Calibrated and completeness-corrected optical stellar density maps of the northern Galactic plane

H. J. Farnhill,[1][*] J. E. Drew,[1][*] G. Barentsen[1,2] and E. A. González-Solares[3]

[1]*School of Physics, Astronomy & Mathematics, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, UK*
[2]*NASA Ames Research Center, Mail Stop 245-30, Moffett Field, CA 94035, USA*
[3]*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 OHA, UK*

ABSTRACT

Following on from the second release of calibrated photometry from IPHAS, the INT/WFC Photometric Hα Survey of the Northern Galactic Plane, we present incompleteness-corrected stellar density maps in the r and i photometric bands. These have been computed to a range of limiting magnitudes reaching to 20th magnitude in r and 19th in i (Vega system), and with different angular resolutions – the highest resolution available being 1 arcmin^2. The maps obtained cover 94 per cent of the 1800 square degree IPHAS footprint, spanning the Galactic latitude range, $-5° < b < +5°$, north of the celestial equator. The corrections for incompleteness, due to confusion and sensitivity loss at the faint limit, have been deduced by the method of artificial source injection. The presentation of this method is preceded by a discussion of other more approximate methods of determining completeness. Our method takes full account of position-dependent seeing and source ellipticity in the survey data base. The application of the star counts to testing reddened Galactic disc models is previewed by a comparison with predicted counts along three constant-longitude cuts at $\ell \simeq 30°$, $90°$ and $175°$: some overprediction of the most heavily reddened $\ell \simeq 30°$ counts is found, alongside good agreement at $\ell \simeq 90°$ and $175°$.

Key words: atlases – dust, extinction – Galaxy: disc – Galaxy: stellar content – Galaxy: structure.

1 INTRODUCTION

The positioning of the Solar system almost in the equatorial plane of the Milky Way places a significant obstacle in the way of understanding the structure of our own Galaxy. Despite the fact that the disc of the Milky Way is effectively the largest object in the night sky, offering vastly superior angular resolution compared to that achievable for any other galaxy, sightlines at low Galactic latitude remain a challenge to decipher because of large and variable amounts of dust extinction. Given that the formation and maintenance of galactic discs is an important problem in galaxy evolution (van der Kruit & Freeman 2011), an improved vision of our own galactic disc is needed. And as our home in the Universe, it is of interest in its own right.

The reliable empirical determination of the three-dimensional (3D) distribution of the Milky Way's interstellar dust, needed to make sense of the disc, is now becoming possible through increasingly sophisticated analyses of comprehensive digital survey data (Drimmel & Spergel 2001; Marshall et al. 2006; Sale et al. 2009, 2014). In the next 5–10 years, these advances will complement the

astrometric harvest being gathered by the *Gaia* mission and usher in a much better, sharper vision of the 3D Milky Way. Nevertheless, the anticipated *Gaia* parallax precision at fainter magnitudes (100 µas at $G \sim 19$) will still leave stars beyond the first 1–2 kpc with distances known to a precision no better than ∼10–20 per cent. Accordingly, it remains useful to supplement our knowledge through the application of other methods that can test predictive Galactic models. One of these, which puts to good use the uniquely detailed view we have of the Galactic disc, is comparison with magnitude-limited star counts. This approach has been applied successfully in the past and continues to be influential in guiding the content of Galactic models (Robin et al. 2003; Czekaj et al. 2014).

So far, deeper optical star-count mapping has only been carried out in the southern Galactic plane (Ruphy et al. 1997). The options to conduct a well-calibrated stellar density mapping of the northern sky are now appearing in the wake of data product releases from wide-area digital imaging surveys (Stoughton et al. 2002; Magnier et al. 2013). The greatest challenge is at low Galactic latitudes, but even this is now becoming tractable as the 1 arcsec angular-resolution IPHAS (INT/WFC Photometric Hα Survey of the Northern Galactic Plane; Drew et al. 2005) and UVEX (Groot et al. 2009) surveys approach completion. In particular, the recent release of IPHAS DR2 (Barentsen et al. 2014), offering uniform photometry of the

[*] E-mail: hywel@farnhill.io (HJF); j.drew@herts.ac.uk (JED)

northern Galactic plane in *r*, *i* (and Hα), provides the basis for a precise and deep stellar density map across almost 1800 square degrees.

In this work, we describe the construction of *r*, *i* stellar density maps at a range of angular resolutions, up to a maximum of 1 arcmin^2, based on IPHAS DR2. In order that this mapping is not hampered by variable observing conditions and correspondingly variable levels of source loss, great care has been taken to make corrections for incompleteness by evaluating the results of artificial source injection tailored to every survey field. In the Galactic plane where confusion can present as a significant issue, it is important to do this reliably. Because literature presentations of source injection and other methods of correction have so far been limited [but see e.g. Harvey et al. (2006) for more than usual detail], we present a reasonably full description of how we arrive at the position- and magnitude-dependent corrections applied. This is presented in Sections 3.3–3.7, after a brief restatement of the main features of the IPHAS survey in Section 2 and some appraisal via two more approximate techniques in Sections 3.1 and 3.2.

The resulting corrected *r*, *i* stellar density maps are described and briefly discussed in Section 4. The maps themselves, with maximum magnitude limits of *r* = 19 and *i* = 18, respectively (Vega system), are presented as plots within Appendix A and are provided within machine-readable online supplementary material. In Section 5, a first purely illustrative comparison is made between counts derived from the *i* < 18 map and Galactic model predictions. This uses two distinct reconstructions of the 3D distribution of interstellar extinction in order to gain a first impression of how influential the choice of extinction map can be. Interestingly, the results from this preliminary exploration are already mixed: we find that at $\ell \simeq 90°$ and 175°, the agreement between the star counts and model prediction is good – but at the lowest Galactic longitude examined, $\ell \simeq 30°$, a significant discrepancy appears.

The paper ends in Section 6 with some concluding remarks, including comment on the best use of these star-count maps.

2 OBSERVATIONS

2.1 IPHAS broad-band photometry

The most recent release, DR2, of the IPHAS was presented by Barentsen et al. (2014), while the basic specification of IPHAS was set down by Drew et al. (2005). This survey, conducted using the Wide Field Camera (WFC) mounted on the Isaac Newton Telescope (INT) in La Palma, provides photometry of the complete northern Galactic plane within the latitude range −5° < *b* < +5°, in narrow-band Hα and broad-band Sloan *r* and *i*. It is the broad-bands that are the focus of this paper. The typical 5σ magnitude limit reached in the survey is 21.2 in *r* and 20.0 in *i*, achieved at a median seeing of 1.1 arcsec. The typical external photometric precision of DR2 is close to 0.03 mag, as judged by comparisons with Sloan Digital Sky Survey (SDSS) DR9 data. In Galactic longitude, this survey spans the range 30° < ℓ < 215°. In time it will be fully complemented by the VPHAS+ survey (Drew et al. 2014) covering the southern Galactic plane and bulge.

A highly characteristic feature of the IPHAS survey, due to its targeting the Galactic plane, is that the stars dominating the faint end of the captured *r*- and *i*-band magnitude distributions are typically either intrinsically red or highly reddened. Consequently, a reasonable generalization to apply is that at *r* = 19, stars captured by the survey are typically brighter in *i* by around 1 mag (in the Vega system, adopted by the survey – see figs 15 to 17 in Barentsen

et al. 2014). Because of this, in the presentation that follows, comparisons between the *r* and *i* density maps obtained are generally drawn between those derived for *r* = 19 and *i* = 18 limits.

Whilst IPHAS is certainly uniform in its execution, the weather at the telescope represented an important variable. The observations were obtained via standard allocations of time, with the result that a wide variety of observing conditions are contained within the survey data base. Thanks to the opportunity to obtain repeats of fields exposed in poor conditions, it was possible to apply quality cuts in the preparation of IPHAS DR2 so as to omit clearly inferior data, whilst still achieving over 92 per cent coverage of the survey footprint. Even so, there is still a broad quality range within DR2 in terms of both measured widths of the point spread function (PSF) and limiting magnitude (for full details, see Barentsen et al. 2014). This variation, along with over a factor of 10 contrast in the typical observed stellar density along the northern plane (see fig. 3 in González-Solares et al. 2008), makes it important to achieve uniformity of outcome through careful position-dependent completeness corrections. This task, along with the development and application of an algorithm that ensures reliable source counting without duplication and proper accounting for small gaps and other irregularities in the data, is described here (see also Farnhill 2015).

The first step in this process is to establish a working definition of a detected 'star' in the survey data.

2.2 On morphological classification and the expected number density of extragalactic sources

A fraction of sources detected in DR2 fields are known to be misclassified as non-stellar towards field edges – this is due to progressive distortion in the PSF with increasing distance from the optical axis (OA) of the INT/WFC system. This pattern of behaviour has a small impact on the final DR2 catalogue, which provides primary detections that are selected according to a criterion of lowest distance to the OA. But when counting stellar sources (morphological classes −1 and −2) from a single field to produce a stellar density map, the issue of progressive misclassification has to be tackled. Fig. 1 shows an example of a field exhibiting this problem.

Assuming that a constant ratio of stellar to non-stellar sources applies in reality across any one field, comparisons between 5 square degree regions at field centres and edges revealed occasional misclassification of several hundred sources. Of course, some of these objects are not misclassified – they can be genuine examples of galaxies presenting as extended objects. To determine the expected number of galaxies detected in a similar area, galaxy counts were taken from Yasuda et al. (2001), down to the 5σ *r*- and *i*-band
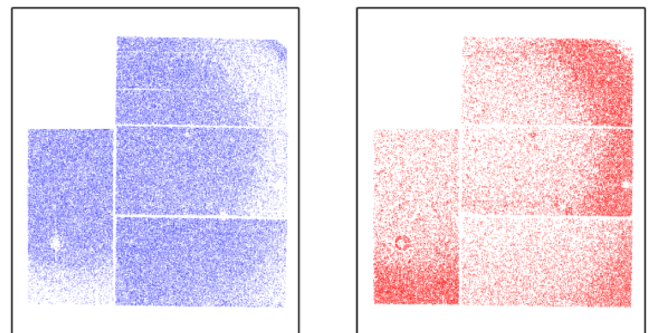

Figure 1. Distributions of sources classified as stellar (*rClass* = −1, blue) and non-stellar (*rClass* = +1, red) in the *r* band for IPHAS field 4975o_aug2004a.

limits of IPHAS, and then extinguished by the median Schlegel, Finkbeiner & Davis (1998) extinction values for all fields [corrected using the Schlafly & Finkbeiner (2011) recalibration]. Based on median limiting magnitudes, ∼1 extragalactic source is predicted per 5 square degree region. Against stellar densities typically orders of magnitude larger, this is very small. It was therefore regarded as safe to include 'non-stellar' +1 classified sources for density mapping purposes – the losses that would be suffered on leaving them out, in fields similar to that depicted in Fig. 1, are much larger than is the likely contribution from galaxies.

3 COMPLETENESS CORRECTION

Genuine astronomical sources falling within IPHAS detection limits can fail to appear in the resulting photometric catalogues for a number of reasons.

Detector issues can prevent sources from being picked up: the WFC is a 4-CCD mosaic leaving gaps between the component detectors, and there are also bad columns and regions of vignetting that will hinder detection. In the majority of such cases (but not quite all), missing sources will be picked up in the offset partner pointing, or in the overlap with a neighbouring field.

More significant and pervasive losses affecting the final star counts are those due either to confusion promoted by high stellar densities or to statistical sensitivity losses at the faint limit for detection. In the Galactic plane at lower longitudes, where the stellar density in IPHAS can reach to more than 100 000 per square degree, confusion will be especially important and liable to determine the effective magnitude limit. Outside the solar circle where the typical stellar density is around 20 000 per square degree, confusion has only marginal significance: in this domain the median 1.1 arcsec PSF implies 70 'beams' per source, to be compared with the rule of thumb confusion threshold of 30 per source (see Hogg 2001). Both kinds of incompleteness are exacerbated by relatively poor seeing. We first consider approximate methods of correction, to gain insight into the role source confusion is playing within the IPHAS catalogues. We then go on to describe our chosen method of incompleteness correction based on artificial source injection.

3.1 Confusion, as estimated from nearest-neighbour analysis

Confusion, the effect of background noise caused by unresolved sources, scales with source density (Condon 1974). For randomly distributed sources, the probability of finding a given number of neighbours within angular distance θ can be described by a Poisson distribution, which leads to the nearest-neighbour distribution

$$n(\theta) = 2\rho^2\Omega\pi\theta e^{-\rho\pi\theta^2}, \tag{1}$$

where ρ is the density of stars per unit solid angle and Ω is the sky solid angle of interest, as presented in Bahcall & Soneira (1983).

Fig. 2 shows this distribution for a sky area matching that of the WFC field of view (0.29 square degrees) containing different numbers of sources. It can be seen that increasing the density of sources increases the number of nearest neighbours at small separations, while lowering the number at greater separations, thereby pulling in the peak separation value, θ_{max}. Equation (1) holds for a population of sources randomly distributed regardless of brightness. A posteriori checks have revealed that, at the ∼1 arcsec resolution characteristic of IPHAS, the angular distribution of stars in the northern Galactic plane conforms quite well to the random case: essentially we find that the Rn statistic or the ratio of observed mean nearest-neighbour separation to the predicted (random) value
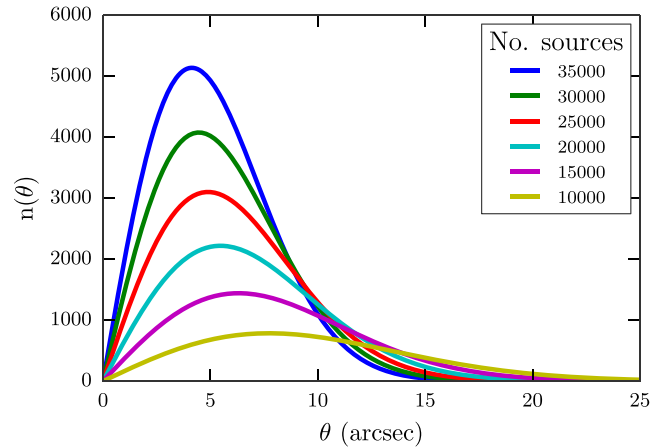

Figure 2. Theoretical distribution (see equation 1) of nearest-neighbour distances for an area with the effective area of a WFC image, for fields containing varying numbers of sources.

is commonly between 1 and 1.1, per field. If clustering were significant, or indeed close pairs due to stellar multiplicity were frequent, Rn would be pulled down well below unity. In this context, it is worth noting the following feature of IPHAS star counts: the magnitude range from ∼13th to 20th begins to sample the Galactic plane from distances of 0.5–1 kpc, at the bright end, out to 4 or more kpc depending on sightline. Accordingly IPHAS star counts say probe the Galactic disc beyond the first kiloparsec, rather than the solar neighbourhood.

In order to reproduce the theoretical distribution (equation 1), all sources in the area Ω would need to be recovered. This will never be the case in dense fields, which will modify the observed distribution in the sense that the smallest separations will be underreported – appearing to relatively boost the proportion of nearest neighbours at larger values of θ.

Three test cases are shown in Fig. 3: these IPHAS fields, located around $\ell \sim 32°$, were chosen because they lie in a region of the survey containing fields spanning a wide range of densities. The $n(\theta)$ distributions were generated by selecting subregions of the CCDs such that a border of width 30 arcsec was excluded – the nearest neighbour would then be identified for each source within the central subregion. This approach eliminates from consideration sources close to CCD edges whose nearest neighbours are likely to be located in the image plane beyond the detector edge, whilst it rightly permits the nearest-neighbour search for each central-region source to extend into the borders. The border width was set at 30 arcsec after obtaining initial nearest-neighbour distributions which showed that the great majority of nearest neighbours occurred at smaller separations than this (see e.g. Fig. 2).

The fields are displayed in order of increasing density. The leftmost field shown contains 3966 sources – a number sufficiently low that confusion would not be expected to make a significant contribution to incompleteness. The fields presented in the middle and right-hand panels, containing 24 291 and 49 150 sources respectively, suffer from increasing levels of confusion. The overplotted theoretical distribution gives an idea of the amount of confusion; at low field densities, the empirical and theoretical distributions agree quite well, while at higher densities the theoretical distributions clearly predict closer nearest neighbours than measured.

While the observed nearest-neighbour distribution is affected by confusion most significantly at smaller ($\lesssim$10 arcsec) separations, the observed number of nearest neighbours at larger separations
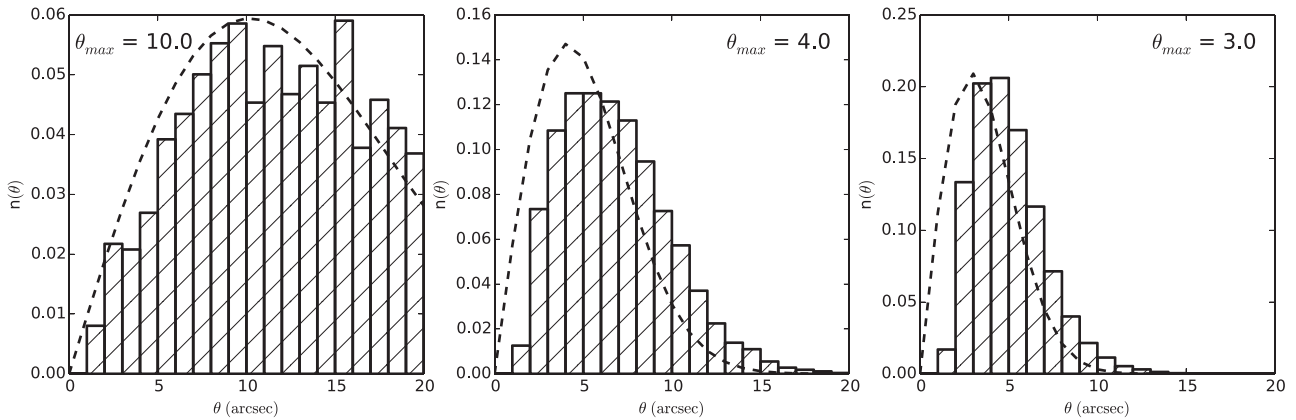
Figure 3. Nearest-neighbour distributions for three IPHAS fields, normalized to allow the overplotting of theoretical distribution (dashed line). Density of fields increases from left to right.

will be less affected. Hence, the theoretical distribution best fitting the tail at large θ values comes close to describing the field as if confusion were not an issue. We have exploited this property as a means to gauge confusion loss at small separations. We have fitted distributions generated from equation (1) to the observed distributions at >10 arcsec separation, varying ρ such that the number of sources in a field ranged from half to twice its observed value. By calculating χ^2 for each fit at large separations, the best-fitting value of ρ was found for each, identifying the correction that should be applied for confusion.

