

A Residual Census of Unidentified X-ray Sources in Omega Centauri after the Modern Pulsar Catalogue

Tim Swanson

The Omega Centauri Society / Post Oak Labs
tim@postoaklabs.com

August 2026

v1.0 (last revised 2026-08-16) — eighth paper of the set (Paper G); companion to the Macro Transcension Hypothesis (Paper A), the inward-migration review (Paper B), the Omega Centauri campaign (Paper C), the migration economics (Paper D), the engineered-IMBH systems paper (Paper E), the joint accretion bound (Paper F), and the mass-tension paper (Paper H); prepared for omegacentauri.me

Abstract

The deepest Chandra catalogue of Omega Centauri ([Henleywillis et al., 2018](#)) lists 233 X-ray sources, 188 of them without an optical identification, and notes that more than 30 of the unidentified sources carry luminosities and X-ray colours resembling millisecond pulsars (MSPs) in other clusters. That catalogue was published when no radio MSP was known in the cluster. Nineteen are known now ([Colom i Bernadich et al., 2026](#)), and the catalogue has never been cross-matched against them as a set. We supply that cross-match and the residual census it leaves behind, running the identical pipeline on 47 Tucanae first as a control. The control recovers 18 of 21 known MSP X-ray counterparts at the 95 per cent match radius, and the three it misses are the three rows the source catalogue itself flags as deviating beyond that radius ([Bhattacharya et al., 2017](#)), so the pipeline reproduces the control’s own flag set rather than an independent error budget. In ω Cen the Chandra astrometric frame carries a measured $+0.42''$ declination offset relative to the radio frame; registering it takes tier-1 matches from three of eight to five of eight and, with the new timing positions, tightens two published associations from $2.65''$ and $2.42''$ to $0.07''$ and $0.01''$. All five registered associations sit between $0.01''$ and $0.35''$ with per-source chance-coincidence probabilities below 2.3×10^{-3} under each of three independent estimators, reported separately and never combined. One published association is discrepant: the timing position of MSP H sits $1.81''$ from source 14c after registration, in the opposite declination sense to the other five anchors, against the $0.73''$ offset quoted for the same pulsar in [Zhao and Heinke \(2023\)](#). We record that as a data-quality finding rather than reconciling it. This is a positional and colour census; variability is deferred and addressed only in the back matter. After subtracting a $\log N - \log S$ background expectation of 146.9 (92.4–215.5) sources from the 159 unidentified sources inside a $480''$ aperture, the residual is 12.1 (–56.5, +66.6), consistent with zero and dominated by the background bracket. Between 39 and 53 unidentified sources fall inside a colour-flux region defined by this cluster’s own confirmed counterparts, a quantity that is not interchangeable with the catalogue’s own “more than 30”. The two published population predictions disagree, predicted totals of 324 and 6 pulsars (a factor of 54), which against 19 known imply 176.6 and zero X-ray-detectable undiscovered pulsars ([Yin et al., 2024](#)). We bound that disagreement against the residual without adjudicating it. The residual is an input to future identification work and is not a detection of anything.

Keywords: globular clusters: individual: NGC 5139 (Omega Centauri); X-rays: binaries; pulsars: general; catalogues; methods: statistical

Contents

1	Introduction	3
2	Data	4
2.1	The X-ray catalogue	4
2.2	The pulsar positions	4
2.3	The published counterparts	4
2.4	Source 24f as a classification-supersession case	4
2.5	The control catalogue	5
3	Methods	5
3.1	Tiered cross-match	5
3.2	Three chance-coincidence estimators	5
3.3	Astrometric frame registration	6
3.4	Membership as a probability	6
3.5	Background accounting	6
3.6	Control first	7
4	Control validation	7
5	Results for Omega Centauri	8
5.1	Frame registration	8
5.2	Matches	8
5.3	The MSP H association is discrepant	9
5.4	Classification census	10
5.5	Background accounting	12
5.6	The residual census	12
5.7	The MSP-like region	12
6	The population disagreement	13
7	Discussion	14
7.1	What would identify the residual	14
7.2	The census as an anomaly screen	14
8	Conclusion	14
A	Caveat checklist	15

