aerospace

MDPI

Article

A Probabilistic Trajectory and Dispersion Study of Martian Tumbleweed Rover Swarms in Global Dust Storm Conditions

Emma C. Belhadfa [1,2,*] and Gabriele M. T. D'Eleuterio [2]

[1] Clarendon Laboratory, Department of Physics, University of Oxford, Oxford OX1 3PU, UK
[2] Institute for Aerospace Studies, University of Toronto, Toronto, ON M3H 5T6, Canada
* Correspondence: emma.belhadfa@physics.ox.ac.uk

Abstract

Global Dust Storms (GDSs) dominate the Martian climate, yet the mechanisms of dust lifting, transport, and deposition remain poorly constrained because in situ observations are limited: landers are stationary, and their Radioisotope Thermoelectric Generators (RTGs) contaminate local measurements. Our solution is a swarm of lightweight, wind-propelled spherical rovers—Dust Rovers (DRs)—that follow natural wind patterns to collect high-spatial-resolution in situ dust data during storms. Here, we quantify the dispersion and communication connectivity of this robotic architecture. Using 625 wind states calculated using the Mars Climate Database (MCD) for the MY34 storm at the Curiosity site, we build a bivariate-normal model of the near-surface wind velocity and characterise its diurnal structure. Rover trajectories and swarm dispersion are propagated with three integrators: the Euler and fourth-order Runge–Kutta schemes for the deterministic dynamics, and a stochastic Ornstein–Uhlenbeck formulation that simulates turbulent gust forcing. Dispersion is controlled by the unmeasured gust correlation time: with a static rolling threshold enforced, the white-gust case ($\tau_g = \Delta t = 1$ s) provides a two-rover separation below 1 m over an hour (0.44 km over a sol), whereas correlation times of 10–60 s yield 6–25 m over an hour (1.4–3.3 km over a sol). We demonstrate that a 20-rover swarm needs only a 1.6–4.1 km per-rover range to remain connected over a sol, and stays connected across the plausible envelope—for gust amplitudes up to twice the nominal value with $\tau_g \leq 60$ s, and at the nominal amplitude up to $\tau_g \approx 150$ s.

Keywords: Mars; swarm robotics; tumbleweed rover; stochastic differential equations; Ornstein–Uhlenbeck process; trajectory modelling; Martian dust storms; passive locomotion

Check for updates

Academic Editor: Deshan Meng

Received: 23 August 2026
Revised: 28 September 2026
Accepted: 28 September 2026
Published: 29 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The interannual variability of Mars' climate is dominated by aerosol forcing, in particular airborne dust. Approximately three times in 10 Mars Years (MYs), a Global Dust Storm (GDS) occurs, suspending and transporting fine dust across the entire planet for weeks or months at a time. Despite their central role in material transport and atmospheric processes, the onset, intensity, and evolution of GDSs remain difficult to forecast because the mechanisms that trigger and sustain them are poorly constrained by observation [1,2]. These storms also pose an existential hazard to surface missions. During the MY34 global storm, the solar-powered Opportunity rover was starved of power and lost [3], whereas the nuclear-powered Curiosity rover survived and returned rare in situ measurements of storm dynamics [4,5]. Curiosity's data, however, were acquired from a single fixed location, and its Radioisotope Thermoelectric Generator (RTG) interfered with the dust environment

it sought to measure. The result is a persistent observational gap: no mission to date has been designed to move with, and thrive within, a Martian dust storm while sampling it densely in space.

Closing this gap requires a platform that is (i) mobile in the strong, variable winds of a storm, (ii) spatially distributed, and (iii) minimally invasive to the dust field being measured. A robotic swarm—a decentralised, self-organising collective of relatively simple agents [6,7]—meets the second and third of these requirements naturally. The first is met by tumbleweed rovers: lightweight, wind-propelled spheres first proposed in the early 2000s [8–11] that harness ambient wind for locomotion and therefore follow, rather than fight, the migration of dust. Because they carry no active propulsion and no RTG, tumbleweeds disturb the local environment far less than a conventional wheeled rover. They can be manufactured on Earth and stowed for flight in numbers.

Previous investigations [12] have deployed similar rovers on Earth, yet the Martian environment reshapes the design problem in ways that do not transfer directly. Foremost is the atmosphere: at roughly 1.6% of Earth's density [13], a given wind speed delivers one to two orders of magnitude less propulsive drag, so a wind-driven vehicle must be far larger and lighter to move—the primary driver behind the 4 m, 20 kg baseline adopted here. Together with Mars' lower gravity, its distinct wind statistics and diurnal cycle, the absence of any positioning or relay infrastructure, and the impossibility of recovery or repair, these conditions require that both the vehicle and the swarm's dispersion behaviour be re-derived from Martian first principles rather than inherited from terrestrial designs.

The platform we analyse is a swarm of wind-propelled vehicles, which we call Dust Rovers (DRs), each carrying six equally spaced photon sensors to measure dust size, shape, and optical properties in situ. Current models cannot predict the precise starting location or onset time of a GDS [1]. Therefore, the central question for a passive swarm is not whether it can be steered to a storm, but whether, once deployed, it will remain a coherent, communicating swarm as the wind carries it. This paper answers that question quantitatively. We characterise the near-surface wind during a real storm from modelled climate data, establish the basic mobility of a DR from first principles, and propagate the resulting rover trajectories and swarm dispersion using both deterministic and stochastic numerical methods. The stochastic treatment, built on the Ornstein–Uhlenbeck (OU) process, is developed in detail because it is the natural framework for the turbulent, gust-driven forcing that ultimately determines whether the swarm holds together.

The broader engineering of the mission is not treated here. Vehicle-level aspects, including locomotion, inflation and deployment, power, and pendulum-based attitude control, have been addressed in tumbleweed concept studies, e.g., [8,9,11]. Entry, descent and landing, thermal survival, telecommunications hardware, autonomous coordination, and planetary protection are challenges common to any Mars surface mission. Earlier work examined tumbleweed group behaviour only under nominal winds [10], and recent field campaigns have advanced the hardware and begun testing swarm-coordination strategies in Mars-analogue terrain [14]—however, neither has quantified the dispersion and connectivity of such a swarm under storm conditions.

The scope of this study is thus confined to swarm dispersion and connectivity under GDS winds, and its contributions are threefold. We first separate the near-surface storm wind into a swarm-common diurnal mean and a rover-specific gust residual, such that the object of analysis is the cohesion of the swarm rather than the displacement of any single rover. We then represent the gust as a coloured (OU) process with an explicit correlation time τ_g and, as this timescale is not constrained by the available data, bound its effect on the dispersion by sweeping it across a plausible range. Finally, we obtain quantitative estimates of the two-rover separation and a network model of the swarm's communication

connectivity—to our knowledge the first graph-theory connectivity treatment for wind-driven Martian rover swarms (cf. the mission-concept discussion of [14])—from which we obtain the critical communication range for connectivity, the network's robustness to rover loss, and the storm duration over which Line-Of-Sight (LOS) relay can hold the swarm together.

2. Materials and Methods

This section presents the maturity of the tumbleweed concept and the proposed swarm, a first-principles analysis of DR mobility, the construction of a multivariate wind model, and the deterministic and stochastic methods used to propagate rover trajectories.

2.1. Platform Background: Tumbleweed Dust Rovers

The proposal is deliberately simple: lightweight, spherical rovers pushed by the wind to follow natural weather patterns [8]. Their passive locomotion and simple structure make them well-suited to mass manufacture and swarm deployment [6], and their low mass permits power and thermal solutions that do not rely on an RTG [9]. Because they naturally follow wind patterns, DRs are ideal candidates to track storms and collect data from within.

2.1.1. Science Traceability

The value of a DR swarm lies in its ability to answer key open science questions about GDSs. No platform is currently able to follow and measure a GDS in situ, across a large spatial region. However, for the swarm to successfully address these questions, it must remain a connected relay network capable of measuring and transmitting data.

Two lines of prior work approach this question without answering it. Southard et al. [10] simulated groups of tumbleweed rovers, but under nominal winds and without a communication model, so neither a storm-driven dispersion rate nor the resulting link requirement was obtained. Kingsnorth et al. [14] advanced the hardware and begun testing swarm-coordination strategies in Mars-analogue terrain, where the atmosphere is two orders of magnitude denser and relay infrastructure is available, so neither the mobility threshold nor the dispersion statistics transfer. Therefore, neither work establishes whether a passive swarm survives a GDS as a connected network.

Table 1 traces the open science questions to the need for swarm connectivity. Every objective that motivates the concept resolves, in its final column, to the same requirement: that the dispersing swarm stay in communication. Therefore, the feasibility of the whole concept of a tumbleweed swarm is gated by the connectivity established in Section 4.3, motivating the need for the analysis and simulations conducted in this study.

Table 1. Science traceability for a DR swarm. Key objectives defined across the literature are outlined, with current gaps described. We show that the ability of a DR swarm—or any other multi-vehicle platform—to deliver on its science objectives hinges on the connectivity analysis conducted here (Section 4.3).

Science Objective	Gap Under GDS Conditions	Dependence on Swarm Connectivity
Spatial structure of near-surface dust lifting and loading through storm onset and growth [15]	Landers fix a single point [4,5]; orbiters lose the surface beneath the dust column; storm timing and location are unpredictable [1,2]	A spatial map exists only if the distributed samples are geolocated and relayed as one dataset—requiring a connected swarm (Section 4.3)

Table 1. *Cont.*

Science Objective	Gap Under GDS Conditions	Dependence on Swarm Connectivity
In situ particle size distribution and shape across the storm, to constrain particle-scattering models [16]	Size and shape retrievals to date are column-integrated (orbital) or single-point (Curiosity cameras) [4]; no distributed in situ microphysics exists for a GDS	Comparison across the swarm needs simultaneous multi-point returns as it disperses, motivating hop-by-hop relay
Dust optical properties (opacity, scattering) mapped across the storm for radiative-transfer modelling [17]	Optical depth is constrained only from orbit or a single lander; dark storm interiors are poorly sampled [5]	The scientific product is the spatial field, realised only if the connected swarm returns its data
Near-surface dust transport, using the rovers as Lagrangian tracers of the storm wind field [18]	No in situ Lagrangian tracers have flown in a GDS; fixed stations yield Eulerian point data only	Trajectories inform only while positions are continuously relayed; lost connectivity truncates the transport record

2.1.2. Vehicle and Payload Configuration

Every in situ mission flown to Mars has used a traditional rover or helicopter architecture, and few have operated through a GDS. To collect the optical dust data needed to constrain particle-scattering models even in dark, extreme storm conditions, a mission must instead be able to image during the storm and from many points at once. A DR carries six equally spaced photon sensors (Figure 1). Previous quasi-static analyses established that such a vehicle can overcome surface friction and travel long distances with no active propulsion [8], and various surface treatments—fabric sails, turbine-like drag surfaces—and stowage schemes that fold or compress the vehicle for transport have been proposed to increase the number carried per unit volume [9,10]. Critically, the DR contains an internal motor–pendulum system [11] capable of generating and storing energy for small attitude corrections, ensuring the sensors remain exposed to the environment. The vehicle is otherwise propelled solely by the wind: the driving forces are the aerodynamic thrust of the wind and the friction of contact with the surface [8].

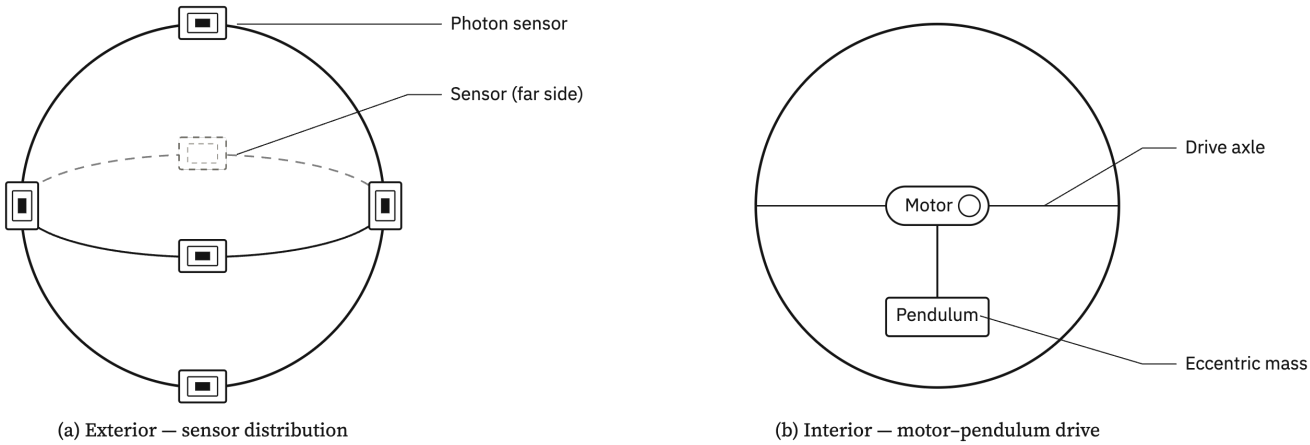

Figure 1. The proposed Dust Rover (DR): external view (**a**) and internal motor–pendulum and payload layout (**b**), adapted from [11].