Fig. 4 shows the effect of confusion on the survey as a whole. The best-fitting $n(\theta)$ distribution (corresponding to a theoretical unconfused source count) was obtained for each IPHAS field, allowing a picture to be built up of the impact of confusion. The statistic chosen to represent the magnitude of the correction is the ratio of $\frac{\text{observed}}{\text{predicted}}$ sources, where the 'predicted' source count is the count corresponding to the best-fitting theoretical $n(\theta)$ distribution. In order to understand the variation in confusion at different limiting magnitudes m_0, fits were performed on sources brighter than, respectively, 18th, 19th and 20th magnitudes in r, in addition to fits across all sources.

In a perfect survey capable of deblending overlapping sources, all fields would exhibit zero confusion, with their curves jumping from zero to 100 per cent of all fields at a ratio of 1.0. In reality, some sources will always be lost to confusion, with the impact increasing – the higher the density of objects. Fainter objects suffer more from confusion losses as their PSFs are more likely to be lost in the vicinity of brighter sources, in addition to their intrinsically higher densities (see Fig. 6 for examples of the magnitude distributions of IPHAS fields). Fig. 4 demonstrates this behaviour as the incompleteness takes hold for a greater number of fields as the cut-off magnitude m_0 is increased.

For a fraction of fields at all values of m_0, the best fit to their nearest-neighbour distributions indicates a lower predicted number of sources than is actually observed. Indeed, the median completeness at brighter m_0 shows signs of converging to 0.98–0.99. However, we notice that the maximum ratio returned is an unphysical ~1.1, for fields which lie in regions of lowest density in the map of the plane. We take this as an indication that the fitting to the $n(\theta)$ distribution tails has an associated uncertainty of ~10 per cent. So large an uncertainty indicates that this method for correcting source counts is not itself sufficiently exact to adopt it for application across the survey. A significant contributing factor to this uncertainty is the adoption of a single nearest-neighbour distribution for each IPHAS
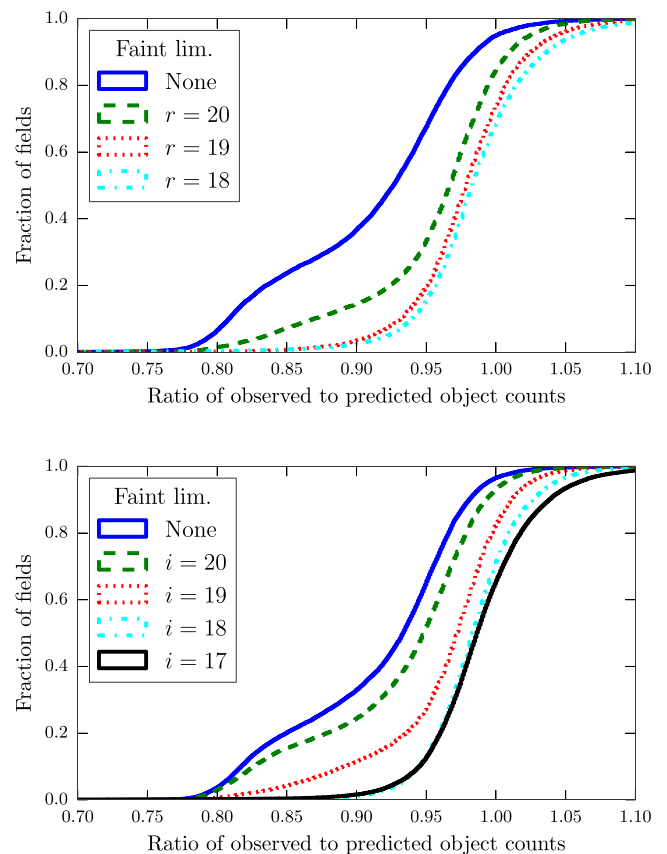

Figure 4. Cumulative histogram of $\frac{\text{observed}}{\text{predicted}}$ source counts in r (upper panel) and i (lower panel), where the predicted counts were estimated by fitting equation (1) to the tail of the nearest-neighbour distribution. The curves in both panels represent the ratio as determined for sources down to different limiting magnitudes, plus an instance where no limit was placed on the magnitude of sources.

field. Marked variations or gradients in stellar density (due particularly to changing extinction) on scales of several arcminutes within fields will introduce error into the corrected count inferred from the tail of the nearest-neighbour distribution – this is behind the unphysical underprediction noted above. Cutting fields up into smaller subregions would bear down on this source of error, at the price of reducing the numbers of objects contributing to each distribution.
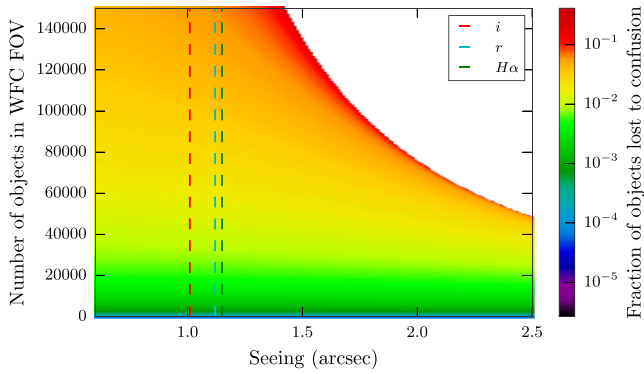
Figure 5. Correction term in equation (2) for the range of $n_{src} = \rho' \times \Omega_{WFC}$ and θ_{FWHM}. Median seeings for each filter are marked by vertical lines. The white region denotes the domain where equation (3) is true and equation (2) breaks down. Note that vertical scale converts to the source count per square degree on dividing by the WFC footprint of 0.29 square degrees.

The more fundamental limitation on the utility of this technique is that it does not easily provide completeness as a function of stellar apparent magnitude.

Nevertheless, Fig. 4 can be regarded as a first demonstration that, as a general rule, at $r < 19$ the impact of confusion in IPHAS is small while, as might be expected, it is close to ubiquitous at $r > 20$. Source counts in the i band are higher thanks to the lower extinction – the same plot for this band shows that confusion becomes a minor consideration once $i < 18$.

3.2 Other approximate measures of completeness

A formula for correcting for the number of sources lost to confusion, that is based upon equation (1), may be written as

$$\rho = -\rho' \frac{\log \left(1 - 4\rho'\pi\theta_{FWHM}^2\right)}{4\pi\theta_{FWHM}^2}, \tag{2}$$

where ρ is the actual source density, ρ' is the observed density and θ_{FWHM} is the seeing in which the field was observed. The expression was derived by Irwin & Trimble (1984), and it was used to correct the observed $n(\theta)$ distributions directly in González-Solares et al. (2008). Clearly, fields with poorer seeing and higher densities will be subject to larger corrections; González-Solares et al. (2008) reported that the IPHAS Initial Data Release fields suffering from greatest confusion were missing 41 per cent of their sources.

However, equation (2) does not hold for the entire $[\rho', \theta_{FWHM}]$ space covered by all IPHAS DR2 fields — it breaks down in dense fields, where the seeing is appreciably worse than the median. Fig. 5 shows the variation of the correction term for the range of parameters relevant to DR2. The white region shows the parameter space in which equation (2) is not applicable, where

$$4\pi\rho\theta^2 > 1. \tag{3}$$

Although this domain is only strictly entered by fewer than 20 fields, the performance of this approximate formalism – well suited to low-density halo fields – degrades as the domain is approached at the high densities encountered away from the mid-plane inside the solar circle. For this reason, its deployment across the survey is not the preferred option.

An approach used in several previous large-area surveys (Ruphy et al. 1997; Cambrésy et al. 2002; Lucas et al. 2008) is to estimate the completeness limit by taking the magnitude distribution of sources and identifying at which magnitude the distribution begins to drop off. As part of a study of stellar populations in the Galactic plane nearly two decades ago, Ruphy et al. (1997) noted that evaluating completeness based on recovery from simulated images is quantitatively the superior option, but for simplicity kept to the alternative of a completeness limit based on star-count histograms of uncrowded fields. This approach may serve well in studies of limited sky regions, or if sources fainter than the deduced completeness limit can be set aside without causing problems. In this same study, confusion due to crowding was recognized as the dominant source of incompleteness, whilst 'slight variations due to the observing conditions' were judged to be relatively unimportant. This last presumption would be unsafe here.

Cambrésy et al. (2002) estimated their completeness from the turnover of 2MASS magnitude distributions, and reason that their density maps would show imprints of individual observations had they overestimated their limiting magnitudes. This would certainly be the case for IPHAS – without any magnitude cut in place, the field-to-field variation of densities is extremely obvious. In applying a similar treatment of incompleteness, Lucas et al. (2008) quote the 90 per cent completeness limit of the UKIRT Infrared Deep Sky Survey (UKIDSS) Galactic Plane Survey, noting that in uncrowded fields the modal depths vary by 0.25 mag due to observing conditions. This method still depends heavily on the turnover in the magnitude distribution, and relies on 'visually extrapolating' the histogram.

A single completeness statistic as provided by Lucas et al. (2008) could be used to apply a uniform completeness correction, although this assumes that all fields in a survey suffer from the same pattern of incompleteness. Even if this were indeed the case, only the count of sources brighter than the given completeness limit could be corrected; it is not possible with this information to correct the count going deeper.

Fig. 6 shows the magnitude distributions for three IPHAS fields. A visual inspection would suggest that incompleteness sets in at ~20th mag or fainter in r and at ~19th in i and Hα. Attempts to automate the determination of turnover magnitude for all fields are hampered by the facts that the mode of these distributions is dependent on the chosen binning, and that the magnitude distributions of many fields plateau before dropping off. But it is already clear from the previous section and Fig. 4 that the full picture is more complicated, with some small residual source loss still affecting magnitudes as bright as 18–19. In short, there is good reason to address correction rigorously on a field-by-field basis.

3.3 Artificial source injection

A thorough treatment of incompleteness in any survey requires measuring its sensitivity to sources over the entire magnitude range of interest – this is best achieved by simulating observations and then processing these simulated frames in the same way as the original data. This method (in addition to the magnitude distribution method discussed in Section 3.2) was used and described by Harvey et al. (2006) as part of a study of interstellar clouds observed by *Spitzer*.

Such an approach is much more powerful, as statistics can be returned on magnitude bins, rather than on the entire magnitude distribution down to a specified faint limit. It entails the simulation of sources of all magnitudes, resembling the real data as closely as possible, whilst avoiding undue cost in terms of both computing power and time.

For the purposes of correcting the density map, the completeness of each field was assessed using the images and catalogues of CCD 4 only. Using only one of the WFC CCDs per field cuts the processing
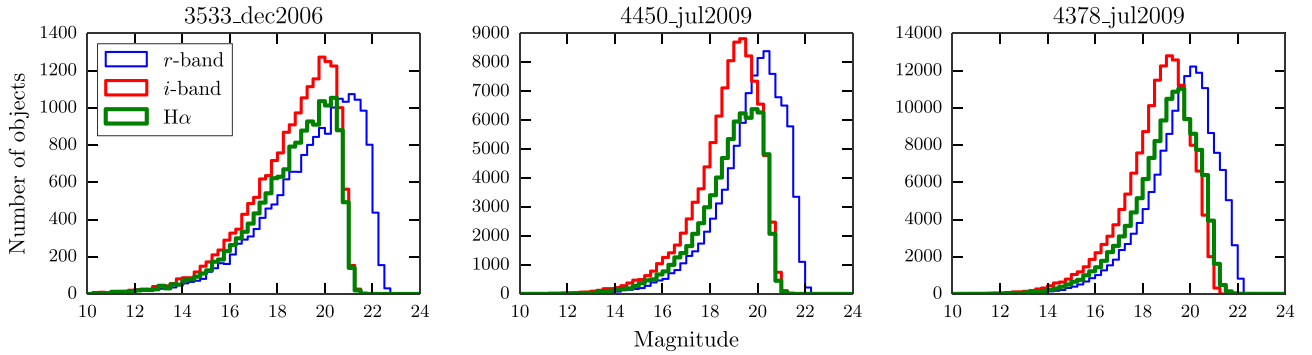
Figure 6. Magnitude distributions for three IPHAS fields of increasing density. From left to right, fields have *r*-band source counts of 18 113, 99 061 and 153 430. A requirement that sources counted here have *nBands* > 1 places the greatest constraint on *i*-band counts, as redder sources are picked up in the *i* band only.

time by a factor of 4, bringing the total time necessary to compute the completeness of the entire survey down to a little over one week (using a high-performance multi-node cluster). Due to the fact that DR2 preferentially selects sources from around field centres (closer to the OA), CCD 4 – the most central of the mosaic – was selected as the chip to represent each field (see Section 3.7 for further discussion).

The main steps of the completeness calculation were as follows. First, the typical properties of an image were determined, and used to generate synthetic sources that closely resemble genuine stellar sources (Section 3.4). The parameters of these synthetic sources were recorded and then catalogues were generated, via the pipeline software, from the new frames with the added synthetic objects. The tables of synthetic source parameters were cross-matched against the newly generated catalogues, and the rate of recovery was measured (Section 3.6). The fraction recovered at different magnitudes will allow a completeness curve to be built up for each field, resulting in tailored corrections that can then be applied across the entire survey.

3.4 Simulating stellar sources

The complexity of simulating stellar sources depends on the accuracy with which the real sources are to be recreated. The simplest approximation of stellar images to two-dimensional Gaussians can be seen to be overly simplistic from a cursory glance at the reduced IPHAS images. Almost all IPHAS frames exhibit some degree of ellipticity (up to $e = 0.2$; see Barentsen et al. 2014), requiring non-symmetric artificial sources. Elliptical 2D Gaussians were therefore chosen as sufficiently realistic profiles for artificial stellar sources (see Section 3.6 for further discussion).

The parameters returned by the aperture photometry performed on the DR2 images indicate the following should be included to accurately recreate a stellar source:

(i) flux,
(ii) FWHM,
(iii) ellipticity,
(iv) position angle and
(v) coordinates (or pixel position on detector).

The choice of each is described below.

3.4.1 Flux

The relation between total flux and magnitude was determined per field by fitting a power law to the photometry returned by IMCORE.

In order to insert a source of a given magnitude m_0, it is necessary to understand what flux in counts per unit time that source should have. The magnitude of an object is determined by

$$m = \mathrm{ZP} - 2.5 \times \log_{10} \frac{c}{t} - \mathrm{APCOR} - \mathrm{PERCORR}, \qquad (4)$$

where ZP is the photometric zero-point of the image as determined via the DR2 uniform calibration (Barentsen et al. 2014), c is the measured counts within the defined aperture, t is exposure time in seconds, while APCOR and PERCORR are small correction terms. APCOR is an aperture correction term calculated by the pipeline software, IMCORE, which uses the curve of growth of stellar sources to determine the correction required to transform the chosen aperture measurement to total flux (Irwin & Lewis 2001). PERCORR is a sky calibration correction, obtained by comparing dark sky regions with the median across each CCD. ZP, APCOR and PERCORR are all taken from the pre-existing catalogue headers.

3.4.2 PSF shape parameters: full width at half-maximum and ellipticity

Initially the parameters of genuine sources in IPHAS frames were determined by fitting elliptical two-dimensional Gaussian profiles to each source detected by IMCORE, building up lists of best-fitting parameters for the purpose of generating realistic artificial sources. Parameters of sources were excluded from these lists if they exhibited one of the following behaviours.

(i) The source position returned is further than 5 pixels from the position reported by IMCORE – these fits are likely to have been disturbed by nearby sources.
(ii) The peak value of the best-fitting Gaussian is >55 000 counts – this is the regime for the WFC where bright sources saturate, distorting the PSF.

The best-fitting full width at half-maximum (FWHM) values returned for the semi-major and semi-minor axes of the sources were gathered for each of the fields entering this exploration, and Gaussians were then fitted to the distribution to recover the distribution mean and σ. The FWHM value representative of each field returned by IMCORE fell within 1σ of the mean, confirming that the pipeline-specified value can be re-used safely for artificial source generation.

Similarly, the measure of field-wide ellipticity returned by IMCORE was found to agree with the typical ellipticities obtained in the

process of fitting Gaussians to each object across the field using the relation

$$e = 1 - \frac{\mathrm{FWHM_{min}}}{\mathrm{FWHM_{maj}}} \qquad (5)$$

(where $\mathrm{FWHM_{min}}$ and $\mathrm{FWHM_{maj}}$ are the full width at half-maxima along sources' semi-minor and semi-major axes, respectively). As with FWHM, the IMCORE ellipticity value per field was used in artificial source generation.

3.4.3 Position angle

The typical distribution of position angle (PA) within a field is quite broad and non-Gaussian in appearance. In practice, there is no reason to expect the choice of PA to significantly influence the ability to recover an artificial source – hence it is set at zero for fields where no obvious peak in the PA distribution was observed, otherwise artificial source PAs were drawn from the best fit to the peak of the PA distribution.

3.4.4 Position

Object x and y pixel positions were drawn randomly from uniform distributions across the 2048×4096 pixels of the WFC CCDs.

