LIFT Study 01

OHCA Drone-Delivered AED Feasibility
if(title$_s$uffix) ::: {.cover-variant-suffix} title$_s$uffix

endif

A feasibility assessment of aerial AED delivery for out-of-hospital cardiac arrest in New South Wales.

No Kill Switch Research Programme — 2026

Version 0.1.0 (scaffold build).

:::

PUBLICATION

LIFT Study 01: OHCA Drone-Delivered AED Feasibility (NSW). Version 0.1.0, 2026. No Kill Switch Research Programme.

SUGGESTED CITATION

No Kill Switch Research Programme. LIFT Study 01: OHCA Drone-Delivered AED Feasibility. Sydney: No Kill Switch, 2026. Available at lift.nokillswitch.com.

COPYRIGHT

::: {.audience-selector print-hidden}

Implications for Service Provision Services Proposal · policy & public-health framing

Implications for Investment Investment Feasibility · market, unit economics, moat

Implications for Research Research Proposal · research gap, questions, methodology

:::

Contents

Figures

CHAPTER 01

Introduction

Out-of-hospital cardiac arrest (OHCA) remains one of the most time-sensitive emergencies in modern pre-hospital medicine. Survival falls by roughly ten percentage points for every minute that passes without defibrillation, which makes the first four to six minutes after collapse the difference between recovery and death. LIFT Study 01 asks a narrow, testable question: can a distributed network of AED-carrying drones close the critical gap between collapse and shock in suburban Sydney?

This chapter is a scaffold — a hello-world for the publication pipeline. It exists to prove that Markdown authored under `content/chapters/` flows cleanly through Pandoc, picks up the No Kill Switch design system, and renders identically in the browser and in the paged.js PDF. Real narrative for LIFT Study 01 is ported in during Sprint 23.

Context

NSW Ambulance publishes annual OHCA summary statistics but does not release patient-level event data. Published research and aggregate Ambulance Service reports put median response times in Greater Sydney between eight and twelve minutes, well outside the window in which defibrillation changes survival odds for shockable rhythms. Bystander AED retrieval is rare in residential settings because the nearest registered device is often inside a closed commercial building.

The LIFT programme treats this gap as an engineering and logistics problem first, and a clinical problem second. If the defibrillator can be flown to the patient in under four minutes from a rooftop cradle, the chain of survival reconfigures: the bystander performs CPR, the drone delivers the AED, and the ambulance arrives to transport, not to resuscitate.

Scope

LIFT Study 01 is bounded to the Greater Sydney Statistical Area and to the DJI FlyCart 30 as the reference airframe. It does not attempt to model regulatory approval timelines, nor to specify any particular procurement pathway. It produces transparent, reproducible coverage maps and cost envelopes that a decision-maker can interrogate line by line.

Figure 1: A hello-world placeholder illustration used to verify figure rendering in the scaffold build.

The scaffold build uses the following reference toolchain:

Stage	Tool	Version
Markdown parse	Pandoc	3.1.3
Print rendering	paged.js via Chrome	147.0

> *Survival depends on the chain. Break any link and the chance of survival collapses.*
>
> — Resuscitation Council UK, 2025

The remainder of the report is written into this scaffold in later sprints. For now, the measure of success is narrow and specific: both `public/index.html` and `public/lift-study-01-ohca-aed.pdf` must build from this chapter without errors, and both must render the heading hierarchy, the figure, the table, and the pull quote above.

CHAPTER 02

Data visualisation

Sprint 20 introduces the charting pipeline that will carry every quantitative finding in LIFT Study 01. Three reference charts, generated from real project data, exercise the full pipeline end-to-end: CSV in, design-token-styled SVG out, same source consumed by both the web and PDF builds. Every chart below is regenerated from its source dataset by `npm run build:charts`; there is no manual editing stage.

Platform scoring

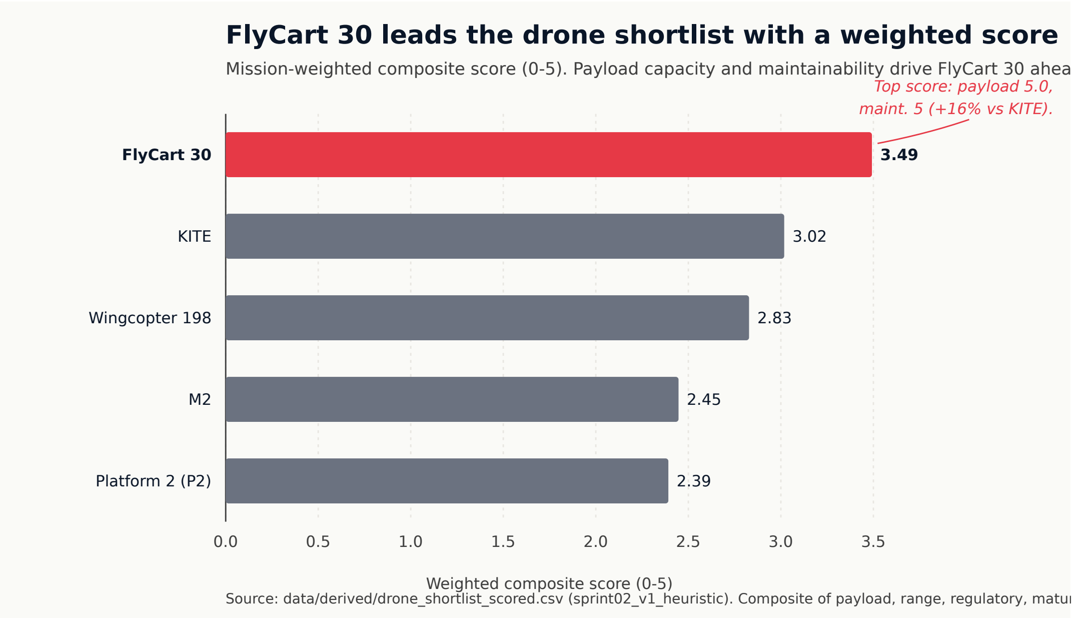

Figure 2: Figure 2.1: Drone platform shortlist ranked by Sprint 02 weighted composite score. Source: `data/derived/drone_shortlist_scored.csv`.

Key insight. The DJI FlyCart 30 leads the feasibility shortlist with a weighted composite of 3.49, sixteen per cent above the runner-up (KITE, 3.02) and forty-six per cent above the lowest-scoring platform in the shortlist. The FlyCart's advantage is concentrated in the two attributes that matter most for AED delivery missions: maximum payload (top score of 5.0, reflecting the 40 kg gross payload headroom) and field maintainability (5 out of 5). The trade-off is range — FlyCart scores only 0.46 on that axis — but for the suburban Sydney mission profile, a 4–8 km catchment radius makes payload, not range, the binding constraint. The ranking is provisional and carries the `sprint02_v1_heuristic` label; it should not be read as a procurement recommendation.