2.1.3. Swarm Architecture

Although a swarm may in principle be heterogeneous, the DR mission calls for a homogeneous architecture (Figure 2). The number of agents is constrained by launch mass and

stowed volume and is assumed to be modest (<20) [8]. The three required tasks—travelling autonomously as a group, collecting in situ dust data, and communicating with the nearest neighbour—demand little dexterity, since sensing is passive. However, failure of any one task ends the mission, so high reliability is paramount: a homogeneous swarm ensures that no single DR is mission-critical, while simplifying design and manufacture and lowering early phase cost [7].

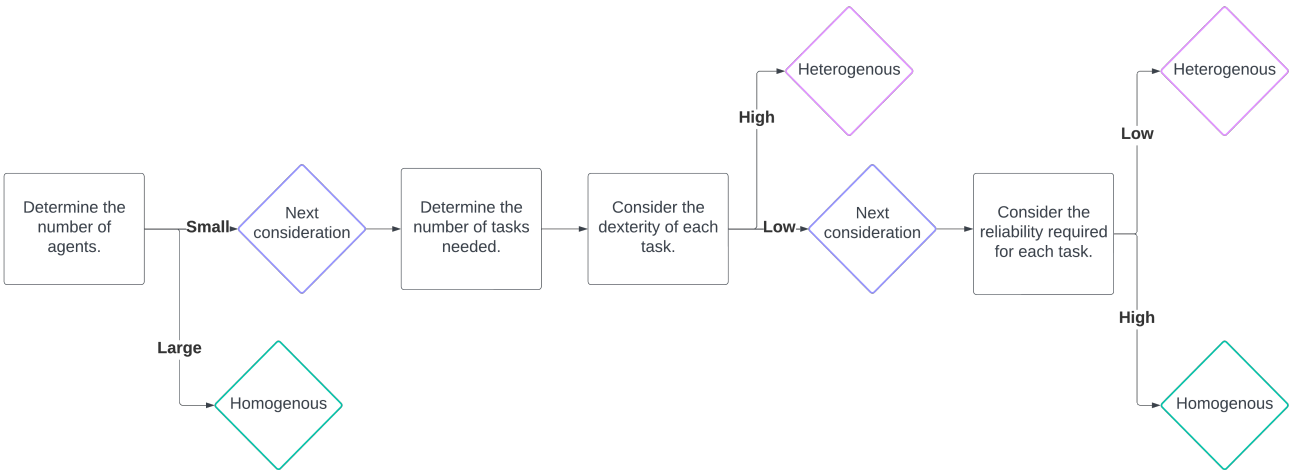

Figure 2. Trade flow for the swarm architecture decision. Rectangular boxes represent design considerations and questions posed to inform architecture selection. This diagram walks through the decisions made to select a homogeneous mission architecture.

2.1.4. Manufacturing and Storage

Previous investigations [8,9] have proposed various manufacturing and storage techniques. The most common, and best suited to this swarm, would be to manufacture the swarm on Earth and store them flat during transit. Upon insertion on Mars' surface, the tumbleweeds would inflate to their full deployed diameter, unfolding from a compact stowed volume into load-bearing spheres. A particularly attractive variant is the "Wedges" concept [8], in which the sphere is formed from several independent inflatable wedge sections arranged around a central instrument package. This format maximises the apparent diameter for a given stowed volume while providing redundancy, since the loss of a single section need not disable a vehicle rolling over the sharp, jagged Martian terrain.

2.2. Dust Rover Dynamics and Mobility

Before modelling the swarm statistically, we establish that a single DR can move and negotiate terrain under Martian conditions. The relevant atmospheric and vehicle parameters are listed in Table 2. Following prior tumbleweed studies [8,9], we model the DR as a smooth, uniform sphere of mass m and radius r, beginning at rest and driven in the direction of the wind, with drag distributed over its cross-section $A = \pi r^2$. The drag resultant is taken to act through the geometric centre, at height r above the contact point; this fixes the moment arms used in Equations (1) and (2).

Table 2. Atmospheric and dust-rover parameters used throughout this study.

Parameter	Symbol	Value	Notes
Gravitational acceleration	g	3.71 m s^{-2}	Mars surface
Atmospheric density	ρ	0.020 kg m^{-3}	Near-surface mean; ∼1.6% of Earth [13]
Drag coefficient	C_d	0.4	Sphere [8]

Table 2. *Cont.*

Parameter	Symbol	Value	Notes
Rolling-resistance coefficient	C_{rr}	0.06	Smooth sphere on regolith [9]
DR mass	m	20 kg	Margin for payload and pendulum [11]
DR diameter	d	4 m	Baseline; $r = d/2$
Cross-sectional area	A	12.57 m^2	πr^2

2.2.1. Minimum Force for Motion

A DR at rest begins to roll when the aerodynamic thrust generates a moment about the contact point exceeding the resisting moment of gravity. Rolling resistance is modelled as follows: contact deformation displaces the normal-force resultant a distance o ahead of the geometric contact point, and the rolling-resistance coefficient is defined by $C_{\mathrm{rr}} = o/r$ [9]. Taking the resultant drag to act through the centre, the DR begins to roll when the driving moment exceeds the restoring moment, $f_{\mathrm{drag}} r > mgo$, so the minimum forward force is

$$f_{\min} = \frac{mg\,o}{r} = C_{\mathrm{rr}}\, mg. \tag{1}$$

For the baseline DR, $f_{\min} = 4.5$ N is independent of diameter and shown by Antol et al. [8] to be readily supplied by the wind. Equivalently, a DR reaches a wind-relative terminal speed, $v_{\mathrm{rel}}^{\mathrm{term}}$, when drag balances rolling resistance, $\frac{1}{2}\rho C_{\mathrm{d}} A\,(v_{\mathrm{rel}}^{\mathrm{term}})^2 = C_{\mathrm{rr}} mg$, yielding $v_{\mathrm{rel}}^{\mathrm{term}} = \sqrt{2C_{\mathrm{rr}} mg/(\rho C_{\mathrm{d}} A)} = 9.4\ \mathrm{m\,s^{-1}}$.

2.2.2. Obstacle and Slope Negotiation

To traverse terrain autonomously, a DR must surmount rocks and slopes. Balancing the drag moment against the gravitational restoring moment about the top edge of a step obstacle yields a maximum surmountable obstacle height, h_{obs},

$$h_{\mathrm{obs}} = r\left(1 - \sqrt{\frac{(mg)^2}{f_{\mathrm{drag}}^2 + (mg)^2}}\right), \qquad f_{\mathrm{drag}} = \tfrac{1}{2}\rho C_{\mathrm{d}} A\, v_{\mathrm{w}}^2, \tag{2}$$

where v_{w} is the free-stream wind speed, parallel to the mean surface.

The same horizontal wind sets the maximum climbable slope. On an incline of angle θ, the drag f_{drag} projects onto the slope as $f_{\mathrm{drag}} \cos\theta$, whereas the rolling resistance acts tangentially along the slope and is proportional to the normal load, $f_{\mathrm{fric}} = C_{\mathrm{rr}} mg \cos\theta$. The rover is on the verge of climbing when the up-slope driving force balances the tangential gravity component, $f_{\mathrm{drag}} \cos\theta - C_{\mathrm{rr}} mg \cos\theta = mg \sin\theta$. Hence,

$$mg \sin\theta_{\max} = (f_{\mathrm{drag}} - C_{\mathrm{rr}} mg)\cos\theta_{\max} \tag{3}$$

which, on dividing by $mg \cos\theta_{\max}$, yields

$$\theta_{\max} = \arctan\left[\frac{\rho C_d A\, v_{\mathrm{w}}^2}{2mg} - C_{\mathrm{rr}}\right]. \tag{4}$$

Both capabilities grow with wind speed and diameter, so a DR is substantially more mobile during storms than in quiescent conditions. The obstacle capability in Equation (2) uses the free-stream wind appropriate to a stationary rover; a rover already near its rolling speed sees a reduced relative wind and therefore a smaller climbing margin. Rock size-frequency distributions measured for the Mars Exploration Rovers [19] show that obstacles below the centimetre scale are common, but that rocks exceeding it occur at metre-to-

decametre spacing at rocky sites, so arrest by rock fields is a credible competing mechanism. We therefore limit the mobility claims to unobstructed terrain.

2.3. Wind Characterisation

The trajectory of a passively driven DR is set entirely by the wind, so a faithful statistical model of the near-surface wind is the foundation of the analysis.

2.3.1. Data Set

Curiosity carries no functioning wind sensor, as the Rover Environmental Monitoring Station (REMS) wind sensor was damaged on landing and returned usable data for only about two MYs [20], before the onset of the MY34 storm. Hence, no in situ wind measurements exist at the site for this event, and a climatology model is the only practical source of a complete diurnal wind field.

Here, the near-surface wind is characterised using the Mars Climate Database v6.1 (MCD) [21], a climatological database compiled from Mars general-circulation-model simulations [22–24]. The MCD supplies model-derived wind fields: the values used here are synthetic and are neither measured nor independently validated by instruments at the surface. Therefore, the dataset is a physically plausible statistical description rather than ground truth, a limitation we return to in Section 4. We adopt the MY34 global dust storm, the most extensively documented recent global event [5], and query the database at 2 m above the Curiosity landing site (4.6° S, 137.4° E), taken as a representative low-latitude location where other instruments onboard Curiosity collected data throughout the storm. Three fields were used: the horizontal wind speed (v) and the meridional (south–north, w) and zonal (west–east, u) components, sampled hourly (25 local-time bins) across the active storm period, $160^\circ \leq L_s \leq 250^\circ$ in 25 steps, yielding $25 \times 25 = 625$ wind states. The three fields are mutually consistent by construction, i.e., $v = \sqrt{u^2 + w^2}$. This window spans the expansion, peak (i.e., planet encirclement near $L_s \approx 193^\circ$) and decay phases of the MY34 event [2]. Because pre- and post-peak states are included, the resulting mean velocity describes the storm-period in its entirety, rather than just the peak-storm; we report it as such throughout.

2.3.2. Wind-Speed Distribution

Wind speed is classically modelled by a two-parameter Weibull distribution [25], with probability density

$$p_W(v; k, \lambda) = \frac{k}{\lambda}\left(\frac{v}{\lambda}\right)^{k-1} \exp\left[-\left(\frac{v}{\lambda}\right)^{k}\right], \qquad v \geq 0, \tag{5}$$

where k is the shape parameter and λ is the scale parameter.

The parameters are estimated by maximum likelihood [26]. For a sample $\{v_i\}_{i=1}^{n}$, the log-likelihood of the two-parameter Weibull is

$$\ln L_w(k, \lambda) = n \ln k - nk \ln \lambda + (k-1) \sum_{i=1}^{n} \ln v_i - \sum_{i=1}^{n} \left(\frac{v_i}{\lambda}\right)^{k}. \tag{6}$$

Setting $\partial \ln L_w / \partial \lambda = 0$ provides the scale in closed form,

$$\hat{\lambda} = \left(\frac{1}{n} \sum_{i=1}^{n} v_i^k\right)^{1/k}, \tag{7}$$

and $\partial \ln L_w / \partial k = 0$, after eliminating λ, leaves a single equation for the shape,

$$\frac{\sum_{i=1}^{n} v_i^k \ln v_i}{\sum_{i=1}^{n} v_i^k} - \frac{1}{k} - \frac{1}{n}\sum_{i=1}^{n} \ln v_i = 0, \tag{8}$$

which is solved numerically for $\hat{k}$, whereupon $\hat{\lambda}$ follows from the closed form above [26]. Applied to the 625 wind states, this calculation yields $k = 2.72$ and $\lambda = 8.52\ \mathrm{m\,s^{-1}}$ (storm-period mean speed $7.60\ \mathrm{m\,s^{-1}}$; diurnal mean peaks at $9.9\ \mathrm{ms^{-1}}$). The direction must, however, be treated jointly with speed, as we now show.

2.3.3. Joint Structure and Diurnal Variability

Modelling horizontal wind speed v and heading ϕ as independent variables is convenient but incorrect: the wind direction on Mars is organised by a strong diurnal cycle [27]. The heading is influenced by diurnal tides: a dominant night-time regime near 120–180° with the strongest speeds ($\sim$9–10 $\mathrm{m\,s^{-1}}$) and a weaker daytime regime near -60–$-90°$. Additional thermal tides, including semi-diurnal wave-2 tides, complicate modelling of all harmonics. Here, we focus on the dominant diurnal tide and treat the system as bimodal. Therefore, speed and heading are statistically dependent, and a single Gaussian placed on the wrapped angle ϕ is a poor representation. Rather than force a distribution onto the circular variable, we model the wind's zonal and meridional components, (u, w), directly. This solution avoids the angle-wrapping pathology entirely, embeds the speed–direction dependence in a single covariance, and yields the wind vector required to force the rover dynamics.