The completeness calculated for the unvignetted CCD 4 of each field was taken to represent the completeness over all four CCDs. This allowed the border zone – which can complicate the appraisal of completeness – to be easily defined: a rectilinear zone 10 arcsec (or 30 pixels) wide inside the CCD edges was imposed such that no artificial source could be placed within it. This made sure that no injected artificial source would lose flux across the CCD edge.

3.5 Practical implementation of artificial source injection

Ideally measuring the completeness using artificial sources would be done by adding a single source at a time, verifying whether or not it is detected. This process would be repeated as many times as required to obtain decent statistics, for the entire magnitude range of interest. This would be the route taken in the absence of any limitations on either time or computing resources. The limiting factor here is the computing time taken for the basic step of catalogue generation (20 s). To generate the number of sources detailed in Table 1, it would take ~ 65 h to process just one field.

To reduce the cost, multiple sources must be inserted into each image simultaneously. In practice, it was necessary to also ensure that too many sources were not inserted at any one time, since too high a number would modify the intrinsic properties of the image. For example, probing the completeness of an originally sparse image with a high number of artificial sources inserted would not return statistics useful for understanding the original image. The value of $\frac{\delta n}{n_{\mathrm{image}}}$ needed to be kept sufficiently low, where δn is the number of artificial sources added to the image and n_{image} is the number of sources present originally. The quality control information available for DR2 fields reports that the most sparsely populated fields contain more than 1000 stars. This is not an extremely constraining limit; a maximum of up to 50 stars was chosen as a value that would keep $\frac{\delta n}{n_{\mathrm{image}}}$ below 0.05 for all fields. Allowing up to 50 sources per artificial source injection increases processing efficiency, but offers too few to sample the pre-existing magnitude distribution of the sources in each image faithfully. Hence, the approach adopted was to split the magnitude range of interest into bins 0.25 mag wide,

Table 1. Number of artificial sources added per *i*-band magnitude bin. 'No. ofsources' denotes total number that will be generated over M runs. N is the number of sources that will be added to each image, which will be repeated M times. A total of 12 300 are added per field, across 410 images containing artificial sources. For the *r* band, the bins were shifted 1 mag fainter.

Magnitude bin				
Start	End	N	M	No. ofsources
11.0	11.25	10	10	100
11.25	11.5	10	10	100
⋮	⋮	⋮	⋮	⋮
13.5	13.75	10	10	100
13.75	13.0	10	10	100
14.0	14.25	20	10	200
⋮	⋮	⋮	⋮	⋮
15.25	15.5	20	10	200
15.5	15.75	30	10	300
⋮	⋮	⋮	⋮	⋮
16.75	17.0	30	10	300
17.0	17.25	40	10	400
⋮	⋮	⋮	⋮	⋮
17.75	18.0	40	10	400
18.0	18.25	50	10	500
18.25	18.5	50	10	500
18.5	18.75	50	15	750
18.75	19.0	50	15	750
19.0	19.25	50	20	1000
⋮	⋮	⋮	⋮	⋮
19.75	20.0	50	20	1000
Total:		410	12 300	

inserting sources from only one bin at a time in each catalogue regeneration.

Table 1 details N, the number of sources inserted per image, and M, the number of artificial images generated for reprocessing and recovery for each magnitude bin. N rises with increasing magnitude and is capped at 50 at the faint end for the reasons discussed above. To accumulate a larger number of sources injected in the faintest bins, and minimize noise where the most significant corrections arise, M increases from 10 to 15 at $m_0 = 19.5$, and then to 20 at $m_0 = 20.0$. This setup requires around 2.2 h (~ 20 s $\times$ 410 IMCORE runs). For the 14 115 fields that make up DR2, this results in a total computing time of $\sim 31\,000$ h. Utilizing 128 CPUs simultaneously, as was done here, brings the total time to estimate completeness for the entire survey in a single band to ~ 10 d.

3.6 Recovery of simulated sources

Cross-matching the recovered sources to the added artificial sources required the imposition of matching thresholds in both position and brightness. While, ideally, a recovered source should lie centred on the exact x,y pixel position where it was added, there is no guarantee that IMCORE will report unchanged centroid coordinates. Background variations in the original image, pixel binning and blending may cause the reported centroid to shift by a small amount; hence, a generous upper limit of 5 pixels (compared to a typical shift of less than 1 pixel) was set on this displacement.

Simply cross-referencing the recorded position of an artificial source with the regenerated catalogue is not sufficient to confirm
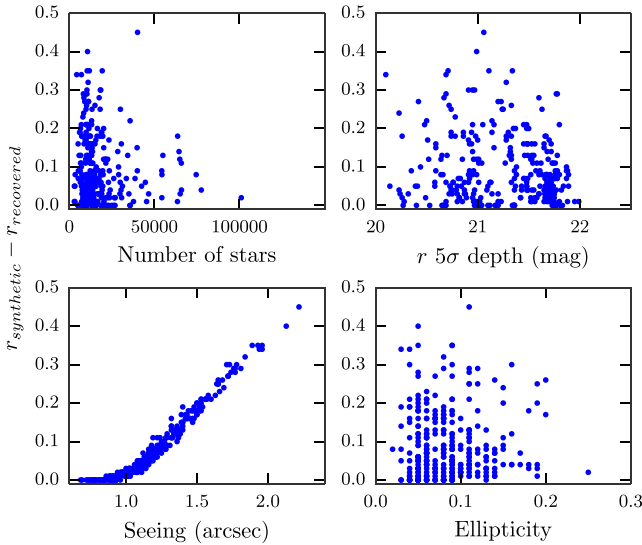
Figure 7. Modal difference between artificial inserted and recovered magnitude per field, plotted against field parameters.
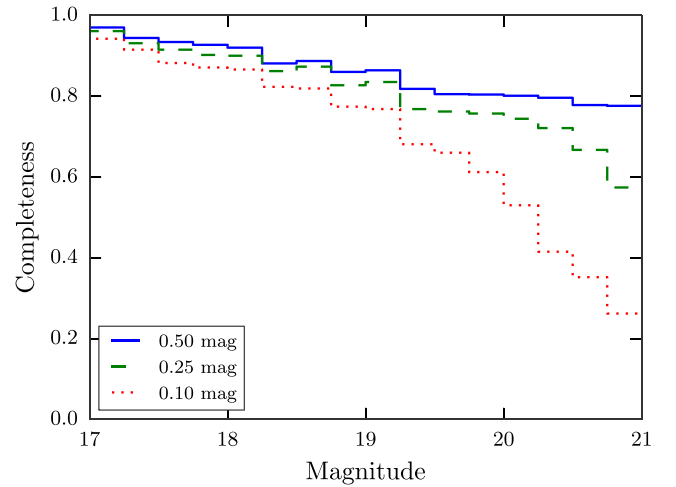

Figure 8. An example of *r*-band completeness fractions for a single field exposure obtained using offset tolerances between inserted and recovered magnitudes of 0.5 (blue solid), 0.25 (green dashed) and 0.1 mag (red dotted). The middle value of 0.25 is the adopted tolerance.

successful recovery. A constraint on the difference between the recovered magnitude $m_{recovered}$ of the source and its insertion value $m_{synthetic}$ is essential, in order to reduce spurious cross-matches between added sources and pre-existing nearby sources (typically of different brightness). Choosing a bound on this offset, $\Delta_m = m_{synthetic} - m_{recovered}$, was more involved than for pixel position. Previous attempts to use artificial photometry to understand incompleteness have faced similar issues. For example, Mateo & Hodge (1986), as part of a study of a globular cluster in the Large Magellanic Cloud, added artificial stars to images using DAOPHOT: they considered a source recovered if it was returned at the same coordinates at a magnitude within 0.5 of the inserted value. Harvey et al. (2006) used this approach in a study of stellar sources in the Serpens molecular cloud, and plotted the modulus of the average difference for artificial sources injected into their data, finding a range from 0.1 mag at ∼10–13 mag to more than 0.5 mag at ∼15 mag.

For each field, the distribution of the modulus of Δ_m was binned up and the modal value determined. For a number of fields, the modal value reached ≈0.4 mag – an effect that would require a Δ_m tolerance of >0.5 mag. Such a large tolerance would likely result in many spurious cross-matches and hence an unduly optimistic estimate of completeness fractions, especially at fainter magnitudes. To better understand this, the modal Δ_m for each field was plotted against a number of field parameters. The resulting plots, seen in Fig. 7, clearly show that this quantity correlates most strongly with the PSF FWHM (the 'seeing'). Indeed, the relationship underlying the tight trend is that the returned magnitudes from a regenerated catalogue are offset by a systematic amount rather than spread widely around the mode – a pattern that persists, up to the highest PSF FWHM encountered in the survey. This finding has motivated the application of a seeing-dependent uniform shift to all recovered magnitudes from an image, applied to bring the median Δ_m of a field back to zero. The origin of this offset is likely to be the difference between the stellar profiles of the artificial and the real sources. In this sense, the empirically determined offset and its correction are the trade for having used Gaussian profiles in place of more complex Moffat profiles.

After taking out the systematic shifts to Δ_m values, a bound on the acceptable corrected difference Δ_m^{cor} still needs to be set.

As illustrated in Fig. 8, the completeness curves remain relatively unchanged at $r \lesssim 19$ as the Δ_m^{cor} limit is varied between 0.5 and 0.1 mag. This limit was set at 0.25 mag, as a compromise between avoiding acceptance of a large number of spurious detections at the faint end of the magnitude distribution, and ensuring that few objects are missed due to the occasional larger Δ_m^{cor} value.

3.7 Completeness fractions

Fig. 9 shows the completeness curves of all DR2 fields; the discrete quantization of completeness is more prominent at the bright end due to fewer artificial sources being injected to assess the incompleteness of bright sources. The upper two cumulative histograms shown in the upper set of panels in Fig. 9 show that incompleteness remains low at bright magnitudes, with a median completeness of 99.5 per cent at $r = 15$. At 18th magnitude the median completeness of IPHAS is 97.7 per cent, falling to 93.1 per cent at 20th magnitude. The pattern in the *i* band (lower panels) is similar, except that completeness declines more rapidly with increasing magnitude such that the data for *i* at 19th magnitude more closely resemble *r* at 20th than *r* at 19th.

In order to apply completeness corrections when generating the density map, each detected source needs to contribute to the map by a factor that takes into account the incompleteness of sources of similar magnitude from its field of origin. In order to characterize the incompleteness fraction of each field as a quantity varying continuously with magnitude (*r* as written below), a function was fitted to the measured completeness fractions of each field, of the form

$$C(r) = \alpha - \gamma \times e^{\frac{r}{\beta}}, \qquad (6)$$

where α, β and γ are parameters allowed to vary to find the best fit for each field. These parameters were collected and formed a lookup table for use in correcting the density map. The curves of this form can be seen plotted for every field in both *r* and *i* in Fig. 10. Median values for the coefficients are $\alpha = 0.996$, $\beta = 1.20$, $\gamma = 8.08 \times 10^{-9}$ for the *i* band, and $\gamma = 0.999$, $\beta = 1.73$, $\gamma = 6.89 \times 10^{-7}$ for the *r* band. The curves corresponding to these median values have
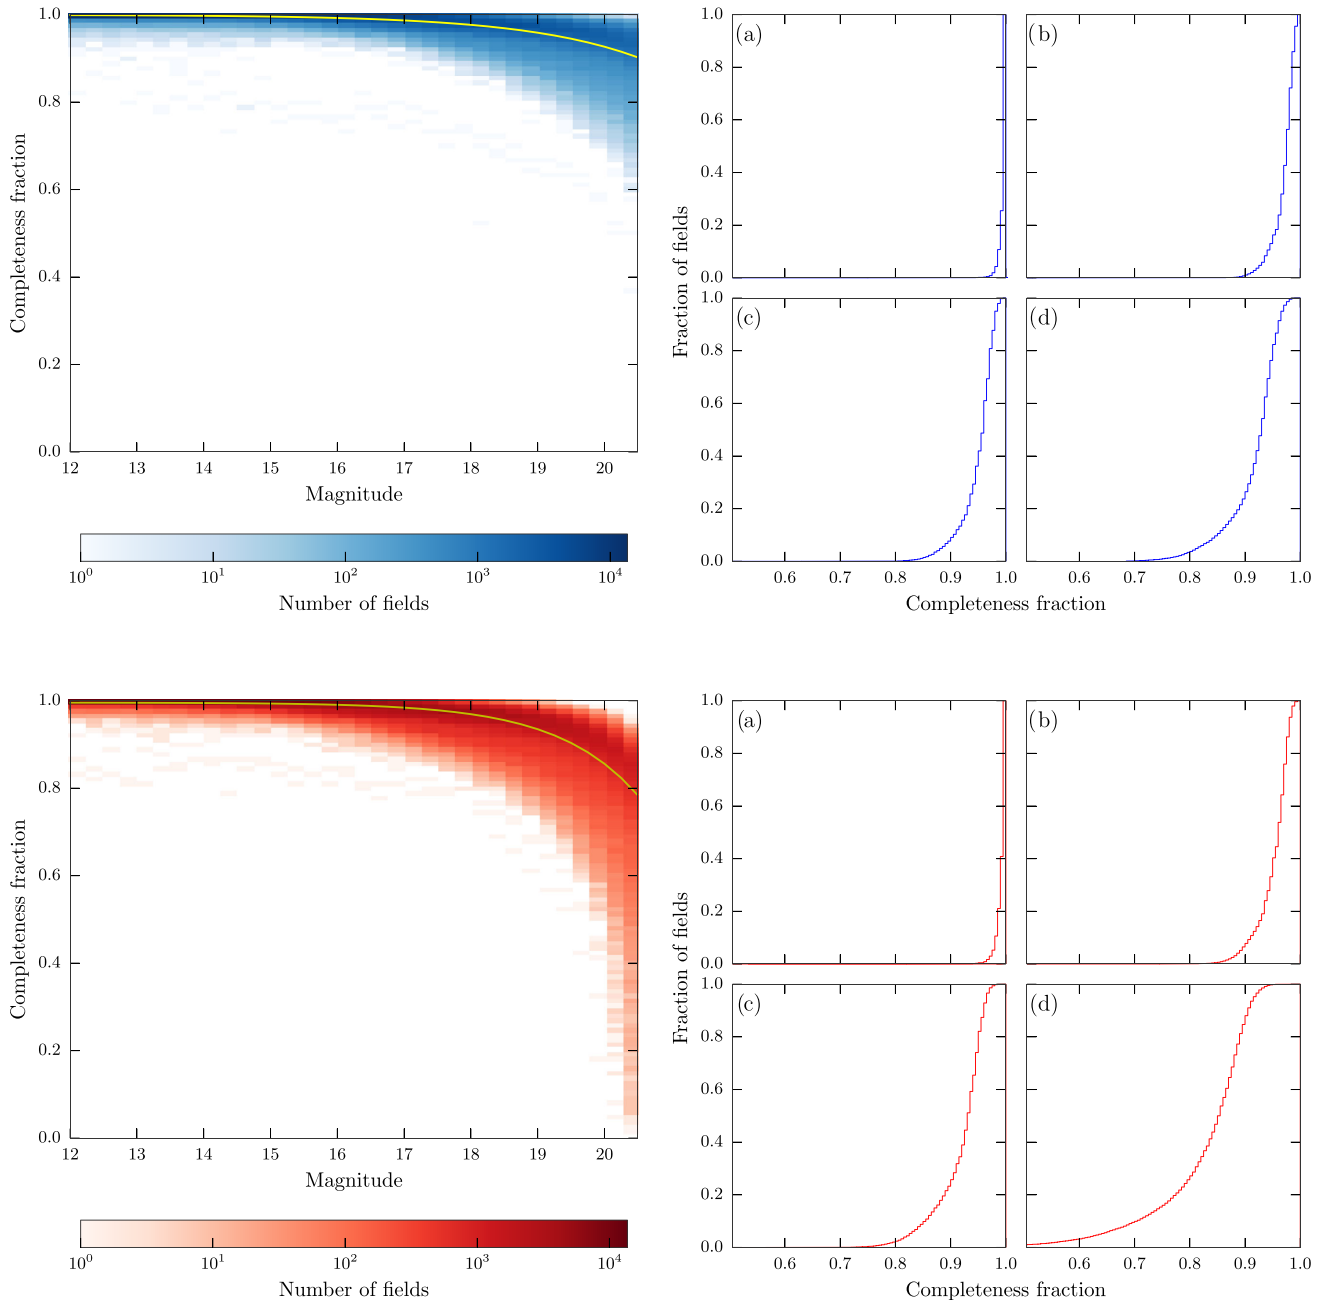
Figure 9. Left: a density plot incorporating every completeness curve in DR2, for magnitude bins of width 0.25 in between 12th and 21st mag in *r* (upper) and *i* (lower). The yellow line superposed in each panel is the fit curve obtained using equation (6) with the median value for each fit parameter. Right: cumulative histograms of field completeness for magnitude bins (a) 15.0–15.25, (b) 18.0–18.25, (c) 19.0–19.25, (d) 20.0–20.25.

been plotted in Fig. 9. A table containing sets of coefficients for all DR2 fields is available online.[1]

As mentioned in Section 3.4, the completeness curves generated from CCD 4 of each field were used to represent those of the field as a whole. To ensure that this approach was sensible, a randomly selected set of fields had completeness curves generated for all four CCDs in order to compare completeness curves and ensure that the variation between chips was acceptably low. For the magnitude range $12 < r < 21$, standard deviations in the completeness corrections were calculated. At $r = 19$, the deviation in completeness

corrections between CCDs reaches as high as 0.025, where a bright star appears in one or more CCDs of the field. But these cases are rare: the median σ at $r = 19$ is 0.005.