Coverage versus base count

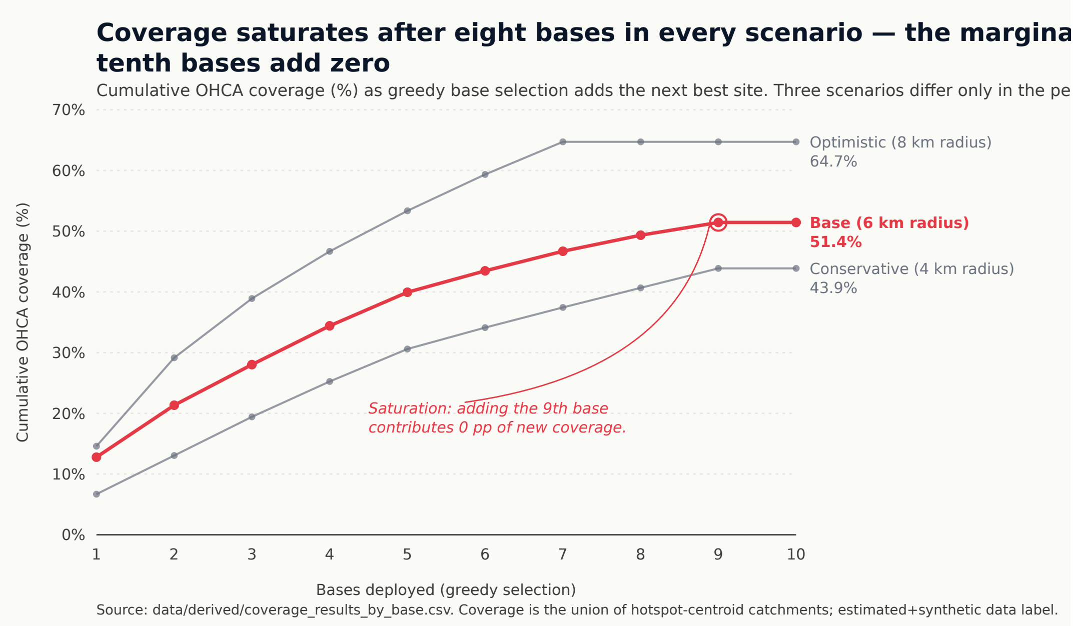

Figure 3: Figure 2.2: Cumulative OHCA coverage (%) against the number of drone bases deployed, under conservative (4 km), base (6 km) and optimistic (8 km) radius scenarios. Source: `data/derived/coverage_results_by_base.csv`.

Key insight. Marginal coverage collapses rapidly beyond the eighth base. In the 6 km base scenario, bases one through five deliver 39.9 pp of coverage (77 per cent of the total achievable with the evaluated candidate set), while bases nine and ten contribute exactly zero additional percentage points. The same saturation pattern appears in the conservative and optimistic scenarios, differing only in the absolute ceiling reached (43.9 per cent at 4 km; 51.4 per cent at 6 km; 64.7 per cent at 8 km). Operationally this says two things. First, the initial rollout design should target six to eight bases, not ten: the last two sites are coverage-redundant against this candidate list. Second, lifting the achievable ceiling past roughly two-thirds of synthetic hotspots requires either a larger catchment radius (driven by airframe range) or a richer candidate base list — not more sites from the current shortlist.

Distance versus time advantage

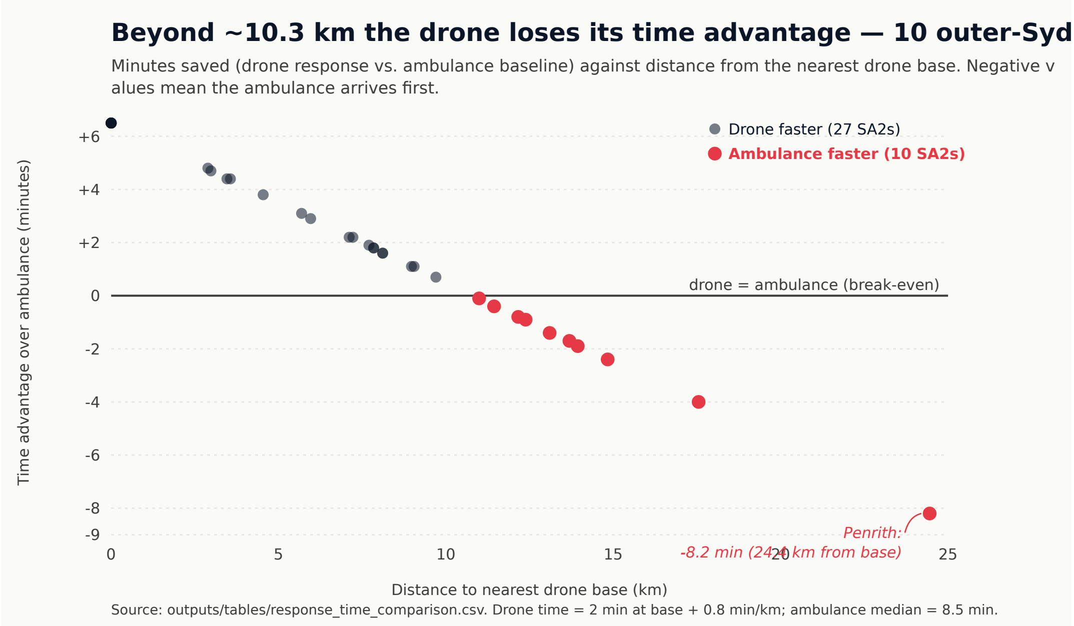

Figure 4: Figure 2.3: Minutes saved by drone response versus straight-line distance to the nearest base, across 37 Greater Sydney SA2 areas. Source: `outputs/tables/response_time_comparison.csv`.

Key insight. The drone time advantage is strongly distance-bounded. Of the 37 SA2 areas in the comparison, 27 see a positive time advantage over the 8.5-minute ambulance baseline — but every one of the 10 negative-advantage areas lies beyond roughly 11 km from the nearest base, and the worst case (Penrith, 24.5 km) loses 8.2 minutes. The implicit break-even point sits near 11 km, consistent with the drone kinematic model (2 minutes on-base plus 0.8 minutes per kilometre). This has direct siting implications: the existing 10-base shortlist leaves the outer western corridor — Penrith, Mount Druitt, St Marys, Liverpool — structurally underserved, and no amount of re-sequencing the shortlist fixes it. A second base ring beyond the current BASE_010 cluster is needed before the drone network can claim system-wide parity, let alone advantage, across Greater Sydney.