2.3.4. Bivariate Normal Component Model

Let $\mathbf{V} = (u, w)^\top$ denote the near-surface wind velocity vector, with u and w its zonal and meridional components. We model its distribution as bivariate normal,

$$\mathbf{V} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \tag{9}$$

and estimate the mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ from the 625 MCD wind states as the sample mean and sample covariance.

This choice is supported by the central limit theorem—the aggregate of many independent atmospheric perturbations tends to normality [28]—and by the analytic tractability of the Gaussian, whose dependence structure is fully captured by $\boldsymbol{\Sigma}$ without recourse to copulas [29], which are computationally expensive. A single bivariate Weibull would require such a copula construction, which lies outside our scope. Samples are drawn efficiently through the Cholesky factor $\mathbf{L}$ of $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^\top$, via $\mathbf{V} = \boldsymbol{\mu} + \mathbf{L}\mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$.

The bivariate normal above pools all 625 states and so blends the diurnal cycle with the turbulent gusts. To isolate the fast, rover-specific gusts that drive dispersion from the slow, swarm-common diurnal trend, we bin the same 625 states by local hour. The states form a 25×25 grid in local time and solar longitude, so each of the 25 local-hour bins holds $n_h = 25$ states ($\sum_h n_h = 625$). For bin h we compute—exactly as for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ above, but per bin—the sample mean $\bar{\mathbf{V}}_h$; together, these form the diurnal mean field. The gust residual is the deviation of each state from its own hourly mean,

$$\boldsymbol{\xi}(t) = \mathbf{V}(t) - \bar{\mathbf{V}}_{h(t)}, \tag{10}$$

which therefore has zero sample mean within every bin, and hence zero mean overall, by construction; only its shape is assumed, namely Gaussian with the pooled residual covariance $\boldsymbol{\Sigma}_g = \mathrm{cov}(\boldsymbol{\xi})$. The wind is then generated, rather than decomposed, as

$$\mathbf{V}(t) = \bar{\mathbf{V}}_{h(t)} + \boldsymbol{\xi}(t), \qquad \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_g), \tag{11}$$

with $\bar{\mathbf{V}}_{h(t)}$ the mean at the local hour $h(t)$ of time t. The two descriptions are consistent views of the same data: averaging the hourly means recovers the global mean, $\boldsymbol{\mu} = \frac{1}{25}\sum_h \bar{\mathbf{V}}_h$, and the total covariance separates into between-hour (diurnal) and within-hour (gust) parts, $\boldsymbol{\Sigma} = \mathrm{cov}(\bar{\mathbf{V}}_h) + \boldsymbol{\Sigma}_g$, so $\boldsymbol{\Sigma}_g$ is the smaller, within-hour component of $\boldsymbol{\Sigma}$. Equation (11) specifies the marginal distribution of the gust at a given instant; its temporal correlation structure is left open here and is specified in Section 2.4.

2.4. Trajectory Models

With a probabilistic wind model in hand, we propagate the two-dimensional rover position $\mathbf{p} = (x, y)$ from the starting position at Curiosity's landing site, $\mathbf{p}_0 = (0, 0)$, over an hour ($\Delta t = 1$ s, $N = 3600$ steps) and a full sol ($N = 88{,}800$). Writing the rover velocity as $\dot{\mathbf{p}}$, Newton's second law provides the deterministic force per unit mass:

$$\mathbf{F}(\mathbf{V}, \dot{\mathbf{p}}) = \frac{1}{m}\left[\tfrac{1}{2}\rho C_d A \,\|\mathbf{V} - \dot{\mathbf{p}}\|(\mathbf{V} - \dot{\mathbf{p}}) - C_{rr} m g \,\frac{\dot{\mathbf{p}}}{\|\dot{\mathbf{p}}\|}\right], \tag{12}$$

so that the rover state $(\mathbf{p}, \dot{\mathbf{p}})$ evolves as a second order equation,

$$\ddot{\mathbf{p}} = \mathbf{F}(\mathbf{V}, \dot{\mathbf{p}}). \tag{13}$$

The rolling resistance direction, $\dot{\mathbf{p}}/\|\dot{\mathbf{p}}\|$, is undefined at zero velocity, so the dynamics are integrated with an explicit rest/motion state: a resting rover stays at rest until $\|F_{\mathrm{drag}}\| > f_{\min}$, after which static friction opposes the drag direction. A moving rover feels kinetic rolling resistance $f_{\min}$ opposing its velocity, with the per-step friction impulse capped at the current speed (the numerical regularisation replacing $\dot{\mathbf{p}}/\|\dot{\mathbf{p}}\|$) so it can, at most, return the rover to rest, at which point it stays if the drag is sub-threshold.

We integrate this system three ways: two deterministic integrators applied as Monte-Carlo ensembles, and one stochastic integrator based on an OU formulation. Our goal is to compare how considering stochastic effects impacts the resulting dispersion and trajectories, informing the ways in which stochastic considerations (i.e., coloured versus white noise) improve upon traditional, deterministic methods.

All simulations were implemented in Python 3.16 (NumPy/SciPy) and executed on a single Apple M3 core; ensembles are fully vectorised over rovers. The code can be found on the corresponding author's GitHub [30]. Because each Monte-Carlo member samples the distribution of a single rover's displacement, swarm dispersion is quantified by the distance between two independently drawn ($\ell = 0$) rover endpoints, $D = \|\mathbf{p}^{(i)} - \mathbf{p}^{(j)}\|$. To fully quantify the range of outcomes, we report the mean and upper percentiles of D, which are independent of ensemble size (unlike the extremal separation over n trials).

Connectivity is a property of the whole swarm rather than of a single pair, so we model the N-rover swarm as a time-varying unit-disk graph: the rovers are nodes, an edge joins any two separated by less than the communication range r_c, and the swarm is connected when the graph is a single component, so that data reach every rover by multi-hop relay. For each of $R = 2000$ Monte-Carlo swarms, we draw N trajectories from the ensemble and evaluate the graph at every recorded step (600 s). Two families of metrics follow. From the end-of-sol (most-dispersed) state, we report the critical range $r_c^\star$—the smallest r_c that connects the swarm, equal to the longest edge of its Euclidean minimum spanning tree [31] and reported as the 95th percentile over the 2000 swarms—and the algebraic connectivity λ_2, the second-smallest Laplacian eigenvalue [32] (Table 5). The time-resolved evaluation provides the fraction of the sol for which the swarm is connected and the outage statistics (Section 3.4).

2.4.1. Deterministic Integration: Euler and Runge–Kutta

The explicit Euler scheme [33] advances the state by

$$\mathbf{p}_{n+1} = \mathbf{p}_n + \Delta t\, \dot{\mathbf{p}}_n, \tag{14}$$

$$\dot{\mathbf{p}}_{n+1} = \dot{\mathbf{p}}_n + \Delta t\, \mathbf{F}(\mathbf{V}_n, \dot{\mathbf{p}}_n), \tag{15}$$

with $\mathbf{V}_n = \bar{\mathbf{V}}_{h(t_n)} + \boldsymbol{\xi}_n$ and $\boldsymbol{\xi}_n \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_g)$ as an independent gust draw at each step, such that the deterministic schemes use the same mean field and gust covariance as the stochastic scheme of Section 2.4.2 and differ from it only in the gust correlation time. Specifically, it treats gusts as uncorrelated (i.e., white) noise.

The fourth-order Runge–Kutta (RK4) scheme evaluates $\mathbf{F}$ at four sub-points per step and combines them as a weighted average [34], giving local error $O(\Delta t^5)$ against Euler's $O(\Delta t^2)$. In both, the spread of the ensemble estimates the swarm dispersion.

To enable a direct comparison across integrators, all methods use identical gust histories (i.e., matched random inputs).

2.4.2. Stochastic Model: The Ornstein–Uhlenbeck Process

The deterministic schemes resample the wind as white noise, but a turbulent gust is not instantaneous; it persists over a finite correlation time. The natural formalism for such coloured, mean-reverting forcing is a stochastic differential equation driven by a Wiener process. A standard, two-dimensional Wiener process $\mathbf{W}_t$ is defined by $\mathbf{W}_0 = \mathbf{0}$, independent increments, Gaussian increments $\mathbf{W}_t - \mathbf{W}_s \sim \mathcal{N}(\mathbf{0}, (t-s)\mathbf{I}_2)$ for $t > s$, and continuous sample paths [35]. We model the gust as an OU process with correlation time τ_g and stationary covariance $\boldsymbol{\Sigma}_g = \mathbf{L}_g \mathbf{L}_g^\top$, driven by the Wiener process $\mathbf{W}_t$,

$$\mathrm{d}\boldsymbol{\xi}_t = -\frac{1}{\tau_g}\, \boldsymbol{\xi}_t\, \mathrm{d}t + \sqrt{\frac{2}{\tau_g}}\, \mathbf{L}_g\, \mathrm{d}\mathbf{W}_t. \tag{16}$$

Each rover is driven by the total wind $\mathbf{V}(t) = \bar{\mathbf{V}}_{h(t)} + \boldsymbol{\xi}_t$, so its state $(\mathbf{p}, \dot{\mathbf{p}})$ obeys

$$\mathrm{d}\mathbf{p}_t = \dot{\mathbf{p}}_t\, \mathrm{d}t, \qquad \mathrm{d}\dot{\mathbf{p}}_t = \mathbf{F}(\bar{\mathbf{V}}_h + \boldsymbol{\xi}_t,\, \dot{\mathbf{p}}_t)\mathrm{d}t, \tag{17}$$

with $\mathbf{F}$ the force per unit mass of Equation (12). The correlation time τ_g is the single parameter that the hourly MCD data cannot constrain; its influence is quantified in Section 4.1. Setting $\tau_g = \Delta t$ does not produce a fully independent gust: under the exact transition of Equation (18), it leaves a lag-one autocorrelation of $e^{-1} \approx 0.37$, and the independent-draw limit is $\tau_g \ll \Delta t$. The deterministic Euler and RK4 integrators of Section 2.4.1 instead draw an independent gust $\boldsymbol{\xi}_n \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_g)$ at every step (Table 3), whereas the $\tau_g = \Delta t = 1$ s row of the OU sweep (Table 4) retains this small residual correlation. This is why the OU white-gust row gives ~ 1 m over an hour, while the fully independent deterministic schemes give 0.38–0.42 m.

Table 3. Two-rover separation D and mean displacement after one hour under mean storm-period wind ($n = 10{,}000$), matched-seed white-gust ($\tau_g = \Delta t = 1$ s) integrators on the static-threshold dynamics.

Method	Runtime [s]	Mean D [m]	99th Pct. [m]	Mean Displ. [m]
Euler	3.6	0.42	1.35	1.30
Runge–Kutta (RK4)	9.0	0.38	1.24	1.35

Table 4. Two-rover separation versus gust correlation time τ_g under the static-threshold baseline (hour case: mean storm wind; sol case: diurnal cycle). The white-gust discrete case ($\tau_g = \Delta t = 1$ s) is shown for reference; it is not a general lower bound but the $\tau_g = 1$ s member of the sweep, whose diffusion scales as $\sqrt{\Delta t}$. The range is extended to 300 s to bracket the loss of one-sol connectivity: the sol 99th percentile crosses the 7.4 km horizon near $\tau_g \approx 30$–35 s and the sol mean near $\tau_g \approx 300$ s. Hour mean values are rounded to the nearest metre.

τ_g [s]	Hour Mean [m]	Hour 99th [m]	Sol Mean [km]	Sol 99th [km]
1	1	2	0.44	1.32
10	6	34	1.37	4.08
20	12	71	1.94	5.80
30	17	100	2.37	7.13
40	20	129	2.74	8.22
60	25	181	3.34	10.08
100	31	275	4.29	13.04
150	35	350	5.23	15.98
200	37	413	6.02	18.48
300	39	512	7.34	22.68

Equations (16) and (17) are integrated with the exact OU transition for the gust and an explicit Euler step for the rover state [36,37]. The update is therefore

$$\boldsymbol{\xi}_{n+1} = e^{-\Delta t/\tau_g}\,\boldsymbol{\xi}_n + \sqrt{1 - e^{-2\Delta t/\tau_g}}\;\mathbf{L}_g \mathbf{z}_n, \qquad \mathbf{z}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2), \tag{18}$$

$$\dot{\mathbf{p}}_{n+1} = \dot{\mathbf{p}}_n + \mathbf{F}(\bar{\mathbf{V}}_h + \boldsymbol{\xi}_n, \dot{\mathbf{p}}_n)\Delta t, \qquad \mathbf{p}_{n+1} = \mathbf{p}_n + \dot{\mathbf{p}}_{n+1}\,\Delta t. \tag{19}$$

Because the gust of Equation (18) is advanced exactly, its stationary covariance $\boldsymbol{\Sigma}_g$ is independent of Δt, and the only discretisation error is the first-order Euler step (19). The full procedure is presented in Algorithm 1.