4 THE DENSITY MAPS

4.1 Source counting and application of the incompleteness corrections

The IPHAS footprint was split into 1 arcmin × 1 arcmin cells, and for each cell, a table identifying the extent of every image contributing to DR2 (at the CCD level, i.e. 4 CCD footprints per WFC
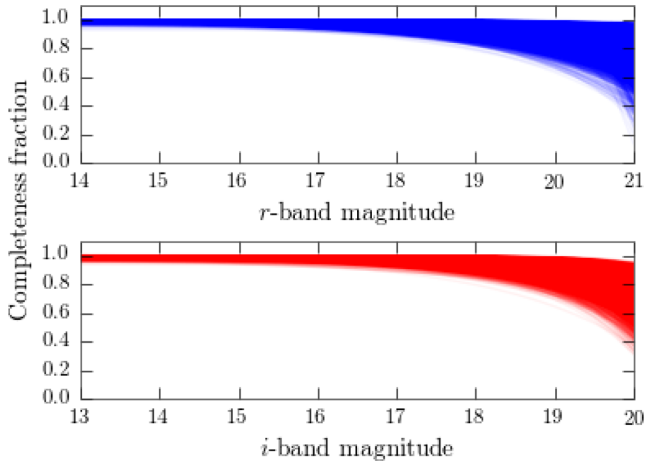
Figure 10. Best fits to completeness curves for both *r* (upper) and *i* (lower) catalogues.

exposure) was queried to identify which intersect the cell (either completely or in part). In order to calculate the effective footprint of each CCD image, the pixel coordinates of its corners were determined, and a customized unusable border region was put in place. This served mainly to exclude sources detected far from the WFC's OA where the counts are less reliable.

For each cell, the relevant files identified in the coverage table were accessed, and all sources falling within the cell boundaries were selected. Sources meeting the following criteria were retained:

(i) morphology classification −1, −2 or +1,
(ii) brighter than the adopted faint limit,
(iii) flag errBits < 64 for band of interest,

where the errBits criterion eliminates sources affected by one or more of: bad pixels within the PSF; a truncated PSF, vignetting. The exclusion of borders already eliminates the majority of such cases, and so mainly leaves to this criterion the job of removing sources affected by issues such as bad columns and hot pixels.

For each selected source, a *corrected contribution* to the number count of the cell was computed from the completeness curves generated in Section 3.7. The table containing all α, β and γ values (equation 6) was read in, so that the values relevant to the current CCD footprint can be identified. Next, the magnitude of the source under consideration was used with these parameters to determine the source's corrected weight. The original (ϕ) and corrected (Φ) source counts were recorded, along with the fraction of the cell overlapped by the CCD footprint from which the source detections are drawn. Each estimate of total corrected cell occupation, Φ' is given by

$$\Phi' = \Phi / f_c, \tag{7}$$

where f_c is the fraction of the cell covered. A formal error estimate would be the Poisson noise of the observed area scaled to the cell:

$$\Phi'_{err} = \sqrt{\Phi}/f_c . \tag{8}$$

This is repeated for each CCD footprint overlapping the specified cell, and for all cells in the density map.

At 1 arcmin × 1 arcmin resolution, the density map contains 6537 051 cells overlapping with IPHAS photometry. Of these, 92.9 per cent overlap two CCD footprints, with 38.3 per cent overlapping three. The cells benefitting from repeated observations provide empirical estimates of the variance in source counts between
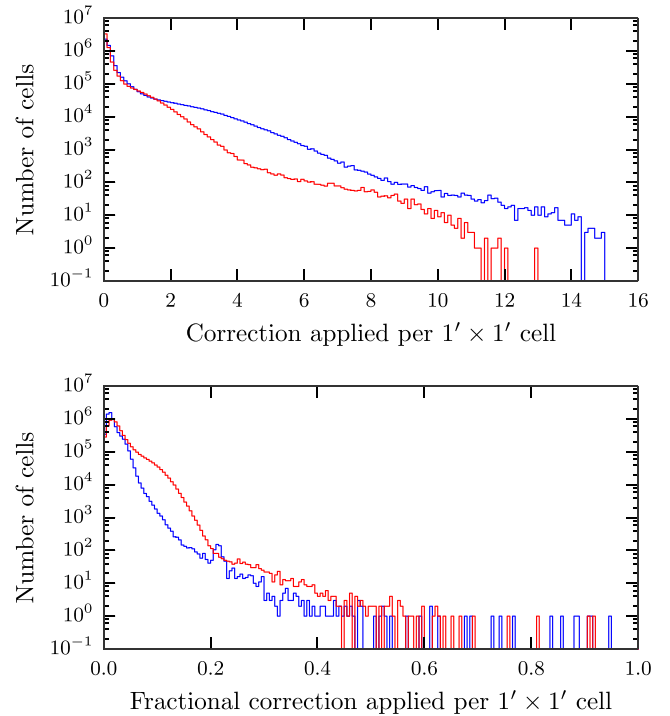
observations, which in turn can be compared with the formal Poisson uncertainty (equation 8), discussed further in Section 4.3. In the cases where multiple exposures and detections are available, the count from the CCD providing the best coverage fraction was adopted, with the observation made under the best conditions chosen from sets of observations providing equal coverage.

Fig. 11 shows the histogram of corrections for the *r* = 19 and *i* = 18 density maps, confirming that most cells require an additive count correction that is well under 1.

Figure 11. Histogram of the incompleteness corrections. Upper: expressed as the number of added sources. Lower: exposed as the fractional correction per cell applied to the 1 arcmin × 1 arcmin *r* = 19 (blue) and *i* = 18 (red) density maps.

4.2 Bright stars

The fact that the *r*- and *i*-band maps are generated independently means that spurious sources detected in one filter only can be included in the map. The majority of such sources are noise and as such are eliminated based on their morphological classification. However, in the regions surrounding bright stars ($V \lesssim 3$, particularly), the scattered light produced can lead to a large number of spurious detections that finish up classified as stellar or extended sources. Cross-matching between bands to eliminate such detections is not an option, as this would eliminate redder sources included in the *i*-band map.

Around fainter, but still bright, stars ($V \lesssim 5$) the scattered light is not so severe as to increase the number of spurious sources; instead, a zone of missing sources is likely due to the bright PSF wings of these stars. This incompleteness is more localized than taken into account by the approach of Section 3. To deal with these issues, cells lying within 5 arcsec of stars brighter than $V = 5$ appearing in the catalogue of Hoffleit & Jaschek (1991) were simply excluded from the density map. The resultant cutouts around 102 stars have resulted in the discarding of 0.2 per cent of the density map cells.
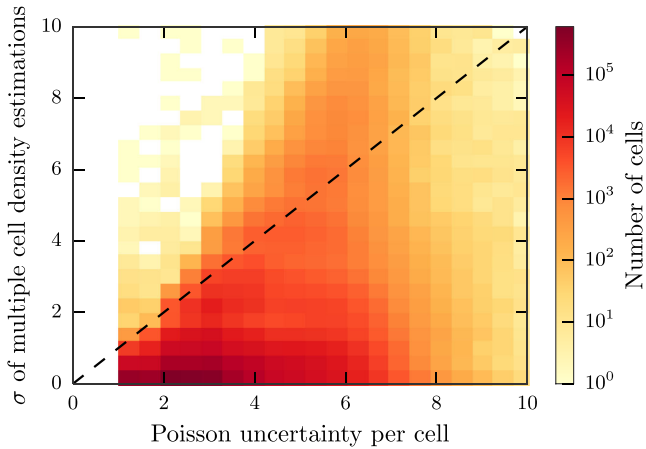
Figure 12. Density plot comparing the deviation between repeated source counts and Poisson uncertainty for ∼4 million cells of the 1 arcmin × 1 arcmin corrected $i = 18$ density map. These cells are those which are covered by multiple IPHAS observations, where the coverage fraction of all observations is greater than 0.5. The equivalent plot for the $r = 19$ density map uncertainty appears indistinguishable from the $i = 18$ – this is not unexpected; the typical $r - i$ value of 1 suggests the two maps are capturing the same population.

4.3 Uncertainties in the star counts

For sky cells captured in more than one exposure, the availability of repeated source counts provides a pessimist's measure of the impact of the observationally driven error, existing in the density maps at the level of the 1 arcmin × 1 arcmin cells (after correction for incompleteness). It is conservative in the sense that variance computed between counts from the different independent and independently corrected observations pays no heed to the relative qualities of the compared observations. It is significant then that this variance is nevertheless, in the vast majority of cases, distinctly lower than the scaled Poisson uncertainty (equation 8) on the adopted cell count – as shown in Fig. 12. More specifically, down to the $r = 19$ limit, for cells covered multiple times with coverage fraction greater than 0.5, the median deviation between cell source counts is 0.03 compared to a median Poisson uncertainty of 2.68. Clearly simple Poisson statistics do not apply. To the same limit, the median source count is ≈6.2 (which scales up to ≈20 000 sources per square degree). These numbers are similar for the $i = 18$ limit, with a median deviation between repeated cell counts of 0.03, and a median Poisson uncertainty of 2.66. The ratio, 0.03/6.2, points to a typical error at maximum resolution of the map of ≲0.5 per cent due to the source count derived from IPHAS observations. Rebinning the map to coarser resolution will bring this typical error down.

Special measures are needed for deciding whether empty cells are reliable or not. First, cells which are fully covered by IPHAS but contain no sources (i.e. genuinely empty cells down to the adopted faint limit) have an uncertainty placed on them equal to the contribution of a single source at the faint limit of the density map. Continuing the example of the r map to 19th magnitude, the median source count of 6.2 in 1 arcmin × 1 arcmin density map cells indicates, on average, that 16 per cent of the area of any one cell of this size needs to be covered in order to provide a single source detection – if the fraction covered is less than this, a null source count should be regarded as unreliable. We use this reasoning to identify and weed out unreliable empty cells: specifically, on multiplying the expected uncertainty for an empty cell (typically ≈1.2 at $r = 19$) by the minimum acceptable cell coverage fraction

Table 2. The FITS extension corresponding to faint limiting magnitudes for both r- and i-band density maps.

	17.0	17.5	18.0	18.5	19.0	19.5	20.0
r	–	–	1	2	3	4	5
i	1	2	3	4	5	–	–

(rounding it to 20 per cent), a limiting scaled error of 6 counts is obtained for it. Any unoccupied cells with computed (Poisson) errors exceeding this level were discarded as being unreliable. This has meant the retention of 22 500 genuinely empty cells, and the rejection of 16 352. Genuinely empty cells are most commonly found in regions associated with high extinctions (e.g. towards the Aquila and Cygnus Rifts). The rejected cells mostly trace out the edges of fields, and are preferentially located at higher Galactic Longitudes, where the majority of gaps in IPHAS DR2 coverage fall.

The main causes of error in the derived stellar density in any one cell (with high coverage fraction) are the uncertainty in the completeness correction and the imprecision in the applied magnitude calibration. The completeness correction errors can be kept as low as computing power permits – here, the many hundreds of artificial sources in the fainter bins (see Table 1) would imply errors in the completeness fraction estimates of the order of 3–5 per cent. How this feeds through to the corrected density maps will scale as the count correction: e.g. a high-end 30 per cent count correction to a cell would add under 2 per cent to the error budget (the fractional count corrections are mapped out in Appendix A). Turning to the magnitude calibration, the external precision of IPHAS DR2 r and i is formally in the region of 0.03 mag (see Barentsen et al. 2014). Viewing this conservatively by allowing that the magnitude cuts applied to the density maps may be wrong by as much as 0.05, an uncertainty in source count of ∼2 per cent would be deduced. These factors together imply a small total error budget in the region of 3 per cent that applies at any density map resolution and to both the r and i bands. This source of error, being greater than the source count errors discussed, is therefore the dominant source of uncertainty in the final density maps. Errors of no more than ∼3–4 per cent should be borne in mind when studying source densities derived from the maps (discussed in Section 5).

4.4 The final stellar density maps

Both r- and i-band density maps are available for a range of depths and resolutions. The resolutions available begin at 1 arcmin × 1 arcmin and increase to cells of 2, 10, 30 arcmin and 1° on a side. At each resolution, the maps are stored as multi-extension FITS files, with the extensions containing maps of increasingly faint limiting magnitude. Table 2 lists which extension corresponds to which faint limit for each band. These files are provided as supplementary materials to this paper, and they will also be accessible from CDS and from the IPHAS survey website.[2] In Appendix A, we show plots of the full final stellar density maps at 2 arcmin × 2 arcmin resolution, for both photometric bands. These are accompanied by plots showing the scale of the incompleteness corrections that had to be applied.

For the purpose of immediate illustration, three cutouts at 2 × 2 arcmin², taken from the i-band density map to sample contrasting Galactic longitudes, are presented in Fig. 13. The inner Galaxy
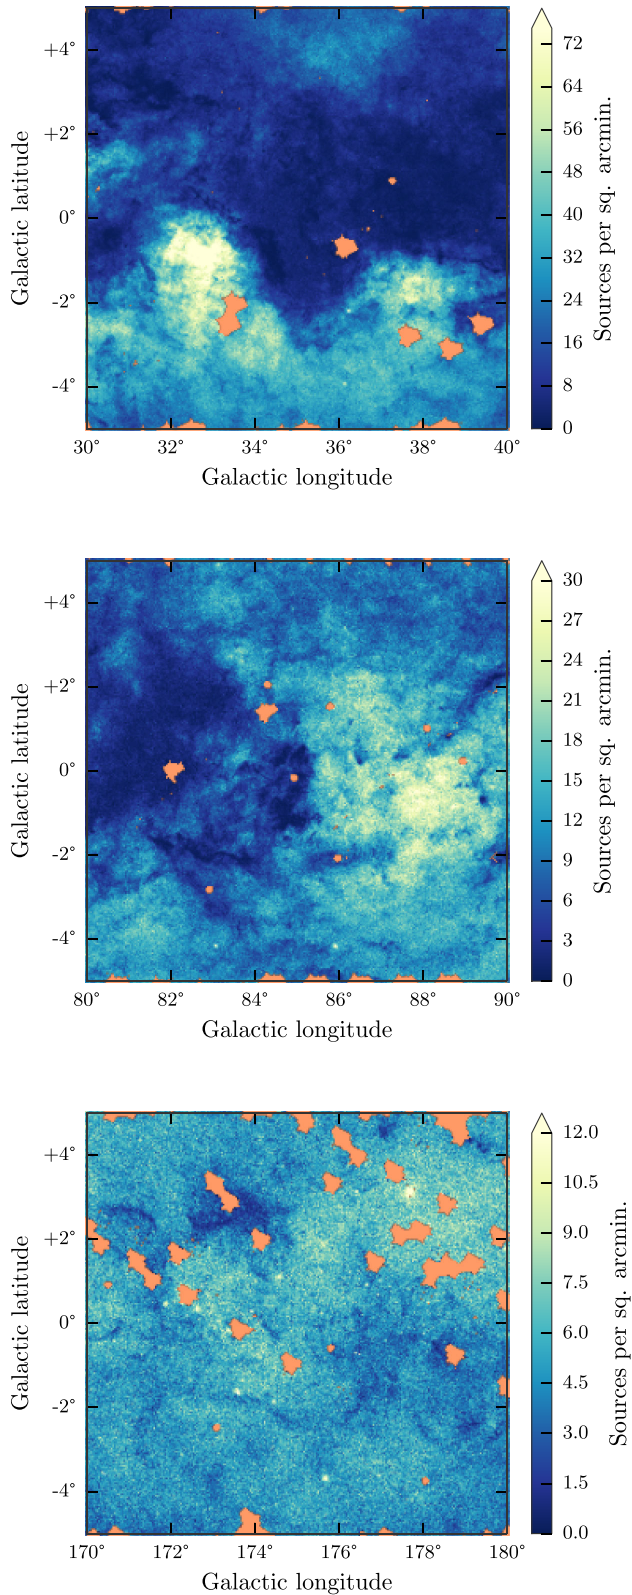
---

[2] http://www.iphas.org/

Figure 13. Cutouts of the $i < 18$ density map with $2\,\mathrm{arcmin} \times 2\,\mathrm{arcmin}$ resolution. Upper: inner Galaxy sightline showing the Aquila Rift. Middle: sightline covering the Cygnus region. Lower: anticentre sightline. The areas shaded in salmon pink are regions where no counts are presently available from the IPHAS DR2 catalogue, or where regions of 5 arcsec have been masked around stars brighter than $V = 5$. Note that the colour scales applied in the different panels are different.

sightlines in the range $30° < \ell < 40°$ (top panel in the figure) contain some of the lowest and the highest density cells in the entire map, owing to the most extreme contrasts between high equatorial extinction and intrinsically high stellar densities away from the mid-plane. These pronounced contrasts ease off with increasing longitude swinging around towards the outer disc, as seen in the $80° < \ell < 90°$ and anticentre sightlines also presented in Fig. 13. The areas left grey in these cutouts are areas for which counts are not yet available: among these, the smaller circular patches will be zones of avoidance around bright stars, while the larger areas with a more ragged outline will be due to no data having been accepted into IPHAS DR2. The latter are much more common in the anticentre region that is overhead in the La Palma winter, when the weather is less clement on average.

At resolutions on the arcminute scale, the density maps betray the fractal nature of the interstellar medium, presenting many fine dark filaments due to obscuring dust – some of these are especially easily picked out in the middle panel of Fig. 13. In experimental comparisons we have made, we have noticed that much of this detail echoes the fine structure seen in the CO maps of Dame, Hartmann & Thaddeus (2001) presented at a spatial resolution of 7.5 arcmin and in the more recent *Planck* maps of Galactic dust at 5 arcmin resolution. Indeed, the effective resolution in the star counts is higher – we particularly recommend the 2 arcmin × 2 arcmin sampling, as the median count per cell is 25 and benefit is gained from the effective interpolation over smaller gaps in survey coverage.

