



A MUNICIPAL DIGITAL TWIN SIMULATOR FOR DRONE-ASSISTED EMERGENCY ACCESS IN RURAL TERRITORIES

JUAN VELADO-ARENAS¹, JORGE CISNEROS-GONZÁLEZ¹, JOSÉ ANTONIO ONDIVIELA GARCÍA², JAVIER SÁNCHEZ-SORIANO¹

¹Advanced Artificial Intelligence Group (A2IG), Escuela Politécnica Superior, Universidad Francisco de Vitoria, Pozuelo de Alarcón, 28223 Madrid, Spain

²Citizen-Centric Intelligent Cities Research Institute, Universidad Francisco de Vitoria, Madrid, Spain

KEYWORDS	ABSTRACT
Municipal Digital Twin Drone-assisted Emergency Services Rural Resilience Virtual Testing Urban Air Mobility Explainable Decision Support Unmanned aerial vehicles	Public services planned around averages can conceal structurally underserved residents in dispersed rural territories. This study develops a municipal digital twin simulator for testing drone-assisted emergency access in a Spanish municipality of nineteen settlements covering 406.71 km ² . The study first characterizes existing road-based emergency access and identifies four peripheral settlements, comprising 46 residents, with disproportionately long response times. The resulting equity indicators make this territorial discontinuity visible and provide a criterion for evaluating intervention scenarios. The simulator integrates the territorial model, emergency resources, deterministic allocation logic, three-dimensional visualization, replayable virtual scenarios and an inspectable decision account. Its prototype was validated in collaboration with the municipal authority through scenario-based sessions using the real territorial and operational context. Municipal feedback was incorporated into successive iterations, and the resulting system was assessed for functional behavior, interpretability and practical usefulness to non-technical decision makers. The prototype therefore represents a pilot-validated simulation environment for municipal planning and is consistent with the TRL5 criterion of a validated prototype in a real-world context. The study shows how digital-twin simulation can support equitable, transparent and citizen-centered emergency planning before physical deployment of drone services.

1. Introduction

Public services are planned around averages, and averages are calibrated to density. Where population is concentrated, the mean experience of a service and the experience of most of its users converge, so a planner who improves the mean improves most lives. Where population is dispersed across a large territory in very small settlements, that convergence fails. The mean describes nobody, and the residents whose service is worst contribute least to the figure by which the service is judged. This paper concerns a municipality of the second kind: some three thousand inhabitants distributed across four hundred square kilometers in nineteen separate settlements, at a density of approximately seven inhabitants per square kilometer.

Emergency medical response is the sharpest instance of the problem, for three reasons. The cost of delay is an outcome rather than an inconvenience. The threshold of adequacy is clinical and externally given, so whether a service arrives on time is not a matter of administrative preference. And provision is indivisible: a settlement of eight residents cannot be assigned a

fraction of an ambulance, and no efficiency criterion will ever justify assigning it a whole one. The three combine to produce a situation in which a small number of people are structurally underserved, the deficit is measurable against an external standard, and the conventional instruments of service planning are guaranteed not to detect it.

Into that situation arrives a technological offer. Municipalities are increasingly invited to adopt digital twins and artificial intelligence as instruments of urban efficiency, and the invitation carries an implicit addressee: an authority with sensor networks, technical staff and a budget for continuous data acquisition. A municipality of three thousand inhabitants receives the vocabulary without the preconditions and is left either to decline the offer or to pursue an artefact it could not operate. We argue that the productive question is not whether such a municipality should have a digital twin, but which level of digital representation can operate, and whether that level is sufficient for the decision in front of it.

This paper makes three contributions. The first is a model of road-based emergency access that applies an emergency-driving speed correction differentiated by road class, in place of the single multiplier commonly used, and shows why the undifferentiated version systematically flatters the incumbent service in territories that combine motorway with local road. The second is a set of equity indicators for emergency access in dispersed municipalities, constructed to identify the worst-served residents rather than to summarize the typical one, and including a measure of the access gap that is independent of where the coverage threshold is placed. The third is a municipal digital twin simulator that integrates territorial representation, emergency-resource modelling, deterministic decision support, replayable virtual scenarios and an inspectable decision account, validated with the municipal authority as a pilot-level decision-support prototype for scenario testing.

In dispersed rural territories, equitable emergency access is also a resilience problem: service continuity depends on a small number of centralized resources, while peripheral settlements remain exposed to disproportionate delays when road-based response is stretched by distance, geography or competing demand.

Importantly, the prototype was not developed as a purely technical demonstration: it was validated with the municipal authority in the real territorial and institutional context for which the system was designed, and municipal feedback was incorporated into successive iterations.

Running through all three is a single claim. In the territory studied here, the residents with the worst access to emergency care number forty-six, slightly over one per cent of the municipality. No criterion that aggregates over the population will ever reach them, because the aggregate is already satisfactory and the cheapest way to improve it is to serve those who are served well already. If they are to be reached at all, the criterion must be equity of access rather than efficiency of provision and if that is the criterion, then the instrument a municipality needs is one that can show it who is left outside, and let it test what would change if it acted.

The paper proceeds as follows. Section 2 situates the work in literature on urban digital twins, artificial intelligence in public services, and unmanned aerial vehicles in emergency response. Section 3 states the objectives and the boundaries of what is reported here. Section 4 describes the territory and its existing provision. Sections 5 and 6 develop the access model and the equity indicators respectively, and Section 7 applies them. Section 8 presents the instrument, Section 9 the framework that situates it, and Section 10 the regulatory and institutional conditions for deployment together with what would be transferred to another municipality. The paper closes with conclusions.

2. Background

Digital twins entered urban practice from manufacturing and aerospace, where the term denoted a synchronized virtual replica of a physical asset, continuously updated by sensor data and capable of informing decisions about the object it mirrors. Its migration to the city has been rapid and terminologically loose. In their synthesis of digital twins in smart city development, Peldon et al. (2024) observes that the label is now attached to artefacts ranging from static

three-dimensional models to fully instrumented urban platforms (a breadth that complicates both comparison between projects and assessment of what any given deployment delivers).

That ambiguity matters for municipalities, because it obscures cost. Haraguchi et al. (2024) confront it directly, proposing a maturity model for city digital twin technology and reading the governance implications of each stage. Their contribution is to treat the digital twin as a spectrum rather than a binary property: each successive level demands additional sensing infrastructure, data-management capacity and institutional commitment, and each confers correspondingly greater operational capability. On this view, a municipality asking whether it "has" a digital twin is asking the wrong question. The useful question is which level it can build, fund and operate, and whether that level is sufficient for the decision to hand.