Algorithm 1 Stochastic integration of the DR trajectory ensemble

1. **Input:** $\mathbf{p}_0$; drift $\mathbf{F}$; gust factor $\mathbf{L}_g$; correlation time τ_g; step Δt; steps N; ensemble size n
2. **for** $j = 1$ to n (rovers, in parallel) **do**
   $\mathbf{p} \leftarrow \mathbf{p}_0, \quad \dot{\mathbf{p}} \leftarrow \mathbf{0}, \quad \boldsymbol{\xi} \leftarrow \mathbf{0}$
   for $i = 0$ to $N - 1$ **do**
   $\bar{\mathbf{V}} \leftarrow$ hourly mean wind at local hour $h(t_i)$
   $\boldsymbol{\xi} \leftarrow e^{-\Delta t/\tau_g}\boldsymbol{\xi} + \sqrt{1 - e^{-2\Delta t/\tau_g}}\;\mathbf{L}_g \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$
   $\dot{\mathbf{p}} \leftarrow \dot{\mathbf{p}} + \mathbf{F}(\bar{\mathbf{V}}_h + \boldsymbol{\xi}, \dot{\mathbf{p}})\,\Delta t; \quad \mathbf{p} \leftarrow \mathbf{p} + \dot{\mathbf{p}}\,\Delta t$
   end for
3. **end for**
4. **Output:** rover endpoints $\{\mathbf{p}^{(j)}\}_{j=1}^{n}$

Rovers close together sample much of the same flow, so we correlate the gust spatially: the forcing covariance between two rovers, a and b, is $C(\delta_{ab})\Sigma_g$, where $C(\delta) = \exp(-\delta/\ell)$ and δ is the rover separation. Co-located rovers therefore share the single-point covariance and the relative forcing vanishes as $\delta \to 0$. The correlation length ℓ is treated as a sensitivity parameter; the independent-gust case ($\ell = 0$) is retained as the conservative bound. In practice, for a swarm of N rovers, we build the $N \times N$ spatial-coherence matrix $C_{ab} = \exp(-\delta_{ab}/\ell)$ from the current pairwise separations δ_{ab}, take its Cholesky factor $\mathbf{C} = \mathbf{L}_s\mathbf{L}_s^\top$, and apply $\mathbf{L}_s$ to the per-rover white-noise draws before the single-point OU update of Equation (18) (whose factor $\mathbf{L}_g$ sets the per-rover covariance); $\mathbf{C}$ is recomputed periodically as the rovers separate. The spatial-correlation results (Table 6, Figure 10) are

therefore produced by joint N-rover swarm simulations. By contrast, the two-rover statistic D defined above is computed from independently drawn ensemble members and thus assumes independent forcing; it is used only for the $\ell = 0$ results and is not applied for $\ell > 0$.

3. Results

3.1. Wind Model

The fitted wind-speed distribution and the bivariate component model are shown in Figure 3. The Weibull fit ($k = 2.72$, $\lambda = 8.52\ \mathrm{m\,s^{-1}}$) reproduces the speed marginal. The (u, w) components are strongly anti-correlated, with a Pearson coefficient $\rho_p = -0.72$ and covariance $\mathbf{\Sigma} = \begin{bmatrix} 19.9 & -18.2 \\ -18.2 & 31.9 \end{bmatrix}\ \mathrm{m^2\,s^{-2}}$. This dependence confirms that an independent speed–direction model would be inappropriate. Removing the hourly mean leaves a gust residual with the much smaller covariance $\mathbf{\Sigma}_g = \begin{bmatrix} 2.7 & -2.2 \\ -2.2 & 4.5 \end{bmatrix}\ \mathrm{m^2\,s^{-2}}$ (standard deviations 1.6 and 2.1 $\mathrm{m\,s^{-1}}$), against a diurnal mean-speed swing from 0.65 to 9.9 $\mathrm{m\,s^{-1}}$ (Figure 4). The wind field is thus dominated by a coherent, swarm-common diurnal cycle, with comparatively small rover-to-rover gusts.

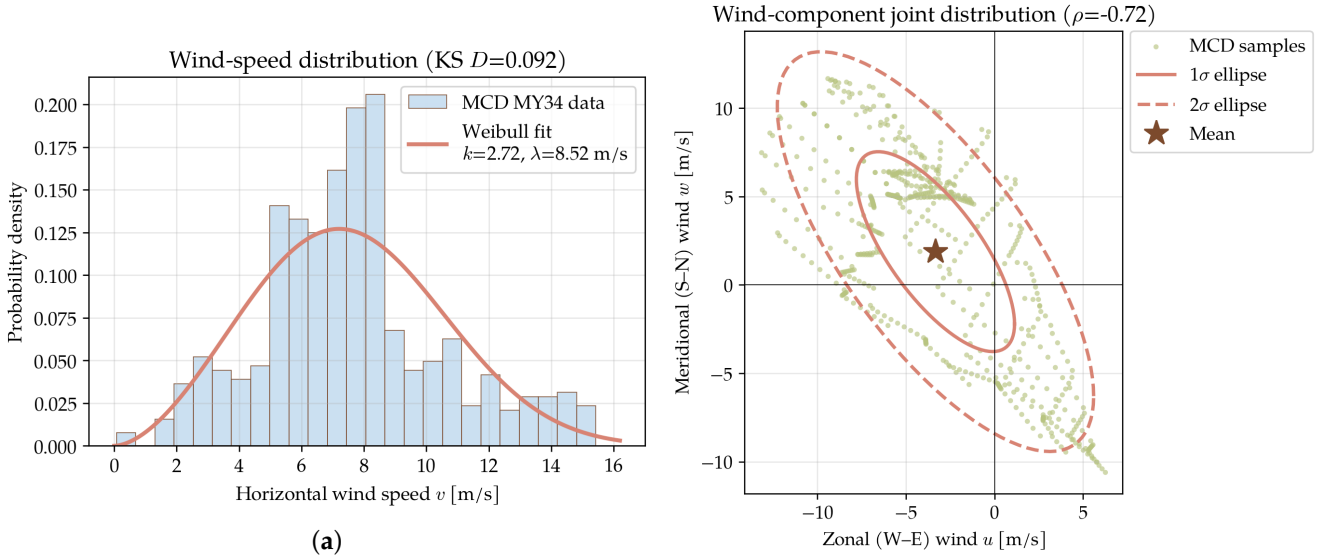

Figure 3. Wind model from MCD MY34 data at the Curiosity site. (**a**) Horizontal wind-speed histogram with maximum-likelihood Weibull fit, with a Kolmogorov–Smirnov statistic KS $D = 0.092$. Because the 625 states form a structured 25 × 25 grid in local time and L_s rather than 625 independent samples, we report D as a goodness-of-fit diagnostic and do not quote an independent and identically distributed (iid) significance level; a block bootstrap over L_s columns yields 95% intervals $k \in [2.30, 3.24]$, $\lambda \in [7.69, 9.22]$, wider than the iid parametric intervals $k \in [2.57, 2.90]$, $\lambda \in [8.27, 8.78]$. (**b**) Joint distribution of the zonal (u) and meridional (w) wind components with 1σ and 2σ bivariate-normal ellipses.

$\mathbf{\Sigma}_g$ is the residual of hourly states about the diurnal mean and is dominated by seasonal and storm-phase variation. We therefore treat $\mathbf{\Sigma}_g$ as a proxy for the unresolved variability in this model, not as a measured turbulent gust variance. Its amplitude is swept as a sensitivity parameter (Section 4.1).

The gust residual ξ departs from Gaussian (Shapiro $p \leq 10^{-16}$; excess kurtosis ≈ -1), so the bivariate-normal gust is a tractability assumption for the OU model rather than a fitted distribution. The bivariate-normal component model, not the Weibull speed marginal,

is used in the dynamics; the two are not identical (a general anisotropic correlated Gaussian does not imply a Weibull speed).

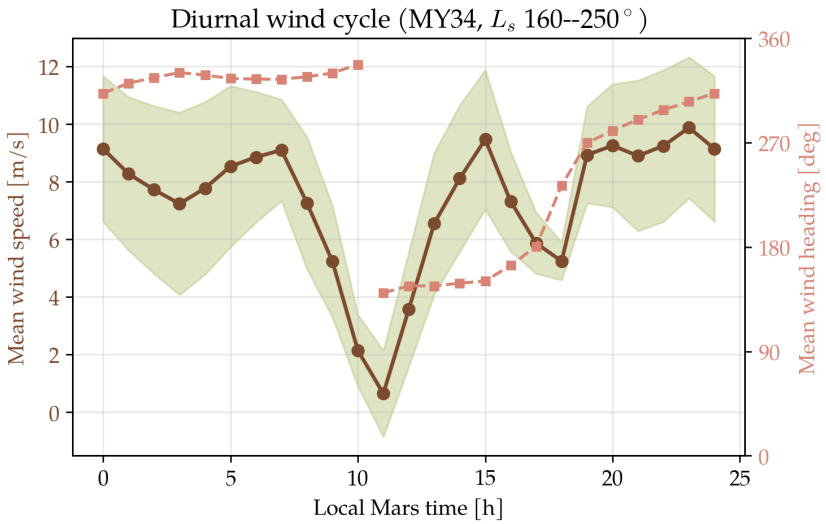

Figure 4. Diurnal cycle of the near-surface wind (MY34, L_s 160–250°): mean speed (brown) and mean heading (pink) versus local Mars time. The shaded band (green) is ±1 standard deviation of speed. The heading line is broken across the pre-noon calm (hour 11, mean speed 0.65 m s^{-1}), where the mean wind is near zero and its direction is poorly defined (i.e., the heading swings by ∼180°).

3.2. Mobility Envelope

Figure 5 evaluates Equations (2) and (4) across a spread of DR diameters (1–6 m) for representative winds. The baseline 4 m DR clears centimetre-scale obstacles—around 1 cm in typical winds and ∼7 cm in storm-strength (20 m s^{-1}) gusts—and climbs slopes of order 10° under strong winds; a larger 6 m rover clears ∼0.4 m and ∼30° at 20 m s^{-1}. Since both capabilities scale with v_w^2, storm conditions significantly enhance mobility. The diameter is also a direct mobility–stowage trade-off, with larger rovers clearing more significant obstacles. We also acknowledge that large-diameter mobility is optimistic, as areal density and inflation-structure mass are not scaled.

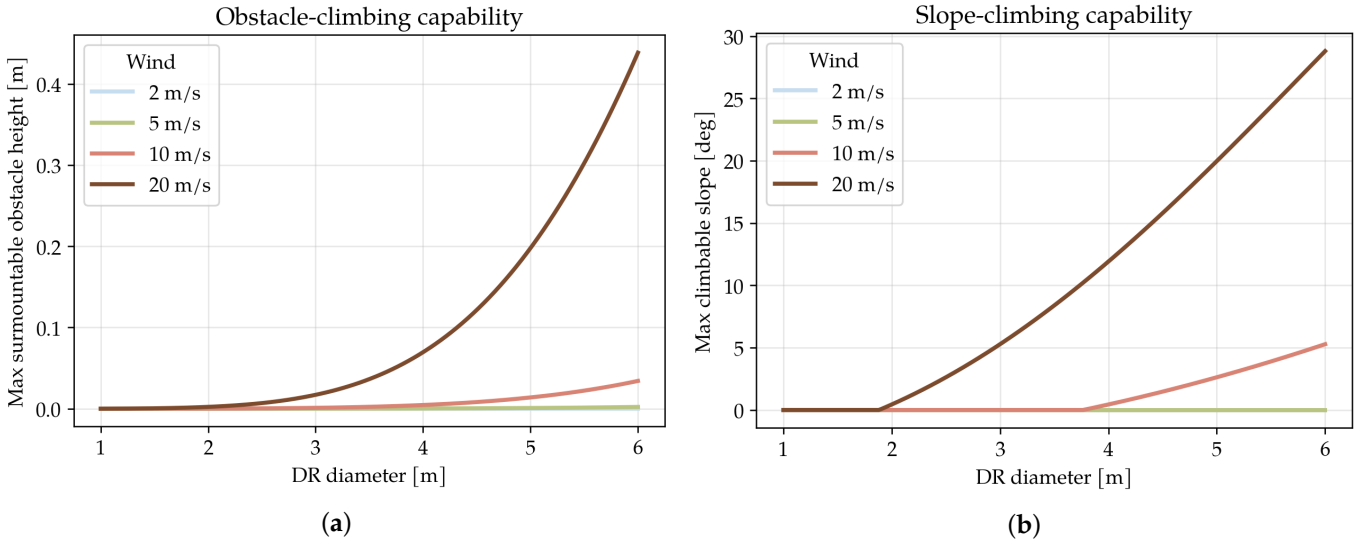

Figure 5. DR mobility versus diameter for winds of 2, 5, 10, and 20 m s^{-1}. (**a**) Maximum surmountable obstacle height. (**b**) Maximum climbable slope.

3.3. Trajectory and Dispersion Results

Hour-long ensembles (n = 10,000) under mean storm-period wind yield the two-rover separation statistics in Table 3 and produce the trajectories shown in Figure 6. The deterministic integrators, which resample the gust as white noise each second, produce mean separations of 0.42 m (Euler) and 0.38 m (RK4). The stochastic separation depends on the gust correlation time τ_g and is quantified in Table 4.

