



Universiteit
Leiden
The Netherlands

Giant galactic outflows and shocks in the cosmic web

Oei, M.S.S.L.

Citation

Oei, M. S. S. L. (2023, December 12). *Giant galactic outflows and shocks in the cosmic web*. Retrieved from <https://hdl.handle.net/1887/3666253>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3666253>

Note: To cite this publication please use the final published version (if applicable).

They are ill discoverers that think there is no land, when they can see nothing but sea.

Francis Bacon, English philosopher, *The Advancement of Learning* (1605)

7

Do luminous giants populate special large-scale environments? Or: the radio luminosity–Cosmic Web density relation for radio galaxies

M. S. S. L. Oei, R. J. van Weeren, M. J. Hardcastle, A. R. D. J. G. I. B. Gast, F. Leclercq, H. J. A. Röttgering, P. Dabhade, T. W. Shimwell, A. Botteon — *Astronomy & Astrophysics*, accepted

Abstract

CONTEXT Giant radio galaxies (GRGs, giant RGs, or giants) are megaparsec-scale, jet-driven outflows from accretion disks of supermassive black holes, and represent the most extreme pathway by which galaxies can impact the Cosmic Web around them. A long-standing but unresolved question is why giants are so much larger than other radio galaxies.

AIMS It has been proposed that, in addition to having higher jet powers than most RGs, giants might live in especially low-density Cosmic Web environments. In this work, we aim to test this hypothesis by pinpointing Local Universe giants and other RGs in physically principled, Bayesian large-scale structure reconstructions.

METHODS More specifically, we localise a LOFAR Two-metre Sky Survey (LoTSS) DR2-dominated sample of luminous ($L_\nu(\nu = 150 \text{ MHz}) \geq 10^{24} \text{ W Hz}^{-1}$) giants and a control sample of LoTSS DR1 RGs, both with spectroscopic redshifts up to $z_{\text{max}} = 0.16$, in the BORG SDSS Cosmic Web reconstructions. We measure the Cosmic Web density on a smoothing scale of $\sim 2.9 \text{ Mpc } h^{-1}$ for each RG; for the control sample, we then quantify the relation between RG radio luminosity and Cosmic Web density. With the BORG SDSS tidal tensor, we also measure for each RG whether the gravitational dynamics of its Cosmic Web environment resemble those of clusters, filaments, sheets, or voids.

RESULTS For both luminous giants and general RGs, the Cosmic Web density distribution is gamma distribution-like. Luminous giants populate large-scale environments that tend to be denser than those of general RGs. This result is corroborated by gravitational dynamics classification and a cluster catalogue cross-matching analysis. We find that the Cosmic Web density around RGs with 150 MHz radio luminosity L_ν is distributed as $1 + \Delta_{\text{RG}} | L_\nu = L_\nu \sim \Gamma(k, \theta)$, where $k = 4.8 + 0.2 \cdot \downarrow$, $\theta = 1.4 + 0.02 \cdot \downarrow$, and $\downarrow := \log_{10}(L_\nu (10^{23} \text{ W Hz}^{-1})^{-1})$.

CONCLUSIONS This work presents more than a thousand inferred Mpc-scale densities around radio galaxies, which may be correct up to a factor of order unity — except in clusters of galaxies, where the densities can be more than an order of magnitude too low. We pave the way to a future in which Mpc-scale densities around RGs are common inferred quantities, which help to better understand their dynamics, morphology, and interaction with the enveloping Cosmic Web. Our data demonstrate that luminous giants inhabit denser environments than general RGs. This underlines that — at least at high jet powers — low-density environments are no prerequisite for giant growth. Using general RGs, we quantify the relation between radio luminosity at 150 MHz and Cosmic Web density on a smoothing scale of $\sim 2.9 \text{ Mpc } h^{-1}$. This positive relation, combined with the discrepancy in radio luminosity between known giants and general RGs, reproduces the discrepancy in Cosmic Web density between known giants and general RGs. Our findings are consistent with the view that giants are regular, rather than mechanistically special, members of the radio galaxy population.

Key words: radio continuum: galaxies – galaxies: active – jets – inter-galactic medium – large-scale structure of Universe

7.1 INTRODUCTION

Supermassive Kerr black holes are key building blocks of the Universe, on galactic and cosmological scales alike. During episodes of baryon accretion, they turn into active galactic nuclei (AGN), launching winds and jets that warm and rarefy the interstellar medium (e.g. [Fabian, 2012](#); [King & Pounds, 2015](#)). These energy flows generally suppress the formation of new stars, although local star formation enhancement can occur within expanding kiloparsec-radius rings (e.g. [Dugan et al., 2017](#)). Meanwhile, jet-mediated AGN outflows — also known as radio galaxies (RGs) — can have a vast, megaparsec-scale reach, protruding from both the galaxy and its circumgalactic medium. The synchrotron radiation from RGs illuminates their jets, lobes, cocoons, and threads (e.g. [Ramatsoku et al., 2020](#)), and dominates the extragalactic radio sky. Because the behaviour of RGs links to the physics of black hole accretion and galactic winds, the pressure field and magnetisation history (e.g. [Vazza et al., 2017](#)) of the warm-hot intergalactic medium (warm-hot IGM, or WHIM), and shocks (e.g. [Nolting et al., 2019a](#)), vorticity (e.g. [Nolting et al., 2019b](#)), and cooling flows ([Fabian et al., 1984](#)) in the intracluster medium, a precise understanding of the RG phenomenon is indispensable to modern astrophysics.

A key goal of the study of RGs is to identify the main factors that determine their dynamics and to formulate models (e.g. [Scheuer, 1974](#); [Turner & Shabala, 2015](#); [Hardcastle, 2018](#)) that describe them quantitatively. One way to investigate the growth of RGs is to search for particularly large examples and to analyse what internal or external traits set them apart from the rest. Following this logic, [Oei et al. \(2022a\)](#) presented and studied *Alcyoneus*, a giant radio galaxy (GRG, giant RG, or simply *giant*) whose proper length component in the plane of the sky $l_p = 4.99 \pm 0.04$ Mpc. In general, giants are members of the RG population for which $l_p \geq l_{p,\text{GRG}}$ — where the latter is some fixed megaparsec-scale threshold — and that therefore rank among the largest RGs in existence. However, despite being one of the largest RGs known, *Alcyoneus* is not particularly luminous, and its host galaxy does not feature a particularly massive central black hole or stellar population — at least, when compared to other giants and their hosts ([Oei et al., 2022a](#)). The question naturally arises whether external properties, rather than those internal to the host galaxy, are the most important drivers of RG growth.

It is well established that the IGM resists the growth of RGs by forcing jets to convert a part of their kinetic energy into work spent to form lobe cavities (e.g. [Hardcastle & Croston, 2020](#)). The work needed to free up a cavity is the product of its volume and the local pressure.¹ Following this line of reasoning, an RG with fixed intrin-

¹As the pressure field in modern large-scale structure is not constant, as predicted by approximate

sic properties should reach a larger end-of-life extent in more tenuous (and colder) Cosmic Web (CW) environments. It has therefore been proposed that the astonishing growth of giants might be explained by their presumptive tendency to reside in tenuous parts of the Cosmic Web.

Several previous works have investigated the role of the enveloping Cosmic Web density field on GRG growth. Notably, in all cases, the authors traced Cosmic Web environments through three-dimensional galaxy positions estimated with photometric or spectroscopic redshifts. The pioneering work of [Subrahmanyam et al. \(2008\)](#) presented a case study of GRG MSH 05-22 and its Cosmic Web environment as traced by 6dF data ([Jones et al., 2004](#)). Using a sample of 12 giants and environments traced by 2dF/AAOmega data ([Sharp et al., 2006](#)), [Malarecki et al. \(2015\)](#) concluded that the lobes of giants grow in directions that avoid denser regions of the Cosmic Web. In the most comprehensive study yet, [Lan & Prochaska \(2021\)](#) used a sample of 110 giants and environments traced by DESI Legacy Imaging Surveys DR9 data ([Dey et al., 2019](#)), and did not find evidence that giants occur in more dilute environments than non-giant RGs. A major limitation of this latest study is the use of photometric redshifts as a probe of the Cosmic Web density field.

In order to determine decisively how GRG growth and the Cosmic Web relate, analyses with both more giants and more accurate Cosmic Web reconstructions seem necessary. In this work, we intend to present major developments in both regards. Firstly, the Low-Frequency Array (LOFAR; [van Haarlem et al., 2013](#)) Two-metre Sky Survey (LoTSS; [Shimwell et al., 2017](#)) has made possible the discovery of thousands of previously unknown giants in its Northern Sky imagery at observing frequency $\nu_{\text{obs}} = 144$ MHz and at resolutions $\theta_{\text{FWHM}} \in \{6'', 20'', 60'', 90''\}$. In particular, the joint search efforts of [Dabhade et al. \(2020b\)](#) in the LoTSS DR1 ([Shimwell et al., 2019](#)), those of [Simonte et al. \(2022\)](#) in the LoTSS Boötes Deep Field, and those of [Oei et al. \(2023a\)](#) in the LoTSS DR2 ([Shimwell et al., 2022](#)) have tripled the total number of known giants, which now stands at $\sim 3 \cdot 10^3$. In particular, the number of known giants in the part of the Local Universe covered by the Sloan Digital Sky Survey (SDSS; [York et al., 2000](#)) DR7 ([Abazajian et al., 2009](#)) has quintupled — a fact whose relevance will become clear in Sect. 7.2.3. Secondly, the last two decades have seen the development of principled, physics-based Bayesian inference techniques through which the three-dimensional total (i.e. baryonic *plus* dark) matter density field of the Cosmic Web can be reconstructed (e.g. [Wandelt et al., 2004](#); [Kitauro & Enßlin, 2008](#); [Jasche et al., 2010b](#); [Jasche & Kitauro, 2010b](#); [Jasche & Wandelt, 2013](#); [Jasche & Lavaux, 2019](#)). These techniques make use of the fact that the statistical

hydrostatic equilibrium, this expression only holds for sufficiently small volumes.

behaviour of the Early Universe’s total matter density field is known, as is the dominant process by which this field has evolved over cosmic time: gravity. By simulating gravity acting on Early Universe density fields and comparing the evolved fields to the observed spatial distribution of galaxies, the late-time density field can be inferred. These late-time density fields subsequently enable megaparsec-scale density measurements, including uncertainties, around individual (giant) radio galaxies. We shall use these measurements as an improved probe of the Cosmic Web density field.

Section 7.2 presents the data we use: radio galaxy observables and late-time density field reconstructions via which we probe the influence of the Cosmic Web on the growth of giants. In Sect. 7.3, we combine these data to determine megaparsec-scale densities and large-scale structure type probability distributions for giants and general RGs. Section 7.4 presents these results, alongside a quantification of the relation between RG radio luminosity and Cosmic Web density. We finally evaluate evidence for the claim that the Cosmic Web affects RG growth. In Sect. 7.5, we test the reliability of RG Cosmic Web density measurements and dynamical classifications and discuss caveats of and promising future extensions to the current work, before we present conclusions in Sect. 7.6.

For consistency with Oei et al. (2022a), Oei et al. (2023a), and Mostert & Oei (2023), we assume a flat, inflationary Λ CDM model with parameters from Planck Collaboration et al. (2020): $b = 0.6766$, $\Omega_{\text{BM},0} = 0.0490$, $\Omega_{\text{M},0} = 0.3111$, and $\Omega_{\Lambda,0} = 0.6889$, where $\Omega_{\text{DM},0} := \Omega_{\text{M},0} - \Omega_{\text{BM},0} = 0.2621$ and $H_0 := b \cdot 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. With ‘Local Universe’, we refer to the spherical region of space observed to have redshift $z < z_{\text{max}} := 0.16$. All reported redshifts are heliocentric. Terminology-wise, we strictly distinguish an RG (a radio-bright structure of plasma and magnetic fields, consisting of a core, jets, hotspots, lobes, a cocoon, and collimated synchrotron threads) from the host galaxy that has generated it. As in our previous work, we define giants to be RGs with projected proper² lengths $l_{\text{p}} \geq l_{\text{p,GRG}} := 0.7 \text{ Mpc}$. We define the spectral index α so that it relates to flux density F_{ν} at frequency ν as $F_{\nu} \propto \nu^{\alpha}$; under this convention, radio spectral indices are typically negative.

²In Cosmic Web filament environments, where giants appear most common (Sect. 7.4), lobes may expand along the Hubble flow, rendering their proper and comoving extents different. To avoid ambiguity, we stress that our projected lengths are proper instead of comoving. A less precise synonym for ‘projected proper length’ often found in the literature is ‘largest linear size’ (LLS).




Figure 7.1: Mollweide view of the sky showing the locations of all giants and LoTSS DR1 RGs in the Local Universe for which we infer Cosmic Web densities and dynamical states. The background shows the Milky Way at 150 MHz (Zheng et al., 2017), on which we overlay the LoTSS DR1 footprint (hatched dark rectangle). RGs are drawn as lemniscates of Bernoulli; we do not attempt to portray realistic morphologies or position angles. The colours represent redshifts $z \in (0, z_{\max} := 0.16)$, whilst the diameters are proportional to projected proper lengths $l_p \in (\sim 1 \text{ kpc}, 4.6 \text{ Mpc})$. Giants are translucent. Upon zooming in, the reader can appreciate the wide variety of sizes that radio galaxies can attain.

7.2 DATA

To compare the Cosmic Web environments of Local Universe giants to those of surrounding smaller radio galaxies, we combined a GRG catalogue, a general RG catalogue, and a Cosmic Web density field reconstruction. All three data sets are publicly available.

7.2.1 GIANT RADIO GALAXIES

As our source of giants we used the catalogue aggregated by [Oei et al. \(2023a\)](#), which contains 3341 giants with $l_p \geq l_{p,\text{GRG}} := 0.7 \text{ Mpc}$.³ This catalogue is intended to be complete up to (and including) September 2022. Thanks to the steradian-scale Northern Sky coverage of the LoTSS, combined with its arcsecond-scale resolution and sensitivity up to degree scales, LoTSS-discovered giants ([Dabhade et al., 2020b](#); [Simonte et al., 2022](#); [Oei et al., 2023a](#)) dominate the catalogue.

In their LoTSS DR2 manual visual search for giants, [Oei et al. \(2023a\)](#) used an angular length threshold of $\frac{1}{2}(\varphi_{\text{max}} - \varphi_{\text{min}}) = 5'$ to limit the duration of their search to a manageable few hundred hours. At the same time, this $5'$ -threshold ensured that the search would yield most giants with sufficient surface brightness in the Local Universe.⁴ This has been by design: [Oei et al. \(2023a\)](#) aimed to build a surface brightness-limited, but otherwise complete census of Local Universe giants in the LoTSS DR2 footprint with the intent of localising them within the Cosmic Web density field reconstructions presented in Sect. 7.2.3.

As elaborated upon in Sect. 7.3.1, we retained 281 giants in the part of the Local Universe where Cosmic Web analysis is possible. Of these, 260 have spectroscopic redshifts; only these giants could be reliably localised. We show their sky locations in Fig. 7.1. LoTSS DR2 discoveries make up 208 of the final 260 giants (80%). Through six example LoTSS DR2 giants, Fig. 7.2 provides the reader a sense of the quality of the radio imagery underpinning this work, from which angular lengths φ have been inferred, a sense of the reliability of our host galaxy identification, from which spectroscopic redshifts z_s have inferred, and a sense of the morphological and surface brightness diversity of the objects under consideration in this study. This publication is the first to contain images of the objects shown.

³[Oei et al. \(2023a\)](#) provide references to all discovery articles.

⁴[Oei et al. \(2023a\)](#) illustrate this point in their Fig. 9.




Figure 7.2: LoTSS DR2 cutouts at central observing frequency $\nu_{\text{obs}} = 144$ MHz and resolution $\theta_{\text{FWHM}} = 6''$, centred around giant-generating BCGs of Local Universe clusters. All giants shown are discoveries of [Oei et al. \(2023a\)](#) for which no previous images have been published. Each cutout covers a solid angle of $15' \times 15'$. Contours signify 3, 5, and 10 sigma-clipped standard deviations above the sigma-clipped median. For scale, we show the stellar Milky Way disk (with a diameter of 50 kpc) generated using the [Ringermacher & Mead \(2009\)](#) formula, alongside a 3 times inflated version. Each DESI Legacy Imaging Surveys DR9 (g, r, z) inset shows the central $3' \times 3'$ region.

7.2.2 GENERAL RADIO GALAXIES

In order to determine whether the Cosmic Web environments of giants are exceptional in any way, we must create a sample of reference Cosmic Web environments. For this reason, we also localised within the Cosmic Web a sample of general RGs, selected without regard for their length.

In particular, our starting point for this sample is the radio-bright active galactic nucleus (RLAGN) sub-sample described by [Hardcastle et al. \(2019\)](#), which contains 23,344 of the 318,520 sources (7%) in the LoTSS DR1 value-added catalogue of [Williams et al. \(2019\)](#). To construct their sub-sample, [Hardcastle et al. \(2019\)](#) combined multiple criteria that separate RLAGN from star-forming galaxies (SFGs). At the low redshifts considered in this work ($z < z_{\text{max}} := 0.16$), most RLAGN are identified through SDSS spectroscopy, and we therefore expect high RLAGN completeness and minimal contamination from SFGs.

As elaborated upon in Sect. 7.3.1, we retained 1870 LoTSS DR1 RGs in the part of the Local Universe where Cosmic Web analysis is possible. Of these, 1443 have spectroscopic redshifts; only these RGs could be reliably localised. We show their sky locations in Fig. 7.1. The LoTSS DR1 image quality is very similar to the LoTSS DR2 image quality, which Fig. 7.2 illustrates.