Already known star clusters frequently stand out. In the $30° < \ell < 40°$ region, the globular cluster NGC 6749 is apparent as a count peak at $\ell = 36°.167$, $b = -2°.178$. In the vicinity of the Cygnus-X region shown in the $80° < \ell < 90°$ cutout, the open clusters Berkeley 54 and NGC 7044 can be noticed as peaks at $\ell = 83°.129$, $b = -4°.143$ and at $\ell = 85°.890$, $b = -4°.150$, respectively. In the generally quieter star-count background of the anticentre region, star clusters become even more evident to inspection: M37 particularly stands out as a broad bright peak at Galactic coordinates, $177°.64, +3°.09$. In addition, Czernik 21 ($171°.89, +0°.45$), NGC 1907 ($172°.62, +0°.31$), NGC 1893 ($173°.59, -1°.68$) and Berkeley 17 ($175°.65, -3°.65$) are all obvious to the eye. A possibility to pursue in the future is to take advantage of the uniformity and red sensitivity of these star counts to eliminate *false* clusters in the literature: this can be achieved by cross-matching the maps with published cluster catalogues to confirm (or not) that clusters identified from blue images are also apparent in the r and i bands.

Fig. 14 shows the variation in source counts per square degree across the Galactic plane, with counts taken directly from the $1° \times 1°$ density map (the coarsest resolution produced). The plot is similar to fig. 3 from González-Solares et al. (2008), with some differences. The older diagram was based simply on counting rows per ~ 0.25 square degree survey field in the IPHAS Intermediate Data Release, whilst the new ones are the result of rigorously counting sources in each of the r and i photometric bands down to the specified calibrated magnitude limits. The imposition of a magnitude limit naturally reduces the source count such that now the typical stellar density in the outer disc is in the vicinity of 20 000 per square degree in both bands (the count being higher in i), as compared to 50 000 per square degree. Otherwise, the general pattern in the dependence of stellar density on Galactic longitude remains much the same, including some of the substructure (e.g. the peak at $\ell \simeq 87°$). As a general rule the stellar surface density in the i band is between 1.5 and 2 times higher than in the r band, thanks to reduced extinction.
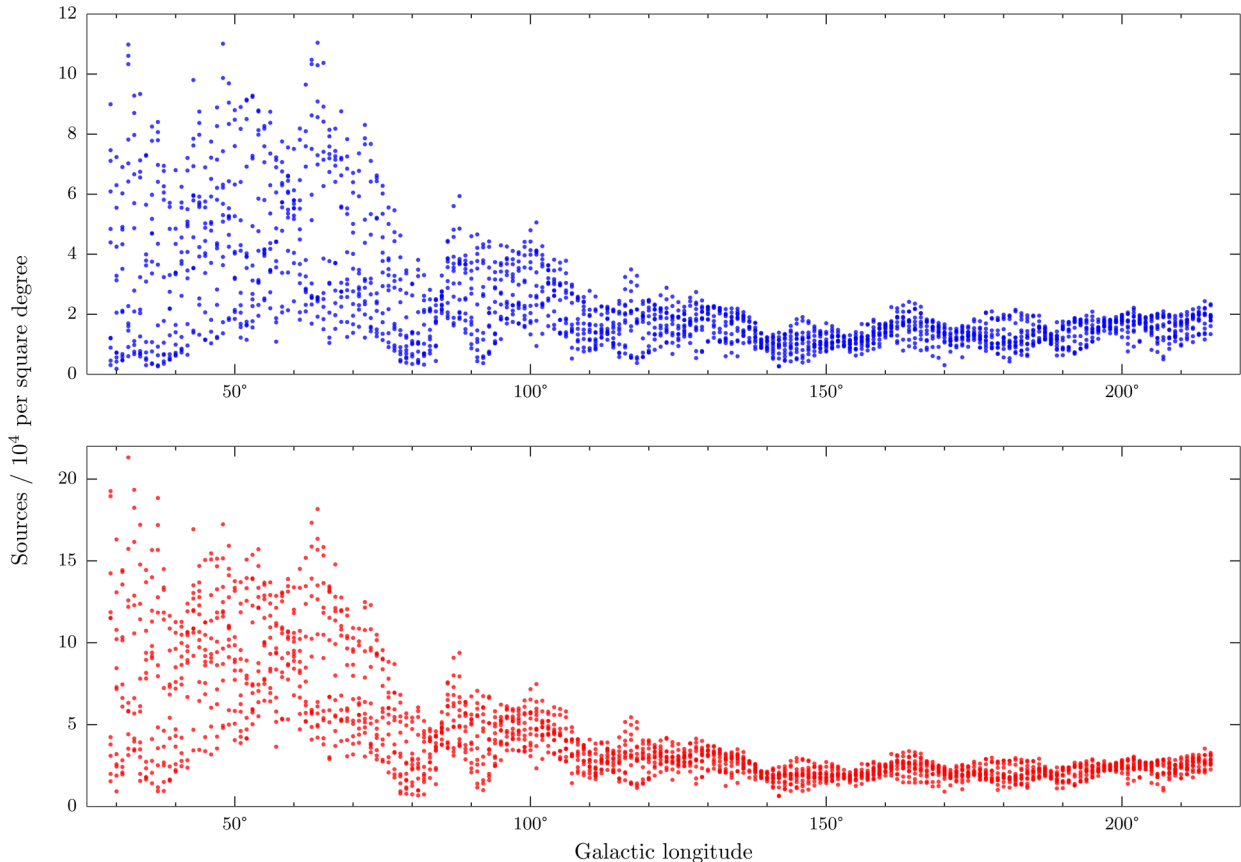
Figure 14. Source counts for $r < 19$ (upper) and $i < 19$ (lower) as a function of Galactic longitude, per 1 square degree cell (per-cell counts taken from the coarsest resolution density maps produced).

5 FIRST COMPARISONS WITH GALACTIC MODEL PREDICTIONS

In this final section of results, we provide a limited demonstration of how these star counts could be deployed as combined tests of Galactic models and 3D extinction maps. It will certainly be appropriate to take such tests further, but these first comparison provides evidence that the recent extinction maps of Sale et al. (2014) are an improvement on the older 2MASS-based maps of Marshall et al. (2006).

5.1 Choice of galactic model

The Besançon model (Robin et al. 2003) has been available to the astronomical community for over a decade now: it is designed to generate synthetic Galactic populations, compiling knowledge of the Milky Way from several sources. The required inputs to the model include stellar density gradients, star formation history, age–metallicity relationships and a stellar initial mass function. It is one of the tools that will be used to interpret *Gaia* data, and is being updated to incorporate the latest developments in studies of Galactic evolution, structure and kinematics (Czekaj et al. 2014). However, since the older 2003 model has already been used extensively (Momany et al. 2006; Ivezić et al. 2008) and continues to offer downloadable synthetic catalogues,[3] we use it here to make

the very first comparisons against IPHAS completeness-corrected star counts.

The synthetic stellar catalogues delivered by the Besançon website can be specified in terms of both heliocentric distance and Galactic coordinates. The parametrization of the overlaid extinction is adjustable, while spectral types, along with absolute and apparent magnitude ranges, can also be specified. Two output photometric systems are available: Johnson–Cousins and CFHTLS-MegaCam (the instrument used for the Canada-France-Hawaii Telescope Legacy Survey). For the present purpose of comparison with IPHAS data, the SDSS-like MegaCam filters are the appropriate choice. The IPHAS/SDSS transformations presented in Barentsen et al. (2014) have been applied here to bring the returned synthetic magnitudes into the IPHAS photometric system. We have retrieved simulated catalogues for three $2° \times 10°$ strips, perpendicular to the Galactic equatorial plane, in order to sample a range of likely outcomes (each one taken from a region plotted in Fig. 13). We have set the diffuse extinction parameter to zero in these requested simulations, so that we may correct the catalogues using more recently published 3D extinction data (see Section 5.3). By limiting our attention to just these three strips, it became practical to collect the full range of stellar populations and to retain the default absolute magnitude range ($-7 < M < 20$). An apparent magnitude range of $8 < r < 25$ was imposed in our selections to exclude any objects falling far outside IPHAS detection limits.

While the Besançon model adopts a thin disc with a hole at its centre alongside a separately specified bulge component, we only need to consider the thin disc component here. The influence of the
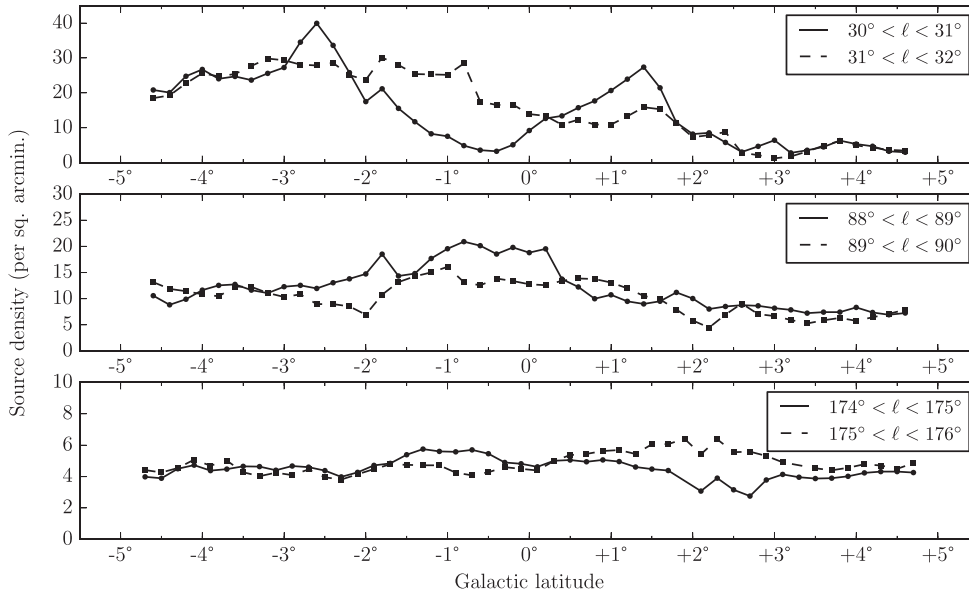
---

[3] http://model.obs-besancon.fr/

Figure 15. Stellar density profiles in Galactic latitude, averaged over 1°-wide strips in Galactic longitude, as given by the 20 arcmin × 20 arcmin resolution IPHAS *i*-band density map, for a magnitude limit of 18. Each panel shows two neighbouring strips for each of the three general directions within the Galactic plane.

central Galaxy specification can be neglected as even at the lowest Galactic longitudes covered by IPHAS, the sightlines towards the inner Galaxy have their closest approach ∼4 kpc from the Galactic Centre – beyond the region affected by the bulge component.

5.2 The selected sightlines

The regions chosen for comparison have been selected to come from contrasting longitudes ($\ell \approx 30°$, 90° and 175°), with the precise choices designed to avoid gaps in IPHAS DR2 coverage. From the Besançon web interface, we have extracted the relevant *i*-band catalogues with no applied extinction.

Fig. 15 shows the variation in stellar densities with Galactic latitude for three pairs of strips, as retrieved from the 20 arcmin × 20 arcmin resolution IPHAS $i < 18$ density map. At each of the three longitudes investigated, the star-count data for two adjacent 1°-wide strips are shown to provide a more representative impression of the degree-scale variations present.

The inner Galaxy sightline shows the largest variation with latitude and the largest variation between the two separated 1°-wide strips. The Aquila region ($\ell \lesssim 45°$) contains both the highest and lowest valued cells of the entire IPHAS density map. A dark cloud at ($\approx 30°$, $-1°$) is responsible for the reduction in stellar density (of the amplitude ≈10–20 sources arcmin^{-2}) in the $30° < \ell < 31°$ profile relative to the $31° < \ell < 32°$ profile. The Aquila Rift, just 200–300 pc away, is responsible for the drop-off in stellar density at higher latitudes in both strips. Taking the spiral arm positions of Vallée (2008), it can be seen (in Fig. 16) that this sightline will pass through the Sagittarius–Carina arm twice by a distance of 10 kpc, as well as follow the tangent of the Scutum–Crux arm: in combination with the local dust, such spiral arm crossings would be expected to produce the complex extinction distribution associated with this direction.

The sightline towards $\ell = 90°$, bordering the Cygnus region (60° $\lesssim \ell \lesssim 90°$), shows similarly pronounced variation with latitude, and between the two strips. Further from the Galactic mid-plane the two
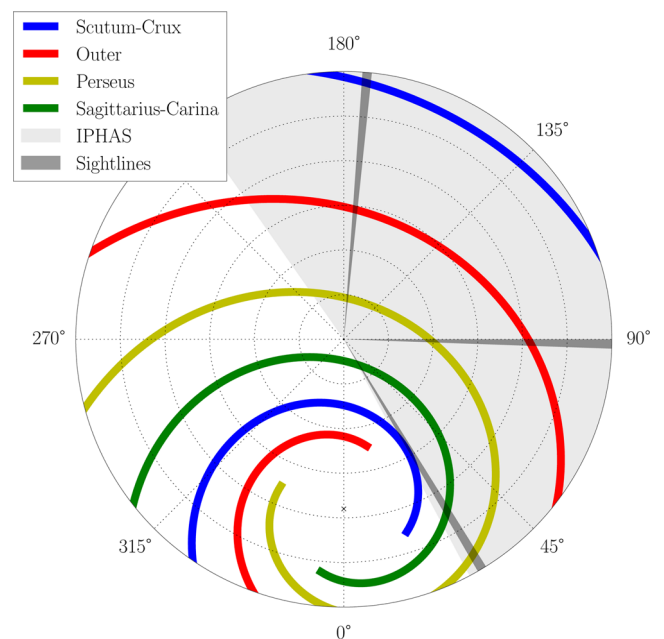

Figure 16. Distribution of spiral arms in Galactic coordinates, with distance from the Sun to the spiral arms based on the parameters obtained by Vallée (2008). Dotted circles denote distances in increments of 2 kpc. The Galactic Centre is shown as a black cross. The region covered by IPHAS is highlighted in grey. The sightlines compared with Besançon model predictions are denoted by darker hatched regions.

seem to be in good agreement. The spiral arm map of Vallée (2008) places the Perseus arm at ≈4 kpc in this direction. In contrast, the sightline passing close to the anticentre shows the least variation: at positive latitudes the difference between the two adjacent strips is no more than ≈2 arcmin^{-2}. The flatness of these distributions highlights the relative lack of structure apparent in the density map at these longitudes; this suggests that the Besançon model that does not attempt to emulate spiral arm structure is likely to perform

well along this sightline – a simple exponentially decreasing stellar density with Galactocentric radius may capture the observed smooth behaviour. Nevertheless, the Perseus Arm, described by Vallée (2008) as lying at a distance of ≈2 kpc at this longitude, also passes through.

5.3 Adding in extinction to the synthetic catalogues

With the diffuse extinction parameter set to 0 mag kpc^{-1}, the synthetic catalogues returned provide an unreddened view of the Milky Way. As the catalogues specify the distance to each generated object, a custom extinction profile can be imposed on the simulated data. This has come from one of two empirically based sources, as detailed below.

5.3.1 Marshall et al extinction map

The 3D extinction map produced by Marshall et al. (2006) has been applied to the catalogues, taking the sightline closest to the synthetic catalogue coordinates – the Marshall et al. (2006) sightlines are binned in 0°.25 increments in Galactic longitude and latitude. This map only covers sightlines with $\ell < 100°$, which means that the anticentre comparison we make here cannot benefit from this model. It is based on analysis of 2MASS near-infrared photometry in combination with constraints from the Besançon model: accordingly, the extinction in this case is specified for the K_s band (A_{K_s}) as a function of distance. For present purposes, this requires the intermediate step of conversion to A_i. To do this, we use extinction law data from Cardelli, Clayton & Mathis (1989), along with conversion factors for the K and I filters, thereby obtaining $A_i = 4.2A_{K_s}$.

5.3.2 Sale et al extinction map

The 3D map of Sale et al. (2014) is the result of combining IPHAS DR2 photometry with the hierarchical Bayesian-inference model developed by Sale (2012), which estimates the distance–extinction relationship along a given sightline, along with estimates of the atmospheric parameters of the stars sampled. A prior adopted in this treatment is that all sightlines are occupied only by thin disc stars, fitting to an exponential scale length of 3 kpc. It is important to note that, although this extinction map was also constructed from IPHAS DR2, it is only the precision on the returned extinction to any given distance that carries a strong dependence on the stellar density distribution deduced from the photometric data in common. In other words, there is no circularity of argument here, as the comparisons below make plain.

The map provides a typical angular resolution of 10 arcmin; the nearest sightline to the Besançon catalogue under consideration was adopted in extinguishing the synthetic photometry. While each sightline comes with a warning of a maximum reliable distance (after which a prior on extinction takes over), holding the extinction constant beyond these limits would risk too large a contribution to the predicted star count from these distant regions. For this reason, the extinction was permitted to continue to rise, even if it is no longer directly empirically constrained. All three longitudes considered here fall into the area treated by Sale et al. (2014).

5.4 The observed and predicted star-count profiles compared

At each longitude, we make a comparison for just one of the two 1°-wide cross cuts at constant longitude. These are shown in Fig. 17.

The strip sampling $\ell = 30°$ to $31°$ is compared to model predictions using both 3D extinction maps. It is striking that, irrespective of the extinction map applied, there is a marked overprediction of the star count – particularly near the Galactic equator, where the discrepancies are on the scale of a factor of 2 or more. However, the Sale et al. (2014) extinction data do a somewhat better job than the K_s-band-based Marshall et al. (2006) data, actually managing to agree at positive Galactic latitudes. The need to convert K_s-band extinction to suit i may partly be behind the weaker performance of the Marshall et al. (2006) modelling, but the differences in the 'predicted' profile shapes also indicate that the two extinction maps describe different line-of-sight distributions of the dust. One obvious difference is that Sale et al. (2014) mapping is better able to place the sharp rise in extinction due to the Aquila Rift, presenting anything from 3 to 7 mag of extinction, within the first kiloparsec where it belongs. Nevertheless, the biggest discrepancies around the Galactic equator coincide with the domain in which the maximum extinctions inferred in this mapping fall well short of the large total extinctions described in the Schlegel et al. (1998) map (see fig. 10 in Sale et al. 2014), raising the possibility of too little in the 3D map. There is a known bias of this sense that affects extinction mapping and it may have greater impact along these especially extinguished sightlines.