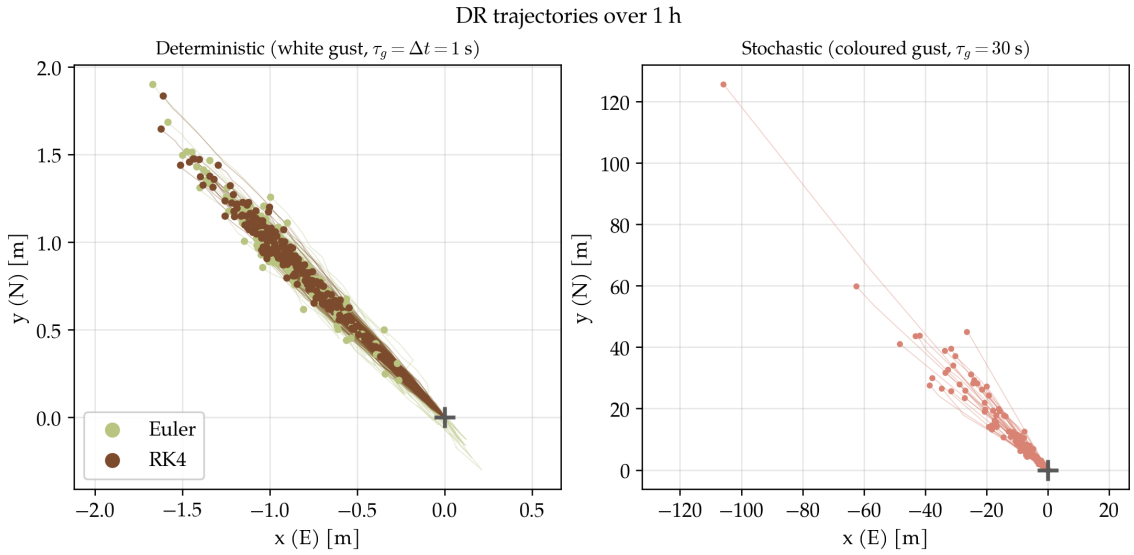

Figure 6. DR trajectory ensembles over one hour under the deterministic white-gust case for both integrators ($\tau_g = \Delta t = 1$ s, **left**) and a coloured gust ($\tau_g = 30$ s, **right**).

Over a full sol with the diurnal wind cycle (Figure 7), the deterministic drift carries the swarm ∼11 km with a characteristic looping path as the day- and night-time regimes reverse. The two-rover separation shows similar patterns as over an hour: 0.44 km on average (99th percentile 1.3 km) under the deterministic, white-gust case, rising to 1.4–3.3 km on average (99th percentile 4.1–10.1 km) for stochastic, gust correlation times of 10–60 s (Table 4).

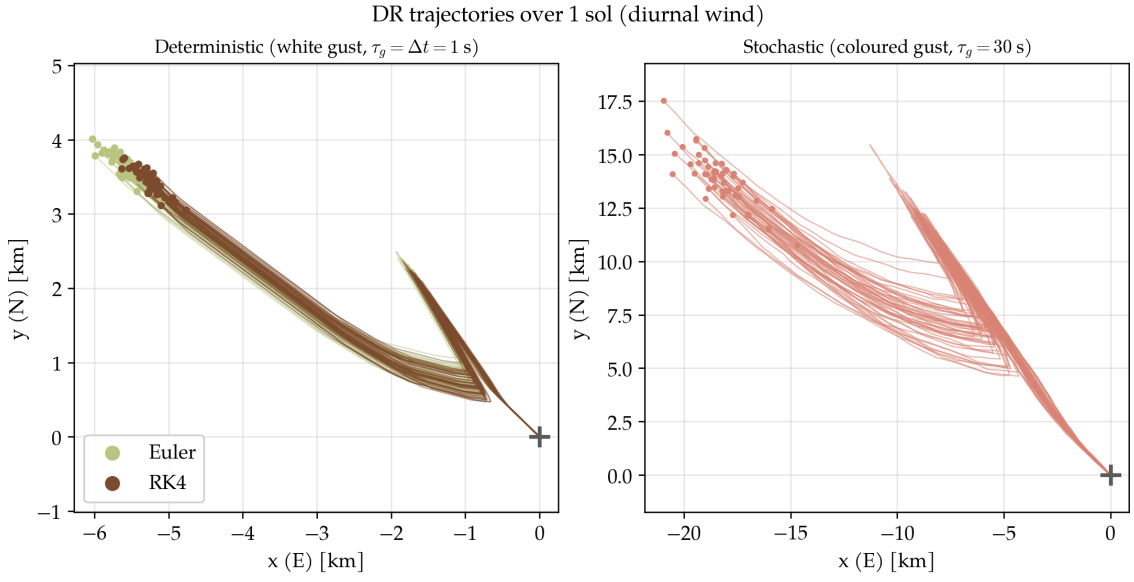

Figure 7. DR trajectories over one sol under the deterministic white-gust case for both integrators ($\tau_g = \Delta t = 1$ s, **left**) and a coloured gust ($\tau_g = 30$ s, **right**), showing the diurnal reversal.

3.4. Communication-Network Connectivity

Beyond the pairwise separation, we evaluate the connectivity of the swarm modelled as a communication graph (Section 2.4), at the most-dispersed end-of-sol state. Figure 8a presents the probability that a 20-rover swarm forms a single connected component as a function of the per-rover communication range r_c; the transition is sharp, and the critical range $r_c^\star$ (the 95th percentile of the longest minimum-spanning-tree edge) rises from 1.6 km at $\tau_g = 10$ s to 4.1 km at $\tau_g = 60$ s (Figure 8b, Table 5). In every case, we find $r_c^\star$ to be below the 7.4 km LOS horizon for one sol. Evaluated through time rather than only at the end of the sol, the swarm is connected at the 7.4 km horizon for essentially the whole sol at these correlation times, and at the 3 km design range for a fraction 1.00, 1.00, 0.99, 0.98, and 0.93 of the sol for $\tau_g = 10, 20, 30, 40$, and 60 s, respectively. At the 3 km design range and $\tau_g = 60$ s (the least-connected plausible case), 42% of swarms remain connected for the entire sol; when a disconnection occurs, the longest single outage is ∼1 h on average (95th percentile ∼6 h), which sets the onboard data-buffering requirement. At the 7.4 km LOS horizon outages are negligible.

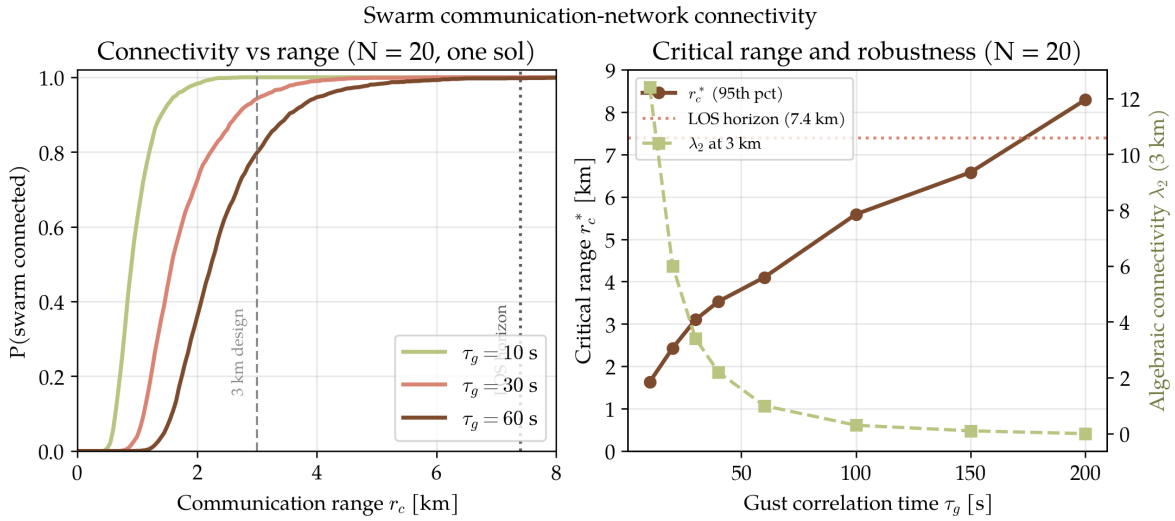

Figure 8. Swarm communication-network connectivity ($N = 20$, end of one sol). (**a**) Probability the swarm forms a single connected component versus per-rover communication range r_c, for gust correlation times $\tau_g = 10, 30, 60$ s; the dotted line marks the 7.4 km LOS horizon. (**b**) Critical range $r_c^\star$ (95th percentile) versus τ_g, against the LOS horizon.

Table 5. Swarm communication-network metrics versus gust correlation time τ_g ($N = 20$, static-threshold baseline, independent gusts $\ell = 0$). $r_c^\star$ is the 95th percentile of the longest minimum-spanning-tree edge (end of sol); λ_2 is the algebraic connectivity evaluated at the 3 km design range (at the LOS range it saturates to the complete-graph value and is uninformative); P_{conn} is the probability the swarm is a single connected component at 7.4 km and 3 km; "frac. sol" is the mean fraction of the sol connected at 3 km. $T^\star = (7.4/r_c^\star)^2$ is a heuristic diffusive extrapolation.

τ_g [s]	$r_c^\star$ [km]	λ_2 (3 km)	P_{conn} (7.4 km)	P_{conn} (3 km)	Frac. Sol (3 km)	$T^\star$ [sols]
10	1.6	12.4	1.00	1.00	1.00	21
20	2.4	6.0	1.00	0.99	1.00	9
30	3.1	3.4	1.00	0.94	0.99	6
40	3.5	2.2	1.00	0.90	0.98	4
60	4.1	1.0	1.00	0.80	0.93	3
100	5.6	0.3	0.99	0.53	0.84	2
150	6.6	0.1	0.97	0.32	0.72	1
200	8.3	0.0	0.91	0.11	0.61	<1

Reported at the 3 km design range, the algebraic connectivity falls from $\lambda_2 = 12.4$ at $\tau_g = 10$ s to 1.0 at $\tau_g = 60$ s, meaning the network thins from a dense mesh towards a sparse but connected graph as τ_g grows (Figure 9). The critical range decreases with swarm size, from 4.4 km at N = 5 to 2.8 km at N = 50 (for $\tau_g = 30$ s). Extending the sweep, $r_c\star$ reaches the 7.4 km horizon near $\tau_g \approx 180 - 200$ s, bracketing the loss of one-sol connectivity. Because $r_c\star$ grows diffusively, a heuristic extrapolation of $r_c\star \propto \sqrt{t}$ provides a connectivity horizon T of ∼21 sols at $\tau_g = 10$ s, falling to ∼3 sols at $\tau_g = 60$ s.

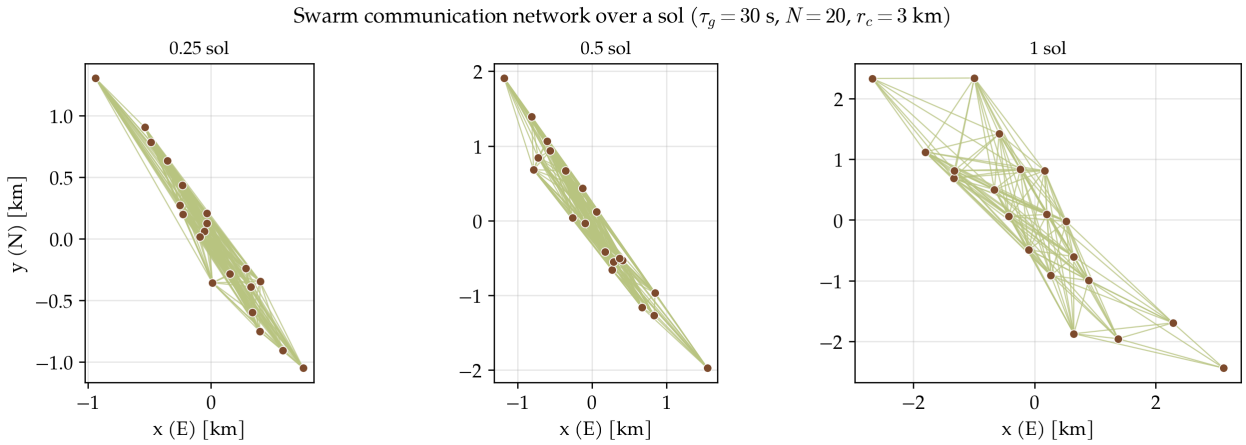

Figure 9. Communication graph of one 20-rover swarm at 0.25, 0.5, and 1 sol ($\tau_g = 30$ s, range $r_c = 3$ km), centred on the swarm centroid. Edges are viable single hops; the relay network stretches but stays connected as the swarm disperses. A 3 km design range limit is marked by a dashed line in the left panel, and λ_2 is overlayed at 3 km in the right panel.