7.2.3 COSMIC WEB LATE-TIME DENSITY FIELD

To localise RGs within the Cosmic Web, we used data products from the Bayesian Origin Reconstruction from Galaxies (BORG; [Jasche & Wandelt, 2013](#)) SDSS run ([Jasche et al., 2015](#)). The BORG SDSS uses second-order Lagrangian perturbation theory (2LPT; [Bouchet et al., 1995](#)) to forward model structure formation and evaluates the plausibility of a proposed structure formation history by comparing its late-time density field to the three-dimensional positions of galaxies in the SDSS DR7 Main Galaxy Sample (MGS; [Abazajian et al., 2009](#)). As a result, the BORG SDSS provides a late-time total⁵ matter density field posterior for the part of the Local Universe covered by the SDSS DR7 footprint. The posterior is represented by a Hamiltonian Monte Carlo (HMC; [Duane et al., 1987](#)) Markov chain of approximately ten thousand samples after the burn-in phase. Each sample covers the same volume of $(750 \text{ Mpc } h^{-1})^3$ extent with a cubical grid of 256^3 voxels. Thus, the side length of a BORG SDSS voxel $L = \frac{1}{256} \cdot 750 \text{ Mpc } h^{-1} \approx 2.9 \text{ Mpc } h^{-1}$. The late-time posterior can be compactly summarised by taking the mean and standard deviation (SD) of the samples on a per-voxel basis. In this work, we will use both the individual samples and

⁵The BORG algorithm does not differentiate between baryonic and dark matter, assuming identical behaviour on the multi-megaparsec scale.

the summary cubes, the latter of which we shall call the BORG SDSS mean and SD. However, by using these summary cubes, information contained in higher moments of single-voxel posteriors and inter-voxel correlations remains unused.

Leclercq et al. (2015) have extended the BORG SDSS by calculating, for each sample, the Cosmic Web classification as stipulated by the T -web definition (Hahn et al., 2007). The result is a probabilistic classification with a marginal distribution for each voxel. As one aspect of a broader information theoretic analysis, Leclercq et al. (2016) demonstrated how these classifications can be used to predict galaxy properties, such as their $g - r$ colours. We now use these classifications to characterise the dynamical environments of luminous giants and general RGs.

Both the BORG SDSS posterior and a posterior from another BORG run, the BORG 2M++ (Jasche & Lavaux, 2019), have been used before to relate properties of active galaxies to the density of the enveloping Cosmic Web (Frank et al., 2016; Porqueres et al., 2018). Our work is the first to relate properties of radio galaxies to BORG (or BORG-like) Cosmic Web reconstructions.

7.3 METHODS

As pointed out in Sect. 7.1, the number of known giants has increased substantially in recent years. For example, within the Local Universe covered by the SDSS DR7, the manual visual search of Oei et al. (2023a) alone has quintupled the number of known giants with spectroscopic redshifts — from 52 to 260. Meanwhile, the BORG SDSS now offers the first physically principled, probabilistic reconstruction of the total matter density field over this entire volume. In this section, we combine both advances by describing the localisation of giants and general RGs within the BORG SDSS.

7.3.1 BORG SDSS LOCALISATION PROCEDURE

For each RG in our two samples (be it a giant or a general RG), we first transformed right ascension, declination, and redshift into a vector with comoving coordinates $\mathbf{r}_i := [x_i, y_i, z_i]^\top$ following the coordinate system convention of Jasche et al. (2015). We note that this transformation is cosmology-dependent. The BORG SDSS adopts the cosmological parameters of Jasche et al. (2015) to convert SDSS DR7 MGS redshifts into radial comoving distances and subsequently infer the structure formation history of the Local Universe. In order to obtain valid localisations, it is imperative that we use the same conversion between redshift and radial comoving distance for

our RGs. We therefore adopted the [Jasche et al. \(2015\)](#) cosmology for this particular procedure. Afterwards, we simply associated each RG to the voxel nearest to $\mathbf{r}_i$.

After localising an RG to a voxel, we verified whether the late-time density field at this location is sufficiently constrained by the BORG SDSS. More precisely, we evaluated whether — at the RG’s voxel — the survey response operator of [Jasche et al. \(2015\)](#)’s lowest r -band absolute magnitude bin ($-21.00 < M_r < -20.33$), R^0 , equals or exceeds some threshold $R_{\min}^0$. If indeed $R^0(\mathbf{r}_i) \geq R_{\min}^0$, we retained the RG for Cosmic Web analysis; if not, we discarded it. In this work, we chose $R_{\min}^0 = 0.1$. Higher choices for $R_{\min}^0$ rid the sample of comparatively uninformative (i.e. more prior-dominated) density and dynamical state measurements, but come at the cost of a reduced sample size.

Next, in order to obtain density distributions for each retained RG, we explored two methods.

FIXED VOXEL METHOD

In the ‘fixed voxel method’, we considered for each RG only the *marginal* posterior density distribution at its voxel — i.e. the distribution for that voxel in isolation — despite the fact that the BORG SDSS provides reconstructions with complex inter-voxel density correlations. Moreover, we summarised the voxel’s marginal distribution through the mean and SD only. This method is the simplest of the two methods we have used.

In Fig. 7.3, we show six example localisations of giants in the Local Universe. The giants in the upper two panels are pinpointed to galaxy clusters, the giants in the middle two panels are pinpointed to filaments, and the giants in the bottom two panels are pinpointed to sheets.

FLEXIBLE VOXEL METHOD

There are at least two issues associated with the fixed voxel method, though. One issue is the fact that the BORG SDSS has inferred the Cosmic Web with a limited set of bright SDSS DR7 galaxies, which (amongst other factors) causes reconstruction uncertainty. In practice, this means that a given cluster or filament may morph and wiggle around in different BORG SDSS samples. If one does not wiggle around the voxel to sample from accordingly (but sticks with the same voxel all the time), one regularly samples outside of the cluster or filament in which the RG of interest resides. This, of course, biases the inferred densities low.

On top of Cosmic Web reconstruction uncertainty, host galaxy peculiar motion uncertainty leads to additional difficulty in the determination of an RG’s Cosmic




Figure 7.3: Example localisations of giant radio galaxies within the large-scale structure of the Local Universe. The top row shows two cluster giants, the middle row shows two filament giants, and the bottom row shows two sheet giants. For each giant, we show a slice of constant Cartesian comoving z through the late-time BORG SDSS posterior mean total matter density field and a LoTSS DR2 $6''$ image at $\nu_{\text{obs}} = 144$ MHz (inset). Outside of the SDSS DR7–constrained volume, the posterior mean tends to the Universe’s late-time mean total matter density $\bar{\rho}_0$. The locations of the giants are marked by white circles.

Web density. In Appendix 7.A1, we show that redshift uncertainties generally lead to sub-voxel localisation uncertainty when spectroscopic redshifts are used, even after taking low-mass galaxy cluster-like peculiar motion into account. For high-mass galaxy clusters however, peculiar motion can cause localisation errors of several voxels. Even more alarmingly, we show that photometric redshift uncertainties cause localisation uncertainties of 10^1 – 10^2 Mpc, or up to tens of voxels, that are unworkably large. Cosmic Web localisation therefore only seems possible for RGs with spectroscopically detected hosts.

To counteract the fixed voxel method’s tendency to sample densities outside of the clusters and filaments in which our RGs truly reside, we propose a flexible voxel method. In this method, the voxel considered hitherto serves as a reference voxel, around which we search for the most likely correct voxel to sample from. More precisely, we iterate over 1000 BORG SDSS samples equally spaced within the MCMC, and consider for each sample all voxels in a sphere of radius 5 Mpc h^{-1} . We then simply adopt the voxel with the highest density as the most likely correct voxel for that sample. This procedure encapsulates our prior knowledge that the massive ellipticals which give rise to radio galaxies are more likely to form in a high-density region than in a low-density region of the same extent.

The search radius of 5 Mpc h^{-1} chosen here is arbitrary to some degree. Clearly, for the flexible voxel method to be any different from the fixed voxel method, this radius must exceed the side length of a single voxel, 2.9 Mpc h^{-1} . For larger radii, we are able to correct for larger peculiar motion errors, and thus provide more accurate densities for RGs located in massive galaxy clusters. At the same time, for larger radii, we are at risk of straying too far from the reference voxel; for example, this could lead to sampling cluster-like densities for RGs that actually reside in an adjacent filament. Because the majority of RGs appears to reside in filaments, and because the gravity solver of the BORG SDSS significantly limits the usefulness of cluster densities anyways (see Sect. 7.5), we chose a relatively ‘small’ search radius of 5 Mpc h^{-1} — less than two voxels in each direction.

In Fig. 7.4, we compare the mean Cosmic Web densities inferred for all RGs considered in this work, with fixed voxel method densities on the horizontal axis, and flexible voxel method densities on the vertical axis. Flexible voxel-based relative densities are typically a factor two higher. The methods agree most often at the lowest redshifts, where the BORG SDSS Cosmic Web reconstructions are least uncertain.




Figure 7.4: Comparison between Cosmic Web densities of giants and LoTSS DR1 RGs in the Local Universe, inferred via two variations of our BORG SDSS-based method. The horizontal axis represents densities as measured through the fixed voxel approach of Sect. 7.3.1, while the vertical axis represents densities as measured through the flexible voxel approach of Sect. 7.3.1.

7.3.2 BORG SDSS LOCALISATION IN PRACTICE

As described in Sect. 7.3.1, for each giant and LoTSS DR1 RG in our samples, we measured the mean and SD of its marginal posterior density RV. Of 281 giants that lie within the part of the BORG SDSS volume where $R^0(\mathbf{r}) \geq R_{\min}^0$, there are 260 with a spectroscopic redshift (93%), which were therefore suitable for Cosmic Web analysis. Of these, 208 (80%) are LoTSS DR2 discoveries (Oei et al., 2023a). In exactly the same way, of the 1870 LoTSS DR1 RGs that lie within the constrained BORG SDSS volume, we retained 1443 specimina with spectroscopic redshifts (77%), which we selected for Cosmic Web analysis.

At the BORG SDSS resolution of $2.9 \text{ Mpc } h^{-1}$ per voxel side, the baryonic matter density field approximately equals the dark matter density field scaled down by a fac-

tor $\frac{\Omega_{\text{BM},0}}{\Omega_{\text{DM},0}}$. The BORG SDSS does not distinguish between these two fields. Instead, it provides the sum of the baryonic and dark matter density field: the total matter density field. In this article, we shall exclusively mention relative total matter densities — total matter densities divided by today’s cosmic mean total matter density $\Omega_{\text{M},0}\rho_{\text{c},0}$, where $\rho_{\text{c},0}$ is today’s critical density. To be less verbose, we shall just refer to ‘relative densities’, and write $1 + \delta$, where δ denotes overdensity in the usual sense.

7.4 RESULTS

In this section, we present the first empirical distributions of the Cosmic Web total matter density and dynamical state around giants and general RGs. We also determine the radio luminosity–Cosmic Web density relation, and explore whether it can cause the density distribution discrepancy between observed giants and general RGs.

7.4.1 COSMIC WEB DENSITY DISTRIBUTIONS

The marginal relative density distributions of individual RGs resemble lognormal distributions, irrespective of whether the fixed or flexible voxel method is used to determine them. We demonstrate this in Fig. 7.5 through example marginals for a typical giant and a typical LoTSS DR1 RG.

Indeed modelling $1 + \Delta_{\text{GRG,obs}} \mid 1 + \Delta_{\text{GRG}} = 1 + \delta_i \sim \text{Lognormal}(\mu, \sigma^2)$, we can succinctly summarise each giant’s measured density distribution with two parameters. We provide these, for 50 out of 260 giants, in Table 7.1. (For access to such data for all giants, and for similar data on LoTSS DR1 giants, see its footnote.)

By aggregating just the means of these marginal relative density distributions, we can analyse the distributions for our observed populations as a whole. In green, Fig. 7.6 shows kernel density estimated (KDE) relative density distributions for both giants (top panels) and the broader population of RGs (bottom panels) in the Local Universe. The panels in the left column represent the fixed voxel method, whilst the panels in the right column represent the flexible voxel method. More precisely, these KDE distributions approximate the distributions of the observed GRG relative density RV $1 + \Delta_{\text{GRG,obs}}$ and the observed RG relative density RV $1 + \Delta_{\text{RG,obs}}$.

We sought to summarise these distributions parametrically. After testing various two-parameter distributions for continuous, non-negative RVs (such as the gamma distribution and the lognormal distribution), the KDE distributions of $1 + \Delta_{\text{RG,obs}}$ in the bottom panels — which are based on all 1443 selected LoTSS DR1 RGs — appeared best approximated by a gamma distribution.⁶ Maximum likelihood estima-

⁶In Appendix 7.A2, we provide an astrophysical–statistical argument that could explain the




Figure 7.5: Distributions for the measured relative total matter density RVs $1 + \Delta_{\text{GRG,obs}} | 1 + \Delta_{\text{GRG}} = 1 + \delta_i$ and $1 + \Delta_{\text{RG,obs}} | 1 + \Delta_{\text{RG}} = 1 + \delta_i$ for an individual giant (top row; solid), and for an individual LoTSS DR1 RG (bottom row; hatched). Fixed voxel method densities (left column; green) are lower than flexible voxel method densities (right column; orange). The MLE-fitted lognormal PDFs (solid curves) demonstrate that the distributions are approximately lognormal.

tion (MLE) suggested $1 + \Delta_{\text{RG,obs}} \sim \Gamma(k = 2.1, \theta = 1.8)$ for the fixed voxel method and $1 + \Delta_{\text{RG,obs}} \sim \Gamma(k = 3.6, \theta = 1.8)$ for the flexible voxel method. Similarly, fitting a gamma distribution to $1 + \Delta_{\text{GRG,obs}}$ through MLE gave $1 + \Delta_{\text{GRG,obs}} \sim \Gamma(k = 2.5, \theta = 1.9)$ for the fixed voxel method and $1 + \Delta_{\text{GRG,obs}} \sim \Gamma(k = 4.7, \theta = 1.6)$ for the flexible voxel method, although the latter does not provide a tight fit. These gamma distributions constitute practical, two-parameter representations of the underlying data and are drawn as solid lines in Fig. 7.6.

The mean of a gamma-distributed RV $1 + \Delta \sim \Gamma(k, \theta)$ is $\mathbb{E}[1 + \Delta] = k\theta$. Thus, for the fixed voxel method, $\mathbb{E}[1 + \Delta_{\text{GRG,obs}}] = 4.8$ and $\mathbb{E}[1 + \Delta_{\text{RG,obs}}] = 3.8$. Similarly, for the flexible voxel method, $\mathbb{E}[1 + \Delta_{\text{GRG,obs}}] = 7.5$ and $\mathbb{E}[1 + \Delta_{\text{RG,obs}}] = 6.5$. Naively, it appears that we can conclude that, in a statistical sense, giants occupy denser regions of the Cosmic Web than radio galaxies in general. Using a two-sample Kolmogorov–Smirnov (KS) test, we formally tested the null hypothesis that the observed giant and general RG relative density distributions of Fig. 7.6 share a common underlying dis-

gamma distribution’s emergence here.




Figure 7.6: Probability density functions (PDFs) of the relative total matter density RV of 260 Local Universe giants (top row), and of 1443 Local Universe RGs (bottom row), determined through the fixed voxel method (left column) and flexible voxel method (right column). These RVs correspond to the density field smoothed to a scale of $2.9 \text{ Mpc } h^{-1}$. We also show PDFs (solid lines) of gamma-distributed RVs with parameters obtained via maximum likelihood estimation (MLE), alongside PDFs (dash-dotted lines) of gamma-distributed RVs with parameters obtained via MLE with a heteroskedasticity correction. We warn that these distributions are affected by surface brightness selection, and thus represent observed populations only.

tribution. For both the fixed and flexible voxel methods, the p -value $p \lesssim \sim 10^{-6}$. The null hypothesis is thus rejected (for typical significance levels).

However, these distributions come with two major caveats.

HETEROSKEDASTICITY

Firstly, BORG SDSS relative density measurements are heteroskedastic: measurements of higher densities have higher errors than measurements of lower densities (see e.g. Fig. 6 of Jasche et al., 2015). Heteroskedasticity causes the observed distributions to differ from the intrinsic distributions; not only by widening them, as also occurs in the more familiar homoskedastic setting, but also by systematically shifting the distributions towards lower densities. Therefore, in Appendix 7.A3, we propose a simple method to infer distributions corrected for this effect. As a result, we distin-

⁷We share a table with all 260 entries, alongside the analogous table for our 1443 selected LoTSS DR1 RGs, in Flexible Image Transport System (FITS) format through the Centre de Données astronomiques de Strasbourg (CDS). One can either use anonymous File Transfer Protocol (FTP) to <ftp://cdsarc.cds.unistra.fr> (130.79.128.5) or visit <https://cdsarc.cds.unistra.fr/cgi-bin/qcat?J/A+A/>.