Applied work has nonetheless concentrated at the upper end of that spectrum and in the largest cities, where the enabling conditions are already present: dense sensor networks, dedicated technical staff, and budgets that absorb continuous data acquisition. The result is a literature whose implicit reader is a metropolitan authority. Very little of it speaks to a municipality of a few thousand inhabitants distributed across hundreds of square kilometers, which is precisely the administrative unit facing the sharpest service-provision problems in depopulating regions (Llorent-Bedmar et al., 2021). Such a municipality does not lack a use case for simulation; it lacks any account of which rung of the maturity ladder is realistic for it, and what that rung would buy.

A parallel strand examines artificial intelligence in the delivery of public services. The review by de Sousa et al. (2019) maps the field and its research agenda, while Gesk and Leyer (2022) investigate the conditions under which citizens accept the use of such systems in public services. This literature is directly relevant here for two reasons. First, emergency response is a high-stakes domain in which an algorithmic recommendation carries consequences that citizens are entitled to interrogate. Second, the municipal officers who would operate such a system are not data scientists, so the intelligibility of a recommendation is not presentational refinement but a condition of the system being used at all. Both considerations push towards decision support that can explain itself, rather than towards automation that merely performs well.

More broadly, distributed logistics problems increasingly rely on optimization-based approaches to balance operational efficiency with multiple service constraints. Recent work has shown how algorithmic optimization can be used to jointly consider routing costs and service preferences, illustrating the value of computational decision support in complex logistics environments (Sánchez-Soriano et al., 2025).

A third relevant strand concerns the use of unmanned aerial vehicles (UAVs) for emergency medical response. Boutilier et al. (2017) established the template by optimizing the placement of a drone network for the delivery of automated external defibrillators, subsequently extended in Boutilier and Chan (2022). Field evidence has followed: Schierbeck et al. (2023) compared drone-delivered defibrillators with ambulance arrival in real-life suspected out-of-hospital cardiac arrests in Sweden, and Jakobsen et al. (2025) reported a Danish feasibility study using a semi-autonomous platform. Related architectural work outside the medical domain has addressed the remote management of UAV fleets for infrastructure maintenance and surveillance, coupling airborne and ground-based sensing within a modular system design (Bemposta Rosende et al., 2020).

Two observations follow from this body of work. The first is that its center of gravity lies in network design (where to place aircraft so that a coverage target is met) and in the clinical validation of the resulting response times. The instrument by which a local authority might itself explore, test and defend such a configuration receives comparatively little attention, even though that instrument, and not the optimization model, is what a council actually encounters. The second is that the field evidence has been generated in comparatively well-served European regions, where an ambulance service already achieves short response times and the drone competes at the margin. The dispersed rural periphery presents a different problem, in which the incumbent service is structurally unable to reach part of the population within any clinically meaningful window, and in which the affected population is small enough to be invisible to

aggregate efficiency criteria: A tension that the accessibility literature has long framed as one of equity rather than of efficiency (Neutens, 2015).

The gap this paper addresses sits at the intersection of these strands. The digital twin literature offers maturity as an organizing concept but assumes metropolitan capacity; the UAV literature offers siting methods but assumes the decision has already been delegated to an optimization model; the public-sector AI literature establishes that intelligibility conditions acceptance but does not treat territorial emergency planning. None of them answers the question a small rural council actually asks: what level of digital twin we can operate, what can we legitimately decide with it, and on what criterion should that decision rest when those who stand to benefit are a numerically negligible minority of our residents. Sections 5 and 6 characterize the access problem and propose the criterion; Sections 8 and 9 present the instrument and situate it on the maturity spectrum.

3. Objectives and scope

The paper pursues three objectives: to characterize the deficit in emergency medical access across a dispersed rural municipality using an explicit model of the existing road-based service; to propose indicators that make that deficit visible where conventional summary statistics conceal it; and to establish what level of digital representation a municipality of this size can operate, and what decisions that level supports.


The study focuses on the simulator as a municipal decision-support instrument rather than on the detailed formulation of the underlying optimization procedures. The paper therefore does not reproduce the complete mathematical formulation of the UAV network-design and allocation algorithms. Instead, it demonstrates how these computational components are integrated into the simulator and how the resulting configurations and operational situations can be examined through replayable scenarios. The emphasis is consequently on municipal virtual testing, territorial equity, interpretability and practical validation rather than on algorithmic novelty in network optimization.

A note on terminology is also owed, because the vocabulary of this field has become elastic. The decision layer uses deterministic combinatorial optimization rather than machine learning or a language model. We therefore do not claim that the simulator is a generative-AI or LLM-based system. Instead, computational intelligence is used in the form of algorithmic decision support: the system evaluates feasible alternatives, applies explicit operational criteria and produces reproducible recommendations. This positioning is consistent with the broader scope of AI for cities considered in this special issue, where AI-driven and algorithmic approaches are considered alongside digital-twin simulation and other decision-support technologies.

4. Study area and the access problem

San Esteban de Gormaz is a municipality in the province of Soria, in the autonomous community of Castilla y León, covering 406.71 km² with 2,962 inhabitants recorded by the national statistics institute at 1 January 2025. Its population density of approximately 7.3 inhabitants per square kilometer places it well inside the range conventionally used to characterize severely depopulated territory in Spain, and among the lowest in western Europe. The municipality extends some 25 kilometers from north to south and fifteen from east to west, across terrain ranging from 839 to 1,155 meters above sea level.

Figure 1. Study area: terrain and settlement pattern.



Source(s): own elaboration from IGN MDT05 elevation data.

What matters for service provision is not the area but the way population is arranged within it. The municipality comprises nineteen inhabited nuclei: a core settlement, also called San Esteban de Gormaz, and eighteen dependent hamlets scattered across the territory at straight-line distances from the core ranging from just over four kilometers to more than eighteen. The distribution of population between them is extremely unequal. On municipal records the core holds roughly 87% of inhabitants and the eighteen hamlets together some 13%, with individual settlements ranging from eight residents to one hundred and twenty-five. These settlement-level figures, taken from municipal administrative records, sum to more than the official census total for the municipality and should be read as an upper reference rather than an exact count; they are used here because no census breakdown at settlement level is available, and because the analysis depends on their relative distribution rather than their absolute values.

Table 1. Dependent settlements: population and straight-line distance from the municipal core.

Settlement	Population	Straight-line distance (km)
Torraño	15	18.26
Quintanas Rubias de Arriba	8	15.51
Torremocha de Ayllón	15	14.30
Quintanas Rubias de Abajo	8	13.92
Morcuera	12	11.75
Villálvaro	125	10.12
Inés	30	8.90
Velilla de San Esteban	20	8.81
Piquera de San Esteban	10	8.10
Rejas de San Esteban	20	7.44
Matanza de Soria	15	7.29
Soto de San Esteban	20	6.39
Olmillos de San Esteban	40	5.85
Peñalba de San Esteban	10	5.80
Pedraja de San Esteban	25	5.41
Aldea de San Esteban	15	4.93
Atauta	50	4.31
Quintanilla de Tres Barrios	15	4.23

Source(s): municipal administrative records; distances computed from settlement coordinates.