The spatial coherence of the gust ($C(\delta) = \exp(-\delta/\ell)$; Figure 10 and Table 6) reduce the two-rover mean separation monotonically from 2.36 km ($\ell = 0$, independent) to 1.31 km ($\ell = 1000$ m) at $\tau_g = 30$ s and never reduce connectivity (P_{conn} at 7.4 km stays 1.00). The feasibility result is therefore robust to plausible spatial correlation: the mean separation falls, while the critical range shows no systematic trend and the connected fraction is unchanged (Table 6).

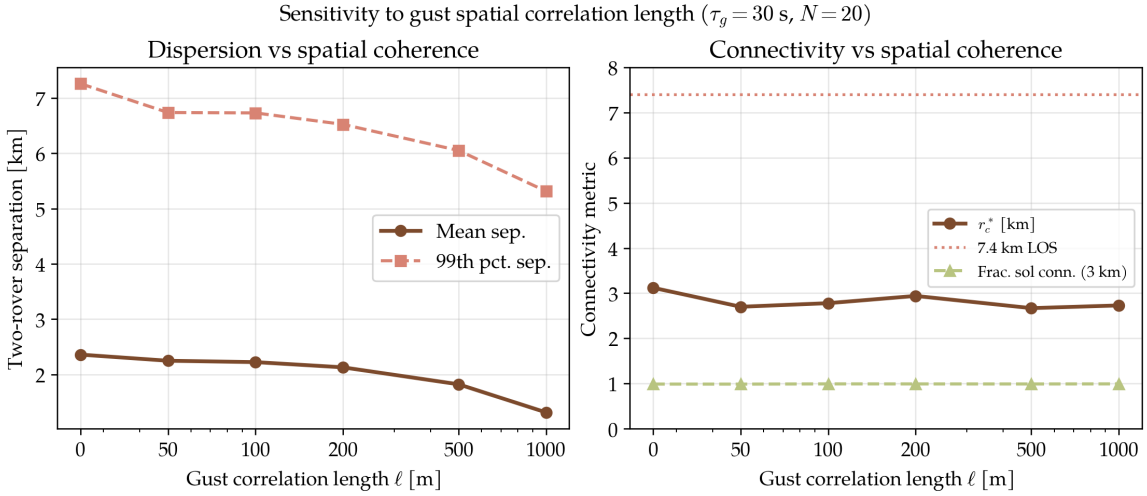

Figure 10. Sensitivity of dispersion and connectivity to the gust spatial correlation length ℓ ($\tau_g = 30$ s, $N = 20$). The gust field is assigned an exponential spatial coherence, $C(\delta) = \exp(-\delta/\ell)$, so that two rovers separated by δ share a fraction of their forcing; $\ell = 0$ is the spatially independent case, retained as the conservative bound. **Left**: Mean and 99th-percentile two-rover separation [km] versus ℓ (logarithmic axis). **Right**: Critical communication range $r_c^\star$ [km], the 7.4 km LOS horizon (dotted), and

the fraction of the sol connected at the 3 km design range, versus ℓ. Greater spatial coherence reduces relative dispersion and does not degrade connectivity (the critical range shows no systematic trend and the connected fraction stays at 0.99), so the independent-gust result is conservative.

Table 6. Sensitivity of one-sol dispersion and connectivity to the gust spatial correlation length ℓ (exponential coherence $C(\delta) = \exp(-\delta/\ell)$; $\tau_g = 30$ s, $N = 20$). $\ell = 0$ is the independent-gust conservative bound. Increasing spatial coherence reduces relative dispersion monotonically and does not degrade connectivity—the critical range shows no systematic trend (2.67–3.12 km) and the connected fraction stays at 0.99—so the feasibility conclusion is robust to plausible spatial correlation.

ℓ [m]	Mean D [km]	99th D [km]	$r_c^\star$ [km]	Frac. Sol (3 km)
0	2.36	7.26	3.12	0.99
50	2.25	6.74	2.70	0.99
100	2.22	6.73	2.78	0.99
200	2.13	6.52	2.94	0.99
500	1.82	6.05	2.67	0.99
1000	1.31	5.31	2.73	0.99

Because the gust amplitude Σ_g is itself uncertain, we report connectivity over a two-parameter envelope (amplitude scale, s, and correlation time, τ_g; Table 7 and Figure 11). The one-sol swarm stays connected within the 7.4 km horizon ($r_c\star < 7.4$ km, $P_{\text{conn}} \geq 0.95$) for all $s \leq 2$ with $\tau_g \leq 60$ s and along the baseline $s = 1$ up to $\tau_g \approx 150$ s. Connectivity is lost only when the amplitude and correlation time are simultaneously large. Tables 5 and 7 are independent Monte-Carlo runs; the 95th-percentile $r_c^\star$ carries a sampling uncertainty of $\approx\pm 0.04$ km at $\tau_g = 10$ s and $\approx\pm 0.15$ km at $\tau_g = 30$ s (95% bootstrap CI over the 2000 swarms), which accounts for the small differences between the tables (e.g., 1.6 vs. 1.7 km at 10 s and 3.1 vs. 2.9 km at 30 s).

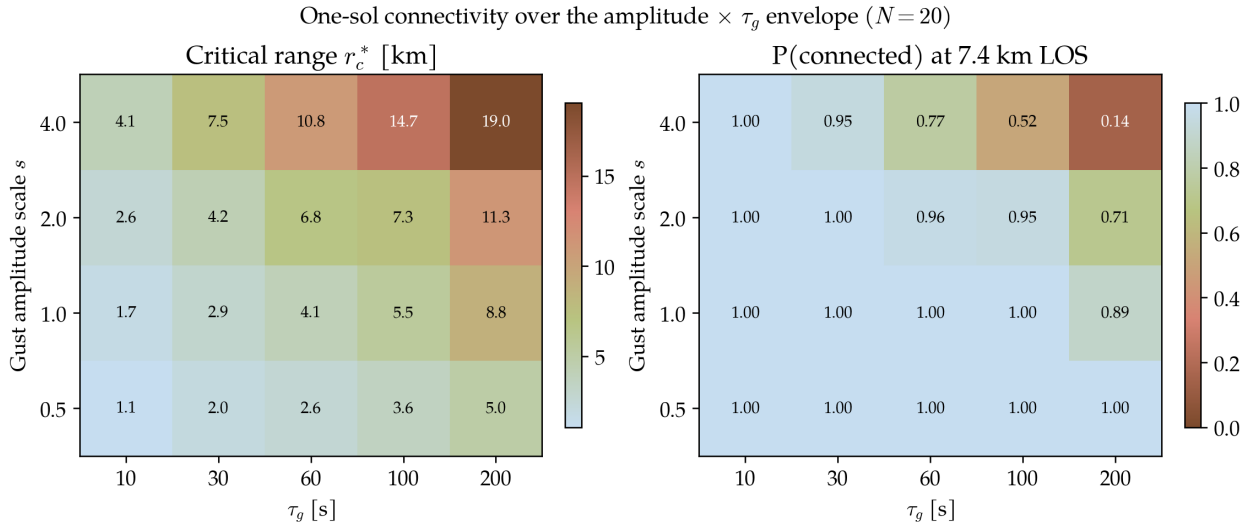

Figure 11. One-sol swarm connectivity over the gust amplitude–correlation-time envelope ($N = 20$). The stationary gust covariance is scaled by a factor s (rows), and the correlation time τ_g is varied (columns), both treated as bounded unknowns. **Left**: Critical communication range $r_c^\star$ [km]—the smallest per-rover range that connects the swarm, taken as the 95th percentile of the longest edge of the swarm's Euclidean minimum spanning tree. **Right**: Probability that the swarm forms a single connected component at the 7.4 km geometric line-of-sight (LOS) horizon. In both panels, darker cells denote poorer connectivity (larger $r_c^\star$, lower P); the swarm remains connected within the LOS horizon across most of the plausible envelope, failing only where the gust amplitude and correlation time are simultaneously large.

Table 7. One-sol connectivity over the two-parameter uncertainty envelope (amplitude scale s on Σ_g × gust correlation time τ_g; $N = 20$, $\ell = 0$). Each cell provides the critical range $r_c^\star$ [km] and, in parentheses, the probability the swarm is connected within the 7.4 km LOS horizon. The swarm stays connected ($r_c^\star < 7.4$ km, $P_{\text{conn}} \geq 0.95$) across the plausible region $s \leq 2$, $\tau_g \leq 60$ s and along $s = 1$ up to $\tau_g \approx 150$ s; connectivity is lost only for simultaneously large amplitude and long correlation time.

$s \backslash \tau_g$ [s]	10	30	60	100	200
0.5	1.1 (1.00)	2.0 (1.00)	2.6 (1.00)	3.6 (1.00)	5.0 (1.00)
1.0	1.7 (1.00)	2.9 (1.00)	4.1 (1.00)	5.5 (1.00)	8.8 (0.89)
2.0	2.6 (1.00)	4.2 (1.00)	6.8 (0.97)	7.3 (0.95)	11.3 (0.71)
4.0	4.1 (1.00)	7.5 (0.95)	10.8 (0.77)	14.7 (0.52)	19.0 (0.14)

4. Discussion

4.1. Numerical Convergence and Sensitivity Analysis

The three integrators play complementary roles. The deterministic Euler and RK4 schemes resample the gust as white noise, so successive gust bursts cancel, and separation stays small (below 1 m over an hour). The stochastic scheme integrates the coloured gust, whose correlation time τ_g is left unconstrained. We verified convergence in the time step by repeating the sol-scale simulation ($\tau_g = 30$ s) at $\Delta t = 4, 1, 0.25$, and 0.0625 s with matched random seeds and 95% Monte-Carlo confidence intervals (Table 8 and Figure 12). Across this 64× refinement, the sol-scale mean separation varies by less than 2% (Table 8). This variation across Δt is far smaller than the variation across τ_g (Table 4, a factor of several over 10–300 s): the dispersion is limited by the physical, unconstrained gust correlation time rather than by numerical resolution, which the hourly, single-point MCD data cannot resolve (Section 4.2).

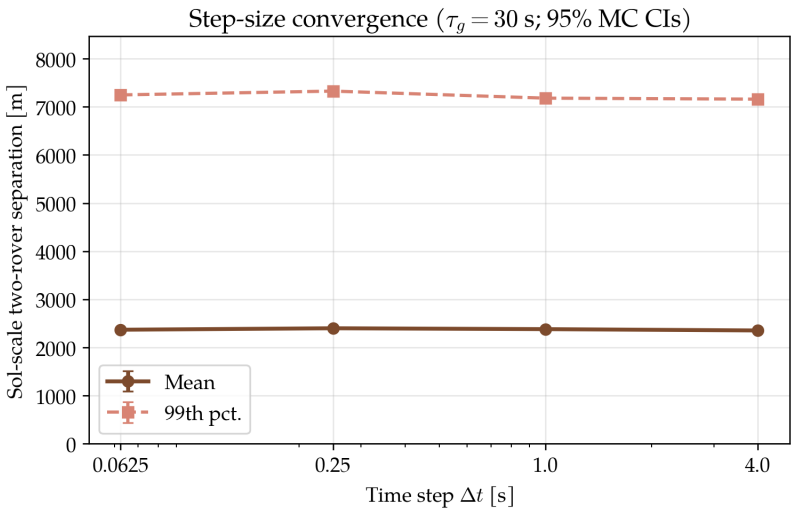

Figure 12. Step-size convergence of the sol-scale separation (coloured gust, $\tau_g = 30$ s).

Table 8. Step-size convergence of the sol-scale two-rover separation ($\tau_g = 30$ s), matched seeds, with 95% Monte-Carlo confidence intervals. Across a 64× refinement the mean varies by <2% and the 99th percentile by ∼2%, comparable to the sampling CIs, so $\Delta t = 1$ s is adequate.

Δt [s]	Steps	Sol Mean [m] (95% CI)	Sol 99th [m] (95% CI)
4.00	22,200	2354 (2348–2359)	7160 (7123–7193)
1.00	88,800	2381 (2376–2387)	7180 (7151–7212)
0.25	355,200	2399 (2394–2405)	7327 (7290–7363)
0.0625	1,420,800	2369 (2363–2375)	7248 (7197–7291)

The gust correlation time τ_g is unconstrained by the hourly data yet governs how far neighbours drift apart. Sweeping it over 10–300 s (Table 4 and Figure 13), the sol-scale two-rover separation grows approximately as $\sqrt{\tau_g}$ (from 1.4 to 3.3 km over $\tau_g = 10$–60 s, a ratio of 2.4 against $\sqrt{6} \approx 2.45$); the hour-scale mean grows faster, from 6 to 25 m (a ratio of ∼4.2), because, at the hour scale, the motion is still threshold-gated and has not reached the diffusive regime. The white-gust case ($\tau_g = \Delta t = 1$ s) gives ∼0.4 m and 0.44 km. The sol-scale 99th-percentile separation stays within the 7.4 km LOS horizon for $\tau_g \lesssim 30$–35 s, and the sol mean up to $\tau_g \approx 300$ s; at the swarm level, $r_{c\star}$ reaches the horizon near $\tau_g \approx 180$–200 s. The physically plausible band is roughly 10–100 s because the rover drag-relaxation time is $m/(\rho C_\mathrm{d} A U) \approx 21$ s at the rolling-threshold wind speed, and Martian daytime boundary-layer eddy turnover is tens of seconds to a few minutes. The 100–300 s is retained as an upper-bounding stress test.