Table 7.1: Cosmic Web properties for 50 of 260 BORG SDSS-constrained giants, sorted by right ascension. We provide MLE parameters of lognormal fits to total matter (i.e. both baryonic and dark matter) density distributions. We report parameters for both the fixed and flexible voxel methods. Densities span the BORG SDSS comoving voxel volume of $(2.9 \text{ Mpc } h^{-1})^3$ and are relative to the Universe’s mean late-time total matter density. The mean and variance of a relative density distribution are $\mathbb{E}[1 + \Delta_{\text{GRG,obs}} | 1 + \Delta_{\text{GRG}} = 1 + \delta_i] = \exp(\mu + \frac{1}{2}\sigma^2)$ and $\mathbb{V}[1 + \Delta_{\text{GRG,obs}} | 1 + \Delta_{\text{GRG}} = 1 + \delta_i] = (\exp(\sigma^2) - 1) \exp(2\mu + \sigma^2)$. The components of the T -web probability vector $\mathbf{p} = (p_1, p_2, p_3, p_4)$ correspond to clusters, filaments, sheets, and voids, respectively; due to rounding, their sum can appear to deviate from 1. Cluster masses are given only for giants generated by BCGs, and stem from cross-matching with [Wen & Han \(2015\)](#).⁷

rank ↓	host coordinates equatorial J2000 (°)	spectroscopic redshift z (1)	Cosmic Web density fixed voxel μ, σ^2 (1)	Cosmic Web density flexible voxel μ, σ^2 (1)	Cosmic Web T -web probabilities $\mathbf{p}$ (1)	cluster mass ($10^{14} M_\odot$)
1	111.5768, 38.6332	0.1539 $\pm$ 3 $\cdot$ 10 ⁻⁵	-0.49, 0.60	0.60, 0.45	(0.0, 0.2, 0.7, 0.1)	-
2	113.7719, 41.9743	0.0873 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.94, 0.06	2.43, 0.03	(0.1, 0.9, 0.0, 0.0)	-
3	114.9887, 43.9833	0.1489 $\pm$ 2 $\cdot$ 10 ⁻⁵	-0.73, 0.52	0.42, 0.42	(0.0, 0.2, 0.6, 0.2)	-
4	115.9636, 28.3578	0.1063 $\pm$ 2 $\cdot$ 10 ⁻⁵	0.88, 0.53	1.87, 0.20	(0.1, 0.8, 0.2, 0.0)	-
5	116.0338, 43.9917	0.1348 $\pm$ 2 $\cdot$ 10 ⁻⁵	0.95, 0.23	2.41, 0.05	(0.0, 0.9, 0.1, 0.0)	-
6	117.5396, 26.7355	0.1304 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.59, 0.17	2.16, 0.06	(0.2, 0.8, 0.0, 0.0)	-
7	118.1447, 35.8398	0.1365 $\pm$ 3 $\cdot$ 10 ⁻⁵	0.95, 0.25	2.09, 0.10	(0.0, 0.8, 0.2, 0.0)	-
8	119.1594, 32.4631	0.1462 $\pm$ 2 $\cdot$ 10 ⁻⁵	0.47, 0.50	1.28, 0.23	(0.1, 0.5, 0.4, 0.0)	-
9	119.2623, 36.6224	0.1395 $\pm$ 3 $\cdot$ 10 ⁻⁵	-0.02, 0.42	0.94, 0.22	(0.0, 0.2, 0.8, 0.0)	-
10	119.2794, 35.9797	0.1330 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.15, 0.28	2.08, 0.07	(0.1, 0.8, 0.1, 0.0)	1.3
11	119.4716, 36.6728	0.1284 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.26, 0.15	1.66, 0.08	(0.1, 0.8, 0.1, 0.0)	0.8
12	120.2556, 13.8312	0.1087 $\pm$ 2 $\cdot$ 10 ⁻⁵	2.46, 0.03	2.58, 0.02	(0.7, 0.3, 0.0, 0.0)	3.2
13	120.3054, 34.6752	0.0827 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.56, 0.08	1.93, 0.04	(0.1, 0.8, 0.0, 0.0)	-
14	120.3832, 47.6045	0.1568 $\pm$ 3 $\cdot$ 10 ⁻⁵	-0.30, 0.56	1.14, 0.32	(0.0, 0.4, 0.6, 0.0)	-
15	120.8052, 51.9326	0.0695 $\pm$ 2 $\cdot$ 10 ⁻⁵	2.00, 0.04	2.21, 0.03	(0.3, 0.7, 0.0, 0.0)	-
16	121.0142, 40.8026	0.1262 $\pm$ 1 $\cdot$ 10 ⁻⁵	0.85, 0.20	1.59, 0.11	(0.1, 0.8, 0.1, 0.0)	2.5
17	121.3279, 28.6242	0.1426 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.18, 0.39	1.78, 0.17	(0.1, 0.6, 0.3, 0.0)	-
18	121.3804, 25.8032	0.1370 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.02, 0.31	1.89, 0.09	(0.1, 0.8, 0.2, 0.0)	-
19	121.4295, 16.2322	0.1000 $\pm$ 2 $\cdot$ 10 ⁻⁵	2.46, 0.03	2.67, 0.02	(0.6, 0.4, 0.0, 0.0)	0.9
20	122.1485, 38.9145	0.0408 $\pm$ 1 $\cdot$ 10 ⁻⁵	1.79, 0.04	2.26, 0.02	(0.1, 0.9, 0.0, 0.0)	-
21	122.2863, 29.6790	0.1257 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.18, 0.18	2.03, 0.06	(0.1, 0.8, 0.1, 0.0)	-
22	122.3128, 41.2890	0.1334 $\pm$ 3 $\cdot$ 10 ⁻⁵	0.86, 0.38	1.48, 0.14	(0.2, 0.6, 0.2, 0.0)	-
23	123.4121, 41.3650	0.0998 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.05, 0.15	2.27, 0.04	(0.0, 0.7, 0.2, 0.0)	0.8
24	123.9160, 50.5406	0.1380 $\pm$ 3 $\cdot$ 10 ⁻⁵	1.16, 0.29	1.85, 0.08	(0.1, 0.6, 0.2, 0.0)	-
25	124.4446, 54.7009	0.1187 $\pm$ 2 $\cdot$ 10 ⁻⁵	2.22, 0.07	2.63, 0.03	(0.2, 0.8, 0.0, 0.0)	-
26	125.7805, 10.5983	0.0660 $\pm$ 0.0001	0.95, 0.14	1.53, 0.05	(0.2, 0.6, 0.2, 0.0)	-
27	126.4742, 41.6025	0.1536 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.38, 0.38	2.03, 0.15	(0.2, 0.7, 0.1, 0.0)	-
28	127.8646, 32.3241	0.0512 $\pm$ 1 $\cdot$ 10 ⁻⁵	0.69, 0.12	1.78, 0.04	(0.0, 1.0, 0.0, 0.0)	-
29	127.9987, 30.6585	0.1070 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.80, 0.08	2.07, 0.04	(0.5, 0.6, 0.0, 0.0)	-
30	128.1421, 4.4100	0.1060 $\pm$ 1 $\cdot$ 10 ⁻⁵	0.47, 0.24	1.34, 0.10	(0.0, 0.2, 0.7, 0.0)	-
31	129.0326, 26.8121	0.0878 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.22, 0.12	1.91, 0.04	(0.1, 0.7, 0.2, 0.0)	-
32	130.1852, 58.6971	0.1439 $\pm$ 3 $\cdot$ 10 ⁻⁵	0.15, 0.40	1.21, 0.19	(0.0, 0.4, 0.6, 0.0)	1.2
33	130.5014, 38.9375	0.1198 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.82, 0.09	2.23, 0.04	(0.2, 0.8, 0.0, 0.0)	-
34	130.5398, 41.9404	0.1257 $\pm$ 2 $\cdot$ 10 ⁻⁵	0.67, 0.31	1.93, 0.08	(0.0, 0.8, 0.2, 0.0)	-
35	131.2433, 42.0774	0.1493 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.57, 0.30	2.40, 0.08	(0.1, 0.8, 0.1, 0.0)	-
36	131.3629, 44.9240	0.1506 $\pm$ 3 $\cdot$ 10 ⁻⁵	1.70, 0.27	2.19, 0.10	(0.3, 0.6, 0.1, 0.0)	-
37	132.7367, 42.8046	0.0924 $\pm$ 1 $\cdot$ 10 ⁻⁵	1.05, 0.17	1.68, 0.05	(0.2, 0.7, 0.0, 0.0)	1.3
38	133.4574, 14.8739	0.0693 $\pm$ 2 $\cdot$ 10 ⁻⁵	2.00, 0.04	2.24, 0.02	(0.8, 0.2, 0.0, 0.0)	0.7
39	133.8097, 49.1933	0.1177 $\pm$ 1 $\cdot$ 10 ⁻⁵	0.99, 0.36	2.14, 0.09	(0.1, 0.8, 0.2, 0.0)	-
40	134.2352, 47.9559	0.1493 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.20, 0.44	2.00, 0.17	(0.1, 0.7, 0.2, 0.0)	-
41	134.5815, 46.3709	0.1171 $\pm$ 1 $\cdot$ 10 ⁻⁵	1.19, 0.24	1.54, 0.11	(0.2, 0.6, 0.2, 0.0)	-
42	135.3674, 55.0446	0.0460	2.15, 0.03	2.35, 0.01	(0.4, 0.6, 0.0, 0.0)	-
43	135.4596, 55.9243	0.1409 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.60, 0.21	2.06, 0.09	(0.4, 0.6, 0.0, 0.0)	-
44	136.0915, 19.7245	0.0995 $\pm$ 2 $\cdot$ 10 ⁻⁵	1.30, 0.10	1.72, 0.06	(0.1, 0.8, 0.0, 0.0)	-
45	137.1353, 18.2803	0.1155 $\pm$ 3 $\cdot$ 10 ⁻⁵	1.17, 0.17	1.96, 0.06	(0.1, 0.8, 0.1, 0.0)	-
46	137.2374, 58.3841	0.1436 $\pm$ 3 $\cdot$ 10 ⁻⁵	0.17, 0.52	1.33, 0.20	(0.0, 0.6, 0.4, 0.0)	-
47	137.8783, 48.5475	0.1086 $\pm$ 2 $\cdot$ 10 ⁻⁵	0.32, 0.30	1.45, 0.09	(0.0, 0.6, 0.4, 0.0)	-
48	139.7475, 31.8613	0.0619 $\pm$ 1 $\cdot$ 10 ⁻⁵	1.65, 0.08	1.78, 0.04	(0.6, 0.4, 0.0, 0.0)	-
49	139.9519, 57.8489	0.1369 $\pm$ 3 $\cdot$ 10 ⁻⁵	1.21, 0.17	1.67, 0.09	(0.1, 0.8, 0.1, 0.0)	-
50	140.3421, 54.8650	0.0447 $\pm$ 1 $\cdot$ 10 ⁻⁵	1.90, 0.02	2.31, 0.02	(0.0, 0.6, 0.4, 0.0)	-

guish between $1 + \Delta_{\text{GRG,obs}}$ and $1 + \Delta_{\text{GRG}}$, and between $1 + \Delta_{\text{RG,obs}}$ and $1 + \Delta_{\text{RG}}$ — in both cases, the latter RV has been corrected for heteroskedasticity.

As the effect appears to induce minor shifts only, we assumed that the heteroskedasticity-free distributions are also well approximated by gamma distributions. We then applied the method of Appendix 7.A3 to both our GRG and RG data. Having no closed-form MLE parameter expressions at our disposal, we obtained the MLE parameters of the heteroskedasticity-free gamma distributions by evaluating the likelihood function over an exhaustive grid of parameter values (k, θ) . For the fixed voxel method, this yielded $1 + \Delta_{\text{GRG}} \sim \Gamma(k = 3.7, \theta = 1.3)$ and $1 + \Delta_{\text{RG}} \sim \Gamma(k = 3.0, \theta = 1.3)$; $\mathbb{E}[1 + \Delta_{\text{GRG}}] = 4.8$ and $\mathbb{E}[1 + \Delta_{\text{RG}}] = 3.9$. For the flexible voxel method, this yielded $1 + \Delta_{\text{GRG}} \sim \Gamma(k = 6.7, \theta = 1.1)$ and $1 + \Delta_{\text{RG}} \sim \Gamma(k = 4.8, \theta = 1.4)$; $\mathbb{E}[1 + \Delta_{\text{GRG}}] = 7.4$ and $\mathbb{E}[1 + \Delta_{\text{RG}}] = 6.7$. We draw the corresponding PDFs as dash-dotted lines in Fig. 7.6.

SURFACE BRIGHTNESS SELECTION

The second — and plausibly most important — caveat is the imprint of selection effects, and in particular the surface brightness selection effect induced by the LoTSS noise level $\sigma_{I_\nu} = 25 \text{ Jy deg}^{-2}$ at $\theta_{\text{FWHM}} = 6''$. This effect causes RGs whose lobe surface brightnesses are below the noise level to evade sample inclusion. Lobe surface brightness depends on the RG’s projected proper length l_p , morphological parameters F_l and F_{l_ν} , radio luminosity L_ν , redshift z , and lobe spectral index α . Equation 40 of [Oei et al. \(2023a\)](#) presents an approximate formula for the mean lobe surface brightness B_ν :

$$B_\nu = \frac{2 \cdot F_{l_\nu} \cdot L_\nu}{\pi^2 \cdot \mathbb{E}[D](\eta(F_l)) \cdot F_l^2 \cdot l_p^2 \cdot (1 + z)^{3-\alpha}}, \quad (7.1)$$

where F_{l_ν} is the fraction of the total radio luminosity that belongs to the lobes, F_l is the fraction of the RG’s axis that lies inside the lobes, $\eta = \frac{F_{l_\nu}}{2 - F_l}$, and $\mathbb{E}[D]$ is the mean deprojection factor as given by Eq. A.29 of [Oei et al. \(2023a\)](#):

$$\mathbb{E}[D] = (1 + \eta) \left(\frac{\pi}{2} - \frac{\eta}{\sqrt{1 - \eta^2}} \ln \left(\frac{\sqrt{1 - \eta^2}}{\eta} + \frac{1}{\eta} \right) \right). \quad (7.2)$$

To explicitly demonstrate that surface brightness selection biases our Local Universe GRG and RG samples towards short lengths and high radio luminosities, we approximated B_ν for all LoTSS DR1 RGs with $z < 0.2$, assuming $F_l = F_{l_\nu} = 0.3$ ([Oei et al., 2022a](#)) and $\alpha = -0.7$.




Figure 7.7: Mean lobe surface brightness estimates B_ν at $\nu_{\text{ref}} = 150$ MHz for LoTSS DR1 RGs at $z < 0.2$, shown as a function of projected proper length l_p . During the majority ($\sim 90\%$) of their lifetime, RGs grow while maintaining a roughly constant radio luminosity. If — additionally — growth is shape-preserving, surface brightness decreases quadratically with l_p . The seven golden lines therefore denote approximate evolutionary tracks of RGs at $z = 0$ with end-of-life radio luminosities $l_\nu \in \{10^{21} \text{ W Hz}^{-1}, \dots, 10^{27} \text{ W Hz}^{-1}\}$. The region of parameter space where RGs can become detectable giants, $l_\nu \geq l_{\nu,\text{min}} := 10^{24} \text{ W Hz}^{-1}$, is shaded gold. The region of parameter space occupied by giants, $l_p \geq l_{p,\text{GRG}} := 0.7 \text{ Mpc}$, is shaded green. The region of parameter space inaccessible to the LoTSS, $B_\nu < \sigma_{I_\nu} = 25 \text{ Jy deg}^{-2}$, is shaded red. The grey lines indicate levels of constant minimum-energy lobe pressure P_{min} .

Figure 7.7 shows the results. As mature RGs grow, they trace out a curve approximately parallel to the golden lines (which represent constant radio luminosities), directed towards the bottom-right of the plot. If their lives last long enough, they therefore reach a length — which depends on their radio luminosity — beyond which they disappear into the noise of the LoTSS. In particular, our calculations suggest that the only giants that can be detected through the LoTSS are those with $l_\nu \geq 10^{24} \text{ W Hz}^{-1}$.⁸ Indeed, the lowest-radio luminosity giant in the LoTSS DR1 GRG sample of [Dabhade et al. \(2020b\)](#) — which consists of 239 giants — has $l_\nu \sim$

⁸In Fig. 7.7, the approximate evolutionary track for $l_\nu = 10^{24} \text{ W Hz}^{-1}$ shows that when RGs with this radio luminosity turn into giants (i.e. when the fourth golden line intersects the green line), they disappear in the LoTSS noise (this golden line intersects the red line).




Figure 7.8: PDFs of the observed RG relative total matter density RV given a radio luminosity at $\nu_{\text{ref}} = 150$ MHz, using MLE parameter values for the model described in Sect. 7.4.2. The black contours denote CDF values. We overplot all 1443 selected LoTSS DR1 RGs (dots), with those above the empirical median density coloured red, and those below coloured blue. We use flexible voxel method densities here. For fixed voxel method densities, see Fig. 7.19.

$1 \cdot 10^{24} \text{ W Hz}^{-1}$. Clearly, if RGs with $l_\nu \sim 10^{23} \text{ W Hz}^{-1}$ can also form giants, then this subpopulation will be missing in LoTSS GRG samples. Equivalently, the giants that *are* included in our LoTSS GRG sample necessarily have radio luminosities $l_\nu \geq 10^{24} \text{ W Hz}^{-1}$. It is overwhelmingly likely that the sample’s radio luminosities are biased high.

Concerningly, Croston et al. (2019) have shown that, for RGs at $z < 0.4$, radio luminosity and environmental richness correlate positively. This gives rise to the possibility that the trend shown in Fig. 7.6, where giants appear to reside in denser Cosmic Web environments than RGs in general, is a result of surface brightness selection: a GRG sample biased high in radio luminosity will be biased high in Cosmic Web density.

7.4.2 RADIO LUMINOSITY–COSMIC WEB DENSITY RELATION

Interestingly, our data allow us to revisit the radio luminosity–environmental richness correlation that Croston et al. (2019) have found. In Fig. 7.8, we plot all 1443




Figure 7.9: Posterior distribution over $k(0)$, $\theta(0)$, $k'(0)$, and $\theta'(0)$, based on selected Local Universe LoTSS DR1 RGs. We show all two-parameter marginals of the likelihood function, with contours enclosing 50% and 90% of total probability. We mark the maximum likelihood estimate (MLE) values (grey dot) and the maximum a posteriori (MAP) values (grey cross). The one-parameter marginals again show the MLE (dash-dotted line), a mean-centred interval of standard deviation-sized half-width (hashed region), and a median-centred 90% credible interval (shaded region). We use flexible voxel method densities here. For fixed voxel method densities, see Fig. 7.20.

selected LoTSS DR1 RGs in Cosmic Web density–radio luminosity parameter space. By eye, it appears that for higher radio luminosities, the Cosmic Web density distribution shifts upwards. To quantify the apparent tendency that more luminous radio galaxies occupy denser regions of the Cosmic Web, we extend the gamma distribution fitting discussed in Sect. 7.4.1. In particular, we assume that the RG relative density

RV at some fixed radio luminosity, $1 + \Delta_{\text{RG}} \mid L_\nu = l_\nu$, is gamma distributed:

$$1 + \Delta_{\text{RG}} \mid L_\nu = l_\nu \sim \Gamma(k(\mathbb{J}), \theta(\mathbb{J})). \quad (7.3)$$

Instead of assuming constant parameters k and θ , we now assume that k and θ depend on radio luminosity:

$$\mathbb{J} = \mathbb{J}(l_\nu) := \log_{10} \left(\frac{l_\nu}{l_{\nu, \text{ref}}} \right). \quad (7.4)$$

The natural minimal extension to assuming that k and θ remain constant as $\mathbb{J}$ changes, is to consider Taylor polynomial approximations of degree one (around $\mathbb{J} = 0$ — i.e. Maclaurin polynomials):

$$k(\mathbb{J}) \approx k(0) + k'(0) \cdot \mathbb{J}, \quad \theta(\mathbb{J}) \approx \theta(0) + \theta'(0) \cdot \mathbb{J}. \quad (7.5)$$

However, we note that the current model does not allow for a direct comparison to our data, which are samples from $1 + \Delta_{\text{RG, obs}} \mid L_\nu = l_\nu$ rather than of $1 + \Delta_{\text{RG}} \mid L_\nu = l_\nu$. Therefore, we must extend the model by applying heteroskedastic measurement errors. In line with Appendix 7.A3, we assume

$$1 + \Delta_{\text{RG, obs}} \mid 1 + \Delta_{\text{RG}} = 1 + \delta \sim \text{Lognormal}(\mu, \sigma^2), \quad (7.6)$$

where the parameters $\mu = \mu(a, b)$ and $\sigma^2 = \sigma^2(a, b)$ are given by

$$\mu = \ln(1 + \delta) - \frac{1}{2} \ln \left(1 + a(1 + \delta)^{b-2} \right); \quad (7.7)$$

$$\sigma^2 = \ln \left(1 + a(1 + \delta)^{b-2} \right). \quad (7.8)$$

Here, a and b are parameters that characterise the heteroskedasticity. By considering the means and SDs of the relative densities of all selected LoTSS DR1 RGs, we found $a = 0.4$ and $b = 1.1$ for the fixed voxel method and $a = 0.4$ and $b = 0.9$ for the flexible voxel method.