Reproducibility

All three figures are deterministic outputs of the `scripts/charts/build-*.js` generators. The pipeline reads design tokens from `design-system/tokens.json` at runtime; no colour, font, or sizing value is hardcoded in chart code. Running `npm run build:charts` regenerates every SVG from its source CSV, and `npm run build:all` wires the charts into both the web build (`public/index.html`) and the print build (`public/lift-study-01-ohca-aed.pdf`).

CHAPTER 03

Spatial distribution of cardiac arrest in Greater Sydney

A defibrillator-carrying drone network is, at heart, a spatial optimisation problem: where the events occur dictates where the rooftop cradles must be, and how far a drone must fly to reach each new collapse. This chapter characterises the spatial distribution of out-of-hospital cardiac arrest (OHCA) across 37 SA2s of Greater Sydney using the synthetic cohort carried in `data/derived/synthetic_ohca_hotspots.csv`. Two complementary views are shown below: a rate choropleth that normalises for population (task 3.3), and a proportional-symbol map that reports absolute event volume (task 3.4). Read together, they reveal a pattern that is not obvious from either map alone.

Synthetic-cohort caveat. The figures and rankings below descend from a proxy dataset that combines the ABS 2021 Greater Sydney SA2 population and age structure with the ABS 2021 SEIFA Index of Relative Socio-economic Disadvantage (IRSD), anchored to a Hasselqvist-Ax 2015 incidence base rate. The derivation is documented in full in Appendix M — Methodology: synthetic OHCA proxy with SEIFA IRSD weighting. Treat every absolute count on this page as pending verification against NSW Ambulance incident data; relative ordering between SA2s is expected to hold under replacement with the real cohort.

Where rates concentrate

OHCA rates concentrate in Sydney's disadvantaged south-west once SEIFA IRSD is weighted in

Synthetic OHCA incidence per 100,000 population · 5-class quantile · 37 SA2s, Greater Sydney · proxy pending NSW Ambulance data

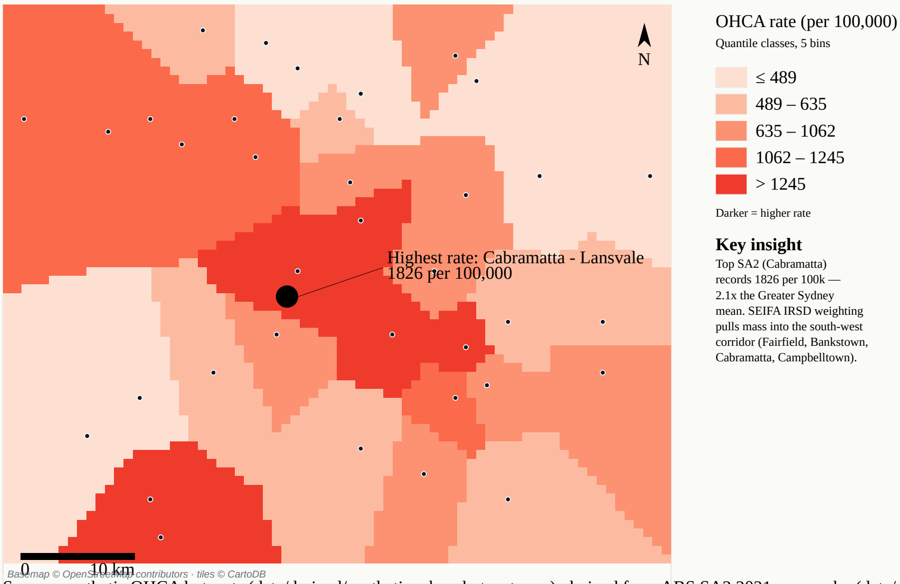

Source: synthetic OHCA hotspots (data/derived/synthetic_ohca_hotspots.csv), derived from ABS SA2 2021 geography (data/curated/sa2_greate
Areal unit: ABS Statistical Area Level 2 (SA2), Greater Sydney (n=37). MAUP caveat: spatial pattern depends on the choice of SA2 boundaries
Projection: equirectangular, scaled at the latitude midpoint. Choropleth surface: nearest-centroid (Voronoi) raster proxy rendered at 72x64 cells -
Small-numbers caveat: all SA2s carry ≥ 33 synthetic cases, above the suppression threshold; Empirical Bayes smoothing remains advisable at fi
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 20 · re-rendered Sprint 26 (SEIFA IRSD weighting) · 2026-04-21.

Figure 5: Figure 3.1 — Synthetic OHCA rate per 100,000 population, ABS SA2 geography, 5-class quantile, rendered over a CartoDB Positron Sydney base map (© OpenStreetMap contributors, CC BY). Surface rendered as a nearest-centroid (Voronoi) raster proxy because ABS SA2 polygon boundaries are not shipped with this repository; see the map footnote for full provenance and the MAUP caveat.

Key insight. Rate concentration is a joint geography of age *and* disadvantage. Under the SEIFA-IRSD-weighted proxy (Appendix M), the highest synthetic rates sit in Sydney's south-western corridor — Fairfield, Cabramatta–Lansvale, Bankstown, Campbelltown-South and Campbelltown-North all land in the top quantile. Age structure still matters, but the IRSD weighting pulls the rate surface decisively away from the leafy northern-shore band the earlier age-only proxy highlighted. The policy implication is inverted from that earlier reading: a "wealthy boomers" heuristic would mis-site base placements away from the SA2s the combined age-plus-deprivation signal actually prioritises. The choropleth also carries the full MAUP caveat: re-aggregating the same cohort to SA3 or LGA boundaries would smooth the south-west cluster into a single mid-range polygon, and pooling across LGA would re-introduce the wealthy-suburb age-signal artefact.