Emergency medical provision follows the same concentration. The municipality is served by a basic life support ambulance stationed permanently in the core and available around the clock, a provision confirmed against the regional health service's published record. For residents of the core, the question addressed in this paper does not arise: the vehicle is already where they are, and no alternative mode of delivery would improve on that. The eighteen-dependent settlements are in a different position. Every emergency in any of them is answered by that same vehicle travelling outward by road, and their access is therefore a function of the road network and of their position on it. The analysis that follows is restricted to those eighteen settlements for this reason, and the core is retained throughout as the reference case in which existing provision is already adequate (a distinction that strengthens rather than weakens the argument, since it concedes at the outset the territory in which no case can be made).

Straight-line distance is a poor guide to what that outward journey costs. It identifies the broad shape of the problem, in that the settlements furthest from the core are also those the ambulance reaches last, but it misorders the middle of the distribution and yields no times at all. Two settlements at comparable straight-line distance can differ appreciably in road response time depending on whether they are reached by dual carriageway or by local road, and the ordering of intermediate settlements by distance and by time does not coincide. Establishing the access problem therefore requires modelling the road network explicitly, which Section 5 does.

Two features of the wider provincial context bear in the argument without forming part of the measurement. The first is that advanced life support is thinly distributed: a single physician-staffed emergency unit serves the whole province from the capital, some seventy kilometers away, and the regional air ambulance operates from a base with a mean response time of the order of thirty-eight minutes. Where the local vehicle is committed, or where the clinical situation exceeds what basic life support can offer, the next tier is distant. The second is that demand is not hypothetical. Provincial civil protection records report in the order of four thousand two hundred medical incidents across Soria in the most recent annual reported period, divided between domestic and public settings. This figure describes the province rather than the municipality and includes many cases requiring patient transport rather than the delivery of material; it is offered as an order of magnitude for the category of demand under discussion, not as a forecast of missions.

5. Modelling road-based emergency access

5.1 Road network and routable graph

The road network was obtained as an OpenStreetMap extract clipped to the comarca (bounding box $-3.50, 41.29$ to $-2.97, 41.80$; approximately $2,511 \text{ km}^2$), retrieved through the BBBike extraction service and used under the Open Database License. The raw extract contains 38,802 nodes and 40,565 way segments. From it a routable graph was built with `pyrosm` and `networkx`, retaining drivable road classes and expanding bidirectional ways into directed edges, which yields 37,789 nodes and 76,362 directed edges. Response times were computed as least-time paths from the ambulance base to each of the eighteen dependent settlements.

No dedicated route engine was used. At the scale of this problem (one origin and eighteen destinations) the computational advantage of a compiled routing service over a graph traversal in Python is immaterial, whereas keeping the network model inside the same analytical environment as the rest of the study removes an interface and, with it, a source of silent inconsistency between components.

The network includes segments of the A-11 dual carriageway, opened in stages between 2020 and 2023, with junctions at San Esteban de Gormaz and at Velilla de San Esteban. This is consequential for the analysis rather than incidental: several settlements along the Duero corridor are connected to the ambulance base by dual carriageway rather than by conventional road, and any speed model treating the network as uniformly rural will misrepresent them. The

OpenStreetMap classification already encodes the distinction, and no manual adjustment is required (provided the speed model reads it).

5.2 Emergency driving speeds

Each edge was assigned a free-flow speed from its OpenStreetMap maxspeed tag where one is present. In this extract only about 13% of rural segments carry the tag, so the remainder were assigned a default by road class following standard Spanish legal limits.

Driving under warning lights and siren is faster than the free-flow speed a routing engine assumes by default, and the correction is conventionally applied as a single multiplier. We argue that a single multiplier is inappropriate for a network of this composition, because the margin of improvement depends strongly on road class.

The available evidence supports this position. Pappinen and Nordquist (2022) analyzed 652,759 GPS positions recorded during urgent ambulance missions against 141,140 recorded during non-urgent missions, disaggregated by the posted speed limit of the road travelled. On roads limited to 80–100 km/h, the measured increase was of the order of 26–33%, substantially below the 45–48% observed on urban roads limited to 40–50 km/h: the gain from emergency driving concentrates where signalized junctions and low limits constrain progress, not on open road. Petzäll et al. (2011) report an absolute speed increase of 23.2 km/h in rural driving against 19.8 km/h in urban driving; because the rural baseline is higher, a comparable absolute gain corresponds to a smaller relative one, which is consistent with Pappinen and Nordquist. The figure of approximately +56%, frequently cited in the wider literature, derives from an exclusively intra-urban sample (Lupa et al., 2021) and does not transfer to rural road.

A factor differentiated by road class was therefore applied, as set out in Table 2.

Table 2. Emergency driving speed correction factors by road class.

OpenStreetMap road class	Factor applied	Supporting evidence
motorway, trunk (and their links)	1.23	Pappinen and Nordquist (2022): +26–33% measured at 80–100 km/h, and less headroom still where the baseline is already 120 km/h. Documented range 1.20–1.26.
primary, secondary (and their links)	1.30	Pappinen and Nordquist (2022) for the 80–100 km/h band; consistent with the relative gain implied by Petzäll et al. (2011) in rural driving. Range 1.26–1.33.
tertiary, unclassified, residential	1.35	Largest relative margin, but applied to an already low baseline; Lupa et al. (2021) measure roughly 28 km/h on local road under siren. Range 1.30–1.40.

Source(s): own elaboration from Pappinen and Nordquist (2022), Petzäll et al. (2011) and Lupa et al. (2021).

Adopting a flat, urban-derived multiplier would inflate ambulance speeds precisely on those segments where the vehicle already travels close to the legal limit, and the inflation would fall disproportionately on the corridor settlements served by the A-11. Since the purpose of this model is to establish how long the incumbent service takes to reach each settlement, an error in that direction flatters the incumbent and understates the access problem. The differentiated factor is thus the conservative choice with respect to the argument advanced in this paper, not merely the more precise one.


5.3 Activation time, threshold handling and sensitivity

Travel time alone understates response time. Three minutes of activation were added to each path (the interval between dispatch and the vehicle beginning to move, conventionally the "wheels rolling" time) taking the NFPA reference standard as a conservative regulatory value. Empirical studies report lower means, which would constitute the optimistic bound of a sensitivity analysis rather than the design assumption.