In order to obtain a fit that is valid over many orders of magnitude in radio luminosity, rather than only over the order of magnitude in which most data are concentrated, we calculated the likelihood in such a way that all decades in radio luminosity that contain at least 20 RGs receive equal weight. We simply determined the likelihood function, and thus the posterior distribution for a flat prior, through a grid search over the four parameters $k(0)$, $\theta(0)$, $k'(0)$, and $\theta'(0)$. We define $l_{\nu, \text{ref}} := 10^{23} \text{ W Hz}^{-1}$. We visualise this flat-prior posterior in Fig. 7.9. In addition, we show percentile scores

of 10%, 50%, 90%, and 95% corresponding to the MAP model in Fig. 7.8. Figures 7.19 and 7.20 demonstrate that, although the densities are lower, the results remain qualitatively the same upon switching to the fixed voxel method. We summarise the posterior for both methods in Table 7.2.

Table 7.2: Maximum a posteriori probability (MAP) and posterior mean and standard deviation (SD) estimates of the free parameters in RG Cosmic Web density–radio luminosity inference. As we assume a flat prior, the MAP parameters are also the MLE parameters.

fixed voxel method:		
parameter	MAP	posterior mean and SD
$k(0)$	2.9	2.9 ± 0.1
$\theta(0)$	1.4	1.4 ± 0.1
$k'(0)$	0.04	0.04 ± 0.08
$\theta'(0)$	0.08	0.08 ± 0.04
flexible voxel method:		
parameter	MAP	posterior mean and SD
$k(0)$	4.8	4.7 ± 0.2
$\theta(0)$	1.4	1.4 ± 0.1
$k'(0)$	0.2	0.2 ± 0.1
$\theta'(0)$	0.02	0.02 ± 0.04

As values of $k'(0)$ and especially $\theta'(0)$ tend to be positive, we find evidence for a positive radio luminosity–environmental richness relation, in line with the results of [Croston et al. \(2019\)](#). More quantitatively, we call the relation positive (at $l_\nu = l_{\nu,\text{ref}}$) when

$$\left. \frac{d\mathbb{E}[1 + \Delta_{\text{RG}} \mid L_\nu = l_\nu]}{d\downarrow} \right|_{\downarrow=0} = k(0)\theta'(0) + k'(0)\theta(0) > 0. \quad (7.9)$$

We combined the posterior with Eq. 7.9 to obtain the posterior probability that the radio luminosity–Cosmic Web density relation is positive — that is to say that more luminous RGs tend to live in denser regions of the Cosmic Web. For both the fixed and flexible voxel methods, this probability exceeds 99%.

SURFACE BRIGHTNESS SELECTION

However, we must again be mindful of a possible surface brightness selection effect. RGs of a given radio luminosity may be able to grow larger in more tenuous Cosmic Web environments, as radio luminosity traces jet power (in close-to-linear fashion;

e.g. Willott et al., 1999; Shabala & Godfrey, 2013; Hardcastle, 2018). If this is indeed the case, then the lobe surface brightnesses of RGs in tenuous environments are likely lower than those of RGs with the same radio luminosity in dense environments. As a result, of all RGs with the same radio luminosity, those in tenuous environments are more likely to be missing from observational samples. This selection effect, that favours higher relative densities, holds for all radio luminosities, and could only explain the trend of Fig. 7.8 if the effect becomes more severe for higher radio luminosities. On the one hand, RGs with higher radio luminosities at a given relative density are expected to grow longer. On the other hand, as is clear from Fig. 7.7, the projected proper length range over which RGs retain detectable lobe surface brightnesses increases with radio luminosity. In other words, at higher radio luminosities, larger RGs remain detectable.

Whether the net result of these counteracting effects is less or more bias towards higher relative densities as radio luminosity increases, depends on the scaling between radio luminosity and growth. Self-similar growth predicts (Kaiser & Alexander, 1997) that an RG's projected proper length scales with jet power Q_{jet} , Cosmic Web density at the host galaxy $1 + \delta_g$, and time since birth t as

$$l_p \propto Q_{\text{jet}}^{\frac{1}{5-\beta}} \cdot (1 + \delta_g)^{\frac{-1}{5-\beta}} \cdot t^{\frac{3}{5-\beta}}, \quad (7.10)$$

where $-\beta$ is the exponent of the local Cosmic Web density profile:

$$(1 + \delta)(r) \propto (1 + \delta_g) \cdot r^{-\beta}. \quad (7.11)$$

Under the additional assumption of equipartition between the relativistic electron–positron kinetic energy density and the magnetic field energy density, and upon neglecting electron–positron energy losses (such as adiabatic, synchrotron, and inverse Compton losses), the RG's radio luminosity obeys (Shabala & Godfrey, 2013)

$$l_\nu \propto Q_{\text{jet}}^{\frac{5+p}{6}} l_p^{3 - \frac{(4+\beta)(5+p)}{12}}. \quad (7.12)$$

Here, $-p$ is the exponent of the initial post-acceleration electron–positron kinetic energy distribution $n(\gamma) \propto \gamma^{-p}$, valid for Lorentz factors between some minimum and maximum values γ_{min} and γ_{max} . Solving for Q_{jet} , substituting the result in Eq. 7.10, and solving for l_p , yields — at fixed $1 + \delta_g$ and t —

$$l_p \propto l_\nu^{\frac{12}{36 + (5+p)(6-\beta)}}. \quad (7.13)$$

In a constant-density environment, $\beta = 0$, and for the strong-shock spectral index $p = 2$, we find $l_p \propto l_\nu^{2/3}$. In an environment where density falls off quadratically with distance, $\beta = 2$, and therefore $l_p \propto l_\nu^{1/3}$.

Meanwhile, Eq. 7.1 predicts that the maximum detectable projected proper length for RGs with a given morphology, at a given redshift, and with given lobe spectral index, scales as $l_{p,\max} \propto \sqrt{l_\nu}$.

These arguments suggest that $l_{p,\max}$ increases more rapidly with l_ν than l_p does. Thus, amongst a family of RGs with varying radio luminosity but seen at the same age (i.e. time after birth) and with the same morphology and Cosmic Web density, the most luminous is most easily detectable. This provides some reassurance that the trend seen in Fig. 7.8 is real: surface brightness selection at a fixed relative density seems to become less severe as radio luminosity increases.

However, there are at least two caveats here. Firstly, more luminous RGs may live longer than less luminous RGs, as jointly suggested by the radio luminosity–stellar mass relation and the stellar mass–radio AGN occurrence relation of [Sabater et al. \(2019\)](#). As a result, in a snapshot of the RG population at a given instant, more luminous RGs may typically be older than less luminous RGs. We illustrate this point in Fig. 7.10.⁹ If more luminous RGs indeed tend to be older than less luminous RGs, then the relation between projected length and radio luminosity at a fixed Cosmic Web density for a snapshot population will be steeper than the one of Eq. 7.13, which

⁹We define

$$F_a := \frac{\mathbb{E}[t_a]}{\mathbb{E}[t_a + t_q]} = \frac{\mathbb{E}[t_a]}{\mathbb{E}[t_a] + \mathbb{E}[t_q]} = \frac{\lambda_a^{-1}}{\lambda_a^{-1} + \lambda_q^{-1}} = \left(1 + \frac{\lambda_a}{\lambda_q}\right)^{-1}. \quad (7.14)$$

We note that the more natural definition

$$F_a := \mathbb{E} \left[\frac{t_a}{t_a + t_q} \right] = \mathbb{E} \left[\left(1 + \frac{t_q}{t_a}\right)^{-1} \right] \quad (7.15)$$

gives, for $r := \frac{\lambda_a}{\lambda_q}$,

$$F_a = \begin{cases} \frac{1}{2} & \text{if } r = 1; \\ \frac{r \ln r - r + 1}{(r-1)^2} & \text{if } r > 0, r \neq 1. \end{cases} \quad (7.16)$$

To arrive at Eq. 7.16, we use standard identities for the PDFs of the ratio distribution and the inverse distribution, and apply l'Hôpital's rule twice. When $r = 1$, that is to say when $\lambda_a = \lambda_q$, both definitions predict the same $F_a = \frac{1}{2}$. For $r < 1$, our definition predicts higher F_a than the natural definition; for $r > 1$, our definition predicts lower F_a . To find λ_a given λ_q and F_a under the natural definition, one must solve Eq. 7.16 numerically. As this procedure is less transparent, we generate Fig. 7.10 using the simpler definition for F_a .




Figure 7.10: Monte Carlo simulation of a population of nine galaxies (thick horizontal bars) with successive active and quiescent phases. We mark the active phases (thin horizontal bars). Each galaxy spends a different fraction of the time F_a in the active phase. We model the time spent in both phases as draws from exponential distributions: $t_a \sim \text{Exp}(\lambda_a)$ and $t_q \sim \text{Exp}(\lambda_q)$. We assume that all galaxies share the same quiescent phase rate parameter $\lambda_q = 10 \text{ Gyr}^{-1}$, although this may not be entirely correct (Turner & Shabala, 2015). This assumption then fixes the active phase rate parameter: $\lambda_a = (F_a^{-1} - 1)\lambda_q$. On human time scales, we can access only a snapshot view of this population. During the snapshot shown here (vertical dash-dotted line), seven galaxies generate RGs, and could be included in an RG sample (thick light-green horizontal bars). RG ages at the snapshot, represented by the hatched parts of the thin green horizontal bars, generally increase with F_a .

assumes equal ages. Secondly, we warn for the possibility — which remains subject of debate (Mingo et al., 2019) — that the distribution of morphologies at a fixed Cosmic Web density may change with radio luminosity.

PREDICTING THE DENSITY DISTRIBUTION OF OBSERVED GIANTS BY TREATING THEM AS LUMINOUS GENERAL RGs

By applying standard PDF identities to the current context, we found

$$\begin{aligned}
 & f_{1+\Delta_{\text{RG}} \mid L_{\nu} \geq l_{\nu, \min}}(1 + \delta) = \\
 & \frac{\int_{l_{\nu, \min}}^{\infty} f_{1+\Delta_{\text{RG}} \mid L_{\nu} = l_{\nu}}(1 + \delta) \cdot f_{L_{\nu}}(l_{\nu}) \, dl_{\nu}}{\int_{l_{\nu, \min}}^{\infty} f_{L_{\nu}}(l_{\nu}) \, dl_{\nu}} \\
 & \frac{\int_{l_{\nu, \min}}^{\infty} f_{1+\Delta_{\text{RG}} \mid L_{\nu} = l_{\nu}}(1 + \delta) \cdot f_{L_{\nu} \mid L_{\nu} \geq l_{\nu, \min}}(l_{\nu}) \, dl_{\nu}}{\int_{l_{\nu, \min}}^{\infty} f_{L_{\nu} \mid L_{\nu} \geq l_{\nu, \min}}(l_{\nu}) \, dl_{\nu}}, \quad (7.17)
 \end{aligned}$$

which one can use to calculate the RG relative total matter density RV for radio luminosities of $l_{\nu,\min}$ and higher. This allowed us to test whether Fig. 7.6's higher densities for observed giants than for general RGs can be fully explained by the radio luminosity–Cosmic Web density relation. Clearly, to proceed, we had to find $f_{L_\nu}(l_\nu)$ for $l_\nu \geq l_{\nu,\min}$ — up to a constant, at least. We modelled $L_\nu | L_\nu \geq l_{\nu,\min} \sim \text{Pareto}(l_{\nu,\min}, \lambda)$, where λ is the tail index or rate parameter of the associated exponential distribution.¹⁰ For both $l_{\nu,\min} = 10^{23} \text{ W Hz}^{-1}$ and $l_{\nu,\min} = 10^{24} \text{ W Hz}^{-1}$, we found $\lambda_{\text{MLE}} = 0.6$ (based on $1 \cdot 10^3$ and $3 \cdot 10^2$ RGs, respectively). For higher $l_{\nu,\min}$, λ_{MLE} increases somewhat, but becomes less reliable because of decreasing sample sizes. We therefore stuck with $\lambda_{\text{MLE}} = 0.6$. We then evaluated Eq. 7.17 for $l_{\nu,\min} = 10^{25} \text{ W Hz}^{-1}$ and added heteroskedastic noise in order to compare our prediction to observations. We propagated the correlated uncertainties on the parameters $k(0)$, $\theta(0)$, $k'(0)$, and $\theta'(0)$, which together determine $f_{1+\Delta_{\text{RG}} | L_\nu=l_\nu}$, by rejection sampling from the joint posterior distribution. We show the resulting mean prediction with standard deviation half-width (pink) in Fig. 7.11. Although the prediction is based solely on observations of general RGs — and so is built without knowledge of the observed GRG distribution (solid green KDE) — it appears to fit it better than a direct MLE gamma fit to the GRG data (solid green curve) does. However, a *better* fit is not necessarily also a *good* fit. Is our sample of Cosmic Web densities for observed giants consistent with the prediction? A Kolmogorov–Smirnov tests suggests an answer in the affirmative, as it yielded a p -value of 0.7.

Section 7.4.1 presented a paradox: known giant radio galaxies reside in denser Cosmic Web environments than their smaller kin, in apparent violation of the long-standing hypothesis that giants should emerge predominantly in the dilute Cosmic Web. The modelling in this section suggests that the higher densities found for known giants than for general RGs can be explained by a combination of these giants' higher radio luminosities and the general relation between RG radio luminosity and Cosmic Web density. Fainter giants — which might inhabit the more dilute Cosmic Web — might exist, but are currently not observationally accessible.

¹⁰Equivalently, one could model

$$\ln \left(\frac{L_\nu | L_\nu \geq l_{\nu,\min}}{l_{\nu,\min}} \right) \sim \text{Exp}(\lambda), \quad (7.18)$$

or

$$\log_{10} \left(\frac{L_\nu | L_\nu \geq l_{\nu,\min}}{l_{\nu,\min}} \right) \sim \text{Exp}(\lambda \cdot \ln 10). \quad (7.19)$$




Figure 7.11: Density distribution predicted for observed giants combining Sect. 7.4.2’s RG radio luminosity–Cosmic Web density relation and a radio luminosity cut-off $l_{\nu, \min} = 10^{25} \text{ W Hz}^{-1}$ (pink); in addition, we show the KDE density distributions for giants with gamma distribution MLE fit (solid green) and for general RGs (hatched green) — all as in the left column of Fig. 7.6.

7.4.3 RADIO GALAXY NUMBER DENSITIES

After accreting baryons, supermassive black holes (SMBHs) can launch jets that give rise to radio galaxies. Although it is clear that SMBH jets and radio galaxies play an important role in galaxy evolution and cosmology, the physics of jet launching is not yet fully understood. Moreover, because SMBHs are astronomical unit-sized, galaxies are kiloparsec-sized (ratio 10^8 – 10^9), and the Cosmic Web is megaparsec-sized (ratio 10^{11} – 10^{12}), it is not yet possible to build cosmological simulations in which a realistic interplay between SMBHs, their host galaxies, and the enveloping Cosmic Web naturally arises. Thus, for now, major advances in our understanding of radio galaxies on the cosmological scale must come from observations instead of from simulations. It is therefore of considerable interest to observationally constrain the occurrence of radio galaxies as a function of Cosmic Web density.

Enticingly, our relative total matter density measurements — in combination with the BORG SDSS HMC Markov chain samples in their entirety — allow us to determine the RG number density as a function of Cosmic Web density: $n_{\text{RG}}(1 + \delta)$. To see why, we first emphasise that the PDF $f_{1+\Delta_{\text{RG}}}$, which we have approximated through observations and visualised in the bottom row of Fig. 7.6, provides a key ingredient for $n_{\text{RG}}(1 + \delta)$. It shows that if a volume of cosmological extent is surveyed, one will obtain more RGs with moderate (filament-like) densities than with low (void-like) or high (cluster-like) densities. We remark that in such a survey sample the number of




Figure 7.12: Number densities of RGs with a given radio luminosity l_ν , as a function of Cosmic Web density at the $2.9 \text{ Mpc } h^{-1}$ scale. The number density functions reflect conditions in the Local Universe, and are given relative to the cosmic mean number density at l_ν . The solid curves reflect the posterior mean, whilst the shaded areas around the mean denote -1 to $+1$ posterior standard deviation ranges. We use flexible voxel method densities here.

RGs whose densities fall within a given interval depends not only on the RG number density at the interval's densities, but also on the prevalence of environments with such densities. As a result, if a sample contains more filament-inhabiting RGs than, say, cluster-inhabiting RGs, one need not necessarily conclude that the RG number density in filaments is higher than in clusters; in fact, the RG number density in the latter environment could well be higher, provided that clusters are sufficiently rare.

Conveniently, the density fields of the BORG SDSS HMC Markov chain samples enable us to quantify the rarity of each Cosmic Web environment with a PDF $f_{1+\Delta_{\text{CW}}}$, which corresponds to the Cosmic Web relative density $\text{RV } 1 + \Delta_{\text{CW}}$. This RV represents the relative total matter density at a randomly chosen point in the Local Universe (and at the BORG SDSS resolution of $2.9 \text{ Mpc } h^{-1}$). We determine the distribution of $1 + \Delta_{\text{CW}}$ simply by binning the density fields of a few hundred Markov chain samples. As a result, $f_{1+\Delta_{\text{CW}}}$ is best determined at low densities, which correspond to the most common environments, and least so at high densities, which

are rare.¹¹ As we expect $f_{1+\Delta_{\text{CW}}}$ to be smooth, we apply a Savitzky–Golay filter (of polynomial order 1) to the binned data; however, this step is not essential.

