integration: mastering continuous change, *Future Gener. Comput. Syst.* 117 (2021) 29–46. <https://doi.org/10.1016/j.future.2020.11.019>
- [89] R. Al-Ali, L. Bulej, J. Kofroň, T. Bureš, A guide to design uncertainty-aware self-adaptive components in cyber-physical systems, *Future Gener. Comput. Syst.* 128 (2022) 466–489. <https://doi.org/10.1016/j.future.2021.10.027>
- [90] S. Mahdavi-Hezavehi, P. Avgeriou, D. Weyns, A Classification Framework of Uncertainty in Architecture-Based Self-Adaptive Systems With Multiple Quality Requirements, Elsevier, 2017, p. 45–77. <https://doi.org/10.1016/b978-0-12-802855-1.00003-4>
- [91] C. Wohlin, P. Runeson, M. Höst, M.C. Ohlsson, B. Regnell, A. Wesslén, *Experimentation in Software Engineering*, Springer, 2012. <https://doi.org/10.1007/978-3-642-29044-2>
- [92] D. Lehner, J. Zhang, J. Pfeiffer, S. Sint, A.-K. Splettstößer, M. Wimmer, A. Wortmann, Model-driven engineering for digital twins: a systematic mapping study, *Soft. Syst. Model.* (2025). <https://doi.org/10.1007/s10270-025-01264-7>
- [93] L. Serena, M. Marzolla, G. D'Angelo, S. Ferretti, Design patterns for multilevel modeling and simulation, in: 2023 IEEE/ACM 27th International Symposium on Distributed Simulation and Real Time Applications (DS-RT), IEEE, 2023, p. 48–55. <https://doi.org/10.1109/ds-rt58998.2023.00015>
- [94] G.V. Bobashev, D.M. Goedecke, F. Yu, J.M. Epstein, A hybrid epidemic model: combining the advantages of agent-based and equation-based approaches, in: 2007 Winter Simulation Conference, IEEE, 2007, p. 1532–1537. <https://doi.org/10.1109/wsc.2007.4419767>
- [95] M. Xiong, M. Lees, W. Cai, S. Zhou, M.Y.H. Low, Hybrid modelling of crowd simulation, *Procedia Comput. Sci.* 1 (1) (2010) 57–65. <https://doi.org/10.1016/j.procs.2010.04.008>
- [96] L. Erazo-Garzon, A. Roman, J. Moyano-Dutan, P. Cedillo, Models@runtime and internet of things: a systematic literature review, in: 2021 Second International Conference on Information Systems and Software Technologies (ICI2ST), IEEE, 2021, p. 128–134. <https://doi.org/10.1109/ici2st51859.2021.00026>
- [97] N. Bencomo, S. Götz, H. Song, Models@run.time: a guided tour of the state of the art and research challenges, *Softw. Syst. Model.* 18 (5) (2019) 3049–3082. <https://doi.org/10.1007/s10270-018-00712-x>
- [98] T. Reiter, K. Altmanninger, W. Retschitzegger, Think Global, Act Local: Implementing Model Management with Domain-Specific Integration Languages, Springer Berlin Heidelberg, 2006, p. 263–276. https://doi.org/10.1007/978-3-540-69489-2_32
- [99] J. Li, L. Wang, S. Tang, B. Zhang, Y. Zhang, Risk-based crowd massing early warning approach for public places: a case study in China, *Saf. Sci.* 89 (2016) 114–128. <https://doi.org/10.1016/j.ssci.2016.06.007>
- [100] M.V. Ramesh, A. Shanmughan, R. Prabha, Context aware ad hoc network for mitigation of crowd disasters, *Ad Hoc Netw.* 18 (2014) 55–70. <https://doi.org/10.1016/j.adhoc.2013.02.006>
- [101] G. Audrito, F. Damiani, G. Torta, Towards Real-Time Aggregate Computing, Springer Nature Switzerland, 2024, p. 49–68. https://doi.org/10.1007/978-3-031-75107-3_4