Handling this constant correctly is less trivial than it appears. Reachability within a threshold must be evaluated against total response time, whereas a shortest-path cut-off operates on edge

weights alone, and those encode driving time only. Applying the threshold directly as a cut-off therefore admits settlements that the ambulance reaches within the threshold while driving, but not once activation is included. In this study the cut-off is set to the threshold minus the activation time, and the resulting reachable set was verified to coincide exactly with a direct count over the computed response times. The distinction is not academic: it changes which settlements are classified as covered, and the error runs in the direction of overstating coverage.

Figure 2. Road network reaches from the ambulance base.



Source(s): own elaboration from OpenStreetMap data (© OpenStreetMap contributors, ODbL).

The choice of factors is consequential, and its consequence is asymmetric. Substituting a flat multiplier of 1.30 (close to the differentiated values themselves) changes the picture almost not at all: the population-weighted mean response time moves by 0.05 minutes, and a single settlement changes coverage classification at the ten-minute threshold. Substituting the frequently cited 1.56 does something else entirely. Every settlement becomes faster to reach, by 1.01 minutes on average, and the gain concentrates on the most peripheral: Torraño, the worst-served of the eighteen, gains 2.77 minutes. At the design threshold the access deficit disappears altogether (no settlement remains outside fifteen minutes, rather than one) and at twelve minutes the excluded group halves, from forty-six residents to twenty-three. The correction is therefore not a refinement of the second decimal place. Applied without differentiation, and with a value drawn from urban driving, it removes the finding.

Three limitations qualify the model. Traffic conditions and weather are not represented, which is defensible in a territory of this density but would not be in a congested one. The sparsity of maxspeed tagging means most speeds are class defaults rather than observed values, so the absolute times should be read as well-founded estimates rather than measurements. Lastly, the model assumes a single ambulance base, which reflects the confirmed provision for seventeen of the eighteen settlements; one settlement lies closer to a neighboring municipality, but no round-the-clock provision there could be confirmed against an official source, and including an unverified base would have biased the comparison in the incumbent's favour.

6. An equity metric for emergency access in dispersed municipalities

Mean response time is the summary statistic that emergency planning reaches by default, and in a dispersed municipality it is close to useless. Averaging over settlements weights a hamlet of eight residents equally with one of a hundred and twenty-five; averaging over residents allows a well-served majority to conceal a badly served remainder. Either way, the settlements whose access is structurally worst (the ones that motivate intervention in the first place) contribute least to the number by which the intervention would be judged. A summary that improves when the periphery is ignored is not a summary of the problem.

The accessibility literature has long framed this as a question of equity rather than of efficiency, and has developed measures accordingly (Neutens, 2015). Not all of them transfer. The two-step floating catchment area family, refined for rural settings by McGrail and Humphreys (2009), measures access to a service whose capacity is shared among competing demanders, and is well suited to primary care, where the binding constraint is provision per head. The constraint here is different in kind: what matters is whether a vehicle arrives before a clinical window closes, and a second simultaneous emergency is rare enough that capacity sharing is not the operative limitation. The relevant measure is therefore time against a threshold, not provision against demand.

That threshold requires justification, and the justification available is narrower than is sometimes assumed. For out-of-hospital cardiac arrest in ventricular fibrillation, survival declines steeply and approximately linearly with the delay to defibrillation, a relationship formalized in the graphic model of Larsen et al. (1993). Spanish registry evidence places the mean ambulance response time at approximately eleven minutes, and the interval beyond which the probability of survival is substantially compromised at around fifteen (Rosell-Ortiz et al., 2017). We adopt fifteen minutes as the design threshold on that basis. Two qualifications must accompany it. The threshold is grounded in one clinical presentation, and no equally established curve exists for hemorrhage, trauma, stroke or anaphylaxis, each with its own therapeutic window. And the survival curve is used here solely to locate a defensible threshold; it is not used to translate response times into survival outcomes, a step that would require assumptions this paper does not make.

Three indicators are computed over the resulting response times, each answering a question the mean cannot. The **uncovered population share** is the proportion of the dependent population whose total response time exceeds the threshold: it answers how many residents the incumbent service cannot reach in time. The **worst-served response time** is the maximum across settlements, reported together with the population it affects: it answers how bad the situation is for those furthest from provision, and reporting the population alongside it prevents the figure from being dismissed as an outlier without the reader knowing how many people the outlier contains. The **access gap** is the difference between the response time of the resident at the ninetieth percentile of the distribution and that of the median resident: it answers how unequally access is distributed, and unlike the other two it is insensitive to where the threshold is set. Taken together the three describe the shape of the distribution rather than its center, which is what a claim about territorial equity requires.

Two choices in the computation warrant statement. The indicators are computed over the eighteen dependent settlements and exclude the municipal core, which hosts the ambulance base and is therefore the reference case rather than a case of deficient access; were it included, it would contribute the large majority of the population weight and would flatten every indicator towards the value observed at the base itself. And the settlement populations are drawn from municipal records that, in aggregate, exceed the official census figure for the municipality, and should be read as an upper reference rather than an exact count. All three indicators are ratios or differences of times computed over the same population figures, so a proportional error in those figures does not propagate to the conclusions; an error in their relative distribution would, and this is noted as a limitation.

7. Results: who is left outside

Applied to the eighteen dependent settlements, the road model returns a population-weighted mean response time of 8.86 minutes, and an unweighted mean across settlements of 9.33 minutes. Both sit comfortably inside any threshold that could reasonably be defended on clinical grounds. A planning exercise that reported either figure and stopped there would conclude that the incumbent service performs adequately and would be wrong: not because the arithmetic is faulty, but because the mean is answering a question about the municipality when the question that matters is about its residents.

Table 3. Road response time to each dependent settlement.


Settlement	Population	Driving (min)	Activation (min)	Total (min)	> 15 min
Torraño	15	13.64	3.0	16.64	yes
Quintanas Rubias de Arriba	8	10.70	3.0	13.70	
Torremocha de Ayllón	15	10.47	3.0	13.47	
Quintanas Rubias de Abajo	8	9.35	3.0	12.35	
Morcuera	12	6.94	3.0	9.94	
Velilla de San Esteban	20	6.89	3.0	9.89	
Rejas de San Esteban	20	6.62	3.0	9.62	
Villálvaro	125	6.61	3.0	9.61	
Piquera de San Esteban	10	6.32	3.0	9.32	
Inés	30	6.03	3.0	9.03	
Matanza de Soria	15	4.79	3.0	7.79	
Peñalba de San Esteban	10	4.35	3.0	7.35	
Olmillos de San Esteban	40	4.03	3.0	7.03	
Soto de San Esteban	20	3.93	3.0	6.93	
Pedraja de San Esteban	25	3.58	3.0	6.58	
Atauta	50	3.41	3.0	6.41	
Quintanilla de Tres Barrios	15	3.30	3.0	6.30	
Aldea de San Esteban	15	3.04	3.0	6.04	

Source(s): own elaboration from OpenStreetMap data (© OpenStreetMap contributors, ODbL) and municipal population records.