Appendix 7.A5 shows that $n_{\text{RG}}(1 + \delta)$, the RG number density as would arise in an environment of constant density $1 + \delta$, obeys

$$\frac{n_{\text{RG}}(1 + \delta)}{\bar{n}_{\text{RG}}} = \frac{f_{1+\Delta_{\text{RG}}}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.20)$$

Here, $\bar{n}_{\text{RG}}$ is the cosmic mean RG number density. Thus, a simple point-wise division of the PDFs of $1 + \Delta_{\text{RG}}$ and $1 + \Delta_{\text{CW}}$ yields the RG number density at a given Cosmic Web density relative to the cosmic mean value.

Unfortunately, as discussed in Sect. 7.4.1, our measurements of $f_{1+\Delta_{\text{RG}}}$ suffer from surface brightness selection. To avoid selection effects, we must formulate a more specific version of Eq. 7.20. In particular, we consider the number density of RGs per unit of radio luminosity as a function of radio luminosity and relative density: $\Upsilon_{\text{RG}}(l_\nu, 1 + \delta)$ — again, relative to a cosmic mean value: $\bar{\Upsilon}_{\text{RG}}(l_\nu)$. Appendix 7.A5 shows that this ratio equals

$$\frac{\Upsilon_{\text{RG}}(l_\nu, 1 + \delta)}{\bar{\Upsilon}_{\text{RG}}(l_\nu)} = \frac{f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.21)$$

This radio luminosity–dependent relative number density is, as before, equal to a point-wise division of PDFs. With respect to the RHS of Eq. 7.20, only the PDF in the numerator has changed. We have modelled this PDF, $f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}$, in Sect. 7.4.2. It depends on the four parameters $k(0)$, $k'(0)$, $\theta(0)$, and $\theta'(0)$, for which the data shown in Fig. 7.8 induce the posterior distribution shown in Fig. 7.9. We evaluate the RHS of Eq. 7.21 by rejection sampling from the posterior, calculating $f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}$ for each sampled quartet $(k(0), k'(0), \theta(0), \theta'(0))$, and dividing the resulting PDFs point-wise by $f_{1+\Delta_{\text{CW}}}$. In Fig. 7.12, for four observationally accessible radio luminosities, we visualise the posterior mean alongside a range that is two standard deviations wide.

¹¹We note that, at first sight, it might be tempting to approximate $f_{1+\Delta_{\text{CW}}}$ by binning the BORG SDSS mean rather than individual samples. This, however, leads to incorrect results. The BORG SDSS mean tends to the cosmic mean wherever the structure of the Cosmic Web is not well constrained by the SDSS DR7 MGS. This tendency to revert to the cosmic mean in the absence of data constraints causes low-density environments to be overrepresented and high-density environments to be underrepresented. As a result, approximating $f_{1+\Delta_{\text{CW}}}$ by binning the BORG SDSS mean would e.g. underestimate its high-density tail. More quantitatively, when using the BORG SDSS mean, the probability densities $f_{1+\Delta_{\text{CW}}}(1 + \delta)$ for $1 + \delta \gtrsim 10$ are up to an order of magnitude lower than predicted by BORG SDSS samples.

As expected, RG number densities generally increase with Cosmic Web density. The more luminous RGs become, the more pronounced is the contrast between their number density in clusters and their cosmic mean number density. Our results indicate that high-luminosity RGs, with $l_\nu = 10^{25} \text{ W Hz}^{-1}$ and $l_\nu = 10^{27} \text{ W Hz}^{-1}$, attain number densities in clusters that are a hundred times higher than average. Interestingly, for radio luminosities $l_\nu = 10^{21} \text{ W Hz}^{-1}$ and $l_\nu = 10^{23} \text{ W Hz}^{-1}$, we obtain statistically significant, preliminary evidence that the RG number density peaks at a cluster-like density that depends positively on l_ν .¹² Beyond this density, these low-luminosity RGs become less prevalent per unit of volume, possibly because the underlying galaxy population starts generating more luminous RGs instead.

We stress that these results are tentative only, as both $f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}$ and $f_{1+\Delta_{\text{CW}}}$ suffer from systematic errors at higher densities. These systematic errors are not reflected in the uncertainties of Fig. 7.12. In the case of $f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}$, the gamma distribution approximation loses validity in the tail: beyond some point, it will consistently under- or overestimate the probability density. Our LoTSS DR1 sample contains only a modest number ($\sim 10^2$) of RGs at relative densities $1+\delta \gtrsim 10$ — especially when selecting those around a given value of l_ν only (in which case just $\sim 10^1$ RGs remain). As a result, in its tail the gamma PDF takes on the character of an extrapolation function rather than of a fitting function, and should consequently be treated with caution. In the future, larger RG samples can resolve this situation. In the case of $f_{1+\Delta_{\text{CW}}}$, the 2LPT model of structure formation used by the BORG SDSS loses validity at cluster densities. As a result, this work’s estimate of $f_{1+\Delta_{\text{CW}}}$ underpredicts the prevalence of cluster-like densities. Newer BORG runs, such as the BORG 2M++, employ an enhanced gravity solver (Jasche & Lavaux, 2019), relieving this limitation.

Although preliminary in nature, the inferences shown in Fig. 7.12 could in principle be used to calibrate cosmological simulations that feature RG feedback. The fact that our number densities are given with respect to a cosmic mean value is not problematic, because such a cosmic mean number density is straightforward to calculate from simulation snapshots. However, apart from making sure that the RG number densities in simulated environments of various densities are correctly proportioned, it is also important to calibrate the *absolute* level of RG feedback in simulations. Interestingly, there are variations on Eq. 7.20 for which an absolute number density can be explicitly calculated with current-day data. Consider, for example, non-giant radio galaxies (i.e. $l_p < l_{p,\text{GRG}} := 0.7 \text{ Mpc}$) in the Local Universe with radio luminosities $l_\nu \geq 10^{24} \text{ W Hz}^{-1}$. Figure 7.7 suggests that this RG subpopulation is already fully accessible by the LoTSS. Calling these RGs luminous non-giant radio galaxies

¹²For radio luminosities $l_\nu = 10^{25} \text{ W Hz}^{-1}$ and $l_\nu = 10^{27} \text{ W Hz}^{-1}$, the uncertainty margins are too large to claim evidence for a peak.




Figure 7.13: Number density of luminous non-giant radio galaxies in the Local Universe, as a function of Cosmic Web density. We define these as RGs with 150 MHz radio luminosities $l_\nu \geq 10^{24} \text{ W Hz}^{-1}$, projected proper lengths $l_p < l_{p,\text{GRG}} := 0.7 \text{ Mpc}$, and redshifts $z < z_{\text{max}} := 0.16$. The Cosmic Web densities encompass baryonic and dark matter, are defined on a $2.9 \text{ Mpc } h^{-1}$ scale, and are given relative to the cosmic mean. We use flexible voxel method densities here. We show the posterior mean with a shaded -1 to $+1$ posterior standard deviation range. We also show five individual realisations (grey).

(LNGRGs), we have

$$n_{\text{LNGRG}}(1 + \delta) = \bar{n}_{\text{LNGRG}} \cdot \frac{f_{1+\Delta_{\text{LNGRG}}}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.22)$$

We calculate $\bar{n}_{\text{LNGRG}}$ by counting all LNGRGs in the LoTSS DR1 footprint up to z_{max} , and find $N_{\text{LNGRG}} = 172$. Because of calibration problems and higher noise levels towards the edges of the footprint, a fraction of the footprint is — artificially — devoid of LNGRGs. We thus apply a correction factor to N_{LNGRG} of 1.2, yielding $N_{\text{LNGRG}} = 200$. Assuming that $N_{\text{LNGRG}} \sim \text{Poisson}(\lambda_{\text{LNGRG}})$, maximum likelihood estimation simply suggests $\hat{\lambda}_{\text{LNGRG,MLE}} = 200$. We combine this MLE Poisson distribution for N_{LNGRG} with the comoving volume $V = 16 \cdot 10^6 \text{ Mpc}^3$ to obtain a probability distribution for $\bar{n}_{\text{LNGRG}}$. We find $\bar{n}_{\text{LNGRG}} = 12 \pm 1 (100 \text{ Mpc})^{-3}$. Next, we approximate $f_{1+\Delta_{\text{LNGRG}}} \approx f_{1+\Delta_{\text{RG}}} | L_\nu \geq 10^{24} \text{ W Hz}^{-1}$, and calculate the latter through

Eq. 7.17 and Sect. 7.4.2’s subsequent procedure. Using the same $f_{1+\Delta_{\text{CW}}}$ as before, we finally obtain a probability distribution over LNGRG number density functions, which we visualise in Fig. 7.13.

As it should, the LNGRG number density function follows the same trend as the functions of Fig. 7.12; the novelty here is the prediction of absolute number densities. We reiterate that systematic errors at cluster-like densities render only the left side of the plot reliable. It predicts that, in filaments, $n_{\text{LNGRG}} \sim 10^{-1} (10 \text{ Mpc})^{-3}$, or $n_{\text{LNGRG}} \sim 10^{-4} \text{ Mpc}^{-3}$. As major filaments have volumes $V \sim 10^2 \text{ Mpc}^3$, only one in ten major filaments harbours an LNGRG.

7.4.4 COSMIC WEB DYNAMICAL STATE DISTRIBUTIONS

For each GRG and RG in the two samples, we obtain a probability distribution over T -web environment classes from the BORG SDSS extensions of [Leclercq et al. \(2015\)](#). The T -web scheme classifies the Cosmic Web at each point on the basis of its local gravitational dynamics. On an intuitive level, the scheme counts the number of dimensions along which the orbit of a test particle released at rest with respect to the environment is stable. More formally, at any location the Hessian of the gravitational potential, or tidal tensor T , has either three, two, one, or zero positive eigenvalues; the environment is interpreted correspondingly as a cluster, filament, sheet, or void.

[Leclercq et al. \(2015\)](#) have determined the T -web environment for each voxel of each BORG SDSS sample. By iterating over samples and counting the frequencies of the four environment classes on a per-voxel basis, one obtains a marginal environment class distribution $\mathbf{p} = (p_1, p_2, p_3, p_4)$ for each voxel. In Fig. 7.14, we show how each of $\mathbf{p}$ ’s components is distributed for Local Universe giants. We note the following.

- For 99% of giants, the most likely environment class is more than 50% probable. For 20% of giants, it is more likely than not that they inhabit a cluster (i.e. $p_1 > 50\%$); for 68% of giants, it is more likely than not that they inhabit a filament (i.e. $p_2 > 50\%$); and for 11% of giants, it is more likely than not that they inhabit a sheet (i.e. $p_3 > 50\%$). Finally, for 0% of giants, it is more likely than not that they inhabit a void (i.e. $p_4 > 50\%$).
- Some giants inhabit sheets (e.g. those in the bottom row of Fig. 7.3), but concluding this in individual cases rarely has a high degree of certainty: for just 0.5% of giants, it is very likely that they inhabit a sheet (i.e. $p_3 > 80\%$).
- A cluster environment is very unlikely (i.e. $p_1 < 20\%$) for a clear majority of 63% of giants, whilst a filament environment is very unlikely (i.e. $p_2 < 20\%$) for only 9% of giants; the filament environment hypothesis can thus seldom




Figure 7.14: Bayesian Cosmic Web reconstructions allow one to determine a probability distribution over the four T -web environment classes — clusters, filaments, sheets, and voids — for each GRG. Here we show a statistical summary of these probability distributions. From top to bottom, we show distributions for the components of vector $\mathbf{p} = (p_1, p_2, p_3, p_4)$, which are probability parameters.




Figure 7.15: Dynamical environment classification of giant radio galaxies (left), radio galaxies in general (centre), and the Local Universe in its entirety (right; from Table 3 of [Leclercq et al., 2015](#)). We define clusters, filaments, sheets, and voids in the T -web sense. These distributions depend on the scale to which the Cosmic Web density field is smoothed; in this case, the smoothing scale is $2.9 \text{ Mpc } h^{-1}$. If giants and RGs were scattered uniformly throughout the Cosmic Web, their distributions would be similar to that of the volume-filling fractions (VFFs). Instead, observed RGs — and observed giants in particular — favour cluster and filament environments.

be excluded. A sheet environment is very unlikely (i.e. $p_3 < 20\%$) for 71% of giants; finally, a void environment is nearly always very unlikely (i.e. $p_4 < 20\%$) — this holds for 99% of giants.

The corresponding distributions for LoTSS DR1 RGs (not shown) are largely similar. This is also clear from Fig. 7.15, in which we compare the environment class distributions of giants, LoTSS DR1 RGs, and the Local Universe in its entirety. Evidently, giants and general RGs are not scattered uniformly throughout the Cosmic Web, but trace cluster and filament environments. Remarkably, nearly a quarter of all giants occurs in clusters, which have a volume-filling fraction (VFF) of just 1–2%. It is easy to derive that the ratio between a GRG environment class probability (left pie) and the corresponding VFF (right pie) yields the GRG number density of that class relative to the cosmic mean GRG number density. Of course, the analogous statement for RGs also holds. In Table 7.3, we present Local Universe relative number densities for giants and general RGs. Clusters are strongly (giant) radio galaxy overdense; voids are strongly (giant) radio galaxy underdense.

Table 7.3: Number densities of observed Local Universe giants and general RGs relative to their cosmic means. These numbers depend on environment class definitions and the scale to which the Cosmic Web density field is smoothed; in this case, we use the T -web classification scheme and a smoothing scale of $2.9 \text{ Mpc } h^{-1}$. Clusters and filaments are (giant) radio galaxy overdense; sheets and voids are underdense.

	cluster	filament	sheet	void
	environment	environment	environment	environment
GRGs	15 ± 2	2.4 ± 0.1	0.29 ± 0.04	0.06 ± 0.04
RGs	10.6 ± 0.6	2.41 ± 0.05	0.39 ± 0.02	0.08 ± 0.02

The left and central pie charts of Fig. 7.15 can be regarded as prior distributions over Cosmic Web dynamical states for luminous giants and RGs in or near the Local Universe. This fact can be useful whenever direct BORG SDSS environment characterisation is not possible. For a randomly picked luminous giant, Fig. 7.15’s left pie chart specifies the prior probability distribution over dynamical states to be $\mathbf{p} = (23\%, 59\%, 17\%, 1\%)$. Interestingly, there is often additional information available to constrain the dynamical state. For example, multiwavelength imagery demonstrates that Alcyoneus (Oei et al., 2022a), one of the longest known radio galaxies, does not inhabit a galaxy cluster; in addition, the number of luminous galaxies in its vicinity excludes the possibility of a void environment.¹³ Combining this information with the prior, we obtain — after rounding — the posterior probability distribution $\mathbf{p} = (0\%, 80\%, 20\%, 0\%)$. We thus estimate the probability that Alcyoneus inhabits a filament to be 80%. Of course, this line of reasoning generalises to any luminous giant outside of galaxy clusters that has massive galactic neighbours in its Mpc-scale vicinity.

7.5 DISCUSSION

7.5.1 COSMIC WEB DENSITY ACCURACY

In this work, we have determined Mpc-scale total matter densities around radio galaxies in the Local Universe, and used these to compare the environments of luminous giants to those of other RGs. A natural question to ask is how accurate these densities are, and to which degree they can be used in other analyses, such as in the fitting of dynamical models to individual radio galaxies (e.g. Hardcastle, 2018). To quan-

¹³Characterising Alcyoneus’s Cosmic Web environment is of scientific interest, as it could help understand how the largest radio galaxies form. However, at $z = 0.25$, it is too far away to probe directly with the BORG SDSS.

tify accuracy, one would ideally know the ground-truth densities that the inferred densities are meant to approach. Problematically, however, no ground-truth densities are known — especially for filaments, the environment type which most radio galaxies seem to inhabit. For Local Universe clusters though, approximate masses are known. We exploit this fact to test the accuracy of our cluster densities. We immediately remark that the accuracy of our densities are likely to vary with density, so that the accuracy of cluster densities might not be informative of the accuracy of filament densities. This is because the 2LPT gravity solver of the BORG SDSS has particularly limited validity in the strongly non-linear cluster regime.

To test the cluster density accuracy, we cross-match our 260 selected giants with the SDSS DR12–based galaxy cluster catalogue of [Wen & Han \(2015\)](#). First, we select all clusters in the Local Universe ($z < z_{\text{max}}$) and discard the rest. To specify each cluster’s location, this catalogue provides the right ascension, declination, and redshift of the brightest cluster galaxy (BCG). Next, we cross-match in an angular sense, and consider a radio galaxy matched with a BCG if their angular separation is less than $3''$. We thus retain 60 BCG giants and 113 BCG LoTSS DR1 RGs. Figure 7.2 shows six members (10%) of this BCG giants sample. The perturbed jet and lobe shapes seen in the radio, and the massive ellipticals and crowded environments seen in the optical, together visually demonstrate the adequacy of our BCG giant identification. We can identify LoTSS DR1 BCG RGs with equal confidence.

As a quality check of the cluster masses, we test whether they are accurate enough to contain fingerprints of known physical effects. In particular, we test whether clusters with RG-generating BCGs are more massive than other Local Universe clusters. The AGN feedback paradigm suggests that this should be the case: in low-mass clusters, AGN feedback converts cool cores to non-cool cores within a gigayear, shutting down further AGN feedback for many gigayears to come; by contrast, high-mass clusters maintain stable AGN feedback cycles with gigayear-scale periods ([Nobels et al., 2022](#)), and thus their BCGs have a higher probability of being observed while generating RGs. The median mass of clusters with RG-generating BCGs is $M_{500} = 1.1 \cdot 10^{14} M_{\odot}$, while for other clusters it is $M_{500} = 1.0 \cdot 10^{14} M_{\odot}$. A Kolmogorov–Smirnov test yields a p -value $p = 7\%$, suggesting to reject the null hypothesis of equality in distribution only for high significance levels (e.g. $\alpha = 10\%$). We also test whether clusters with *giant*-generating BCGs are more massive than other clusters. The median mass of clusters with giant-generating BCGs is $M_{500} = 1.2 \cdot 10^{14} M_{\odot}$, while for other clusters it is $M_{500} = 1.0 \cdot 10^{14} M_{\odot}$. A Kolmogorov–Smirnov test yields $p = 5\%$, suggesting to reject equality in distribution only for rather high significance levels (e.g. $\alpha = 5\%$ or $\alpha = 10\%$).