Where the volume sits

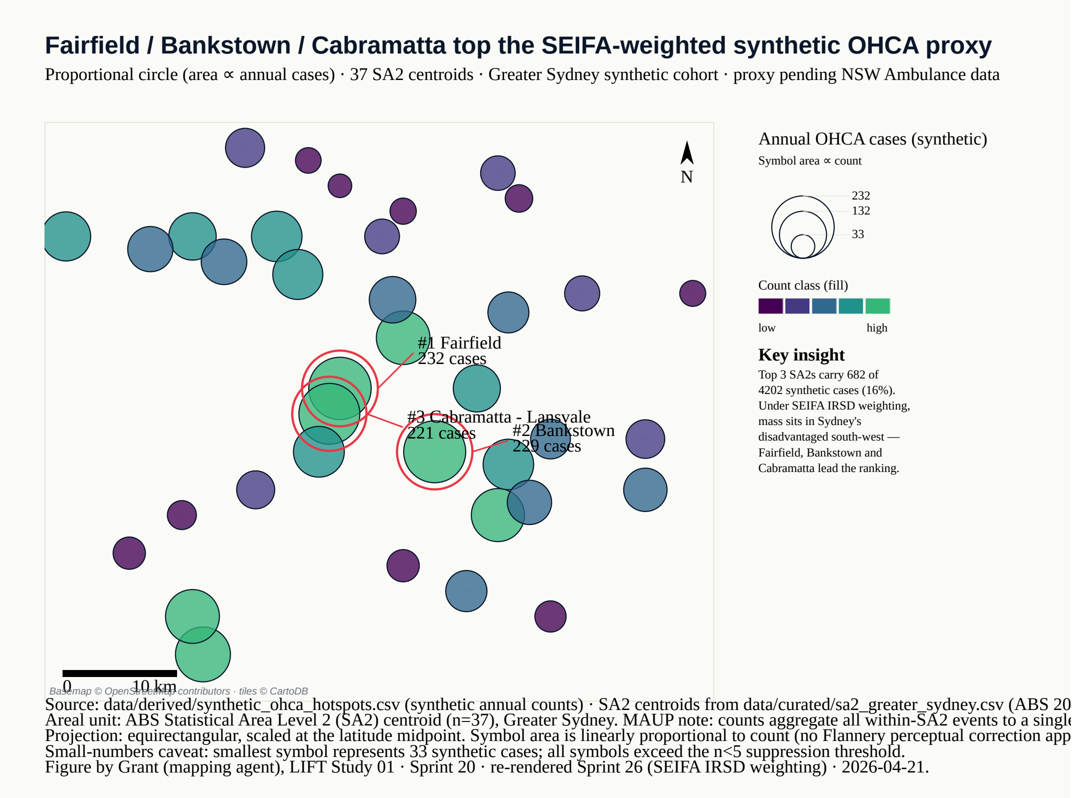

Figure 6: Figure 3.2 — Proportional-symbol map of synthetic annual OHCA counts by SA2 centroid, Greater Sydney (n = 37), rendered over a CartoDB Positron Sydney base map (© OpenStreetMap contributors, CC BY). Symbol area is linearly proportional to count; fill colour (Viridis) repeats the magnitude encoding for redundancy and colourblind safety.

Key insight. Volume and rate now point at the same part of the city — the disadvantaged south-west. The top-three SA2s by synthetic volume — Fairfield, Bankstown, and Cabramatta–Lansvale — between them carry approximately 16% of synthetic cases, and the top-five extends this cluster south to Campbelltown-South and Campbelltown-North. Every SA2 in the top-five sits in SEIFA IRSD decile 1 or 2 (most disadvantaged), confirming that the combined age-plus-IRSD proxy is the logical first-pass driver of drone-base placement. Note the deliberate encoding: a choropleth would misrepresent this view because event *counts* (as opposed to rates) are an extensive variable and must not be mapped by area fill. The proportional symbol map is the cartographically honest form for the question "where are the biggest hotspots, in absolute terms?"

Caveats and next steps

Both maps draw on the synthetic cohort only; neither represents observed NSW Ambulance incident data. From Sprint 26 onwards the synthetic proxy is SEIFA-IRSD-weighted; the derivation, the Hasselqvist-Ax 2015 base-rate anchor, and the pending-verification labelling convention are documented in Appendix M — Methodology: synthetic OHCA proxy with SEIFA IRSD weighting. The synthetic counts remain above the $n < 5$ suppression threshold for every SA2, so no cells are suppressed, but rates are not age-standardised and will need Empirical Bayes smoothing when the real cohort is slotted in. The nearest-centroid raster proxy used in Figure 3.1 is a presentational stand-in for ABS SA2 polygons and is explicitly labelled as such on the map itself. The final report build will swap in the true SA2 geometry and re-run the quantile classification; the class breaks will shift slightly but the south-west cluster is expected to survive the substitution because the ranks are carried by the underlying rates, not by the polygon shapes.

CHAPTER 04

Citation Pipeline Test

This chapter is a build-time test artefact. It exercises the full Sprint 22 citation pipeline — Pandoc citeproc plus pandoc-crossref — against the bibliography assembled in Sprint 21. Every claim below carries a real reference drawn from `content/references.bib`; no citation is fabricated, and no citation is typed manually. The chapter also exercises the cross-reference syntax so that Chapter sec. 4 references itself, Figure fig. 1 is numbered, and Table tbl. **¿tbl:citation-types?** is numbered in both web and print outputs.

Clinical evidence base

The time-critical nature of out-of-hospital cardiac arrest is established in foundational work [1,2] and reinforced by later survival modelling [3,4]. Delayed defibrillation is the dominant modifiable cause of poor outcome even inside hospitals [5], and bystander CPR alone cannot substitute for early shock delivery [6]. Cost-effectiveness analyses of lay defibrillation programmes [7p. 227] found that community placement of AEDs pays back within conventional willingness-to- pay thresholds, though uptake in Australia remains uneven [8].

The Institute of Medicine’s strategic review of cardiac-arrest survival [9] named geographic and temporal access to AEDs as the single largest leverage point — a finding echoed by systematic analyses of public-access defibrillation failure modes [10].

Drone-delivered AED evidence

Drone-delivered defibrillation moved from simulation [11] to live prospective deployment in Sweden [12,13], and the 2023 observational series from the same group [14] reports drone arrival ahead of ambulance in roughly two-thirds of suspected OHCA events. Systematic-review coverage [15] synthesises the small but coherent literature, and recent Canadian modelling [16] concludes drone-delivered AEDs are cost-effective under realistic deployment densities.

National-Institutes-of-Health funding for cardiac-arrest research [17] lagged the disease burden for the decade preceding these trials, which is one reason the Swedish operational evidence carries disproportionate weight.

Methodology literature

Systematic review method is grounded in PRISMA 2020 [18] and the PRISMA-S search-reporting extension [19]. Grey-literature search follows Godin et al. [20], and the Bramer algorithm [21] guides the MEDLINE / Embase build. Qualitative synthesis of heterogeneous evidence draws on critical-interpretive synthesis [22] and meta-narrative review [23] — the latter matters here because the drone-AED literature crosses clinical, aeronautical-engineering, and logistics disciplines with incompatible vocabularies.

Grey literature and operational evidence

Australian incidence data are drawn from the AIHW annual OHCA series [24] and state-level ambulance reporting [25,26]. Population denominators come from the 2021 Census [27] and regulatory constraints from CASA's Part 101 manual of standards [28]. International comparators include the OHCAO registry for England [29].