The distribution in Table 3 is not a gradient. Fourteen settlements fall between 6.04 and 9.94 minutes; the remaining four fall between 12.35 and 16.64. Between those two groups lies an empty band of almost two and a half minutes in which no settlement sits at all. The territory does not contain a continuum of gradually worsening access running outward from the base; it contains two populations with a discontinuity between them. This structural break is the central finding of the diagnosis, and it is what makes the choice of a single threshold less consequential than it first appears: any value placed inside the empty band separates the same two groups.

The three indicators quantify what the mean conceals. At the fifteen-minute threshold the uncovered population share is 3.3%, comprising the fifteen residents of a single settlement. At twelve minutes it rises to 10.2%, comprising forty-six residents across four settlements, and it does not rise further at ten minutes, because the same four settlements are the only ones outside in both cases. The worst-served response time is 16.64 minutes, at a settlement of fifteen residents, against 6.04 minutes at the best-served: a ratio of 2.75 between the fastest and slowest access within one municipality. The access gap, which is independent of where the threshold is placed, is 3.03 minutes, the residents at the ninetieth percentile of the distribution waiting that much longer than the median resident. That percentile already falls inside the peripheral group, which is another way of stating the same discontinuity: the worst decile of residents is not the tail of a distribution but a separate cluster.

Figure 3. Road response time by settlement, sorted, against the design threshold.



Source(s): own elaboration.

The four settlements outside the twelve-minute threshold are Torraño, Torremocha de Ayllón and Quintanas Rubias de Arriba and de Abajo. Together they hold forty-six residents: 10.2% of the population dependent on the base, and 1.3% of the municipality once the core is counted. Their disadvantage is a function of position rather than of size, and the point is made most clearly by the settlement that is not among them. Villálvaro is by some distance the largest of the eighteen, with a hundred and twenty-five residents, and it sits inside the threshold at 9.61 minutes. Being small does not place a settlement outside; being peripheral does, and the two happen to coincide in this territory only because the peripheral settlements are also the smallest.

That coincidence is precisely what makes the case difficult to argue on conventional grounds. Forty-six people distributed across four hamlets will not register in any criterion that aggregates over the municipality: they are a rounding error against a population of some three thousand, and an intervention justified by improvement in mean response time would never reach them, because the mean is already satisfactory and moving it would be achieved most cheaply by serving the settlements that are already well served. The equity indicators are not a more sensitive way of measuring the same thing as the mean. They measure a different thing, and they are the only ones of the three that identify this group at all.


Two features of the method reinforce rather than qualify this result. The differentiated speed correction of Section 5 assigns the incumbent service its most favorable defensible speeds on the dual carriageway segments serving the corridor settlements, and the threshold handling described in the same section excludes settlements that would have been counted as covered had activation time been omitted from the cut-off. Both corrections run in the direction of crediting the existing provision. The access deficit reported here is therefore a lower bound on the deficit that a less careful model would have found, and the four settlements identified would not fall inside the threshold under any assumption this study considered defensible.

8. The instrument: a municipal digital twin simulator for virtual testing

8.1 Architecture of the simulator

The instrument is organized in three layers with a unidirectional flow of information, each holding a single authoritative representation of its own domain. The territorial layer holds the static description of the terrain and of travel between settlements, derived from a high-resolution digital elevation model. The decision layer consumes that description as fixed input and, given an incoming request, determines which resource attends it and when. The presentation layer is a three-dimensional viewer that renders the territory, the resources and the resulting mission over real terrain. The separation is strict: the viewer executes and displays, it does not decide, and no operational judgement is formed in it.

Figure 4. Three-layer architecture of the instrument.



Source(s): own elaboration.

Two design commitments follow from this arrangement and are worth stating because they are what make the instrument usable as a test bench rather than as a demonstration. The first is that the same decision layer serves the interactive application and the automated tests, so a behavior observed on screen and a behavior asserted in a test are produced by identical code. The second is that each quantity has one origin: the travel description consumed by the viewer is the one computed from the elevation model, not a second approximation maintained alongside it. Both commitments are unremarkable as software engineering and consequential as method, because they are what allows a result seen by a council officer to be reproduced exactly by an analyst.

Architectures for the remote management of unmanned aerial fleets have been described for infrastructure applications, coupling airborne and ground-based sensing within a modular design (Bemposta Rosende et al., 2020). The structural problem is comparable (a control layer decoupled from the vehicles it directs) but the operational problem is not. There, missions are planned inspections of a fixed asset, known in advance and schedulable. Here, demand is stochastic and geographically dispersed, and the instrument's purpose is not to execute a plan but to let a decision-maker examine what would happen under circumstances that have not yet occurred.

8.2 Repeatable virtual testing

Examining what has not yet occurred requires control of time, and the instrument provides it through a session clock that substitutes for real time while active. Advancing that clock does not simply jump the world forward: every event that would fall within the interval is processed individually, in the chronological order in which it would occur. A resource that becomes

available midway through a large advance therefore becomes available at that instant and not at the end of the jump. The distinction is not cosmetic. Processing intervening events in bulk would distort precisely the response times that the exercise exists to measure and would do so in a way that is invisible on screen.

On that foundation the instrument supports scripted episodes: declarative files specifying a starting state of the world and a time-ordered sequence of incoming requests. An episode is replayed against the running system, which yields two properties that a live demonstration cannot offer. Replay is deterministic, so the same script produces the same outcome, and a council can be shown a situation rather than talked about it. And each episode in the catalogue carries an automated regression test that replays the same script against the operating interface and pins the outcome it is meant to exhibit, so the catalogue is verified rather than merely curated.

Table 4. Catalogue of scripted situations.

Situation put to the test	What the municipality is shown	Regression test
The nearest resource is free	The baseline case: immediate assignment, no alternatives to weigh.	yes
The settlement needing help is itself a base	Material is already on site, yet delivery still takes time; a second call close behind must wait rather than escalate.	yes
Every resource is committed	A further request joins a queue and is served as soon as one returns.	yes
Every resource is committed and waiting is unreasonable	The request is escalated to road transport instead of waiting.	yes
The nearest resource is busy but returning shortly / returning late	A matched pair: one circumstance changes and the outcome reverses.	yes
Two successive requests depend on the same resource	The second inherits the full cycle of the first, not merely its current flight.	yes
A more serious request arrives behind a less serious one	Priority does not inherit the waiting time already accrued by the other.	yes
Two requests arrive at the same moment, considered separately / together	A matched pair: the allocation differs depending on whether they are weighed as a group.	yes
The nearest resource is recharging, with sufficient / insufficient charge	A matched pair: the charge accumulated decides whether recharging is interrupted.	yes


Source(s): own elaboration.