As a first quality check of our densities, we test whether BCG giants have been




Figure 7.16: Comparison between environmental estimates for giants (large dots) and LoTSS DR1 RGs (small dots) generated by BCGs in the Local Universe. We compare cluster masses M_{500} from [Wen & Han \(2015\)](#) with relative densities $1 + \delta$ from the BORG SDSS, and colour the dots based on their cluster probability p_1 (relative to the average cluster probability $\bar{p}_1$ for giants and LoTSS DR1 RGs). For three randomly chosen giants and three randomly chosen LoTSS DR1 RGs, we visualise uncertainties. Generally, as expected, radio galaxy-generating BCGs claimed to be in more massive clusters by [Wen & Han \(2015\)](#) have higher BORG SDSS densities and T -web cluster probabilities. The contours indicate fractions of the density expected through M_{500} , and show that our relative densities fall short significantly in the cluster regime — by a factor that ranges from less than two to more than an order of magnitude. Densities of low-mass clusters are less biased than densities of high-mass clusters. We use flexible voxel method densities here.

assigned higher densities than non-BCG giants. The median relative densities for BCG giants and non-BCG giants are $1 + \delta = 8.5$ and $1 + \delta = 7.2$, respectively. A Kolmogorov–Smirnov test yields $p = 0.2\%$, indeed suggesting to reject the null hypothesis of equality in distribution. Similarly, we test whether our LoTSS DR1 BCG RGs have been assigned higher densities than our LoTSS DR1 non-BCG RGs. The median relative densities for BCG RGs and non-BCG RGs are $1 + \delta = 8.4$ and $1 + \delta = 5.9$, respectively. A Kolmogorov–Smirnov test yields $p \ll 0.1\%$, again suggesting to reject equality in distribution.

Next, we show the relation between cluster mass and total matter density in Fig. 7.16. Although it is qualitatively clear that BCG RGs in more massive clusters have higher

inferred densities, the densities are all lower than one would expect from spreading out the cluster mass over a BORG SDSS voxel. (For reference, a cluster of mass $M \sim 10^{14} M_{\odot}$ should have a $2.9 \text{ Mpc } h^{-1}$ -scale relative density $1 + \delta \sim 40$.) This systematic discrepancy cannot be attributed to errors in the cluster masses, as [Wen & Han \(2015\)](#) have made sure that their estimates are unbiased. The alternative is to conclude that our densities are biased low in the cluster regime. This is unsurprising, as the assumptions of the BORG SDSS’s 2LPT gravity solver are violated in clusters. Our densities are least accurate for high-mass clusters, where they can be off by an order of magnitude, and most accurate for low-mass clusters, where they can be less than a factor two too low.

It is unclear whether, and if so how, these accuracy characterisations can be extrapolated to the filament density regime. The increased applicability of 2LPT to the filament density regime suggests that the accuracy will increase. We therefore expect our filament densities to be correct up to a small factor of order unity.

Until Bayesian Cosmic Web reconstructions are further developed, we urge caution in using our inferred densities as precision estimates: in filament environments, they should be trusted up to a factor of order unity only; in cluster environments, for the moment, other sources (such as cluster catalogues) appear to provide more accurate density estimates. If one would insist in using the cluster densities, then Fig. 7.16 suggests that one should correct them by multiplying by a factor of ~ 5 .

7.5.2 COSMIC WEB DYNAMICAL STATE ACCURACY

Similar questions of accuracy can be raised about the Cosmic Web dynamical state distributions for luminous giants and RGs presented in the left and central pie charts of Fig. 7.15. To test the cluster GRG and cluster RG percentages, we can again use the cluster catalogue of [Wen & Han \(2015\)](#).

Cluster giants can be subdivided into BCG giants and non-BCG cluster giants. In Sect. 7.5.1, we determined that 60 out of 260 giants (23%) are BCG giants.¹⁴ We define non-BCG cluster giants as giants generated by non-BCG galaxies for which the comoving distance to the BCG is less than some Mpc-scale threshold. We obtain approximate comoving coordinates from right ascensions, declinations, and spectroscopic redshifts, assuming vanishing peculiar motion. In order to allow for a fair comparison to the BORG SDSS, we set the Mpc-scale threshold to a BORG SDSS voxel side length: $2.9 \text{ Mpc } h^{-1}$. Although much larger than a galaxy cluster virial

¹⁴In comparison, [Dabhade et al. \(2020b\)](#) found that only 20 out of 128 LoTSS DR1 giants with $z \leq 0.55$ (16%) were BCG giants. We hypothesise that we find a higher occurrence of BCG giants because our study is restricted to the Local Universe ($z < z_{\text{max}} := 0.16$), where cluster catalogues are generally more complete.

radius, this threshold does ensure that cluster galaxies with significant peculiar motion — a standard deviation $\sigma_{v_p} = 10^2\text{--}10^3 \text{ km s}^{-1}$ is typical — have a reasonable chance to be linked to their cluster. For this choice of threshold, we find 21 non-BCG cluster giants. We thus find a total of 81 cluster giants (31%), which is more than expected based on the BORG SDSS measurement ($23^{+3}_{-3}\%$) shown in Fig. 7.15’s GRG pie chart. This suggests that the BORG SDSS underpredicts the occurrence of cluster giants. Indeed, the strongly underpredicted densities of galaxy clusters discussed in Sect. 7.5.1 also leave their mark on the accuracy of our dynamical states. For example, the giant shown in the top-left panel of Fig. 7.2 and tabulated on the tenth row of Table 7.1, clearly occupies a cluster-like environment with an estimated $M_{500} = 1.3 \cdot 10^{14} M_{\odot}$, but is erroneously assigned an 80% chance to inhabit a filament.

Similarly, cluster RGs can be subdivided into BCG RGs and non-BCG cluster RGs. In Sect. 7.5.1, we determined that 113 out of 1443 selected LoTSS DR1 RGs (8%) are BCG RGs. By defining non-BCG cluster RGs in a way analogous to non-BCG cluster giants, we find 109 non-BCG cluster RGs.¹⁵ We thus find a total of 222 cluster RGs (15%), which is consistent with the BORG SDSS measurement ($16^{+1}_{-1}\%$) shown in Fig. 7.15’s RG pie chart.

To assess the accuracy of the filament GRG and RG fractions shown in Fig. 7.15, we cross-match our giants and selected LoTSS DR1 RGs with the galaxy group and cluster catalogue of Tempel et al. (2017), which provides estimates of M_{200} . Wen & Han (2015) define clusters through $M_{500} \geq 0.6 \cdot 10^{14} M_{\odot}$, which translates to $M_{200} \geq 0.9 \cdot 10^{14} M_{\odot}$: a typical conversion is $M_{500} = 0.7 M_{200}$ (e.g. Pierpaoli et al., 2003). To separate groups from clusters, we therefore define groups through $M_{200} < 0.9 \cdot 10^{14} M_{\odot}$. As with our cluster tests, we associate RGs to groups on the basis of estimated comoving coordinates, assuming vanishing peculiar motion. In the case of groups, this assumption is more reasonable. We again use a $2.9 \text{ Mpc } h^{-1}$ association threshold. We associate 155 out of 260 giants (60%) and 903 out of 1443 RGs (63%) with galaxy groups. These percentages provide lower bounds to the occurrence of filament giants and RGs, as there are also filament giants and RGs that do not reside in galaxy groups. It is therefore likely that the filament GRG and RG fractions of Fig. 7.15 ($59^{+3}_{-3}\%$ and $60^{+1}_{-1}\%$, respectively) are underestimates.

Finally, as a probe of sheet and void giants and RGs, we explore the occurrence of giants and RGs for which the closest galaxy group or cluster is more than 10 Mpc away. Thus, we identify 18 out of 260 giants (7%) and 115 out of 1443 RGs (8%)

¹⁵Whereas three-quarters of all cluster giants are BCG giants, only half of all cluster RGs are BCG RGs. Thus, whenever a giant is found in a cluster, it is most likely generated by the BCG; by contrast, whenever a general RG is found in a cluster, a BCG origin is just as likely as a non-BCG origin.

as occupying sheets and voids. This suggests that the combined sheet and void GRG and RG fractions of Fig. 7.15 ($18^{+2}_{-2}\%$ and $24^{+10}_{-1}\%$, respectively) are significant overestimates.

Rounded to an accuracy of 10%, our ancillary environment analysis suggests an environment class probability distribution for luminous giants of $\mathbf{p} = (30\%, 60\%, 10\%)$, representing clusters, filaments, and a combination of sheets and voids, respectively. For RGs, it suggests $\mathbf{p} = (20\%, 70\%, 10\%)$, again representing clusters, filaments, and a combination of sheets and voids. Both distributions are in fair agreement with the BORG SDSS measurement $\mathbf{p} = (20\%, 60\%, 20\%)$. The ancillary analysis reinforces the idea that, in the Local Universe, luminous giants occupy denser environments than general RGs detectable at LoTSS-like image qualities.

7.5.3 COSMIC WEB DENSITY RESOLUTION

All Cosmic Web density and dynamical state measurements presented in this work correspond to a density field with a resolution, or *smoothing scale*, of $2.9 \text{ Mpc } h^{-1}$. In fact, given the granular nature of matter, one must always adopt a smoothing scale in order to define an informative notion of density.¹⁶ Ideally, this scale is chosen such that the resulting density field varies smoothly over the physical scales of interest — avoiding both unwanted stochastic variations over space and time (as would occur when the smoothing scale is too small), and excessive blurring of the physical phenomena under study (as would occur when the smoothing scale is too large). In the case of this work, the smoothing scale of the Cosmic Web density field could be considered too large. It is set by the limitations of the BORG SDSS 2LPT gravity solver, the ambition to cover the entire Local Universe in the direction of the SDSS MGS footprint, and — most practically — the relation between the number of voxels and the numerical cost of generating the HMC Markov chain. As galaxy clusters have typical radii of $\sim 1 \text{ Mpc } h^{-1}$, they are not resolved by the BORG SDSS; the same therefore certainly holds for galaxy groups. Thus, even if the gravitational dynamics of the reconstructions were exact, the utility of the BORG SDSS in improving dynamical model fits to individual RGs is limited: the smoothing scale of the density field is larger than the scale required to resolve (beta model-like) density profiles within filaments and clusters ($0.1\text{--}1 \text{ Mpc } h^{-1}$). At best, our densities reflect the correct multi-Mpc-scale averages, but these are still much lower than the, say, central 1 Mpc-scale densities. This is because the massive galaxies that usually host RGs tend to occupy the bottom of their local gravitational potential wells, where the large-scale density is typically

¹⁶As there are many situations in which there exists a natural smoothing scale, the smoothing scale is often left unmentioned.

higher than everywhere in the vicinity. Relying on simulations, [Oei et al. \(2023a\)](#) have estimated the relation between the 2.9 Mpc h^{-1} -scale relative total matter density $1 + \delta$ (as used throughout this work), and the 1 Mpc-scale relative IGM density $1 + \delta_{\text{IGM}}$. For galaxies with a stellar mass $M_{\star} = 10^{11} M_{\odot}$, their Fig. 11 suggests $1 + \delta_{\text{IGM}} = 10 \cdot (1 + \delta)^{0.75}$. For luminous giants, the mode of the 2.9 Mpc h^{-1} -scale relative total matter density RV is $1 + \delta = 3.5$ for the fixed voxel method and $1 + \delta = 6.3$ for the flexible voxel method (see Fig. 7.6); the corresponding host-centred 1 Mpc-scale relative IGM density is $1 + \delta_{\text{IGM}} = 23$ for the fixed voxel method and $1 + \delta_{\text{IGM}} = 40$ for the flexible voxel method. This relation in principle allows one to perform Bayesian inference of the distributions of the 1 Mpc-scale Cosmic Web density around RG hosts — possibly even conditioned on RG radio luminosity, as in Sect. 7.4.2 — by converting parametrised high-resolution distributions to the lower BORG SDSS resolution, applying heteroskedastic measurement errors, and finally comparing the resulting distribution to data. We leave this to future work.

Concerning resolution, we finally warn that care must be taken when calculating (AGN feedback-related) quantities that relate quadratically to density, such as the local bremsstrahlung emissivity: these tend to be underestimated if computed from low-resolution densities. To see why, we consider ρ , the density at some natural smoothing scale, alongside $\langle \rho \rangle$, the density averaged over some larger-than-natural smoothing scale. We assume that only the latter is known. Next, we pick a quantity $y = a\rho^2$, which we aim to compute at the larger smoothing scale: $\langle y \rangle = \langle a\rho^2 \rangle = a\langle \rho^2 \rangle$. Naively, one would be tempted to proceed in calculating $\langle y \rangle$ by assuming $\langle \rho^2 \rangle \approx \langle \rho \rangle^2$, as to use $\langle y \rangle \approx a\langle \rho \rangle^2$. However, by doing so, one risks underestimating $\langle y \rangle$, because $\langle \rho^2 \rangle \geq \langle \rho \rangle^2$ by Jensen’s inequality. In Appendix 7.A6, we explicitly calculate the ratio between $\langle \rho^2 \rangle$ and $\langle \rho \rangle^2$ for sheets and voids, filaments, and clusters. For sheets and voids, this ratio is at most $\frac{4}{3}$; for filaments, it is at most a number of order unity; for clusters, it can exceed a thousand.

7.5.4 SPECTROSCOPIC REDSHIFT SELECTION

To localise RGs in the Cosmic Web and subsequently measure their relative total matter densities, we require them to have a spectroscopic redshift. By selecting on spectroscopic redshift availability, we implicitly select on both Cosmic Web density and radio luminosity. The key reason is that spectroscopic redshifts are preferentially available for galaxies with high optical flux densities — so galaxies that are nearby,¹⁷

¹⁷There is an exception to this rule: physically large galaxies at low redshifts, such as ellipticals in nearby galaxy clusters, can be angularly too extended for SDSS fibers and are therefore sometimes excluded from SDSS’s spectroscopic survey. Excluding the RGs associated to these galaxies from our




Figure 7.17: Observations underpinning the RG radio luminosity–Cosmic Web density relation of Fig. 7.8, but with Pan-STARRS r -band Kron flux densities F_ν of the host galaxies indicated. We show uncertainty ellipses whose semi-major and semi-minor axes represent one standard deviation along each dimension. In the top panel, we show KDE Cosmic Web density distributions for $F_\nu < 1$ nJy, and for $F_\nu \geq 1$ nJy. These suggest that spectroscopic redshift selection biases the observed Cosmic Web density distributions high.

and galaxies that have high optical luminosities. Galaxies with high optical luminosities tend to inhabit dense environments and, when active, tend to generate RGs with high radio luminosities.

To demonstrate that spectroscopic redshift selection indeed affects our RG radio luminosity–Cosmic Web density relation, Fig. 7.17 revisits the same Local Universe LoTSS DR1 RG data as shown in Fig. 7.8, but now with the Pan-STARRS r -band Kron flux densities F_ν of the host galaxies indicated. The top panel of Fig. 7.17 shows

samples introduces a bias against high density environments. This artificially vacates the right sides of Figs. 7.5 and 7.8, and leads towards an underestimation of the occurrence of cluster RGs in our environment analysis.




Figure 7.18: Overview of radio luminosity selection resulting from spectroscopic redshift selection, radio surface brightness limitations, and the correlation between radio and optical luminosity. The top panel compares radio luminosity distributions at $\nu_{\text{ref}} = 150$ MHz of LoTSS DR1 RGs in the constrained part of the BORG SDSS, distinguishing between those with spectroscopic redshifts (green) and those without (grey). (We perform KDE on samples of $\log_{10}(L_{\nu}(\nu_{\text{ref}}) \cdot (\text{W Hz}^{-1})^{-1})$ rather than of $L_{\nu}(\nu_{\text{ref}})$, and thus the probability densities on the vertical axis are dimensionless.) Radio galaxies with spectroscopic redshifts dominate (77%) the population. The bottom panel shows the two subpopulations in radio luminosity–redshift parameter space, with RGs without spectroscopic redshifts marked (grey circles). We show uncertainty ellipses whose semi-major and semi-minor axes represent half a standard deviation along each dimension. One retrieves the top panel by collapsing the data on the horizontal axis.

KDEs of $1 + \Delta_{\text{RG,obs}}$ for hosts with $F_{\nu} < 1$ nJy, and for hosts with $F_{\nu} \geq 1$ nJy. Hosts with low optical flux densities tend to inhabit more tenuous environments than hosts with high optical flux densities. By requiring spectroscopic redshifts, we have discarded RGs with only photometric redshifts, whose hosts almost exclusively have $F_{\nu} < 1$ nJy. As a result, our estimates of $1 + \Delta_{\text{RG,obs}}$, $1 + \Delta_{\text{RG}}$, and $1 + \Delta_{\text{RG}} | L_{\nu} = l_{\nu}$ are biased high. The severity of the bias depends on the percentage of RGs that have been

selected out. As discussed in Sect. 7.3.2, 23% of all LoTSS DR1 RGs within the constrained BORG SDSS volume lack a spectroscopic redshift. Our sample of luminous giants is much less affected: only 7% of the population within the constrained BORG SDSS volume lack spectroscopy. Thus, correcting for spectroscopic redshift selection would shift the Cosmic Web density distribution for general RGs to somewhat lower densities, while it would not appreciably shift the corresponding distribution for luminous giants. The result would be an even bigger discrepancy between the Cosmic Web density distributions of luminous giants and general RGs than established in this work, strengthening our results.