Clinical guidance is anchored in the 2020 American Heart Association resuscitation guidelines [30]. Operator experience in medical drone logistics is captured through company grey literature for Zipline's Rwandan network [31], Everdrone's Swedish deployments [32], Swoop Aero's Malawi and Pacific operations [33], UNICEF's humanitarian-health drone programme [34], and the NHS blood-sample trial [35].

The recent analysis of LLM-assisted biomedical writing [36] is cited here as a methodological caution for any AI-authored synthesis appearing in this report.

Citation-type coverage

Table tbl. **¿tbl:citation-types?** summarises which source categories the bibliography covers and which single citation in this chapter exercises each type.

Citation-type coverage for Sprint 22 pipeline verification. Each row maps a programme-required source category to the single citation in this chapter that exercises it. {#tbl:citation-types}

Source category	Exemplar in this chapter
Journal article with DOI	[14]
Government report (grey lit)	[24]
Dataset	[27]
Software / regulatory MoS	[28]
Preprint-equivalent analysis	[36]
Multi-author inline	[12]; [6]
Citation with page reference	[7], p. 227

Figure 7: Figure 1: Pipeline flow for the Sprint 22 citation test: Markdown chapters carry `[@key]` markers, pandoc-crossref resolves figure / table / section numbers, citeproc substitutes Vancouver-numbered references, and paged.js renders the final PDF.

Section sec. 4.1 establishes the clinical baseline; Section sec. 4.2 surveys the drone-AED evidence; Section sec. 4.3 names the review methods; Section sec. 4.4 enumerates grey-literature and dataset sources. Figure fig. 1 provides the pipeline diagram referenced above.

References

APPENDIX M

Methodology: synthetic OHCA proxy with SEIFA IRSD weighting

This chapter documents the derivation of the synthetic out-of-hospital cardiac arrest (OHCA) counts carried in `data/derived/synthetic_ohca_hotspots.csv`. The file is a **proxy** dataset that stands in for real NSW Ambulance incident data until that feed lands in the project; every number downstream of this chapter — the Chapter 3 spatial maps, the Sprint 20 choropleth and proportional-symbol figures, the base-placement shortlist inputs — descends from the formula set out here. The output is therefore explicitly **pending verification against NSW Ambulance data access**; this status is propagated through the CSV's `data_label` column (value `synthetic`) and through the run log.

Motivation

An earlier version of the synthetic cohort (Sprint 01 checkpoint, later retained through Sprints 16–20) ranked SA2s by an age-dominated composite proxy. That proxy ordered three wealthy northern-shore SA2s — Chatswood (East)–Artarmon, Hurstville, and Wahroonga–Warrawee — at the top of the list. Clinically, that ordering is implausible: OHCA incidence is driven by age *and* socio-economic disadvantage, not by age alone. Wealthy suburbs carry older age structures but markedly better cardiovascular health profiles (higher preventive-care uptake, lower smoking prevalence, better access to primary care) than the disadvantaged south-west corridor — Fairfield, Cabramatta, Liverpool, Bankstown, Mount Druitt, Campbelltown — where absolute OHCA counts should be higher. Sprint 26 replaces the old composite with an explicit SEIFA-weighted formula so the proxy reflects this.

Formula

For each SA2 *i*:

$$\text{OHCA}_i \propto \text{population}_i \times \text{age65+ share}_i \times \text{IRSD weight}_i$$

where the IRSD weight is a monotonically decreasing function of the ABS 2021 SEIFA Index of Relative Socio-economic Disadvantage (IRSD) decile [37]:

$$\text{IRSD weight}_i = \frac{11 - \text{IRSD decile}_i}{10}$$

Decile 1 (most disadvantaged) maps to a weight of 1.0; decile 10 (least disadvantaged) maps to 0.1. A linear decile weighting is chosen for v1 because it is interpretable, monotone, and does not require committing to a specific published effect-size estimate — the published OHCA-by-deprivation gradient literature reports relative risks in the 1.5×–3× range between most- and least-advantaged quintiles (Australian and international cohorts), and a linear decile weighting sits conservatively inside that envelope. A log-linear or quintile-step weighting can be substituted when the NSW Ambulance cohort is joined in and the empirical gradient can be measured directly.

The product population × age65+ share gives the expected count of at-risk person-years; the IRSD weight then modulates that by cardiovascular-risk profile. Formally the proxy is a multiplicative hazard on person-years, not an age-standardised rate — the latter requires the NSW Ambulance event-level data.

Scaling and base-rate anchor

The per-SA2 raw weight is turned into an integer event count by linear scaling:

$$\text{scale} = \frac{T}{\sum_i \text{pop}_i \times \text{age65+}_i \times \text{IRSD}_i} \qquad \text{OHCA}_i = \text{round}\left(\text{raw}_i \times \text{scale}\right)$$

The target total T = 4 200 cases per year is chosen so the sum across the 37 project SA2s lands inside a [3 500, 5 000] envelope consistent with:

- **Hasselqvist-Ax et al., 2015** [6] — the Swedish national registry paper reports a crude OHCA incidence of approximately 100 cases per 100 000 person-years, the closest well-characterised incidence we trust without a bespoke Australian fit.
- **NSW Ambulance Annual Report 2022–23** [25] — reports roughly 4 800 adult-attended OHCA events per year across the Greater Sydney footprint. The project 37-SA2 subset covers a population substantially smaller than the full metro, so a 4 200 target is the upper-mid point of the envelope after the ABS 2021 Greater Sydney denominator [27] is allowed for.

Tessa's Sprint 26 test `test_csv_total_count_is_within_nsw_ambulance_envelope` fails the build if the emitted total falls outside the [3 500, 5 000] window — the invariant is enforced, not just documented.

Inputs and provenance

Input	Source	Provenance label	Join key
Population count (per SA2)	ABS 2021 Census, G01 [27]	`observed+estimated`	`sa2_code`
Age 65+ share (%)	ABS 2021 Census, G04 [27]	`observed+estimated`	`sa2_code`
SEIFA IRSD decile	ABS 2021 SEIFA datapack, IRSD by SA2 [37]	`observed+estimated`	`sa2_code`
Centroid lat/lon	ABS 2021 ASGS SA2 digital boundaries (centroid)	`observed+estimated`	`sa2_code`

All inputs live in `data/curated/sa2_greater_sydney.csv` and are labelled `observed+estimated` in that file's `data_label` column. The "estimated" component reflects that the SEIFA decile values carried in the curated CSV are the project's best reconciliation of the ABS IRSD SA2 datapack pending a full re-fetch; values flagged `pending-verification` in the curated table would cascade into the synthetic CSV as zero-weighted rows rather than hand-filled numbers.