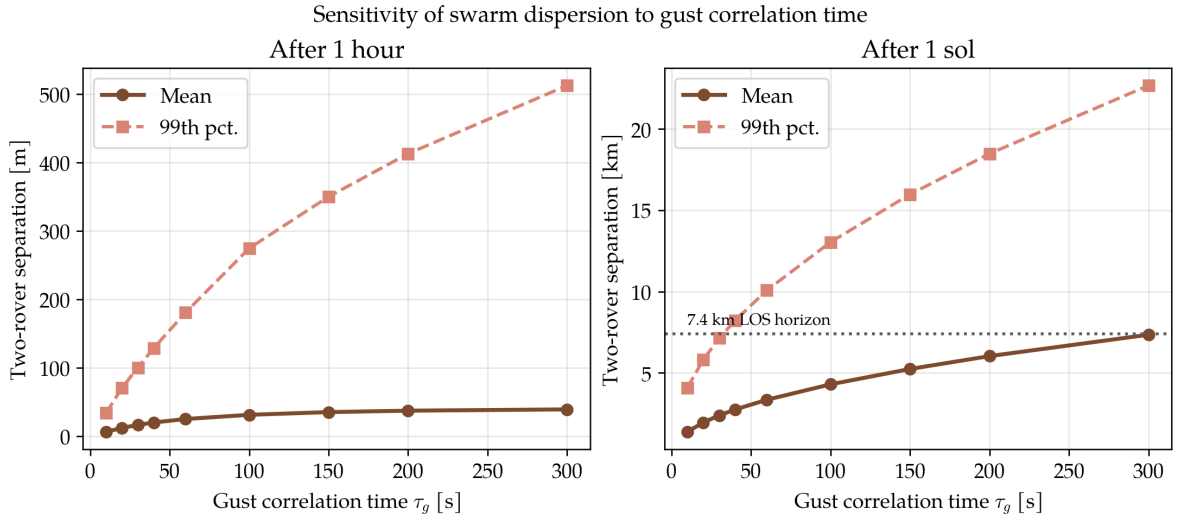

Figure 13. Sensitivity of two-rover separation to the gust correlation time τ_g, after one hour (**left**) and one sol (**right**; dotted line is the 7.4 km LOS horizon). Solid: mean; dashed: 99th percentile.

4.2. Gaussian vs. Weibull Forcing and the Limits of the OU Model

Adopting a bivariate normal distribution for the wind components, rather than an independent or copula-coupled Weibull for (v, ϕ), is what makes the OU formulation tractable: a classical OU process presumes Gaussian (Brownian) increments, so a Weibull-based forcing would require either transforming the variables into an approximately Gaussian space or generalising Equation (17) to non-Gaussian, Lévy-driven noise [38]. Both routes complicate the drift and diffusion and raise the computational burden [39], and a skewed marginal would additionally displace the mean-reversion target from the Gaussian centre. Embedding a realistic speed–direction dependence in that setting would further demand a copula [29], whereas in the Gaussian component model the dependence is carried, exactly and for free, by the off-diagonal terms of $\boldsymbol{\Sigma}$. The component representation also sidesteps the bimodality of the wrapped heading, which no single Gaussian on (v, ϕ) can capture.

The simplification that remains lies within the calibration of the model. The gust is integrated as a coloured, mean-reverting process with a finite correlation time τ_g, and the exact OU transition keeps its stationary covariance independent of the step size. The hourly MCD cannot supply a value for τ_g: it constrains the mean field $\bar{\mathbf{V}}_{h(t)}$ and the residual covariance Σ_g but not the timescale on which the residual decorrelates. We therefore treat τ_g as a bounded unknown and sweep it (Section 4.1) rather than fixing it; calibrating it

would require sub-hourly in situ wind data, such as the 1 Hz REMS record at Gale before the sensor failure [20].

4.3. Communication Feasibility

The mission requires only that each DR remain within radio range of its nearest neighbour, so that data can be relayed hop-by-hop across the swarm and onward to a lander or orbiter. In the white-gust case ($\tau_g = \Delta t = 1$ s), the sol-scale two-rover separation is negligible for this purpose (99th percentile $\sim$1.3 km; Table 4). For the plausible correlation times $\tau_g = 10$–60 s, it grows to 1.4–3.3 km on average (99th percentile 4.1–10.1 km; Table 4, Figure 13), approaching the direct LOS horizon over a single sol and, scaled diffusively over a 30–50 sol storm, reaching tens of kilometres. That horizon is $2\sqrt{2Rh_a} \approx 7.4$ km for two antennas at height $h_a = 2$ m on a sphere model of Mars, i.e., radius $R = 3390$ km.

The connectivity results of Section 3.4 show that the swarm is held together not by a single long link but by multi-hop relay, and with a wide margin: the critical range $r_c^\star = 1.6$–4.1 km is a factor of 1.8–4.6 below the 7.4 km geometric horizon over one sol, and the network is a redundant mesh. We frame this as a geometric connectivity requirement: the 7.4 km is a LOS horizon, not a demonstrated radio range. Whether a rolling, unstabilised, power-limited DR closes a few-kilometre link depends on antenna orientation, losses, and a link budget beyond this paper's scope. Terrestrial long-range radio [40] establishes only that few-kilometre ranges are attainable in principle. Terrain occlusions between rovers would further reduce the effective range below the geometric horizon. Within these caveats, the analysis here indicates that a passive swarm meets the geometric connectivity requirement over a sol with no active station-keeping.

4.4. Deployment, Non-Storm Conditions, and Collisions

While the DR swarm is designed to operate during a storm, little forecasting exists to guarantee that a landing coincides with storm onset, raising the question of how the swarm behaves before the storm arrives. A GDS grows over days to weeks from a regional core before encircling the planet [1,2], so its onset is detectable by existing orbiters well before it reaches a given site. The rovers can therefore be held stowed and deflated—immobile, with negligible aerodynamic cross-section—until a storm is detected, then inflated to deploy into the advancing front; an uninflated rover does not disperse. Quantifying non-storm dispersion directly—for instance, by repeating this analysis at the same solar longitude in a quiescent Mars year from the MCD and testing how often the winds exceed the 9.4 m s^{-1} motion threshold—would improve the deployment-timing trade and is left to future work.

Another consideration for deployment is the risk of collision between rovers. The present model treats each rover as an independent point mass and does not resolve rover–rover contact. Because neighbouring rovers share the large-scale wind effects, their relative velocities are set only by the gust differential. Simulated with spatially correlated gusts, the relative speed between neighbouring rovers near deployment is $\sim$0.07 ms^{-1} on average (95th percentile 0.36 ms^{-1}, maximum $\sim$2.1 ms^{-1}), giving impact energies below $\sim$0.7 J for 95% of contacts, assuming a reduced mass of 10 kg (two 20 kg rovers). Such low-energy contacts are unlikely to threaten swarm cohesion; their effect on dispersal is not modelled here. Beyond the first minutes, the areal density of a $\leq$20-rover swarm is low enough that contacts are rare. To further decrease the risks of collisions, a staggered release with initial spacing exceeding one rover diameter at deployment could be utilised. Resolving contact dynamics explicitly is left to future work.

4.5. Generalising MY34 GDS to Other Storms and Locations

All the wind datasets and the resulting dispersion are derived from the MY34 GDS, specifically at the Curiosity landing site (Section 2.3). Martian GDSs exhibit strong regional

wind differences across latitudes, elevations, and surface thermal inertia, so the specific separation values reported here should not be read as universal.

Two considerations nonetheless support the broader relevance of the result. First, the most recent planet-encircling storms share a common large-scale structure: Wolkenberg et al. [2] found that the MY25, MY28, and MY34 events have a similar expansion-phase duration and that the two equinoctial storms (MY25 and MY34) encircled the planet by the same season ($L_s \approx 193°$). The near-surface wind envelope at the storm peak—the regime that drives the DRs—is therefore broadly representative across equinoctial GDSs rather than specific to MY34.

Second, the differences outlined in [2] (e.g., the post-perihelion onset and longer decay of MY28, regionally distinct dust-lifting centres, and the correlation of dust opacity with atmospheric and night-time surface temperatures) imply that the thermal forcing, and hence, the surface winds, vary from site to site. These variations would shift the mean wind and gust statistics, and therefore the exact separation numbers, at a different location. Hence, a storm over a high-thermal-contrast region, or wind channelling by local topography, could plausibly amplify dispersion.

Crucially, the connectivity conclusion does not rest on the precise wind values but on two site-independent facts: (i) storm-strength winds reach tens of $\mathrm{m\,s^{-1}}$ in any GDS, well above the $\sim$9.4 $\mathrm{m\,s^{-1}}$ threshold for DR motion (Section 2.2); (ii) that threshold, and the hop-by-hop relay on which swarm connectivity depends (Section 4.3), are properties of the vehicle and the swarm architecture, not of the observing site. Because connectivity is maintained through nearest-neighbour relay rather than a single long link, a regional increase in dispersion would enlarge the swarm's overall extent without necessarily breaking that connectivity. We therefore expect the qualitative conclusion that a passive, wind-driven swarm can remain connected over a multi-sol scale to hold across sites. Repeating the analysis at a contrasting location—the northern-lowland Jezero crater, or a southern-highland site—is a natural next step once the trajectory model is coupled to site-specific MCD queries.

4.6. Limitations

Three simplifications bound the scope of these results. First, the terrain is taken as flat; real topography, including boulder fields, craters, and ridges [19], would deflect trajectories and could either concentrate or disperse the swarm. Therefore, coupling the mobility envelope of Section 2.2 to a digital elevation model is a clear next step. Previous work [41] explores terrain modelling for a similar rover construction, considering collisions with a randomised Martian rock field. While it demonstrates a tumbleweed's ability to traverse long-distances within a realistic Martian terrain model, it only focuses on a single rover. Additional work would be required to explore the dispersion and travel of a swarm within Mars' terrain. Second, the wind is sampled from the storm-peak distribution and, for the hour-long runs, from its overall mean rather than each hourly regime. We overcome this limitation as the sol-length simulation uses the full diurnal cycle. Third, although outputs are reported to metre precision for consistency with the MCD, these solutions do not imply metre-level accuracy: the MCD is itself a model constrained by limited remote sensing, without a direct in situ ground truth for these fields [22–24], so the true uncertainty is larger. None of these caveats overturns the central, order-of-magnitude finding that gust-driven dispersion remains well within the communication horizon.

As we set out in Section 1, this study addresses swarm dispersion and connectivity, rather than vehicle subsystems, e.g., delivery and landing, power, thermal control, and attitude control. Those decisions are investigated in other literature, e.g., [6,8–11].

5. Conclusions

We have quantified the dispersion and communication connectivity of a passively driven tumbleweed rover swarm under Martian GDS conditions, combining a first-principles mobility analysis, a multivariate wind model, and deterministic and stochastic trajectory simulations. A 4 m, 20 kg DR is mobile in storm winds, clearing centimetre-scale obstacles and modest slopes. Building the wind model on the bivariate-normal distribution of the velocity components—rather than an independent Weibull in speed and heading—captures the true speed–direction dependence and the diurnal structure while keeping the stochastic dynamics analytically tractable through the OU process. The deterministic and stochastic integrators describe the swarm dispersion: with a static rolling threshold enforced, two rovers separate by under 1 m over an hour (0.44 km over a sol) in the discrete white-gust case and by 6–25 m over an hour (1.4–3.3 km over a sol) for gust correlation times of 10–60 s. Modelled as a communication network, the swarm requires only a 1.6–4.1 km per-rover range to remain connected over a sol (well within the 7.4 km geometric LOS horizon) and forms a mesh that stays connected across the plausible envelope—for gust amplitudes up to twice the nominal value with $\tau_g \leq 60$ s, and at the nominal amplitude up to $\tau_g \approx 150$ s (Table 7)—and across gust spatial correlation lengths. A heuristic $\sqrt{t}$ extrapolation indicates LOS relay could suffice for ∼3–21 sols depending on the gust correlation time. Overall, this study indicates that a passive swarm can meet the geometric connectivity requirement over a sol, subject to the assumed gust correlation time ($\tau_g \lesssim 60$ s), gust amplitude, unobstructed terrain, and a radio link closing the few-kilometre critical range.

Future work could build on these results by integrating surface topography into the trajectory model, calibrating the gust correlation time from sub-hourly wind measurements to exploit the full coloured-noise capability of the OU framework, adopting higher-order stochastic schemes such as the Milstein method, and extending the propagation across the entire multi-sol life of a storm. Coupled with maturation of the DR flight system—stowage, power, thermal, and payload—further development of our model would advance the tumbleweed swarm from a promising concept towards a deployable instrument for resilient, distributed Martian science.