Spectroscopic redshift selection also leaves an imprint on the sample’s radio luminosities. The top panel of Fig. 7.18 shows that the radio luminosities of the 1870 LoTSS DR1 RGs in the constrained BORG SDSS volume are distributed differently between the subpopulations with and without spectroscopic redshifts. In particular, the subpopulation with spectroscopic redshifts is more spread out over radio luminosity compared to the subpopulation without: both very low and very high values are more probable.¹⁸ To understand why, we turn to the bottom panel of Fig. 7.18. The LoTSS 6'' noise level, the inverse square law, and cosmological surface brightness dimming together determine an RG radio luminosity detection threshold that increases with redshift. This threshold induces the hatched region, which is inaccessible to LoTSS-like radio observations. As a straightforward consequence, RGs with low radio luminosities can be detected only at low redshifts. It is not immediately clear whether the host galaxies of low-radio luminosity RGs at low redshifts have low or high optical flux densities. On the one hand, low radio luminosities suggest low optical luminosities: both forms of spectral luminosity correlate positively with stellar mass (as evinced by e.g. Figs. 2 of Mahajan et al., 2018; Sabater et al., 2019) and therefore¹⁹ positively with each other. However, one must again be cautious of selection effects, which can inject strong but spurious luminosity–luminosity correlations in observations (Singal et al., 2019). On the other hand, observing low-redshift galaxies is easier, again because of the inverse square law and cosmological surface brightness dimming. The bottom panel of Fig. 7.18, and in particular its lower left corner, shows that the correlation between radio and optical luminosity is not strong enough to deny the observational advantages of low redshifts: low-radio luminosity RGs at low redshifts have high optical flux densities. As a result, spectroscopic redshifts are more common for such RGs. Similarly, the correlation between radio and optical luminosity explains that spectroscopic redshifts are more common for high-radio lu-

¹⁸The distribution is not only more dispersed, but is also shifted rightwards: the median and mean radio luminosities of the spectroscopic subpopulation are 1.3 and 3.2 times larger, respectively.

¹⁹There exist cases in which this argument does not hold: correlation is not necessarily transitive.

minosity RGs.

This indirect radio luminosity selection effect is not expected to change the radio luminosity–Cosmic Web density relation inferred in this work and shown in e.g. Figs. 7.8 and 7.9. This is because the inferred relation concerns the Cosmic Web density distribution at given radio luminosity. Changes in the number of RGs available in each decade of radio luminosity — and in particular an increase in the number of RGs with $10^{22} \text{ W Hz}^{-1} < L_\nu(\nu = 150 \text{ MHz}) < 10^{23} \text{ W Hz}^{-1}$ — do not materially affect our analysis, as we already perform data rebalancing that downweights our sample’s most populated decades.²⁰

7.5.5 FUTURE OUTLOOK

Excitingly, there appear to be many opportunities to improve our understanding of the large-scale environments of radio galaxies through Cosmic Web reconstructions, both with current and near-future data.

In this work, we have used LoTSS DR1 Local Universe RLAGN with spectroscopic redshifts from [Hardcastle et al. \(2019\)](#). With the release of the LoTSS DR2 radio data ([Shimwell et al., 2022](#)) and the associated optical cross-matching results ([Hardcastle et al., 2023](#)), it will become possible to redo our analysis on a footprint that is more than 13 times larger. Doing so would increase the size of the general RG sample by an order of magnitude, from $\sim 10^3$ to $\sim 10^4$. In turn, this would allow us to revisit Sect. 7.4.2’s RG radio luminosity–Cosmic Web density relation, but with all of Fig. 7.8’s six radio luminosity decades fully evenly weighted. If the positive correlation between radio luminosity and Cosmic Web density is genuine, it will emerge more clearly upon weighting evenly. Such a LoTSS DR2 sample might even be large enough to meaningfully investigate possible differences in the radio luminosity–Cosmic Web density relation for different types of RGs, such as those generated by quasars, and those generated by other AGN. In the near future, the WEAVE–LOFAR Survey ([Smith et al., 2016](#)) will generate $\sim 10^6$ spectra of LOFAR-selected galaxies, which will help alleviate the effects of spectroscopic redshift selection described in Sect. 7.5.4. Finally, the availability of optical spectra for all AGN in the current analysis and for those in similar follow-ups means that a comprehensive classification into HERGs and LERGs, and into HEGRGs and LEGRGs, appears possible. In turn, this classification would allow for studies into the relationships

²⁰The hatched region of Fig. 7.18 shows that radio surface brightness selection causes our Local Universe RG sample to be strongly incomplete in the radio luminosity decades 10^{21} – 10^{22} and 10^{22} – $10^{23} \text{ W Hz}^{-1}$. Extending our sample in the lowest of these decades upweights this decade’s importance, likely revealing that the positive scaling between RG radio luminosity and Cosmic Web density is stronger than we suggest in Sect. 7.4.2.

between AGN accretion mode, RG growth, and the surrounding Cosmic Web. As these AGN reside in the Local Universe, many could be further characterised by cross-matching with X-ray catalogues.

Another promising extension of our work would be to use Cosmic Web reconstructions that are complementary to or improve upon the BORG SDSS. Notably, the BORG 2M++ (Jasche & Lavaux, 2019) offers one such possibility: it provides density reconstructions for the entire sky and at a higher spatial resolution, and features more accurate selection functions, bias modelling, and gravitational dynamics than the BORG SDSS. However, its reconstructions are limited to $z_{\text{max}} = 0.1$, whereas those of the BORG SDSS extend to $z_{\text{max}} = 0.16$. Using the BORG SDSS and the BORG 2M++ in conjunction would allow for an extension of the number of giants for which a density can be determined: the catalogue of giants by Oei et al. (2023a), introduced in Sect. 7.2, contains 83 giants with spectroscopic redshifts $z_s < 0.1$ that lie outside of the constrained BORG SDSS volume. Adding these to the 260 giants for which we have already determined densities would increase the number of BORG giants by 32%. Perhaps more importantly, there are 89 giants for which both BORG SDSS and BORG 2M++ density estimates are possible. Comparing these densities allows for a quality assessment of the BORG SDSS density measurements at large. In particular, given its better gravity solver, we expect BORG 2M++ cluster densities to perform much better in Sect. 7.5.1’s Cosmic Web density–cluster mass test. In addition to the BORG 2M++, additional BORG runs are on their way. Besides, there exist comparable Cosmic Web reconstruction frameworks, such as COSMIC BIRTH (e.g. Kitaura et al., 2021). In any case, thanks to the combination of new spectroscopic galaxy surveys and advances in computing power, high-fidelity Bayesian Cosmic Web reconstructions will become available for an increasingly large fraction of the Local Universe and beyond. We expect that Cosmic Web density $1 + \delta$ will therefore become a standard RG observable, alongside properties such as projected proper length l_p and radio luminosity l_ν .

By building on the work of Leclercq et al. (2015), we have presented probabilistic Cosmic Web environment classifications, based on the *T*-web definition, for 260 giants and 1443 LoTSS DR1 RGs. However, Leclercq et al. (2015) provides additional environmental classifications for the BORG SDSS volume that allow us to measure and understand the Cosmic Web environments of giants and other radio galaxies in complementary ways. Future work could explore and compare RG environments in the *T*-web, DIVA, ORIGAMI, and LICH sense (Leclercq et al., 2016, 2017).

7.6 CONCLUSION

Using Bayesian Cosmic Web reconstructions of the Local Universe, we compare the large-scale environments of giant radio galaxies with those of the radio galaxy population in general. In particular, we measure multi-Mpc-scale Cosmic Web densities around the hosts of $\sim 10^2$ luminous giants and $\sim 10^3$ general RGs. This reveals that the currently observable population of giants inhabits denser regions of the Cosmic Web than general RGs — contradicting the popular hypothesis that giants should primarily form in the dilute Cosmic Web. Our interpretation is that high jet powers, as implied by the high radio luminosities of the known population of giants, enable these RGs to overcome the IGM’s resistance to their Mpc-scale growth.

Next, we quantify the relation between radio luminosity, a proxy for jet power, and Cosmic Web density; to our knowledge, this is the first time this has been done. Our radio luminosity–conditioned Cosmic Web density distributions reinforce the idea that AGN that generate weak jets primarily reside in the dilute Cosmic Web, while AGN that generate powerful jets primarily reside in the dense Cosmic Web. We subsequently show that the radio luminosity–Cosmic Web density relation, which we have inferred using our sample of general RGs only, can accurately explain the density distribution observed for luminous giants.

Additional evidence corroborates that luminous giants form more often in the dense Cosmic Web than radio galaxies in general. Classifying Cosmic Web environments as clusters, filaments, sheets, or voids on the basis of the local gravitational field, we find that luminous giants occur more often in clusters than radio galaxies in general. We also find evidence that clusters with giant-generating BCGs are more massive than other clusters, while the evidence for the analogous claim for clusters with RG-generating BCGs is weaker.

We now summarise our main findings in more detail.

1. Using the BORG SDSS, we have performed the most physically principled measurements yet of the Cosmic Web environments of giant radio galaxies and radio galaxies in general. Our characterisations require spectroscopic redshifts, and are currently confined to the Local Universe ($z < z_{\text{max}} := 0.16$). In particular, we have determined $2.9 \text{ Mpc } h^{-1}$ –scale total matter densities and gravitational environment classifications for 260 giants, of which 208 (80%) have recently been discovered through the LoTSS DR2, and for 1443 LoTSS DR1 RGs.
2. While the marginal probability distributions of the Cosmic Web density around individual RGs are approximately lognormal, the probability distribution of

the Cosmic Web density for the RG population as a whole is approximately gamma. We provide a physical argument that motivates why the gamma distribution could have arisen. The probability distribution of the Cosmic Web density for luminous giants also appears roughly gamma, but favours higher densities. We show that heteroskedastic measurement errors, which are a general feature of our density determination method, spuriously shift the observed density distributions towards lower densities. We demonstrate a forward modelling method to correct for heteroskedasticity.

3. We quantify the radio luminosity–Cosmic Web density relation for general radio galaxies. We find that the $2.9 \text{ Mpc } h^{-1}$ –scale Cosmic Web density distribution at a given 150 MHz radio luminosity is well described by $1 + \Delta_{\text{RG}} | L_\nu = l_\nu \sim \Gamma(k, \theta)$, where $k = 4.8 + 0.2 \cdot \mathbb{J}$, $\theta = 1.4 + 0.02 \cdot \mathbb{J}$, and $\mathbb{J} := \log_{10}(l_\nu \cdot (10^{23} \text{ W Hz}^{-1})^{-1})$. This result shows that more luminous RGs tend to live in denser regions of the Cosmic Web. Treating giants as ordinary radio galaxies with $l_\nu \geq l_{\nu, \text{min}} := 10^{25} \text{ W Hz}^{-1}$, we use this relation to predict the Cosmic Web density distribution for luminous giants. The prediction is consistent with the *observed* distribution for giants: a Kolmogorov–Smirnov test yields a p -value of 0.7. Whether less luminous giants — assuming they exist — also form more often in the dense Cosmic Web, remains to be seen when surveys with increased surface brightness sensitivity commence. If such giants obey the same radio luminosity–Cosmic Web density relation as other RGs, this will not be the case.
4. We show that our methodology enables the inference of radio galaxy number densities as a function of Cosmic Web density. Whether or not the RG number density is a strong function of Cosmic Web density, depends on the radio luminosity considered. In clusters, the number densities of high-luminosity RGs (e.g. 10^{25} or $10^{27} \text{ W Hz}^{-1}$) are $\sim 10^2$ times higher than on average; in the same environment type, the number densities of low-luminosity RGs (e.g. 10^{21} or $10^{23} \text{ W Hz}^{-1}$) are just $\sim 10^1$ times higher than on average. Furthermore, we obtain tentative evidence that the number densities of low-luminosity RGs peak in low-mass clusters, before decreasing again. A possible explanation could be that, as cluster mass increases, the underlying population of active galaxies starts favouring the generation of high-luminosity RGs over low-luminosity RGs.
5. We use the BORG SDSS T -web classification to generate probability distributions over four Cosmic Web structure types — clusters, filaments, sheets, and

voids — for both luminous giants and general RGs. Luminous giants inhabit clusters more often ($23^{+30}_{-3}\%$ versus $16^{+10}_{-1}\%$). Independently, by cross-matching our radio galaxies with a cluster catalogue, we find that the former percentage presumably is an underestimate, while confirming the latter percentage. The same cross-matching procedure reveals tentative evidence ($p = 5\%$) that clusters with giant-generating BCGs are more massive than other clusters. At the same time, evidence that clusters with RG-generating BCGs are more massive than other clusters is weaker ($p = 7\%$), even though the sample involved is almost twice as big.

Although RG environment characterisations with Bayesian Cosmic Web reconstructions are still far from perfect, we have demonstrated that the current generation of reconstructions already allows for addressing open questions in radio galaxy research. Given the pitfalls of the BORG SDSS, we advise to treat every measured RG density with skepticism, and in particular warn that the densities of cluster RGs are strongly underestimated. Nevertheless, given the ever-increasing number of galaxies with spectroscopic redshifts, a growing list of technical improvements, and the widespread availability of computing power, exciting contemporary and near-future opportunities exist to expand and improve upon the results of this work. It is well possible that Cosmic Web density will soon become a staple quantity that helps to unravel the physics of radio galaxies.

M.S.S.L. Oei, R.J. van Weeren, and A. Botteon acknowledge support from the VIDI research programme with project number 639.042.729, which is financed by the Dutch Research Council (NWO). M.S.S.L. Oei warmly thanks Rafaël Mostert, Erik Osinga, and Wendy Williams for helpful discussions and the sharing of data. LOFAR data products were provided by the LOFAR Surveys Key Science project (LSKSP; <https://lofar-surveys.org/>) and were derived from observations with the International LOFAR Telescope (ILT). LOFAR (van Haarlem et al., 2013) is the Low Frequency Array designed and constructed by ASTRON. It has observing, data processing, and data storage facilities in several countries, which are owned by various parties (each with their own funding sources), and which are collectively operated by the ILT foundation under a joint scientific policy. The efforts of the LSKSP have benefited from funding from the European Research Council, NOVA, NWO, CNRS-INSU, the SURF Co-operative, the UK Science and Technology Funding Council and the Jülich Supercomputing Centre.

DATA AVAILABILITY

Oei et al. (2023a)’s LoTSS DR2 GRG data are available at the Centre de Données astronomiques de Strasbourg (CDS) via <https://cdsarc.cds.unistra.fr/cgi-bin/qcat?J/A+A/>. The LoTSS DR1 RG data with optical identifications from the value-added catalogue of Williams et al. (2019) are available at https://lofar-surveys.org/dr1_release.html; those interested in Hardcastle et al. (2019)’s RLAGN subsample should contact Martin J. Hardcastle. The BORG SDSS data release is publicly available at https://github.com/florent-leclercq/borg_sdss_data_release. In-

dividual BORG SDSS HMC Markov chain samples are not available in the public domain; those interested should contact Florent Leclercq. The cluster catalogue of [Wen & Han \(2015\)](#) is available at the CDS.

7.A1 COSMIC WEB LOCALISATION ACCURACY WITH SPECTROSCOPIC AND PHOTOMETRIC REDSHIFTS

In this appendix, we show that RGs with spectroscopic redshifts can be reliably localised within Cosmic Web reconstructions, while RGs with only photometric redshifts generally cannot.

To determine the accuracy of our radio galaxy localisations in the Cosmic Web, we Monte Carlo simulate a radial comoving distance distribution for each RG and compare its dispersion to the BORG SDSS voxel length. Each RG will have a peculiar velocity v_p with respect to us. We treat v_p as a zero-mean Gaussian random variable (RV): $v_p \sim \mathcal{N}(0, \sigma_{v_p})$, and choose a standard deviation σ_{v_p} representative of conditions in low-mass galaxy clusters: $\sigma_{v_p} = 100 \text{ km s}^{-1}$. Similarly, we treat the measured redshift z_m as an RV, and again assume Gaussianity: $z_m \sim \mathcal{N}(\mu_{z_m}, \sigma_{z_m})$. Our catalogue provides the parameters μ_{z_m} and σ_{z_m} for each RG. Using Eqs. 7.23 from left to right, we calculate the relative peculiar velocity RV β_p , the peculiar velocity redshift RV z_p , and the cosmological redshift RV z_c :

$$\beta_p := \frac{v_p}{c}; \quad z_p = \sqrt{\frac{1 + \beta_p}{1 - \beta_p}} - 1; \quad z_c = \frac{1 + z_m}{1 + z_p} - 1. \quad (7.23)$$

Finally, we calculate the radial comoving distance RV $r = r(z_c) = \|\mathbf{r}\|_2$. The distributions of r are approximately Gaussian: for each of the 248 giants in the BORG SDSS volume with a spectroscopic redshift and an associated error, we perform the Shapiro–Wilk test on 1000 Monte Carlo samples and reject the null hypothesis that r is Gaussian for one at significance level $\alpha = 0.01$.²¹ For these giants, 95% of the standard deviations σ_r are between $1.00 \text{ Mpc } h^{-1}$ and $1.12 \text{ Mpc } h^{-1}$; the sample median is $1.06 \text{ Mpc } h^{-1}$. By contrast, the 21 giants in the BORG SDSS volume with only a photometric redshift and an associated error have a σ_r that ranges from $20 \text{ Mpc } h^{-1}$ to $110 \text{ Mpc } h^{-1}$; for details, see Table 7.4. To localise a giant in the BORG SDSS, precision up to the scale of a voxel, which are approximately $2.9 \text{ Mpc } h^{-1}$ long, is

²¹ If one uses sufficiently many Monte Carlo samples and a fixed significance level, any deviation from Gaussianity will become significant. Indirectly, the sample size and significance level together specify a degree of non-Gaussianity identifiable with such Shapiro–Wilk tests.

necessary — but sub-voxel localisation is redundant. We conclude that giants with a spectroscopic redshift can be tied to an individual voxel, even when taking peculiar velocity into account; on the other hand, giants with only a photometric redshift can be mislocalised more than a typical filament length, and are therefore not subjectable to a precise environment analysis. Of course, these conclusions extend to any radio galaxy with an associated host, or indeed to any galaxy.²²

As only few (i.e. $\sim 10^2$) giants are currently known in the volume reconstructed by the BORG SDSS, it is worthwhile to perform host galaxy spectroscopic follow-up for those with photometric redshifts only. We provide a list with targets in Table 7.4.

7.A2 COSMIC WEB DENSITY DISTRIBUTION: THE GAMMA ANSATZ

In this work, we have modelled several RG Cosmic Web density RVs, such as $1 + \Delta_{\text{RG,obs}}$, $1 + \Delta_{\text{RG}}$, and $1 + \Delta_{\text{RG}} | L_\nu = l_\nu$, as gamma variates. This choice is largely driven by data (and a preference for simplicity) rather than by theory, although the requirements of continuity and a strictly positive support have a physical basis. Gamma distributions adequately fit $1 + \Delta_{\text{RG,obs}}$ (see the bottom row of Fig. 7.6), and their use in modelling $1 + \Delta_{\text{RG}} | L_\nu = l_\nu$ is justified by the RG prediction of $1 + \Delta_{\text{RG,obs}}$, which matches observations (see Fig. 7.11). Here, we post hoc theorise why the gamma distribution arises in the current context.