The second and third comparisons in Fig. 17 at $\ell \simeq 90°$ and $\ell \simeq 175°$ respectively use the $\ell = 89°$–$90°$ and the $\ell = 174°$–$175°$ strips presented in Fig. 15. Here the comparisons are encouraging, especially in the case of the Sale et al. (2014) extinction data, where the discrepancies are limited to under ∼1 star arcmin^{-2} in both cross cuts. Certainly the agreement achieved at $\ell = 90°$ is within likely error.

The less compelling performance at $\ell = 30°$ compared to the agreement at the larger longitudes supports the assertion in Section 5.3.2 that the shared IPHAS origin of the Sale et al. (2014) extinction data and the star counts does *not* guarantee a good match. There can be two causes for the divergence between them – either the dependence of extinction on distance is incorrectly represented (already considered above) or the mock star catalogue created from the Besançon model is too well populated. It is not our aim in this limited discussion to try to conclude on the origin of the discrepancy as this would require the examination of a much more comprehensive set of comparisons. Suffice it to say that, in the light of the good behaviour at $\ell \simeq 90°$ and at $\ell = 175°$, it would be premature to insist the cause is all in the 3D extinction mapping.

For now, the possibility remains open that the problem is linked to the distinctive character of this inner Galaxy sightline: climbing up the stellar density gradient, it captures large numbers of more intrinsically luminous highly extinguished stars distributed along the pencil beam. Interrogation of the mock catalogues shows us that the median distance to objects at $\ell \simeq 30°$ included in the $i < 18$ star counts is 5–6 kpc, as compared with 2–3 kpc at $\ell \simeq 175°$. In other words, at $\ell = 30°$, the IPHAS DR2 star counts are, of necessity, made up from more luminous stars than those at $\ell = 90°$ and $175°$. Too many intrinsically brighter stars are equivalent to a stellar luminosity function too well populated at the high end, and hence may imply a star formation history weighted too strongly to recent times. Czekaj et al. (2014) have encountered a similar difficulty in comparing the Tycho-2 catalogue (limited to $V < 11$, and the solar neighbourhood) with predictions based on an updated version of the Besançon model: they also predict too high a star count in the same part of the Galactic plane here (see their fig. 15) – even after imposing a declining star formation rate over the past 10 Gyr. Here, the comparison is with a Galactic model assuming constant
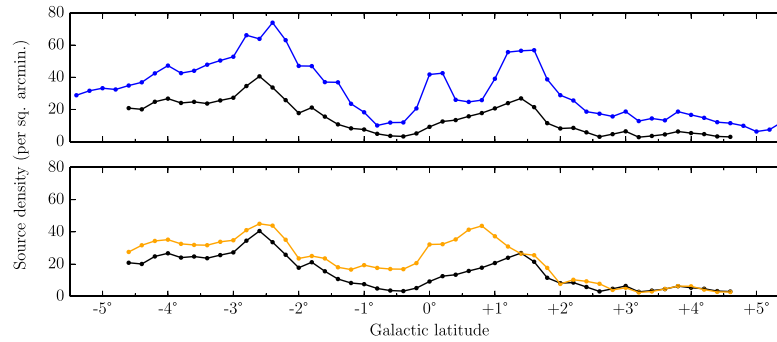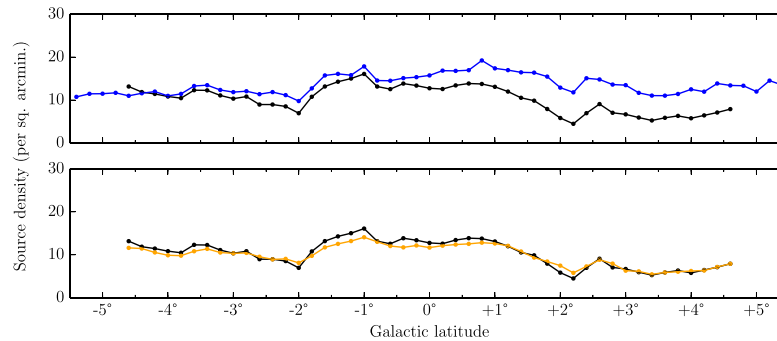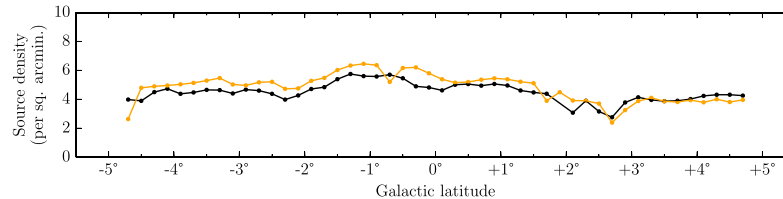
Figure 17. Comparisons between incompleteness-corrected IPHAS star counts to $i = 18$ (black) and Besançon predictions, reddened by the extinction curves of Marshall et al. (2006, blue) and Sale et al. (2014, yellow). From top to bottom, comparisons are for sightlines at $\ell=30°$, 90° and 175°. The errors in the empirical star-count data are 3–4 per cent.

star formation. It will be interesting to examine these issues more closely, taking advantage of the much larger Galactic disc volume these newly provided fainter IPHAS counts can probe.

6 CONCLUDING REMARKS

Our main goal has been to present new IPHAS-based faint r and i star counts for the northern Galactic plane, the first deep optical counts for the northern Milky Way. As the conditions were out of the observers' control, the variety of conditions essentially fed themselves into IPHAS DR2 rather than being fed in by observers/data processing. Here, that scale is set by the camera footprint of ∼0.29 square degrees.

Although it is relatively computation-intensive to use artificial source injection to determine these corrections, our presentation of the alternatives confirms that it is the best available route to placing these faint star counts within a precise quantitative framework. At the same time, the more approximate methods provide insight and supporting evidence that incompleteness at Vega magnitudes brighter than $i = 18$ and $r = 19$ is small. Application of the preferred artificial source method across 94 per cent of the IPHAS 1800 square

degree footprint has traced the appearance of the Galactic disc down to angular resolutions as fine as 1 arcmin^2. The i-band counts are commended for use down to 19th magnitude, and those in r down to 20th. At these limits, the computed completeness fractions are generally above 0.8.

We described in Section 4.4 how the stellar density data are available for re-use in the form of digital maps to a range of limiting magnitudes and at a range of angular resolutions. These are provided with this paper as supplementary material and also may be downloaded from the survey website (see http://www.iphas.org/data/).

There is a limitation to point out. The method discussed in Section 3 takes the average completeness of the central CCD of the WFC as being representative of the complete (4-CCD) IPHAS field. This takes proper account of the prevailing observing conditions, but does not trace local density variations due to e.g. compact star clusters. In regions with small average fractional corrections (under 10 per cent), the corrected density maps will nevertheless remain accurate even on the ∼arcminute scale, since the directly measured count will still be reliable. In contrast, in the minority of sky regions with higher fractional corrections and strong stellar density variations within the WFC footprint, underestimation is liable to affect comparisons of

star counts at scales below ∼10 arcmin × 10 arcmin. Putting this point another way, the star counts presented here should not be assumed to yield accurate peak stellar densities in compact inner Galaxy clusters. At the same time, the higher resolution maps are well suited to the uniform detection of star clusters, creating an opportunity to re-evaluate the content of current optically based cluster catalogues.

Over the next decade, *Gaia* photometry based on the mission's customised system of bandpasses will no doubt become the basis for important work in this area. In this context, it is useful to note that broad-band *Gaia* $G \simeq 20$ roughly corresponds to IPHAS $i \simeq$ 19 (Vega) for the reddened populations dominating the faint star counts, while a similar cross-comparison of the magnitude scales for the *Gaia* blue and red prisms indicates $G_{BP} \simeq 20$ matching to IPHAS $r \simeq 19$. The G_{RP} and IPHAS i scales are similar. These transformations are considered in more detail by Farnhill (2015) and are based on the work of Jordi et al. (2006). The star counts presented here thus reach to the same or slightly greater depth as those likely to ultimately emerge as the *Gaia* mission completes. The important difference is that to brighter magnitude limits, out to a distance of ∼ 1 kpc, *Gaia* counts will benefit from the high angular resolution achieved in space that – together with the advanced astrometric capability – will expose 60 per cent of binaries to 250 pc and 35 per cent of binaries to ∼1 kpc[4] (a star like the Sun will be $r \sim 16$ at 1 kpc). An exhaustive accounting of this volume will be achieved, and stellar multiplicities will be exposed. These specific advantages decay away on the one-to-a-few kiloparsec scale to which the source counts presented here are most sensitive.

To illustrate the potential application of the IPHAS maps, we have presented a first example of comparison between the measured *i* counts and the predictions of the Robin et al. (2003) Galactic model (Section 5.4). These star counts, down to 18th magnitude, are capable of sampling the Galactic disc well beyond the solar neighbourhood, offering a combined test of inferences on the 3D distribution of interstellar dust and of assumptions shaping the Galactic stellar luminosity function. In these first limited tests, the 3D extinction map of Sale et al. (2014) appears to be faring very well along sightlines passing through larger Galactocentric radii, while first signs of a challenge appear for sightlines plunging inwards of the solar circle. The origin of the emerging differences deserves further closer examination, that can draw on a wider selection of count comparisons and also on the information encoded within IPHAS DR2 colour–magnitude data. This represents a particularly appropriate use of IPHAS survey photometry in that both the available Galactic models and these star counts best describe behaviours on angular scales of 0.2–1 deg, averaging over individual clouds and clusters. In these angular and depth domains, these new IPHAS-based stellar density maps are ready to play a validating role, now, in the quest to map out the relative 3D distributions of interstellar dust and the stars in the Galactic plane. Work towards a further paper applying these new counts in a broad examination of the predictions of current Galactic models is now underway.

ACKNOWLEDGEMENTS

REFERENCES

Astropy Collaboration, 2013, A&A, 558, A33
Bahcall N. A., Soneira R. M., 1983, ApJ, 270, 20
Barentsen G. et al., 2014, MNRAS, 444, 3230
Cambrésy L., Beichman C. A., Jarrett T. H., Cutri R. M., 2002, AJ, 123, 2559
Cardelli J. A., Clayton G. C., Mathis J. S., 1989, ApJ, 345, 245
Condon J. J., 1974, ApJ, 188, 279
Czekaj M. A., Robin A. C., Figueras F., Luri X., Haywood M., 2014, A&A, 564, A102
Dame T. M., Hartmann D., Thaddeus P., 2001, ApJ, 547, 792
Drew J. E. et al., 2005, MNRAS, 362, 753
Drew J. E. et al., 2014, MNRAS, 440, 2036
Drimmel R., Spergel D. N., 2001, ApJ, 556, 181
Farnhill H., 2015, PhD thesis, Univ. Hertfordshire
González-Solares E. A. et al., 2008, MNRAS, 388, 89
Groot P. J. et al., 2009, MNRAS, 399, 323
Harvey P. M. et al., 2006, ApJ, 644, 307
Hoffleit D., Jaschek C. V., eds, 1991, The Bright Star Catalogue, 5th edn. Yale University Observatory, New Haven, CT
Hogg D. W., 2011, AJ, 121, 1207
Irwin M., Lewis J., 2001, New Astron. Rev., 45, 105
Irwin M. J., Trimble V., 1984, AJ, 89, 83
Ivezić v. et al., 2008, ApJ, 684, 287
Jordi C. et al., 2006, MNRAS, 367, 290
Lucas P. W. et al., 2008, MNRAS, 391, 136
Magnier E. A. et al., 2013, ApJS, 205, 20
Marshall D. J., Robin A. C., Reylé C., Schultheis M., Picaud S., 2006, A&A, 453, 635
Mateo M., Hodge P., 1986, ApJS, 60, 893
Momany Y., Zaggia S., Gilmore G., Piotto G., Carraro G., Bedin L. R., de Angeli F., 2006, A&A, 451, 515
Robin A. C., Reylé C., Derriere S., Picaud S., 2003, A&A, 409, 523
Ruphy S. et al., 1997, A&A, 326, 597
Sale S. E., 2012, MNRAS, 427, 2119
Sale S. E. et al., 2009, MNRAS, 392, 497
Sale S. E. et al., 2014, MNRAS, 443, 2907
Schlafly E. F., Finkbeiner D. P., 2011, ApJ, 737, 103
Schlegel D. J., Finkbeiner D., Davis M., 1998, ApJ, 500, 525
Stoughton C. et al., 2002, AJ, 123, 485
Taylor M. B., 2005, in Shopbell P., Britton M., Ebert R., eds, ASP Conf. Ser. Vol. 347, Astronomical Data Analysis Software and Systems XIV. Astron. Soc. Pac., San Francisco, p. 29
Vallée J. P., 2008, AJ, 135, 1301
van der Kruit P. C., Freeman K. C., 2011, ARA&A, 49, 301
Yasuda N. et al., 2001, AJ, 122, 1104

[4] http://sci.esa.int/gaia/31441-binary-stars/

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

suppl_data
(http://www.mnras.oxfordjournals.org/lookup/suppl/doi:10.1093/mnras/stv2994/-/DC1).

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the paper.

APPENDIX A: PLOTS OF THE 2 ARCMIN × 2 ARCMIN DENSITY MAPS

We present plots of the *r* and *i* density maps, down to limiting magnitudes of 19th and 18th, respectively (in the Vega system) and at a resolution of 2×2 arcmin^2. The two maps are broken into 30 deg sections, and are shown alongside the corresponding maps of fractional completeness correction so as to highlight where high stellar densities, sometimes in combination with worse-than-median seeing, lead to a need for larger corrections. The reader is reminded that median seeing is 1.1 arcsec in *r*, and 1.0 arcsec in *i* (see Barentsen et al. 2014). These plots also show where there are gaps in the coverage of the northern Galactic plane: the overall fraction of the northern plane covered is 0.94, with more gaps appearing in the outer disc. How the coverage changes with Galactic longitude is laid out in Table A1.

Table A1. Coverage fraction for each segment of the density maps as shown in Appendix A, inside latitude range $-5° < b < +5°$.