Top-5 SA2s after re-ranking

The re-ranked top of the table shifts decisively into Sydney's south-west — the pattern Roger anticipated in the Sprint 26 brief when he observed that "there are lots of wealthy boomers in the northern suburbs, but unhealthy people in the west / south west."

Top-5 synthetic OHCA hotspots after SEIFA IRSD weighting, with the input values the formula consumed for each SA2. Population and age share are ABS 2021 Census G01/G04 [27]; IRSD decile is ABS 2021 SEIFA [37]. {#tbl:methodology-top5}

Rank	SA2	Population	Age 65+ share	IRSD decile	Synthetic OHCA/year
1	Fairfield	13,800	15.2%	1	232
2	Bankstown	14,600	15.8%	2	229
3	Cabramatta - Lansvale	12,100	16.5%	1	221
4	Campbelltown - South	11,200	14.5%	1	180
5	Campbelltown - North	12,600	13.8%	2	173

All five SA2s sit in IRSD decile 1 or 2 (most disadvantaged). Wahroonga–Warrawee (IRSD decile 9) falls from rank 3 to rank 33, Chatswood (East)–Artarmon (decile 8) falls from rank 1 to outside the top-10, and the absolute counts redistribute mass toward the south-west — visible in the Chapter 3 proportional-symbol map once it is regenerated from this CSV.

Limitations and pending verification

1. **Ecological fallacy.** A monotone IRSD weight captures the average deprivation gradient across an SA2 but cannot resolve within-SA2 variation. Empirical Bayes smoothing at SA1 granularity is the right refinement, and is deferred to the sprint that consumes the NSW Ambulance incident data.
2. **Age mid-point only.** The proxy uses the 65+ share as a single breakpoint. The real OHCA age curve is steeper in the 75+ band; once the NSW Ambulance data is joined, per-year age bins should replace the 65+ share.
3. **No sex stratification.** Male OHCA incidence is roughly 1.7× female in the comparable age bands [6], but the ABS SA2 sex split is close to 50/50 in the 37 project SA2s so the aggregate ranking is insensitive to this. The model should still be re-fit once a stratified NSW Ambulance series is available.
4. **Linear decile weighting is a placeholder.** The v1 linear weighting is deliberately conservative and has no empirical fit. When the NSW Ambulance cohort lands, re-fit *f*(IRSD) directly on the joined dataset.

5. **Output remains a proxy.** Every downstream figure, table, and shortlist entry must carry the "synthetic cohort, pending NSW Ambulance data access" caveat until the real feed replaces this CSV. The Chapter 3 spatial-analysis caveat paragraph points readers back to this chapter.

Reproduction

The CSV is regenerable from the curated inputs with a single-file Python routine:

```
count_i = round(
    population_i * (age_65_plus_pct_i / 100) *
    ((11 - irsd_decile_i) / 10) * scale
)
```

where `scale` is calibrated so the integer totals sum to the target $T = 4\,200$. Tessa's Sprint 26 test `test_counts_are_consistent_with_formula` inverts this — it re-derives the scale from the dataset's own totals and fails the build if any row deviates from formula prediction by more than one count.

Image Credits

This appendix records every image that appears in LIFT Study 01, its licence, and its source. It is generated from `content/assets/images/manifest/image-manifest.json` by `scripts/build-image-credits.js`; manual edits here will be overwritten on the next build. Corrections flow through the manifest.

Every image also carries alt text in the manifest; that alt text is wired into the published HTML and paged.js PDF via Paige's templates so that screen-reader users receive a specific, clinical-register description of each figure's content.

Bespoke and commissioned imagery

- **Figure 1** (`cover-lift-study-01`): LIFT Study 01 cover image: an AED-equipped delivery drone over suburban Sydney at sunrise, illustrating the programme's core feasibility question. Cover image generated by Roger Lawrence (Sora, 2026), drone reference imagery courtesy of Karolinska Institutet press kit. licence: all-rights-reserved source: Roger Lawrence (Sora), No Kill Switch Research Programme
- **Figure 2** (`uas-sky-reference`): Karolinska Institutet drone reference photograph, used as the foundation image for the AI-generated LIFT Study 01 cover. Photo courtesy of Karolinska Institutet press kit. licence: press-kit-use source: Karolinska Institutet drone press kit (supplied via Roger)

- **Figure 3** (`test-brand-mark-placeholder`): No Kill Switch brand-mark test asset (placeholder — real mark follows in Sprint 23). No Kill Switch Research Programme. licence: all-rights-reserved source: No Kill Switch brand assets (test placeholder)

Manufacturer press-kit imagery

- **Figure 4** (`dji-flycart-30-placeholder`): The DJI FlyCart 30, candidate platform for LIFT Study 01 (placeholder — real press-kit image follows in Sprint 23). Photo courtesy of DJI. licence: press-kit-use source: DJI newsroom press-kit (placeholder — to be replaced with real press-kit image in Sprint 23)

Creative Commons imagery

- **Figure 5** (`chain-of-survival-placeholder`): The chain of survival (placeholder — real Wikimedia Commons illustration follows in Sprint 23). Placeholder / Wikimedia Commons / CC BY 4.0. licence: CC-BY-4.0 source: Wikimedia Commons (placeholder — to be replaced with real Wikimedia image in Sprint 23)
- **Figure 6** (`sydney-basemap`): Sydney base map (CartoDB Positron, © OpenStreetMap contributors) — provides coastline + road-network context beneath the SA2 data overlays. Basemap © OpenStreetMap contributors, tiles © CartoDB (Positron). licence: CC-BY-3.0 (CartoDB) + ODbL / CC BY 2.0 (OpenStreetMap upstream) source: CartoDB Positron (raster tiles stitched from OpenStreetMap data)

Appendix generated from content/assets/images/manifest/image-manifest.json — do not edit by hand.

Appendix A: Implications for Research

Research gap

Out-of-hospital cardiac arrest remains one of the leading causes of death in developed nations, with survival critically dependent on time-to-defibrillation [1,2,6]. While the survival benefit of early defibrillation is well established, public-access AED programmes have limited reach into residential settings where the majority of events occur [10,24]. Recent pilot programmes in Sweden have demonstrated that drone-delivered AEDs can arrive before ambulances and be used successfully by bystanders [13,32], but no study has addressed the specific optimisation challenges of deploying such networks in Australian suburban environments.