- [102] J.H. Saleh, E.A. Saltmarsh, F.M. Favarò, L. Brevault, Accident precursors, near misses, and warning signs: critical review and formal definitions within the framework of discrete event systems, *Reliab. Eng. Syst. Saf.* 114 (2013) 148–154. <https://doi.org/10.1016/j.ress.2013.01.006>
- [103] L.-M. Messmer, C. Reich, D.O. Abdeslam, Context-Aware Machine Learning: A Survey, Springer Nature Switzerland, 2024, p. 252–272. https://doi.org/10.1007/978-3-031-73110-5_17
- [104] R. Wang, Y. Hao, L. Hu, X. Li, M. Chen, Y. Miao, I. Humar, Efficient crowd counting via dual knowledge distillation, *IEEE Trans. Image Process.* 33 (2024) 569–583. <https://doi.org/10.1109/tip.2023.3343609>
- [105] C. Chen, Y. Liu, S. Kreiss, A. Alahi, Crowd-robot interaction: crowd-aware robot navigation with attention-based deep reinforcement learning, in: 2019 Int. Conf. on Robotics and Automation, IEEE, 2019, p. 6015–6022. <https://doi.org/10.1109/icra.2019.8794134>
- [106] D. Wolinski, M.C. Lin, J. Pettré, Warpdriver: context-aware probabilistic motion prediction for crowd simulation, *ACM Trans. Graph.* 35 (6) (2016). <https://doi.org/10.1145/2980179.2982442>
- [107] P. Trautman, K. Patel, Real time crowd navigation from first principles of probability theory, *Proc. Int. Conf. Autom. Plann. Schedul.* 30 (2020) 459–467. <https://doi.org/10.1609/icaps.v30i1.6741>
- [108] S. Huang, R. Golman, S.B. Broomell, Combining the aggregated forecasts: an efficient method for improving accuracy by stacking multiple weighting models, *Decision* 11 (4) (2024) 668–683. <https://doi.org/10.1037/dec0000245>
- [109] I.G. Georgoudas, G.C. Sirakoulis, I.T. Andreadis, An anticipative crowd management system preventing clogging in exits during pedestrian evacuation processes, *IEEE Syst. J.* 5 (1) (2011) 129–141. <https://doi.org/10.1109/jsyst.2010.2090400>
- [110] C. Deng, Y. Li, Q. Liu, X. Zheng, K. Sun, Quantitative risk assessment for autonomous vehicles: integrating functional resonance analysis method and bayesian network, *Qual. Reliab. Eng. Int.* 41 (3) (2024) 970–991. <https://doi.org/10.1002/qre.3697>
- [111] D. Prabhu, M. Dhivya, A. Gowridurga, M. Dinesh Babu, T.P. Anish, R. Siva Subramanian, Visual analytics for smart decision-making: techniques, applications, and future directions, in: 2025 International Conference on Visual Analytics and Data Visualization (ICVADV), IEEE, 2025, p. 37–43. <https://doi.org/10.1109/icvadv63329.2025.10961125>
- [112] E.F. Langås, M.H. Zafar, F. Sanfilippo, Exploring the synergy of human-robot teaming, digital twins, and machine learning in industry 5.0: a step towards sustainable manufacturing, *J. Intell. Manuf.* (2025). <https://doi.org/10.1007/s10845-025-02580-x>
- [113] M. Krupas, E. Kajati, C. Liu, I. Zolotova, Towards a human-centric digital twin for human-machine collaboration: a review on enabling technologies and methods, *Sensors* 24 (7) (2024) 2232. <https://doi.org/10.3390/s24072232>
- [114] B. Han, H.D. Schotten, Multi-sensory HMI for human-centric industrial digital twins: a 6G vision of future industry, in: 2022 IEEE Symposium on Computers and Communications (ISCC), IEEE, 2022, p. 1–7. <https://doi.org/10.1109/iscc55528.2022.9912932>