The late-time density in a voxel is, of course, equal to the late-time mass in the voxel divided by its (fixed) volume of $(2.9 \text{ Mpc } h^{-1})^3$. Therefore, up to a dimensionful constant, the late-time density and the late-time mass have the same probability distribution. The late-time mass M in the voxel equals the sum of the mass aggregated over cosmic time. If we are to consider a simplistic treatment of mass aggregation, or structure formation, which environment is most relevant? The central pie chart of Fig. 7.15 suggests that most observed RGs have filament environments. In addition, [Pasini et al. \(2021\)](#) suggest that within filaments, most RGs inhabit galaxy groups, which grow by merging with other groups. Therefore, we consider a proto-filament — i.e. a massive structure that extends essentially along a single dimension — on which we consider a Poisson point process with constant spatial rate. The Poisson points represent peaks in the density field where galaxy groups arise. As with any homogeneous Poisson process, the distances between the points are exponentially distributed. It is now helpful to partition the filament in cells, each associated to a single

²²By jointly inferring the large-scale density field and the radial distances to the galaxies used to constrain the former, it is possible to reduce individual photometric redshift uncertainties ([Jasche & Wandelt, 2012](#); [Tsaprazi et al., 2023](#)), but these reductions are not extensive enough to alter our conclusions.

Table 7.4: Overview of all giants in the constrained BORG SDSS volume without a spectroscopic redshift. Roughly speaking, this volume corresponds to the SDSS DR7 MGS footprint up to $z_{\text{max}} = 0.16$. The total matter density of the ambient Cosmic Web can only be reliably determined after spectroscopic follow-up. The giants are sorted (in descending order) by their radial comoving distance standard deviation σ_r , obtained through Monte Carlo simulation, taking into account photometric redshift errors and a random, zero-centred peculiar velocity component with standard deviation $\sigma_{v_p} = 100 \text{ km s}^{-1}$. After spectroscopic follow-up, the projected proper length l_p must be revised.

rank ↓	SDSS DR12 name	photometric redshift (1)	σ_r (Mpc h^{-1})	l_p (Mpc)
1	J122009.84+194833.0	0.15 ± 0.04	110	0.8
2	J155503.00+280430.9	0.10 ± 0.03	100	1.3
3	J134211.92+565839.3	0.11 ± 0.03	70	0.7
4	J130444.37+511119.0	0.09 ± 0.02	70	0.8
5	J122129.95+662644.0	0.10 ± 0.02	70	1.1
6	J152151.93+570635.0	0.07 ± 0.02	50	0.8
7	J112033.54+233559.0	0.12 ± 0.02	50	0.9
8	J140046.38+301900.0	0.06 ± 0.01	40	0.8
9	J090640.80+142522.9	0.13 ± 0.01	40	1.2
10	J112248.98+565243.5	0.08 ± 0.01	30	0.8
11	J085022.99+383547.2	0.14 ± 0.01	30	0.7
12	J145102.13+301227.0	0.16 ± 0.01	30	0.8
13	J152229.23+281911.9	0.13 ± 0.01	30	0.8
14	J140044.26+125219.8	0.11 ± 0.01	30	0.7
15	J090128.15+145158.1	0.14 ± 0.01	30	1.2
16	J145827.21+331312.0	0.11 ± 0.01	20	1.0
17	J152024.13+310557.6	0.06 ± 0.01	20	0.9
18	J102135.20+420022.3	0.13 ± 0.01	20	1.5
19	J172715.85+585220.5	0.15 ± 0.01	20	2.1
20	J141033.40+405932.4	0.14 ± 0.01	20	1.0
21	J134339.20+195301.7	0.15 ± 0.01	20	0.9

point. Assuming that the filament has a constant cross-sectional area and an approximately constant density, the volumes and masses of the cells are simply proportional to the distances between the points, and are thus also exponentially distributed. Finally, if a fixed fraction of the mass in each cell collapses into the cell's galaxy group, then the group masses within the filament — at early times — are, again, exponentially distributed. A group that exists at late times has built up its mass $\mathcal{M}$ by aggregating the masses of early-time groups. Assuming a constant temporal rate of early-time group aggregation, there exists a typical number N of early-time groups that con-

tribute to the mass of a given late-time group: $M = \sum_{i=1}^N M_i$. Now we invoke the fact that the sum of a fixed number of independent and identically distributed RVs with an exponential distribution is an RV with a gamma distribution. Thus, as the M_i are exponentially distributed, M is gamma distributed. This, in turn, implies that the late-time densities of the voxels in which late-time groups fall, are gamma distributed also. As late-time RGs trace late-time groups (and especially the more massive ones), the late-time RG density distribution should be approximately gamma.

In the future, by cross-matching RG catalogues with galaxy group catalogues with accurate masses, the mass distribution of RG-hosting groups in the Local Universe can be sampled and tested against the gamma distribution hypothesis.

7.A3 MODELLING RELATIVE BARYON DENSITY MEASUREMENT HETEROSKEDASTICITY

We must describe the measured RG relative baryon density given a true RG relative baryon density with a continuous probability distribution that has the positive half-line as its support; preferably, we use a simple analytical prescription.

Motivated by the quality of the fits shown in Fig. 7.5, we assume $1 + \Delta_{\text{RG,obs}} \mid 1 + \Delta_{\text{RG}} = 1 + \delta_i \sim \text{Lognormal}(\mu, \sigma^2)$, and likewise for giants. The parameters μ and σ^2 are determined by assumptions on the RV's mean and variance. In particular, we assume that measured RG relative baryon densities are unbiased estimators of the underlying true RG relative baryon densities and that their variances are the same when the underlying true RG relative baryon densities are the same:

$$\mathbb{E} [1 + \Delta_{\text{RG,obs}} \mid 1 + \Delta_{\text{RG}} = 1 + \delta_i] = 1 + \delta_i; \quad (7.24)$$

$$\mathbb{V} [1 + \Delta_{\text{RG,obs}} \mid 1 + \Delta_{\text{RG}} = 1 + \delta_i] = g(1 + \delta_i), \quad (7.25)$$

where $g(1 + \delta)$ is some non-negative function characterising the heteroskedasticity. For the current work, we adopt a two-parameter power-law variance model $g(1 + \delta) := a(1 + \delta)^b$. From basic identities of the lognormal distribution, we find for this choice of g

$$\mu = \ln(1 + \delta) - \frac{1}{2} \ln \left(1 + a(1 + \delta)^{b-2} \right); \quad (7.26)$$

$$\sigma^2 = \ln \left(1 + a(1 + \delta)^{b-2} \right). \quad (7.27)$$

Using the fact that

$$f_{1+\Delta_{\text{RG,obs}}}(1+\delta) = \int_0^\infty f_{1+\Delta_{\text{RG,obs}} | 1+\Delta_{\text{RG}}=x}(1+\delta) \cdot f_{1+\Delta_{\text{RG}}}(x) dx, \quad (7.28)$$

we can deduce the distribution of $1 + \Delta_{\text{RG}}$ given some parametrisation of it and a procedure like MLE. We choose $1 + \Delta_{\text{RG}} \sim \Gamma(k, \theta)$ on the basis that $1 + \Delta_{\text{RG,obs}}$ appears to be much better described by a gamma distribution than by a lognormal distribution and that the effect of heteroskedasticity appears minor. We determine a and b from data.²³

7.A4 RADIO LUMINOSITY–COSMIC WEB DENSITY RELATION: FIXED VOXEL METHOD

In Sect. 7.4.2, we have presented a quantification of the relation between RG radio luminosity and Cosmic Web density. In particular, we infer a probability distribution for the 2.9 Mpc h^{-1} –scale total matter density around the hosts of Local Universe radio galaxies with a given 150 MHz radio luminosity — which can range from 10^{21} to 10^{27} W Hz $^{-1}$. The MLE model shown in Fig. 7.8 and the one- and two-dimensional posterior marginals shown in Fig. 7.9, are based on densities measured through Sect. 7.3.1’s flexible voxel method. In this appendix, through Figs. 7.19 and 7.20, we present the analogous results based on densities measured through Sect. 7.3.1’s fixed voxel method. Both methods reveal a positive scaling between RG radio luminosity and Cosmic Web density.

7.A5 RELATIVE NUMBER DENSITY DERIVATION

We let the RV $1 + \Delta_{\text{CW}}$ represent the 2.9 Mpc h^{-1} –scale relative density at a randomly chosen point in the contemporary Cosmic Web, and let $f_{1+\Delta_{\text{CW}}}$ be its PDF. Similarly, we let the RV $1 + \Delta_{\text{RG}}$ represent the 2.9 Mpc h^{-1} –scale relative density at a randomly chosen RG in the contemporary Cosmic Web, and let $f_{1+\Delta_{\text{RG}}}$ be its PDF. We consider a cosmologically sized, comoving volume of extent V , in which a total of N_{RG} radio galaxies exist. The subvolume in which the relative density is between $1 + \delta$ and $1 +$

²³For the fixed voxel method, we find $a = 0.4$ and $b = 1.1$; for the flexible voxel method, we find $a = 0.4$ and $b = 0.9$. We stick to a simple prescription here, but there appears to be enough BORG SDSS data to describe the heteroskedasticity with a more accurate (though more complex) model while still avoiding overfitting.




Figure 7.19: PDFs of the observed RG relative total matter density RV given a radio luminosity at $\nu_{\text{ref}} = 150$ MHz, using MLE parameter values for the model described in Sect. 7.4.2. The black contours denote CDF values. We overplot all 1443 selected LoTSS DR1 RGs (dots), with those above the empirical median density coloured red, and those below coloured blue. We use fixed voxel method densities here. For flexible voxel method densities, see Fig. 7.8.

$\delta + d\delta$ has extent

$$dV = V f_{1+\Delta_{\text{CW}}}(1 + \delta) d\delta. \quad (7.29)$$

The number of RGs in this subvolume is

$$dN_{\text{RG}} = n_{\text{RG}}(1 + \delta) dV, \quad (7.30)$$

where $n_{\text{RG}}(1 + \delta)$ is the number density of RGs that would arise in an environment of constant density $1 + \delta$. In practice, one never encounters environments of constant density; the average RG number density in the part of the Cosmic Web where the relative density is between $1 + \delta_1$ and $1 + \delta_2$ equals

$$\bar{n}_{\text{RG}}(1 + \delta_1, 1 + \delta_2) = \frac{\int_{\delta_1}^{\delta_2} n_{\text{RG}}(1 + \delta) f_{1+\Delta_{\text{CW}}}(1 + \delta) d\delta}{\int_{\delta_1}^{\delta_2} f_{1+\Delta_{\text{CW}}}(1 + \delta) d\delta}. \quad (7.31)$$




Figure 7.20: Posterior distribution over $k(0)$, $\theta(0)$, $k'(0)$, and $\theta'(0)$, based on selected Local Universe LoTSS DR1 RGs. We show all two-parameter marginals of the likelihood function, with contours enclosing 50% and 90% of total probability. We mark the maximum likelihood estimate (MLE) values (grey dot) and the maximum a posteriori (MAP) values (grey cross). The one-parameter marginals again show the MLE (dash-dotted line), a mean-centred interval of standard deviation-sized half-width (hashed region), and a median-centred 90% credible interval (shaded region). We use fixed voxel method densities here. For flexible voxel method densities, see Fig. 7.9.

As a result, the cosmic mean RG number density $\bar{n}_{\text{RG}}$ is given by

$$\bar{n}_{\text{RG}} := \int_{-1}^{\infty} n_{\text{RG}}(1 + \delta) f_{1+\Delta_{\text{CW}}}(1 + \delta) d\delta. \quad (7.32)$$

The probability that an RG has a relative density between $1 + \delta$ and $1 + \delta + d\delta$, $f_{1+\Delta_{\text{RG}}}(1 + \delta) d\delta$, is

$$f_{1+\Delta_{\text{RG}}}(1 + \delta) d\delta = \frac{dN_{\text{RG}}}{N_{\text{RG}}} = \frac{n_{\text{RG}}(1 + \delta) V f_{1+\Delta_{\text{CW}}}(1 + \delta) d\delta}{N_{\text{RG}}}. \quad (7.33)$$

We note that

$$N_{\text{RG}} = \bar{n}_{\text{RG}} V, \quad (7.34)$$

and combine the last two equations to find

$$\frac{n_{\text{RG}}(1 + \delta)}{\bar{n}_{\text{RG}}} = \frac{f_{1+\Delta_{\text{RG}}}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.35)$$

In other words, through point-wise division of the PDFs of $1 + \Delta_{\text{RG}}$ and $1 + \Delta_{\text{CW}}$, both of which we determine in this work, we obtain the radio galaxy number density that would arise at a given density, up to a constant. This constant is (the reciprocal of) the cosmic mean RG number density; the LHS of Eq. 7.35 can thus be interpreted as the relative RG number density at $1 + \delta$ — that is, the number density relative to the cosmic mean value. Fully analogously, the relative GRG number density at $1 + \delta$ is given by

$$\frac{n_{\text{GRG}}(1 + \delta)}{\bar{n}_{\text{GRG}}} = \frac{f_{1+\Delta_{\text{GRG}}}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.36)$$

By multiplying both sides of Eq. 7.35 by $f_{L_\nu|1+\Delta_{\text{RG}}=1+\delta}(l_\nu)$, using the identity $f_{X|Y=y}(x)f_Y(y) = f_{Y|X=x}(y)f_X(x)$, and dividing both sides by $f_{L_\nu}(l_\nu)$, we obtain

$$\frac{n_{\text{RG}}(1 + \delta) \cdot f_{L_\nu|1+\Delta_{\text{RG}}=1+\delta}(l_\nu)}{\bar{n}_{\text{RG}} \cdot f_{L_\nu}(l_\nu)} = \frac{f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.37)$$

The numerator and denominator at the left-hand side are both number density *densities*: they denote radio galaxy numbers per unit of comoving volume and unit of radio luminosity. As their physical dimensions are distinct from those of ordinary number densities, we will denote them differently. Calling the numerator $\mathfrak{I}_{\text{RG}}(l_\nu, 1 + \delta)$ and




Figure 7.21: Multiplicative discrepancy between the mean squared density $\langle \rho^2 \rangle$ and the squared mean density $\langle \rho \rangle^2$ in voxelised Cosmic Web reconstructions, for density profiles of clusters (red), filaments (orange), and sheets or voids (blue). The length scale $R \sim L$, where L is the voxel side length; here, $L = 2.9 \text{ Mpc } h^{-1}$. The asymptote (grey) corresponds to value $\frac{4}{3}$.

the denominator $\tilde{\Gamma}_{\text{RG}}(l_\nu)$, we write^{2.4}

$$\frac{\tilde{\Gamma}_{\text{RG}}(l_\nu, 1 + \delta)}{\tilde{\Gamma}_{\text{RG}}(l_\nu)} = \frac{f_{1+\Delta_{\text{RG}}|L_\nu=l_\nu}(1 + \delta)}{f_{1+\Delta_{\text{CW}}}(1 + \delta)}. \quad (7.38)$$

7.A6 DISPARITY BETWEEN MEAN SQUARED DENSITY AND SQUARED MEAN DENSITY

Several physical quantities that relate to AGN and RGs, such as the bremsstrahlung emissivity around the host galaxy, scale with the square of IGM density ρ . In Sect. 7.5.3, we have argued qualitatively that such quantities are likely to be underestimated if calculated using low-resolution densities $\langle \rho \rangle$, such as the ones offered by the BORG SDSS. In this appendix, we explicitly calculate how much bigger $\langle \rho^2 \rangle$ is than $\langle \rho \rangle^2$ — in the sense of a multiplicative Jensen’s gap. Our calculations demonstrate a strong

^{2.4} As all n -like symbols in the Roman and Greek alphabets are already in use, we use the Phoenician root $\tilde{\Gamma}$.

dependence on Cosmic Web environment.

When the true density field varies over scales much larger than a voxel side length L , such as in sheets or voids, the density field variation within a voxel resembles a gradient from $(1 - \varepsilon) \langle \rho \rangle$ to $(1 + \varepsilon) \langle \rho \rangle$, where $\varepsilon \in [0, 1]$. If the planes of constant density are parallel to two of the voxel's faces, we can derive

$$\frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2} = \frac{\varepsilon^2}{3} + 1, \text{ so that } 1 \leq \frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2} \leq \frac{4}{3}. \quad (7.39)$$

We show the multiplicative Jensen's gap of Eq. 7.39 in Fig. 7.21 (blue curve).

For a filament whose baryon density is modelled with an isothermal single- β model (Cavaliere & Fusco-Femiano, 1976, 1978) with $\beta = \frac{2}{3}$ and core radius r_c , we have

$$\frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2} = \frac{\left(\frac{R}{r_c}\right)^4}{\left(1 + \left(\frac{R}{r_c}\right)^2\right) \ln^2 \left(1 + \left(\frac{R}{r_c}\right)^2\right)}, \quad (7.40)$$

where we consider a cylinder with length L and radius R , where R is such that the area of a cylindrical section perpendicular to the axis equals L^2 , the area of a voxel face. Thus, $R = \frac{1}{\sqrt{\pi}} L$. We show the multiplicative Jensen's gap of Eq. 7.40 in Fig. 7.21 (orange curve).

When the true density field varies over scales much smaller than a voxel side length L , such as in clusters, the ratio between the mean of the squared density and the square of the mean density can be much larger than 1. In particular, for a Navarro–Frenk–White (NFW) profile (Navarro et al., 1996) with scale radius r_s ,

$$\frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2} = \frac{1 - \left(\frac{r_s}{r_s + R}\right)^3}{9 \left(\frac{r_s}{R}\right)^3 \left(\ln \frac{r_s}{r_s + R} + 1 - \frac{r_s}{r_s + R}\right)^2}. \quad (7.41)$$

where we consider a sphere with radius R , where R is such that the volume of the sphere is equal to the volume of the voxel: $R = \left(\frac{3}{4\pi}\right)^{\frac{1}{3}} L$. We show the multiplicative Jensen's gap of Eq. 7.41 in Fig. 7.21 (red curve). In the limit $\frac{r_s}{R} \rightarrow \infty$, $\frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2} \rightarrow \frac{4}{3}$.