The catalogue is deliberately constructed as a set of minimal contrastive pairs. Four situations exist twice, differing in a single circumstance that reverses the outcome: the resource that returns in seconds against the one that returns in minutes, the battery that has been charging long enough against the one that has not. Presenting a decision in isolation invites the observer to accept it; presenting two nearly identical situations that resolve differently obliges them to identify the circumstance that mattered. For an audience that must scrutinize rather than merely approve an automated recommendation, that is the more useful construction, and it is a property of the test bench rather than of the decision procedure it exercises.

8.3 Making the decision inspectable

A municipal officer is not the reader of an optimization model, and a recommendation that cannot be interrogated is unlikely to be trusted or, ultimately, used. The literature on public-sector artificial intelligence has established that acceptance is conditioned by the intelligibility of the system rather than by its accuracy alone (de Sousa et al., 2019; Gesk and Leyer, 2022). Accordingly, every decision the instrument produces can be opened as a narrative account in ordinary language, accompanied by the table of alternatives that were considered and the reason each was or was not selected. The officer sees not only which resource was assigned, but what the assignment was chosen over.

Figure 5. The viewer with the decision account opens.



Source(s): own elaboration. Terrain and imagery © Cesium ion and contributors; map data © OpenStreetMap contributors, ODbL.

This is the point at which the instrument stops being an analytical device and becomes a deliberative one. A council does not need to reproduce the reasoning to hold it to account; it needs to see that the reasoning has a shape, that alternatives existed, and that the choice among them followed a stated criterion. The same account also serves the transferability argument of Section 10, since it is the artefact through which a different municipality would evaluate whether the system's judgements match its own priorities.

8.4 Pilot validation in the municipal context

The prototype was validated in collaboration with the municipal authority of the study territory. Validation was conducted in the real municipal context using the actual settlement structure, existing emergency provision and the operational assumptions represented in the simulator. Municipal stakeholders were presented with representative scenarios covering resource availability, competing requests, waiting and escalation, changing resource status, simultaneous requests and battery-related constraints.

The validation had three complementary objectives. First, functional validity: the system had to reproduce the expected behavior of the decision logic across predefined scenarios. Second, interpretability: municipal users had to be able to understand the assignment produced, the alternatives considered and the circumstances that determined the outcome. Third, practical usefulness: the simulator had to support the type of counterfactual questions a local authority faces before committing to a new emergency-access configuration.

The validation sessions generated feedback that was incorporated into successive iterations of the prototype, particularly in the presentation of decisions and the organization of scenario information. The resulting version was subsequently re-tested using the same scripted cases,

preserving deterministic behavior and regression-test coverage. This validation does not constitute clinical validation of drone-delivered emergency care or operational certification of the aircraft. It validates the simulator as a municipal decision-support prototype under the real institutional and territorial conditions for which it was designed.

The combination of a working prototype, real municipal data, validation with the intended institutional users, scenario-based testing and iterative refinement places the system at a pilot-validated stage consistent with the TRL5 requirement of this special issue.

8.5 Scope and level of digital-twin integration

The current prototype is designed as a simulation-based municipal digital twin rather than as a fully live-synchronized operational twin. No live telemetry from deployed aircraft is currently connected to the system, and the position of a physical UAV is not propagated into the model in real time. The simulator instead represents the territorial environment, emergency resources, mission states and decision processes required for counterfactual scenario testing.

This distinction does not limit the intended contribution. The planning question addressed by the municipality is inherently counterfactual: whether to deploy drone-assisted emergency resources, where to locate them and under which service conditions. A live operational twin becomes relevant once such a service exists; before deployment, a simulation environment capable of replaying alternative configurations provides the more appropriate decision-support capability. In this sense, the prototype occupies the replayable simulation level of the maturity framework developed in Section 9.

9. How much twin does a rural municipality need?

Section 8 characterized the instrument as a simulation-based municipal digital twin rather than a fully live-synchronized operational twin. Haraguchi et al. (2024) provides a useful maturity perspective for interpreting this distinction, because digital twin capability develops through successive stages with different technical and governance requirements. Their model is addressed to cities with the capacity to occupy any of those stages. What a municipality of three thousand inhabitants needs is the same logic applied to a different constraint: not which stage is most advanced, but which stage answers the question being asked, and what obligations that stage transfers to the authority that operates it. Table 5 sets out the resulting ladder.

Table 5. Levels of digital representation for municipal decision-making.

Level	What question it answers	What it requires, and who operates it	Nature of the recurring burden
L0 Static representation	Where things are.	Existing public cartography. Anyone.	None.
L1 Offline analytical model	What is the case, and what would be under a stated configuration.	Data and technical expertise. An analyst, one commissioned question at a time.	None, but every new question is a new commission.
L2 Replayable simulation	What unfolds over time, and why this outcome rather than another.	L1 plus an execution environment and a curated catalogue of situations. Municipal staff, unassisted.	Occasional revision.
L3 Live-synchronized twin	What is happening now.	Instrumented assets, continuous connectivity, a data pipeline. An operator on duty.	A permanent operational duty.
L4 Closed loop	What should now be done, without being asked.	L3 plus regulatory authorization and an assignment of liability. An accountable service.	Legal responsibility for outcomes.

Source(s): own elaboration, extending Haraguchi et al. (2024).

The transitions between levels are uneven, and the steep one is not where intuition places it. Moving from L0 to L1 costs technical effort. Moving from L1 to L2 costs software engineering, which is real but bounded and once paid, is paid. Moving from L2 to L3 changes the nature of the cost rather than its magnitude: continuous data acquisition, connectivity that must not fail, and above all somebody whose responsibility it is to watch. A municipality can commission a model; it cannot as easily create a post. The literature's implicit assumption of metropolitan capacity is largely an assumption that this post already exists.

The question a council faces before any deployment is a planning question (whether to act, where to place resources, against which threshold) and planning questions are answered at L2. Operational questions, which is what L3 exists to answer, only arise once a service is running. Adopting the rhetoric of the digital twin at the point of deciding whether to build anything therefore invites a municipality to acquire the instrument appropriate to a situation it has not yet entered, at a recurring cost it must sustain before receiving any benefit.