- [115] Y. Lin, L. Chen, A. Ali, C. Nugent, I. Cleland, R. Li, J. Ding, H. Ning, Human digital twin: a survey, *J. Cloud Comput.* 13 (1) (2024). <https://doi.org/10.1186/s13677-024-00691-z>
- [116] A. Correia, A. Grover, D. Schneider, A.P. Pimentel, R. Chaves, M.A. de Almeida, B. Fonseca, Designing for hybrid intelligence: a taxonomy and survey of crowd-machine interaction, *Appl. Sci.* 13 (4) (2023) 2198. <https://doi.org/10.3390/app13042198>

- [117] M.K. Traore, S. Gorecki, Y. Ducq, A Formal Framework for Digital Twin Modeling, Verification, and Validation, Springer Nature, 2024, p. 119–143. https://doi.org/10.1007/978-3-031-69107-2_6
- [118] D. Trentesaux, S. Karnouskos, Engineering ethical behaviors in autonomous industrial cyber-physical human systems, *Cognit. Technol. Work* 24 (1) (2021) 113–126. <https://doi.org/10.1007/s10111-020-00657-6>
- [119] S.A. Mostafa, M.S. Ahmad, A. Mustapha, Adjustable autonomy: a systematic literature review, *Artif. Intell. Rev.* 51 (2) (2019) 149–186. <https://doi.org/10.1007/s10462-017-9560-8>
- [120] E. Veltmeijer, C. Gerritsen, Legal and ethical implications of AI-based crowd analysis: the AI Act and beyond, *AI Ethics* (2025). <https://doi.org/10.1007/s43681-024-00644-x>
- [121] F. Mureddu, A. Paciaroni, T. Pavelka, A. Pemberton, L.A. Remotti, Rights and Responsibilities: Legal and Ethical Considerations in Adopting Local Digital Twin Technology, Springer Nature, 2025, p. 291–317. https://doi.org/10.1007/978-3-031-81451-8_11
- [122] K. Martin, Ethical implications and accountability of algorithms, *J. Bus. Ethics* 160 (4) (2018) 835–850. <https://doi.org/10.1007/s10551-018-3921-3>
- [123] C. Cancela-Outeda, The EU's AI act: a framework for collaborative governance, *Internet Things* 27 (2024) 101291. <https://doi.org/10.1016/j.iot.2024.101291>
- [124] B. Townsend, C. Paterson, T.T. Arvind, G. Nemirovsky, R. Calinescu, A. Cavalcanti, I. Habli, A. Thomas, From pluralistic normative principles to autonomous-agent rules, *Minds Mach.* 32 (4) (2022) 683–715. <https://doi.org/10.1007/s11023-022-09614-w>
- [125] G. Boella, L. van der Torre, H. Verhagen, Introduction to the special issue on normative multiagent systems, *Auton. Agent Multi Agent Syst.* 17 (1) (2008) 1–10. <https://doi.org/10.1007/s10458-008-9047-8>
- [126] K. Josifovska, E. Yigitbas, G. Engels, Reference framework for digital twins within cyber-physical systems, in: *IEEE/ACM 5th Int. Workshop on Software Engineering for Smart Cyber-Physical Systems (SEsCPS)*, IEEE, 2019. <https://doi.org/10.1109/secps.2019.00012>
- [127] D. Sharma, A.P. Bhondekar, A.K. Shukla, C. Ghanshyam, A review on technological advancements in crowd management, *J. Ambient Intell. Humaniz. Comput.* 9 (3) (2016) 485–495. <https://doi.org/10.1007/s12652-016-0432-x>
- [128] H.-T. Dang, B. Gaudou, N. Verstaevl, A literature review of dense crowd simulation, *Simul. Model. Pract. Theory* 134 (2024) 102955. <https://doi.org/10.1016/j.simpat.2024.102955>
- [129] S.P. Hoogendoorn, P.H.L. Bovy, Pedestrian route-choice and activity scheduling theory and models, *Transp. Res. Part B Methodol.* 38 (2) (2004) 169–190. [https://doi.org/10.1016/s0191-2615\(03\)00007-9](https://doi.org/10.1016/s0191-2615(03)00007-9)