Author Contributions: Conceptualisation, E.C.B. and G.M.T.D.; methodology, E.C.B.; software, E.C.B.; validation, E.C.B. and G.M.T.D.; formal analysis, E.C.B.; investigation, E.C.B.; writing—original draft preparation, E.C.B.; writing—review and editing, E.C.B. and G.M.T.D.; visualisation, E.C.B.; supervision, G.M.T.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported in part by a UKRI STFC DTP Studentship and an Oxford University Clarendon Scholarship.

Data Availability Statement: The wind datasets analysed in this study were retrieved from the Mars Climate Database (https://www-mars.lmd.jussieu.fr/mars/access.html (accessed on 1 September 2026)). The simulation code and derived data generated for this study are available on the corresponding author's GitHub (@emmabelhadfa) [30].

Acknowledgments: The authors thank Lisa Romkey for her support of this work.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DR	Dust Rover
GDS	Global Dust Storm

LOS	Line-Of-Sight
MCD	Mars Climate Database
MY	Mars Year
OU	Ornstein–Uhlenbeck
REMS	Rover Environmental Monitoring Station
RK4	Fourth-Order Runge–Kutta
RTG	Radioisotope Thermoelectric Generator

References

1. Barnes, J.R. Initiation and Spread of Martian Dust Storms. In *Dust in the Atmosphere of Mars and Its Impact on Human Exploration*; Cambridge Scholars Publishing: Newcastle upon Tyne, UK, 2020.
2. Wolkenberg, P.; Giuranna, M.; Smith, M.D.; Grassi, D.; Amoroso, M. Similarities and differences of global dust storms in MY 25, 28, and 34. *J. Geophys. Res. Planets* **2020**, *125*, e2019JE006104. [CrossRef]
3. Staab, M.S.; Herman, J.A.; Reich, K.; Sridhar, V.; Nelson, R.W. MER Opportunity Dust-Storm Recovery Operations and Implications for Future Mars Surface Missions. In *Proceedings of the 2020 IEEE Aerospace Conference, Big Sky, MT, USA, 7–14 March 2020* ; IEEE: New York, NY, USA, 2020; pp. 1–20. [CrossRef]
4. Chen-Chen, H. ; Pérez-Hoyos, S.; Sánchez-Lavega, A. Dust particle size and optical depth on Mars retrieved by the MSL navigation cameras. *Icarus* **2019**, *319*, 43–57. [CrossRef]
5. Lemmon, M.T.; Guzewich, S.D.; McConnochie, T.; de Vicente-Retortillo, A.; Martínez, G.; Smith, M.D.; Bell, J.F.; Wellington, D.; Jacob, S. Large dust aerosol sizes seen during the 2018 Martian global dust event by the Curiosity rover. *Geophys. Res. Lett.* **2019**, *46*, 9448–9456. [CrossRef]
6. Iglesias, A.; Gálvez, A.; Suárez, P. Swarm robotics—A case study: Bat robotics. In *Nature-Inspired Computation and Swarm Intelligence*; Academic Press: Cambridge, MA, USA, 2020; pp. 273–302. [CrossRef]
7. Brambilla, M.; Ferrante, E.; Birattari, M.; Dorigo, M. Swarm robotics: A review from the swarm engineering perspective. *Swarm Intell.* **2013**, *7*, 1–41. [CrossRef]
8. Antol, J.; Calhoun, P.; Flick, J.; Hajos, G.A.; Kolacinski, R.; Minton, D.; Owens, R.; Parker, J. *Low Cost Mars Surface Exploration: The Mars Tumbleweed*; Technical Report NASA/TM-2003-212411; NASA Langley Research Center: Hampton, VA, USA, 2003.
9. Ylikorpi, T. A Biologically Inspired Rolling Robot for Planetary Surface Exploration. Ph.D. Thesis, Helsinki University of Technology, Espoo, Finland, 2005.
10. Southard, L.; Hoeg, T.M.; Palmer, D.W.; Antol, J.; Kolacinski, R.M.; Quinn, R.D. Exploring Mars using a group of tumbleweed rovers. In *Proceedings of the 2007 IEEE International Conference on Robotics and Automation (ICRA), Rome, Italy, 10–14 April 2007*; IEEE: New York, NY, USA, 2007; pp. 775–780. [CrossRef]
11. Forbes, J.R.; Barfoot, T.D.; Damaren, C.J. Dynamic modeling and stability analysis of a power-generating tumbleweed rover. *Multibody Syst. Dyn.* **2010**, *24*, 413–439. [CrossRef]
12. Behar, A.; Carsey, F.; Matthews, J.; Jones, J. An antarctic deployment of the NASA/JPL tumbleweed polar rover. In *Proceedings of the World Automation Congress, Seville, Spain, 28 June–1 July 2004*; IEEE: New York, NY, USA, 2004; Volume 15, pp. 453–460.
13. Williams, D.R. *Mars Fact Sheet*; NASA Space Science Data Coordinated Archive; NASA Goddard Space Flight Center: Greenbelt, MD, USA, 2024. Available online: https://nssdc.gsfc.nasa.gov/planetary/factsheet/marsfact.html (accessed on 3 August 2026).
14. Kingsnorth, J.; de Pinto Balsemão, M.; Shanbhag, A.; Pikulić, L.; Moisuc, C.; Bounova, G.; Peterson, M.; Rothenbuchner, J. A Swarm of Wind-Driven Tumbleweed Rovers for in-situ Mapping of Radiation, Water-Equivalent Hydrogen and Magnetic Fields on Mars. In Proceedings of the EPSC-DPS Joint Meeting 2025, Helsinki, Finland, 7–12 September 2025; Abstract EPSC-DPS2025-1779. [CrossRef]
15. He, F.; Rong, Z.; Wu, Z.; Gao, J.; Fan, K.; Zhou, X.; Yan, L.; Wang, Y.; Wei, Y. Martian Dust Storms: Reviews and Perspective for the Tianwen-3 Mars Sample Return Mission. *Remote Sens.* **2024**, *16*, 2613. [CrossRef]
16. Porto, C.; Cortecchia, F.; Cozzolino, F.; Franzese, G.; Mongelluzzo, G.; Esposito, F.; Cosimo Ruggeri, A.; Molfese, C.; Silvestro, S.; Ionut Popa, C.; et al. Optical design of "MicroMED", an optical particle counter to characterize Martian airborne dust. *Measurement* **2024**, *234*, 114778. [CrossRef]
17. Martikainen, J.; Muñoz, O.; Gómez Martín, J.C.; Jardiel, T.; Peiteado, M.; Caballero, A.C.; Pérez-Hoyos, S.; Sánchez Lavega, A.; Becker, T.; Wurm, G.; et al. Experimental Scattering Matrices of Martian Dust Aerosols with Narrow Particle-size Distributions. *Astrophys. J. Suppl. Ser.* **2024**, *273*, 28. [CrossRef]
18. Jackson, B.; Fenton, L.; Brown, T.; Munguira, A.; Martinez, G.; Newman, C.; Viúdez-Moreiras, D.; Golombek, M.; Lorenz, R.; Paton, M.D.; et al. Profiling Near-surface Winds on Mars Using Attitude Data from Mars 2020 Ingenuity. *Planet. Sci. J.* **2025**, *6*, 21. [CrossRef]

19. Golombek, M.P.; Haldemann, A.F.C.; Forsberg-Taylor, N.K.; DiMaggio, E.N.; Schroeder, R.D.; Jakosky, B.M.; Mellon, M.T.; Matijevic, J.R. Rock size-frequency distributions on Mars and implications for Mars Exploration Rover landing safety and operations. *J. Geophys. Res. Planets* **2003**, *108*, 8086. [CrossRef]
20. Viúdez-Moreiras, D.; Gómez-Elvira, J.; Newman, C.E.; Navarro, S.; Marín, M.; Torres, J.; de la Torre-Juárez, M. Gale surface wind characterization based on the Mars Science Laboratory REMS dataset. Part I: Wind retrieval and Gale's wind speeds and directions. *Icarus* **2019**, *319*, 909–925. [CrossRef]
21. Millour, E.; Forget, F.; Spiga, A.; Pierron, T.; Bierjon, A.; Montabone, L.; Vals, M.; Lefèvre, F.; Chaufray, J.Y.; Lopez-Valverde, M.; et al. The Mars Climate Database (Version 6.1). In Proceedings of the Europlanet Science Congress 2022, Granada, Spain, 18–23 September 2022; p. EPSC2022-786. [CrossRef]
22. Forget, F.; Hourdin, F.; Fournier, R.; Hourdin, C.; Talagrand, O.; Collins, M.; Lewis, S.R.; Read, P.L.; Huot, J.P. Improved general circulation models of the Martian atmosphere from the surface to above 80 km. *J. Geophys. Res. Planets* **1999**, *104*, 24155–24175. [CrossRef]
23. Millour, E.; Forget, F.; Spiga, A.; Vals, M.; Zakharov, V.; Montabone, L.; Lefèvre, F.; Montmessin, F.; Chaufray, J.Y.; López-Valverde, M.A.; et al. The Mars Climate Database (Version 5.3). In Proceedings of the From Mars Express to ExoMars, Madrid, Spain, 27–28 February 2018.
24. Colaïtis, A.; Spiga, A.; Hourdin, F.; Rio, C.; Forget, F.; Millour, E. A thermal plume model for the Martian convective boundary layer. *J. Geophys. Res. Planets* **2013**, *118*, 1468–1487. [CrossRef]
25. Tuller, S.E.; Brett, A.C. The characteristics of wind velocity that favor the fitting of a Weibull distribution in wind speed analysis. *J. Clim. Appl. Meteorol.* **1984**, *23*, 124–134. [CrossRef]
26. Cohen, A.C. Maximum likelihood estimation in the Weibull distribution based on complete and on censored samples. *Technometrics* **1965**, *7*, 579–588. [CrossRef]
27. Viúdez-Moreiras, D.; Lemmon, M.; Newman, C.E.; Guzewich, S.; Mischna, M.; Gómez-Elvira, J.; Herkenhoff, K.; Sánchez-Lavega, A.; de la Torre, M.; Rodríguez-Manfredi, J.A.; et al. Winds at the Mars 2020 landing site: 1. Near-surface wind patterns at Jezero crater. *J. Geophys. Res. Planets* **2022**, *127*, e2022JE007522. [CrossRef]
28. Gut, A. *An Intermediate Course in Probability*, 2nd ed.; Springer: New York, NY, USA, 2009. [CrossRef]
29. Krupskii, P.; Joe, H. Factor copula models for multivariate data. *J. Multivar. Anal.* **2013**, *120*, 85–101. [CrossRef]
30. Belhadfa, E. EmmabelhadfaTumbleweedTrajectory: Tumbleweed Trajectory Analysis, 2026. Code and Extracted MCD Wind States. Available online: https://github.com/emmabelhadfa/TumbleweedTrajectory (accessed on 27 September 2026).
31. Penrose, M. *Random Geometric Graphs*; Oxford Studies in Probability; Oxford University Press: Oxford, UK, 2003. [CrossRef]
32. Fiedler, M. Algebraic connectivity of graphs. *Czechoslov. Math. J.* **1973**, *23*, 298–305. [CrossRef]
33. Hemami, H.; Weimer, F.C. Some aspects of Euler–Newton equations of motion. *Ing.-Arch.* **1982**, *52*, 167–176. [CrossRef]
34. Cartwright, J.H.E.; Piro, O. The dynamics of Runge–Kutta methods. *Int. J. Bifurc. Chaos* **1992**, *2*, 427–449. [CrossRef]
35. Malliaris, A.G. Wiener process. In *Time Series and Statistics*; Eatwell, J., Milgate, M., Newman, P., Eds.; Palgrave Macmillan: London, UK, 1990; pp. 316–318. [CrossRef]
36. Kloeden, P.E.; Platen, E. *Numerical Solution of Stochastic Differential Equations*; Applications of Mathematics; Springer: Berlin/Heidelberg, Germany, 1992; Volume 23. [CrossRef]
37. Higham, D.J. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Rev.* **2001**, *43*, 525–546. [CrossRef]
38. Barndorff-Nielsen, O.E.; Shephard, N. Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *J. R. Stat. Soc. Ser. B* **2001**, *63*, 167–241. [CrossRef]
39. Maller, R.A.; Müller, G.; Szimayer, A. Ornstein–Uhlenbeck processes and extensions. In *Handbook of Financial Time Series*; Mikosch, T., Kreiß, J.P., Davis, R.A., Andersen, T.G., Eds.; Springer: Berlin/Heidelberg, Germany, 2009; pp. 421–437. [CrossRef]
40. Foubert, B.; Mitton, N. Long-range wireless radio technologies: A survey. *Future Internet* **2020**, *12*, 13. [CrossRef]
41. Hartl, A.E.; Mazzoleni, A.P. Terrain Modeling and Simulation of a Tumbleweed Rover Traversing Martian Rock Fields. *J. Spacecr. Rocket.* **2012**, *49*, 401–412. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.