Longitude section	Coverage (per cent)
$30° < \ell < 60°$	96.7
$60° < \ell < 90°$	98.4
$90° < \ell < 120°$	98.1
$120° < \ell < 150°$	92.0
$150° < \ell < 180°$	94.4
$180° < \ell < 210°$	90.3
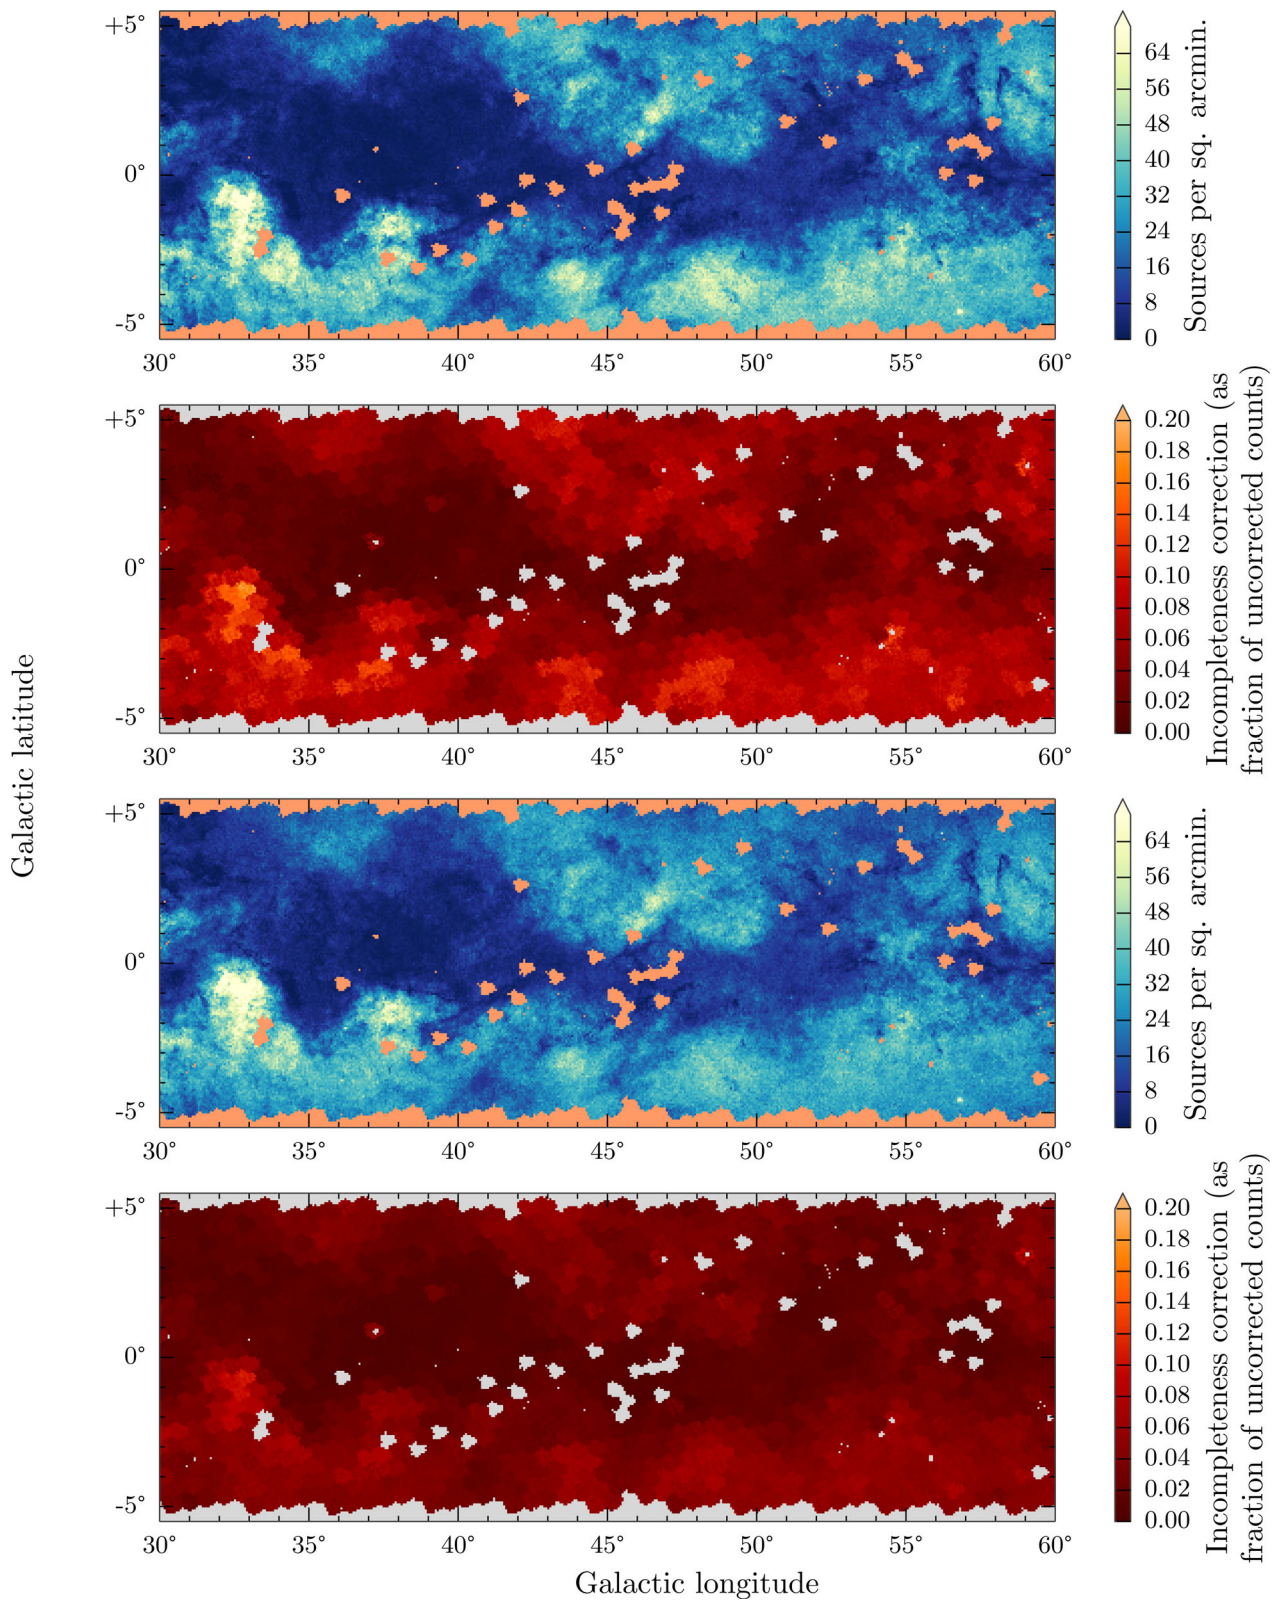
Figure A1. Cutouts of the 2 arcmin × 2 arcmin resolution density maps for the region $30° < \ell < 60°$. (a) *r*-band density map of sources down to 19th mag, (b) *r*-band corrections down to 19th mag, (c) *i*-band density map down to 18th mag, (d) *i*-band corrections down to 18th mag. In (a) and (c), the areas shaded salmon pink are locations without included data presently.
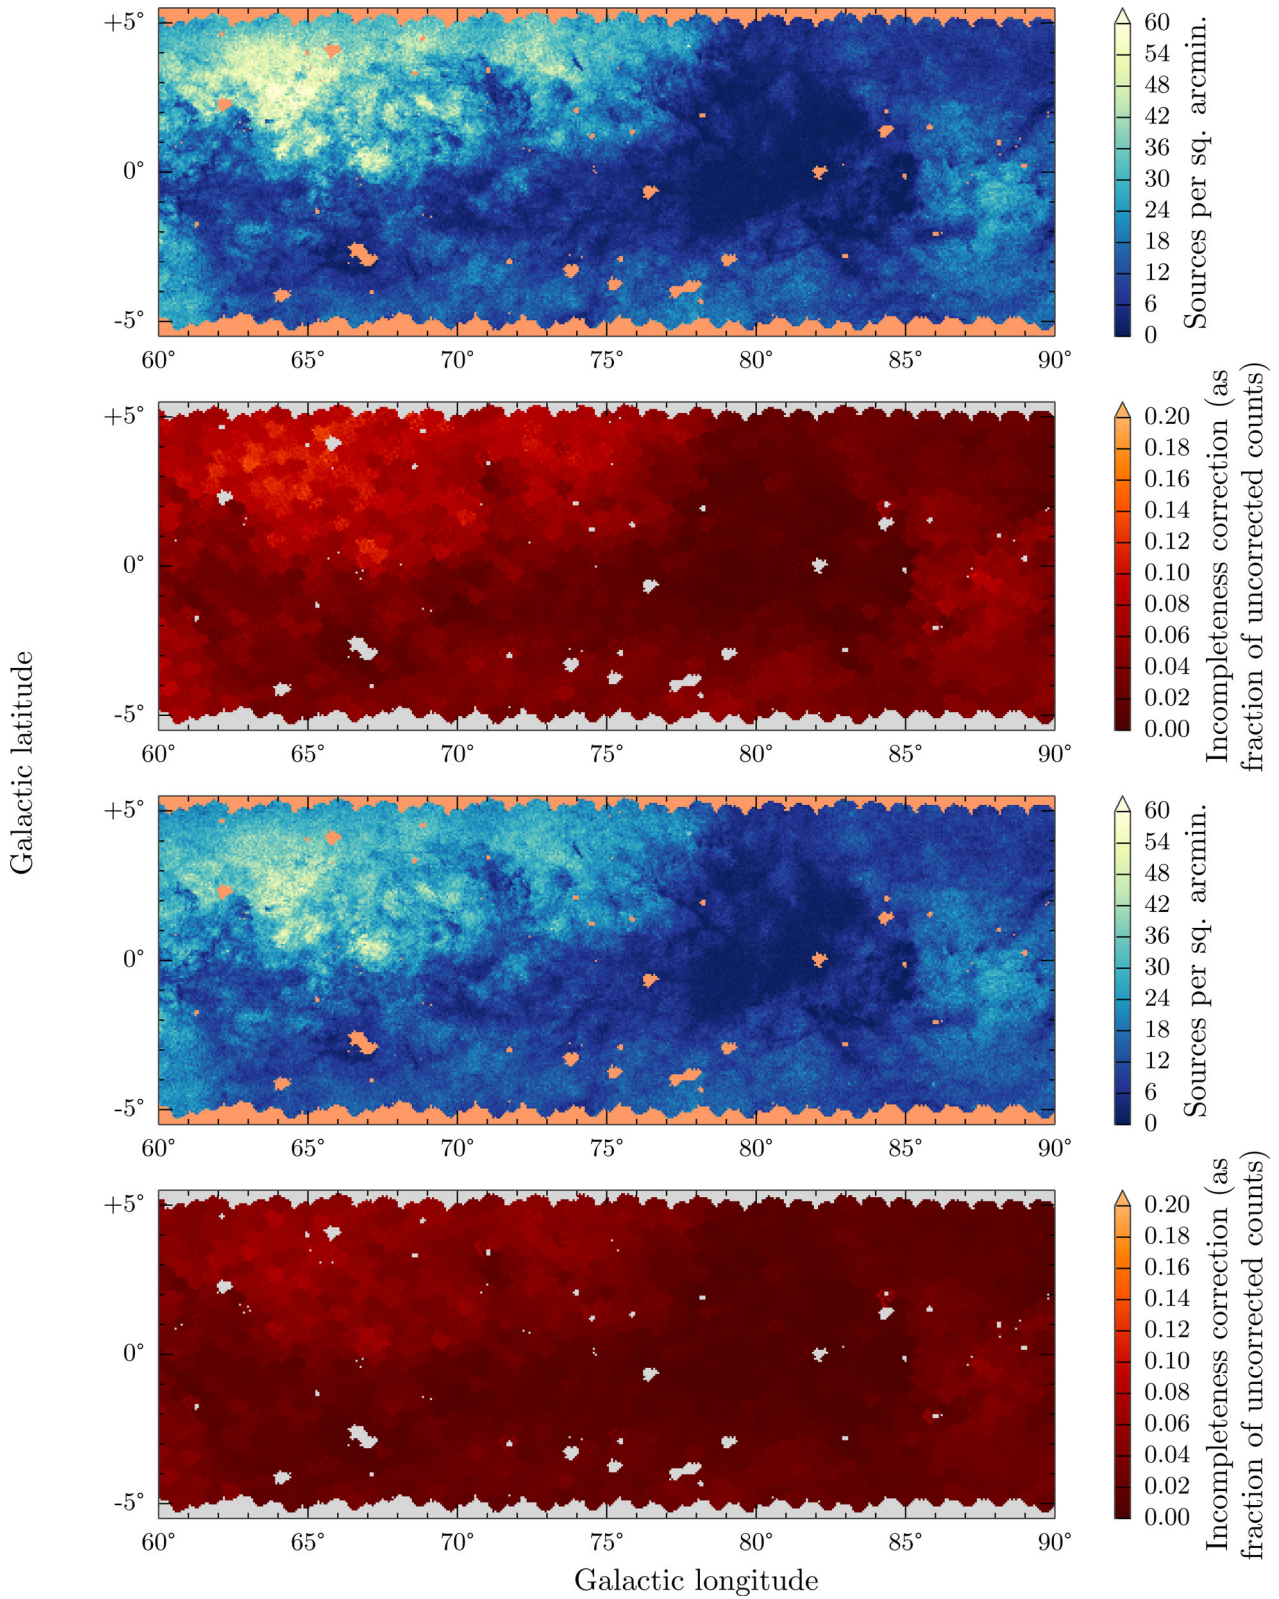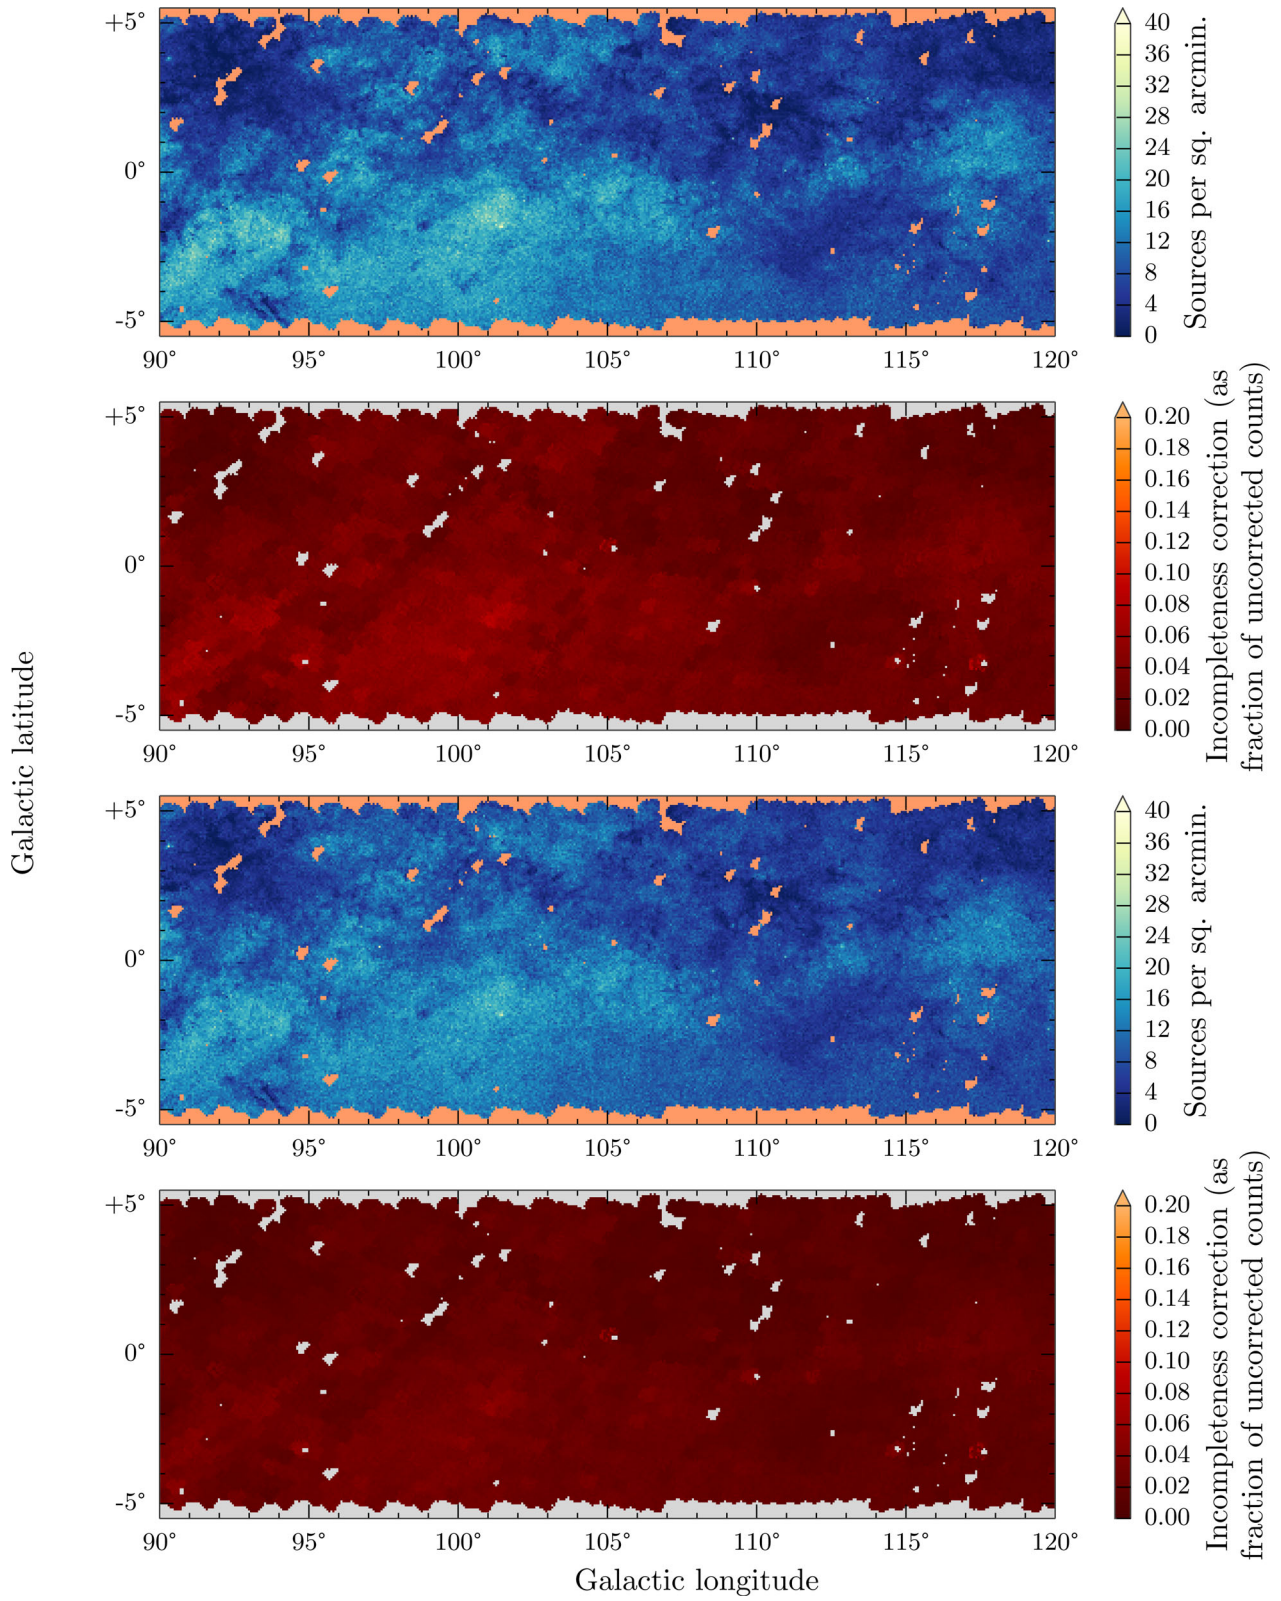
Figure A2. As Fig. A1, only for $60° < \ell < 90°$.