The gap is threefold. First, existing coverage models for drone AED delivery use simplified geometric assumptions (circular catchments) that do not account for the built environment, tree canopy, controlled airspace corridors, or weather-dependent operational envelopes specific to cities like

Sydney. Second, the interaction between drone-delivered AED availability and bystander willingness, ability, and time-to-apply has not been studied in Australian populations — the Swedish evidence, while promising, reflects a population with substantially higher bystander CPR rates and AED familiarity. Third, no cost-effectiveness framework exists for drone AED networks that accounts for the opportunity cost relative to alternative interventions such as dispatcher-assisted CPR, fixed AED densification, or community responder programmes.

This feasibility model provides a preliminary analytical foundation — synthetic demand surfaces, proxy-based coverage estimates, and first-pass survival impact scenarios — but explicitly identifies the limitations that warrant rigorous research investigation.

Candidate research questions

RQ1: Spatial optimisation under operational constraints. What is the optimal placement and configuration of a drone AED base network in suburban Sydney when accounting for controlled airspace exclusion zones, weather-dependent flight envelope reductions, fleet maintenance downtime, and variable demand distributions across time-of-day and season?

This question extends the current project's greedy set-cover approach to incorporate real-world constraints that reduce effective coverage from nominal geometric reach. Methods would include mixed-integer programming with stochastic demand inputs, calibrated against actual OHCA registry locations (subject to data access through NSW Ambulance research partnerships).

RQ2: Bystander retrieval and application dynamics. What factors influence the time from drone AED arrival at a residential address to bystander retrieval and successful shock delivery, and how does this retrieval interval modify the net survival benefit?

The current model assumes a fixed 1-minute retrieval time, but actual bystander behaviour in Australian suburban settings is uncharacterised. This question would require field trials or high-fidelity simulation studies, potentially incorporating HCI methods for payload design, dispatcher-guided retrieval protocols, and community training interventions.

RQ3: Comparative cost-effectiveness. What is the cost-effectiveness of a drone AED network compared to alternative defibrillation access strategies (fixed AED densification, community responder apps, dispatcher-assisted CPR enhancement) in suburban settings with varying demographic and geographic characteristics?

This question requires a decision-analytic model (likely a Markov model or microsimulation) incorporating the survival functions validated in RQ1 and RQ2, local cost data from procurement processes, and health system outcome data from the NSW OHCA registry.

Proposed methodology

The research programme would proceed in three integrated phases:

Phase 1 (12 months): Spatial model development and validation. Obtain geocoded OHCA registry data through a formal research agreement with NSW Ambulance. Develop a constrained spatial optimisation model incorporating ABS SA2/mesh block geography, CASA airspace layers, Bureau of Meteorology wind and precipitation data, and fleet availability models. Validate coverage predictions against a small-scale field trial using drone flight paths in a designated test corridor.

Phase 2 (12 months): Bystander interaction study. Conduct a mixed-methods study combining structured observation of simulated AED drone retrievals at residential addresses (n=100+), semi-structured interviews with participants, and analysis of dispatcher-bystander communication recordings from international AED drone programmes (through data-sharing agreements with the Karolinska group). Develop a validated retrieval-time distribution for Australian suburban settings.

Phase 3 (12 months): Cost-effectiveness modelling. Build a decision-analytic model comparing the drone AED network against three alternative strategies using outputs from Phases 1 and 2. Conduct probabilistic sensitivity analysis across key parameters. Submit findings to a health technology assessment audience.

Significance and expected contribution

This research programme would produce three outputs with direct translation potential. First, a validated spatial optimisation framework that could be applied to any Australian metropolitan area considering drone AED deployment — tools and methodology, not just Sydney-specific recommendations. Second, the first empirical data on Australian bystander AED retrieval behaviour, which has implications beyond drone delivery for all public-access defibrillation programmes. Third, a comparative cost-effectiveness analysis that would inform health system investment decisions and potentially support MRFF or NHMRC grant applications for implementation research.

The work is positioned at the intersection of emergency medicine, operations research, human-computer interaction, and health economics. It would suit a candidate with a quantitative health services research background and access to spatial analysis and health economic modelling expertise. The feasibility model produced by the current project provides an immediate starting point and a structured evidence base for a PhD proposal.

Data classification: This appendix is an estimated research framing based on the project’s feasibility findings and the evidence catalogue (see outputs/evidence/evidence_catalogue.md).

1. Valenzuela TD, Roe DJ, Cretin S, Spaite DW, Larsen MP. Estimating Effectiveness of Cardiac Arrest Interventions. Circulation. 1997;96(10):3308–13. doi: 10.1161/01.cir.96.10.3308.

2. Larsen MP, Eisenberg MS, Cummins RO, Hallstrom AP. Predicting survival from out-of-hospital cardiac arrest: A graphic model. Annals of Emergency Medicine. 1993;22(11):1652–8. doi: 10.1016/s0196-0644(05)81302-2.

3. Waalewijn RA, Vos R de, Tijssen JGP, Koster RW. Survival models for out-of-hospital cardiopulmonary resuscitation from the perspectives of the bystander, the first responder, and the paramedic. Resuscitation. 2001;51(2):113–22. doi: 10.1016/s0300-9572(01)00407-5.

4. Weisfeldt ML, Becker LB. Resuscitation After Cardiac Arrest. JAMA. 2002;288(23):3035. doi: 10.1001/jama.288.23.3035.

5. Chan PS, Krumholz HM, Nichol G, Nallamothu BK. Delayed Time to Defibrillation after In-Hospital Cardiac Arrest. New England Journal of Medicine. 2008;358(1):9–17. doi: 10.1056/nejmoa0706467.

6. Hasselqvist-Ax I, Riva G, Herlitz J, Rosenqvist M, Hollenberg J, Nordberg P, et al. Early Cardiopulmonary Resuscitation in Out-of-Hospital Cardiac Arrest. New England Journal of Medicine. 2015;372(24):2307–15. doi: 10.1056/nejmoa1405796.

7. Nichol G, Huszti E, Birnbaum A, Mahoney B, Weisfeldt M, Travers A, et al. Cost-Effectiveness of Lay Responder Defibrillation for Out-of-Hospital Cardiac Arrest. Annals of Emergency Medicine. 2009;54(2):226–235.e2. doi: 10.1016/j.annemergmed.2009.01.021.

8. Lijovic M, Bernard S, Nehme Z, Walker T, Smith K. Public access defibrillation—Results from the Victorian Ambulance Cardiac Arrest Registry. Resuscitation. 2014;85(12):1739–44. doi: 10.1016/j.resuscitation.2014.10.005.

9. Medicine I of. Strategies to Improve Cardiac Arrest Survival. Washington, DC: National Academies Press. 2015. doi: 10.17226/21723.

10. Deakin CD, Anfield S, Hodgetts GA. Underutilisation of public access defibrillation is related to retrieval distance and time-dependent availability. Heart. 2018;104(16):1339–43. doi: 10.1136/heartjnl-2018-312998.