- [130] S.D. Okegbile, J. Cai, C. Yi, D. Niyato, Human digital twin for personalized healthcare: vision, architecture and future directions, *IEEE Netw.* (2022) 1–7. <https://doi.org/10.1109/MNET.118.2200071>
- [131] B. Schleich, N. Anwer, L. Mathieu, S. Wartzack, Shaping the digital twin for design and production engineering, *CIRP Ann.* 66 (1) (2017) 141–144. <https://doi.org/10.1016/j.cirp.2017.04.040>
- [132] Z. Cheng, Human digital twin with applications, 7th International Digital Human Modeling Symposium (2022). <https://doi.org/10.17077/dhm.31783>
- [133] E. Pretel, E. Navarro, V.C. Pujol, S. Dustdar, Digital twins and artificial collective intelligence: synergies for the future, *IEEE Internet Comput.* 29 (1) (2025) 75–85. <https://doi.org/10.1109/mic.2024.3521607>
- [134] A. Lombardo, G. Morabito, S. Quattropiani, C. Ricci, Sociality-as-a-service: a new platform for networked digital twins, in: *61st Int. Congress Future Telecommunications: Infrastructure & Sustainability*, IEEE, 2022. <https://doi.org/10.23919/fitce56290.2022.9934689>
- [135] R. Casadei, D. Pianini, M. Viroli, D. Weyns, Digital twins, virtual devices, and augmentations for self-organising cyber-physical collectives, *Appl. Sci.* 12 (1) (2022). <https://doi.org/10.3390/app12010349>
- [136] M.E. Khansari, S. Sharifian, S.A. Motamedi, Virtual sensor as a service: a new multi-criteria QoS-aware cloud service composition for IoT applications, *J. Supercomput.* 74 (10) (2018) 5485–5512. <https://doi.org/10.1007/s11227-018-2454-y>
- [137] M. Brown, S. Gilbert, N.A. Lynch, C.C. Newport, T. Nolte, M. Spindel, The virtual node layer: a programming abstraction for wireless sensor networks, *SIGBED Rev.* 4 (3) (2007) 7–12. <https://doi.org/10.1145/1317103.1317105>
- [138] R. Bose, A. Helal, Distributed mechanisms for enabling virtual sensors in service oriented intelligent environments, in: *4th Int. Conf. on Intelligent Environments*, IEEE, 2008, pp. 2B5. <https://doi.org/10.1049/cp:20081121>
- [139] S. Chatterjee, S. Misra, Optimal composition of a virtual sensor for efficient virtualization within sensor-cloud, in: *Int. Conf. on Communications*, IEEE, 2015, pp. 448–453. <https://doi.org/10.1109/ICC.2015.7248362>
- [140] L. Zhou, S. Leng, Q. Wang, A federated digital twin framework for UAVs-based mobile scenarios, *IEEE Trans. Mob. Comput.* (2024) 1–16. <https://doi.org/10.1109/tmc.2023.3335386>
- [141] L. Lei, G. Shen, L. Zhang, Z. Li, Toward intelligent cooperation of UAV swarms: when machine learning meets digital twin, *IEEE Netw.* 35 (1) (2021) 386–392. <https://doi.org/10.1109/MNET.011.2000388>
- [142] Z. Wang, K. Han, P. Tiwari, Digital twin-assisted cooperative driving at non-signalized intersections, *IEEE Trans. Intell. Veh.* 7 (2) (2022) 198–209. <https://doi.org/10.1109/TIV.2021.3100465>
- [143] L.R. Salinas, G. Tzoumas, L. Pitonakova, S. Hauert, Digital twin technology for wild-fire monitoring using UAV swarms, in: *Int. Conf. on Unmanned Aircraft Systems*, IEEE, 2023. <https://doi.org/10.1109/icuas57906.2023.10155819>