Figure A3. As Fig. A1, only for $90° < \ell < 120°$.
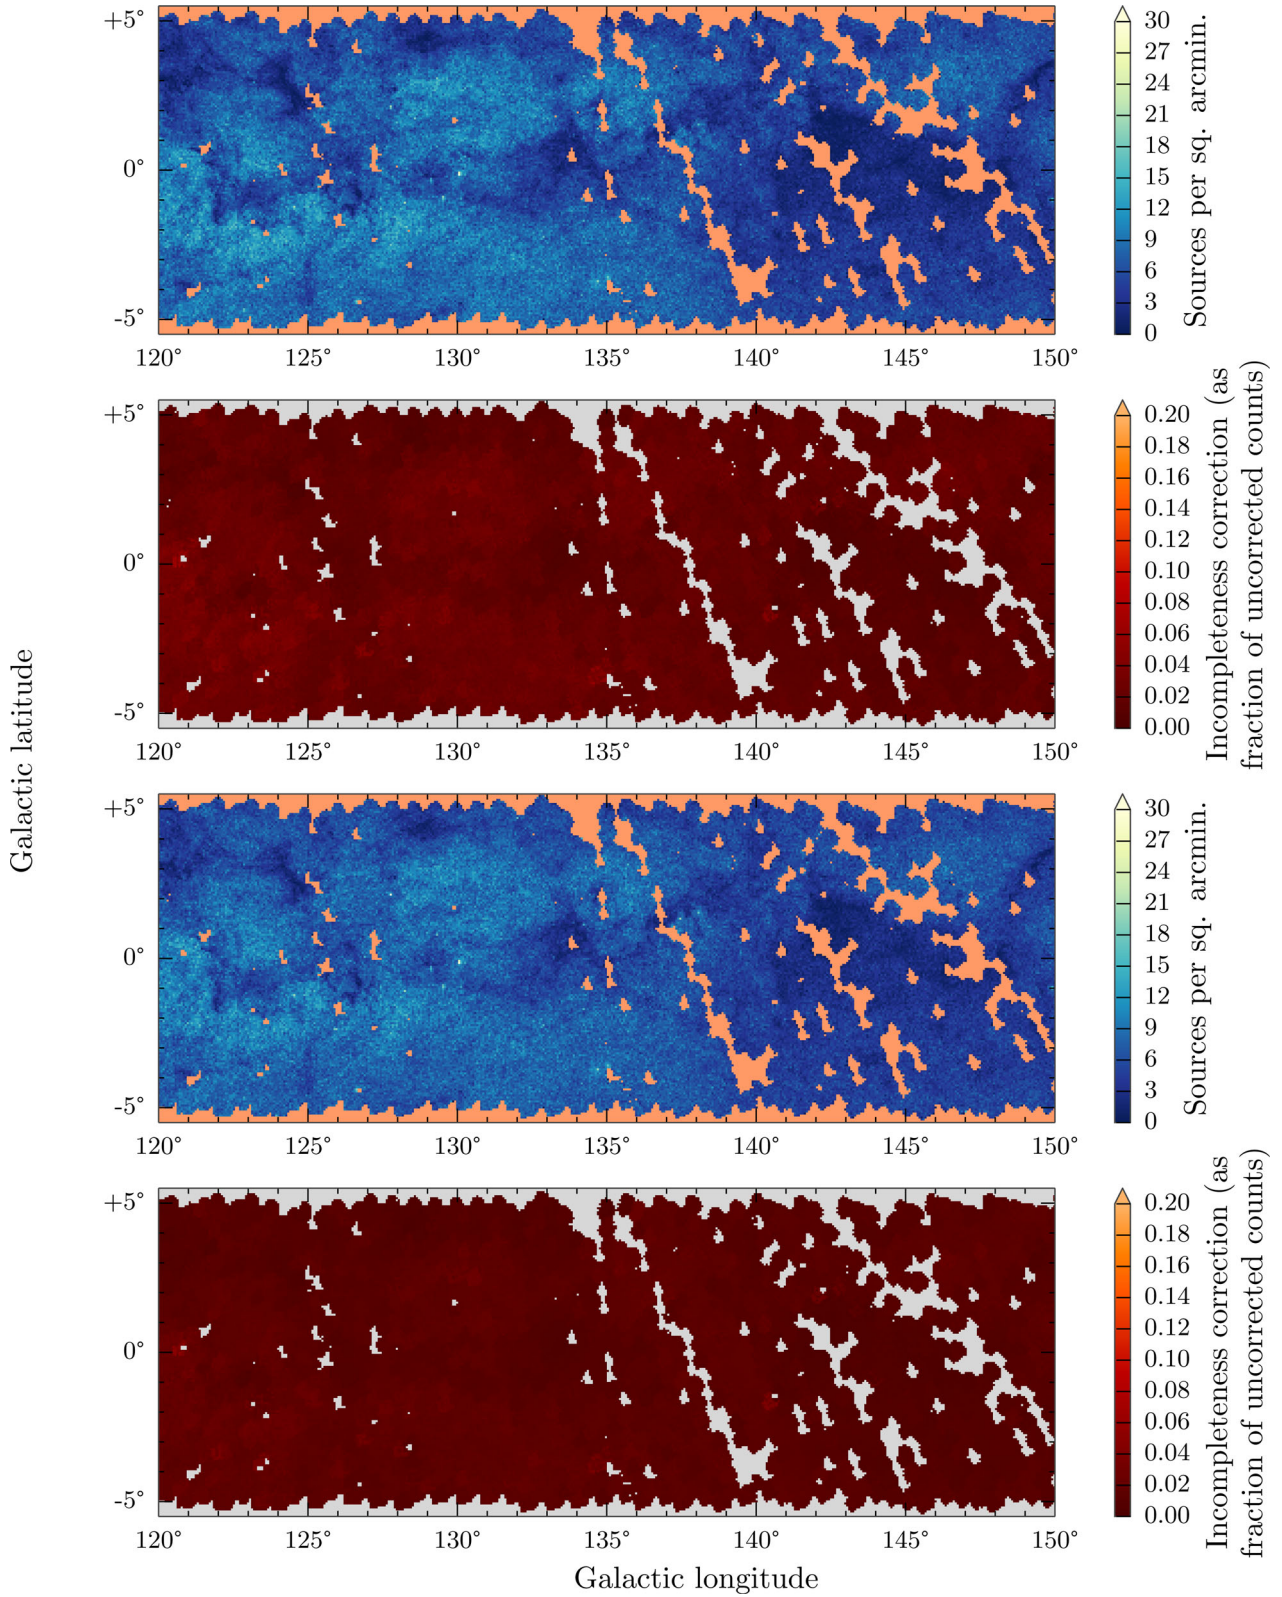
Figure A4. As Fig. A1, only for $120° < \ell < 150°$.
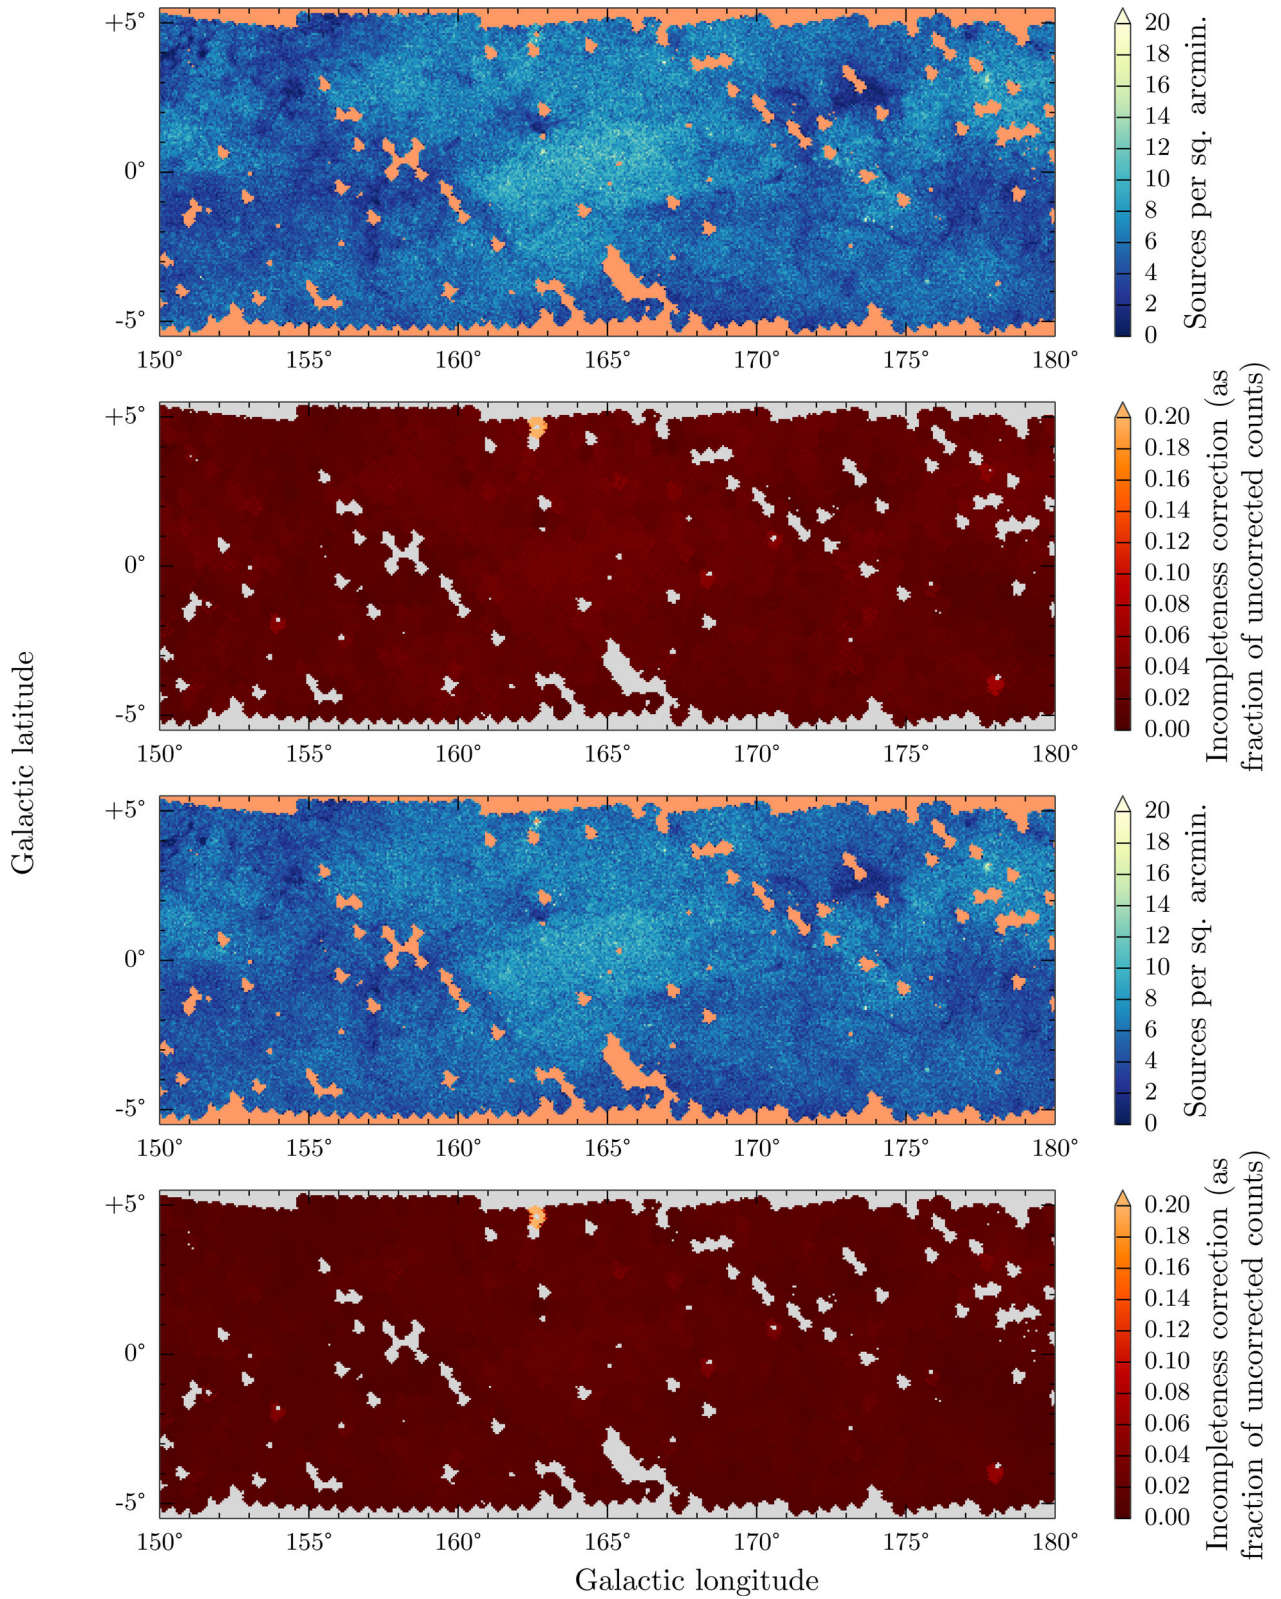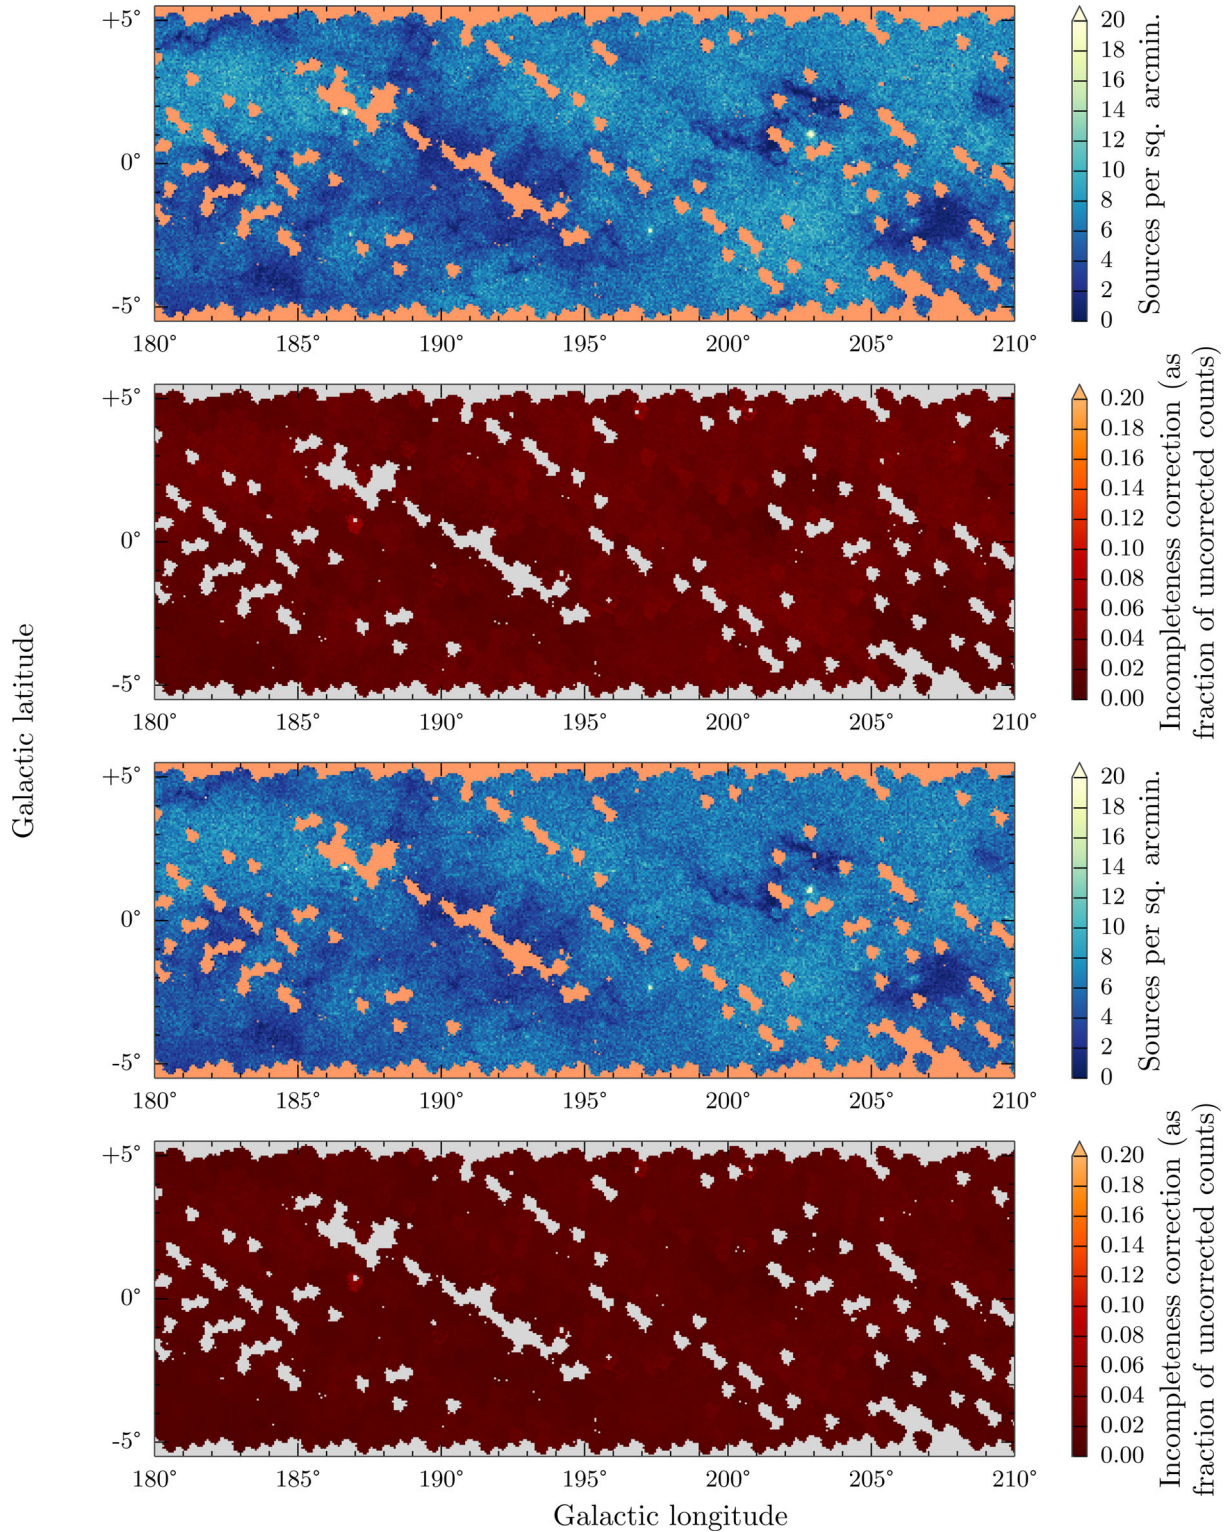
Figure A5. As Fig. A1, only for $150° < \ell < 180°$.

Figure A6. As Fig. A1, only for $180° < \ell < 210°$.

This paper has been typeset from a TEX/LATEX file prepared by the author.