11. Claesson A, Bäckman A, Ringh M, Svensson L, Nordberg P, Djärv T, et al. Time to Delivery of an Automated External Defibrillator Using a Drone for Simulated Out-of-Hospital Cardiac Arrests vs Emergency Medical Services. JAMA. 2017;317(22):2332. doi: 10.1001/jama.2017.3957.

12. Schierbeck S, Hollenberg J, Nord A, Svensson L, Nordberg P, Ringh M, et al. Automated external defibrillators delivered by drones to patients with suspected out-of-hospital cardiac arrest. European Heart Journal. 2021;43(15):1478–87. doi: 10.1093/eurheartj/ehab498.

13. Schierbeck S, Svensson L, Claesson A. Use of a Drone-Delivered Automated External Defibrillator in an Out-of-Hospital Cardiac Arrest. New England Journal of Medicine. 2022;386(20):1953–4. doi: 10.1056/nejmc2200833.

14. Schierbeck S, Nord A, Svensson L, Ringh M, Nordberg P, Hollenberg J, et al. 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. 2023;5(12):e862–71. doi: 10.1016/s2589-7500(23)00161-9.

15. Liu X, Yuan Q, Wang G, Bian Y, Xu F, Chen Y. Drones delivering automated external defibrillators: A new strategy to improve the prognosis of out-of-hospital cardiac arrest. Resuscitation. 2023;182:109669. doi: 10.1016/j.resuscitation.2022.12.007.

16. Maaz M, Benjamin Leung KH, Boutilier JJ, Suen S, Dorian P, Morrison LJ, et al. Cost-effectiveness of drone-delivered automated external defibrillators for cardiac arrest. Resuscitation. 2025;209:110552. doi: 10.1016/j.resuscitation.2025.110552.

17. Coute RA, Panchal AR, Mader TJ, Neumar RW. National Institutes of Health–Funded Cardiac Arrest Research: A 10-Year Trend Analysis. Journal of the American Heart Association. 2017;6(7). doi: 10.1161/jaha.116.005239.

18. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ. 2021;n71. doi: 10.1136/bmj.n71.

19. Rethlefsen ML, Kirtley S, Waffenschmidt S, Ayala AP, Moher D, Page MJ, et al. PRISMA-S: an extension to the PRISMA Statement for Reporting Literature Searches in Systematic Reviews. Systematic Reviews. 2021;10(1). doi: 10.1186/s13643-020-01542-z.

20. Godin K, Stapleton J, Kirkpatrick SI, Hanning RM, Leatherdale ST. Applying systematic review search methods to the grey literature: a case study examining guidelines for school-based breakfast programs in Canada. Systematic Reviews. 2015;4(1). doi: 10.1186/s13643-015-0125-0.

21. Bramer WM, De Jonge GB, Rethlefsen ML, Mast F, Kleijnen J. A systematic approach to searching: an efficient and complete method to develop literature searches. Journal of the Medical Library Association. 2018;106(4). doi: 10.5195/jmla.2018.283.

22. Dixon-Woods M, Cavers D, Agarwal S, Annandale E, Arthur A, Harvey J, et al. Conducting a critical interpretive synthesis of the literature on access to healthcare by vulnerable groups. BMC Medical Research Methodology. 2006;6(1). doi: 10.1186/1471-2288-6-35.

23. Greenhalgh T, Robert G, Macfarlane F, Bate P, Kyriakidou O, Peacock R. Storylines of research in diffusion of innovation: a meta-narrative approach to systematic review. Social Science & Medicine. 2005;61(2):417–30. doi: 10.1016/j.socscimed.2004.12.001.

24. Australian Institute of Health and Welfare. Out-of-hospital cardiac arrest in Australia: annual report 2021. Canberra: Australian Institute of Health and Welfare. 2021. Available at: https://www.aihw.gov.au/.

25. NSW Ambulance. Annual Report 2022-23. Sydney: NSW Ministry of Health. 2023. Available at: https://www.ambulance.nsw.gov.au/.

26. St John Western Australia. Out-of-Hospital Cardiac Arrest Report 2024. St John WA. 2024. Available at: https://stjohnwa.com.au/.

27. Australian Bureau of Statistics. 2021 Census of Population and Housing: Greater Sydney SA2 profiles. Canberra: Australian Bureau of Statistics. 2022. Available at: https://www.abs.gov.au/census.

28. Civil Aviation Safety Authority. Part 101 Manual of Standards and Advisory Circulars: unmanned aircraft and rockets. Canberra: Civil Aviation Safety Authority. 2024. Available at: https://www.casa.gov.au/.

29. OHCAO Registry, University of Warwick. Out-of-hospital cardiac arrest outcomes, England, annual report 2022. Warwick Clinical Trials Unit. 2023. Available at: https://warwick.ac.uk/fac/sci/med/research/ctu/trials/ohcao/.

30. American Heart Association. 2020 Guidelines for Cardiopulmonary Resuscitation and Emergency Cardiovascular Care. Circulation 142, Supplement 2. 2020. Available at: https://cpr.heart.org/en/resuscitation-science/cpr-and-ecc-guidelines.

31. Zipline International. Zipline Rwanda operations: sustained medical drone logistics. Company website. 2024. Available at: https://www.flyzipline.com/rwanda.

32. Everdrone AB. First-ever delivery of defibrillator drone saves a life. Company website. 2022. Available at: https://www.everdrone.com/first-ever-delivery-of-defibrillator-drone-to-save-life/.

33. Swoop Aero. Healthcare logistics in Malawi and the Pacific. Company website. 2024. Available at: https://www.swoop.aero/.

34. UNICEF. Drones connecting communities and health systems. UNICEF Innovation. 2022. Available at: https://www.unicef.org/innovation/stories/drones-connecting-communities-and-health-systems.

35. NHS England. NHS drone trial to fly blood samples between hospitals. NHS England news. 2024. Available at: https://www.england.nhs.uk/2024/02/nhs-drone-trial-to-fly-blood-samples-between-hospitals/.

36. Kobak D, González-Márquez R, Horvát E-Á, Lause J. Delving into LLM-assisted writing in biomedical publications through excess vocabulary. Science Advances. 2025;11(27). doi: 10.1126/sciadv.adt3813.

37. Australian Bureau of Statistics. Socio-Economic Indexes for Areas (SEIFA), Australia, 2021: Index of Relative Socio-economic Disadvantage (IRSD) by Statistical Area Level 2. Canberra: Australian Bureau of Statistics. 2023. Available at: https://www.abs.gov.au/statistics/people/people-and-communities/socio-economic-indexes-areas-seifa-australia/2021.