There is a stronger argument than affordability, and it inverts the ladder. A live twin reports the state of a system that exists. It cannot report the behavior of a configuration that has not been built, because there is nothing to synchronize with. For an authority deciding whether to deploy at all, and in what arrangement, the entire object of interest is counterfactual, and the instrument that reaches it is the one with a clock that can be wound to situations that have never occurred. On this reading L2 is not a compromise between ambition and budget. It is the level at which the relevant question is answerable, and L3 is the level at which that question has already been settled.

The governance reading follows. Each level transfers a different kind of obligation: L1 an analytical dependency on whoever holds the expertise, L2 an internal capacity that the authority itself exercises, L3 a continuous operational duty, L4 responsibility for outcomes produced without human intervention. These are not increments of the same commitment but commitments of different kinds, and a council that ascends the ladder acquires them whether it intended to. The recommendation this paper draws from its own case is accordingly modest and, we think, generalizable: a municipality should occupy the lowest digital-twin maturity level that answers its current planning question and should treat any move upward as the assumption of a new technical and governance obligation.

10. Governance, regulation and transferability

What separates the instrument described here from a service in operation is not principally technical. Flights of this kind fall under the specific category of the European regulatory framework, where authorization proceeds through a Specific Operations Risk Assessment, and a service operating routinely over inhabited territory would additionally engage the U-space provisions for airspace management and the question of handover between the automated system and the emergency services that would receive its output. None of this is engineering work, and none of it is addressed in this paper. It is named here because a proposal that omits it invites the reasonable suspicion that its authors have not looked, and because the sequencing matters: the regulatory case is built on a configuration that has been chosen, which is the decision the preceding sections exist to inform.

The institutional condition is more easily met and more often overlooked. The work reported here was carried out in collaboration with the municipal authority of the territory it studies, which is what distinguishes an exercise conducted on a real jurisdiction from one conducted on a case study. That collaboration supplied the settlement records, the confirmation of existing provision and, less tangibly, the constraint that the output had to be intelligible to the people who would act on it: A requirement that shaped Section 8 more than any technical consideration did.

The claim of transferability advanced in Section 9 should be stated in terms of what a second municipality would actually need. The territorial layer requires a digital elevation model, which national mapping agencies across Europe publish openly. The road model requires an extract of

the open street network, freely obtainable for any bounding box. The demand description requires a list of settlements with coordinates and population, which municipal records hold. The evaluation requires one clinically justified time threshold and the location of existing emergency provision. All five inputs are public, and none is specific to this territory: the instrument is portable in the sense that matters, which is that reproducing it elsewhere is a matter of substitution rather than of reconstruction.

What does not transfer is the conclusion. The threshold adopted here rests on one clinical presentation and would have to be re-argued for a service with a different purpose. The speed corrections are calibrated on driving studies conducted in northern Europe and applied through Spanish road classifications and legal limits, and both would require revision elsewhere. Most importantly, the discontinuity reported in Section 7 (the empty band separating a well-served group of settlements from a peripheral one) is a property of this territory's geography and the position of its single base. Another municipality with the same population and area might present a smooth gradient, in which case the equity indicators would identify no distinct group and the case for intervention would have to be made differently, or not at all.

That asymmetry is the practical argument for the position taken in Section 9. If the instrument transfers and the conclusion do not, then what a municipality needs is not a published recommendation derived from somewhere else, but the modest capacity to pose the question about its own territory and see the answer. That capacity is what the second level of the ladder provides, and it is within reach of authorities considerably smaller than the digital twin literature has so far addressed.

11. Conclusions

Road-based emergency access in the municipality studied here does not degrade smoothly with distance from the single ambulance base. It divides. Fourteen settlements are reached within ten minutes of dispatch and four are not reached within twelve, and between the two groups lies a band of more than two minutes in which no settlement falls at all. Forty-six people live in the second group: a tenth of the population that depends on that vehicle, and around one per cent of the municipality once its core is counted. They are not the tail of a distribution but a distinct population, and that they are also the least numerous is precisely what has kept them from view.

Their invisibility is a property of the measures ordinarily used rather than a lapse of attention. The population-weighted mean response time across the same eighteen settlements is under nine minutes, comfortably inside any threshold that could be defended on clinical grounds, and a planning exercise that reported it would conclude that provision is adequate. The indicators proposed here are not a more sensitive version of that measure; they answer a different question, about the shape of the distribution rather than its center, and they are the only ones that identify the group at all. Both methodological corrections introduced along the way (the emergency speed factor differentiated by road class, and the treatment of activation time within the coverage threshold) run in the direction of crediting the existing service, so the deficit reported here is a lower bound rather than an advocate's estimate.

What follows a diagnosis is a decision, and a decision of this kind is counterfactual: whether to deploy at all, in what arrangement, against which threshold. That is why the level of digital representation capable of answering it is not the most advanced one available. A live-synchronized twin reports the state of a system that exists and can say nothing whatever about a configuration that has not been built. The repayable simulation described in this paper (a controllable clock, a verified catalogue of contrastive situations and a decision account legible to non-engineers) was validated with the municipal authority in the real territorial and institutional context of the study. The resulting prototype therefore demonstrates pilot-level readiness for municipal scenario testing: it is sufficiently mature to be exercised by its intended users, while remaining deliberately below the complexity of a continuously synchronized operational twin.

The inputs it requires are public across most of Europe, so the instrument transfers. The finding does not. The discontinuity reported here is a property of this territory's geography and of the position of its single base; another municipality of comparable size might present a smooth gradient, in which case these indicators would identify no distinct group and the case for intervention would have to be made differently, or would not exist. That asymmetry is the practical argument for equipping authorities with the capacity to ask rather than supplying them with conclusions reached elsewhere.

The case for acting in this territory cannot be made in the language of efficiency, and little is gained by attempting it. Forty-six people distributed among four hamlets are a rounding error against any municipal aggregate, and an intervention justified by improvement in the mean would never reach them, because the mean is already satisfactory and the cheapest way to move it is to serve those who are served well already. A service that reaches them will not be efficient and cannot be made so. It would be fair. Which of those two words governs the provision of emergency care in depopulated territory is a political question rather than a technical one, and the contribution of an instrument such as this is not to answer it, but to ensure that it is answered deliberately, with the people it concerns visible on the screen instead of dissolved into an average.

The simulator consequently contributes not only to service efficiency but to territorial resilience, by allowing a municipality to examine how alternative emergency-access configurations would affect its most vulnerable peripheral settlements before committing physical resources.

12. Acknowledgements

The authors would like to thank Universidad Francisco de Vitoria for its support. Also, to City Council of San Esteban de Gormaz (Soria, Spain) for its collaboration and for access to settlement records, and the regional emergency services

13. Funding

This research was funded through the Specific Collaboration Agreement between Universidad Francisco de Vitoria (UFV) and the City Council of San Esteban de Gormaz (Soria, Spain), signed on December 9, 2025. The funding supported the research activities conducted in this study but had no role in the collection of the dataset, the analysis of the results, or the decision to publish.

14. References

- Bemposta Rosende, S., Sánchez-Soriano, J., Gómez Muñoz, C. Q., & Fernández Andrés, J. (2020). Remote management architecture of UAV fleets for maintenance, surveillance, and security tasks in solar power plants. *Energies*, 13(21), 5712. <https://doi.org/10.3390/en13215712>
- Boutillier, J. J., Brooks, S. C., Janmohamed, A., Byers, A., Buick, J. E., Zhan, C., Schoellig, A. P., Cheskes, S., Morrison, L. J., & Chan, T. C. Y. (2017). Optimizing a drone network to deliver automated external defibrillators. *Circulation*, 135(25), 2454–2465. <https://doi.org/10.1161/CIRCULATIONAHA.116.026318>
- Boutillier, J. J., & Chan, T. C. Y. (2022). Drone network design for cardiac arrest response. *Manufacturing & Service Operations Management*. <https://doi.org/10.1287/msom.2022.1092>
- de Sousa, W. G., de Melo, E. R. P., Bermejo, P. H. D. S., Farias, R. A. S., & Gomes, A. O. (2019). How and where is artificial intelligence in the public sector going? A literature review and research agenda. *Government Information Quarterly*, 36(4), 101392. <https://doi.org/10.1016/j.giq.2019.07.004>

- Gesk, T. S., & Leyer, M. (2022). Artificial intelligence in public services: When and why citizens accept its usage. *Government Information Quarterly*, 39(3), 101704. <https://doi.org/10.1016/j.giq.2022.101704>
- Haraguchi, M., Funahashi, T., & Biljecki, F. (2024). Assessing governance implications of city digital twin technology: A maturity model approach. *Technological Forecasting and Social Change*, 204, 123409. <https://doi.org/10.1016/j.techfore.2024.123409>
- Jakobsen, L. K., Gram, J. K. B., Grabmayr, A. J., Højen, A., Hansen, C. M., Rostgaard-Knudsen, M., Claesson, A., & Folke, F. (2025). Semi-autonomous drone delivering automated external defibrillators for real out-of-hospital cardiac arrest: A Danish feasibility study. *Resuscitation*, 208, 110544. <https://doi.org/10.1016/j.resuscitation.2025.110544>
- Larsen, M. P., Eisenberg, M. S., Cummins, R. O., & Hallstrom, A. P. (1993). Predicting survival from out-of-hospital cardiac arrest: A graphic model. *Annals of Emergency Medicine*, 22(11), 1652–1658. [https://doi.org/10.1016/S0196-0644\(05\)81302-2](https://doi.org/10.1016/S0196-0644(05)81302-2)
- Llorent-Bedmar, V., Cobano-Delgado Palma, V. C., & Navarro-Granados, M. (2021). The rural exodus of young people from empty Spain. Socio-educational aspects. *Journal of Rural Studies*, 82, 303–314. <https://doi.org/10.1016/j.jrurstud.2021.01.014>
- Lupa, M., Chuchro, M., Sarlej, W., & Adamek, K. (2021). Emergency ambulance speed characteristics: A case study of Lesser Poland voivodeship, southern Poland. *GeoInformatica*, 25, 775–796. <https://doi.org/10.1007/s10707-021-00447-w>
- McGrail, M. R., & Humphreys, J. S. (2009). Measuring spatial accessibility to primary care in rural areas: Improving the effectiveness of the two-step floating catchment area method. *Applied Geography*, 29(4), 533–541. <https://doi.org/10.1016/j.apgeog.2008.12.003>
- Neutens, T. (2015). Accessibility, equity and health care: Review and research directions for transport geographers. *Journal of Transport Geography*, 43, 14–27. <https://doi.org/10.1016/j.jtrangeo.2014.12.006>
- Pappinen, J., & Nordquist, H. (2022). Driving speeds in urgent and non-urgent ambulance missions during normal and reduced winter speed limit periods—A descriptive study. *Nursing Reports*, 12(1), 50–58. <https://doi.org/10.3390/nursrep12010006>
- Peldon, D., Banihashemi, S., LeNguyen, K., & Derrible, S. (2024). Navigating urban complexity: The transformative role of digital twins in smart city development. *Sustainable Cities and Society*, 111, 105583. <https://doi.org/10.1016/j.scs.2024.105583>
- Petzäll, K., Petzäll, J., Jansson, J., & Nordström, G. (2011). Time saved with high speed driving of ambulances. *Accident Analysis & Prevention*, 43(3), 818–822. <https://doi.org/10.1016/j.aap.2010.10.032>
- Rosell-Ortiz, F., Escalada-Roig, X., Fernández del Valle, P., Sánchez-Santos, L., Navalpotro-Pascual, J. M., Echarri-Sucunza, A., Adsuar-Quesada, J. M., Cenicerros-Rozalén, I., Ruiz-Azpiazu, J. I., Ibarguren-Olalde, K., López-Cabeza, N., Mier-Ruiz, M. V., Martín-Sánchez, E., Martínez del Valle, M., Inza-Muñoz, G., Cordero Torres, J. A., García-Ochoa, M. J., Cortés-Ramas, J. A., Canabal-Berlanga, R., Zoyo López-Navarro, R., López-Messa, J. B., García del Águila, J., Alonso-Moreno, D., Pozo-Pérez, C., Bravo-Castello, J., Ramos-García, N., Gómez-Larrosa, I., & Mellado-Vergel, F. J. (2017). Out-of-hospital cardiac arrest (OHCA) attended by mobile emergency teams with a physician on board: Results of the Spanish OHCA Registry (OSHCAR). *Resuscitation*, 113, 90–95. <https://doi.org/10.1016/j.resuscitation.2017.01.029>
- Sánchez-Soriano, J., Verdín-Urgal, G., & Gordo-Herrera, N. (2025). Optimizing last-mile deliveries: Addressing customer absence through genetic algorithm. *Technologies*, 13(3), 115. <https://doi.org/10.3390/technologies13030115>
- Schierbeck, S., Nord, A., Svensson, L., Ringh, M., Nordberg, P., Hollenberg, J., Lundgren, P., Folke, F., Jonsson, M., Forsberg, S., & Claesson, A. (2023). Drone delivery of automated external defibrillators compared with ambulance arrival in real-life suspected out-of-hospital cardiac arrests: A prospective observational study in Sweden. *The Lancet Digital Health*, 5(12), e862–e871. [https://doi.org/10.1016/S2589-7500\(23\)00161-9](https://doi.org/10.1016/S2589-7500(23)00161-9)