1 Introduction

Henleywillis et al. (2018) identify 233 X-ray sources in a 222 ks Chandra ACIS-I exposure of Omega Centauri (two epochs, 173.7 and 48.5 ks, per the source paper’s own description; distinct from the deeper 290.9 ks, four-pointing reduction Zhao and Heinke (2023) use for their own X-ray analysis of the same field), 95 of them new, reaching a limiting unabsorbed flux in the core of $f_X(0.5\text{--}6.0\text{ keV}) \approx 3 \times 10^{-16} \text{ erg s}^{-1} \text{ cm}^{-2}$. Forty-five of those sources carry a likely optical identification. The remaining 188 do not. The catalogue’s own summary of that unidentified population is the sentence this paper takes as its starting point: no radio millisecond pulsar was then known in the cluster, and more than 30 unidentified sources had luminosities and X-ray colours like those of MSPs found elsewhere.

The situation has since inverted. MeerKAT found 13 pulsars in the cluster (Chen et al., 2023), and the TRAPUM and Parkes programme now reports a timing set of 19 (Colom i Bernadich et al., 2026). Zhao and Heinke (2023) identified X-ray counterparts for 11 of the 18 MSPs known at the time of their work, nine of them new. The catalogue of unidentified X-ray sources and the catalogue of radio pulsars therefore both exist, and neither has been confronted with the other as a set. The confrontations that have been made are partial and local. Colom i Bernadich et al. (2026) note a marginal positional consistency between pulsar G and Chandra source 24f through an ATCA position and identify 24f as G’s counterpart in their conclusions, and they note pulsar L as positionally consistent at roughly 3σ with source 32d. Those two statements are the prior art this paper extends, and they are credited as such.

Two things follow from doing the cross-match systematically. The first is a set of associations with stated false-association rates, which the partial treatments do not carry. The second is what is left over: a residual population of unidentified sources that survives every known counterpart class and the expected extragalactic background. The residual is the object of interest, and it is worth stating at the outset what kind of object it is. It is a target list. A count of unidentified sources in excess of a background model is a statement about the completeness of identification work in a crowded, unevenly exposed field, and it constrains nothing on its own.

The referee-facing risks in a paper of this kind are well known and each is addressed mechanically rather than rhetorically. Positional cross-matching in a crowded core without a false-association calibration is uninterpretable, so every probe carries a chance-coincidence probability from three independent estimators (Section 3.2). Colour-box classification without membership priors is arithmetic without content, so membership enters as a probability under stated external normalisations and never as a binary label (Section 3.4). Extragalactic contamination dominates a field this deep, so the background enters through a published $\log N\text{--}\log S$ relation with its parameter uncertainties propagated as a bracket (Section 3.5). Fluxes near the detection limit are biased high, which we state as a direction of error rather than correct. And the whole pipeline is validated on a control cluster before it is allowed to produce a number for ω Cen (Section 4).

Section 2 fixes the catalogues and their provenance. Section 3 gives the cross-match, the three chance estimators, the frame registration, the membership treatment, and the background accounting. Section 4 reports the control validation, which is a result and not a preamble. Section 5 gives the ω Cen cross-match, the discrepant association, and the residual census. Section 6 carries the population-model disagreement. Section 7 states what would identify the residual. Appendix A collects the caveats as a checklist.

Epistemic status: Everything in Sections 2–6 is a count, a measured separation, or an expectation from a cited external relation. No detection is claimed anywhere in this paper, and the residual population is presented as an input to future identification work.

2 Data

2.1 The X-ray catalogue

The source catalogue is [Henleywillis et al. \(2018\)](#), obtained as the machine-readable table from VizieR (J/MNRAS/479/2834) and parsed by byte-column specification rather than transcribed. The parse reproduces the paper’s own abstract totals exactly: 233 sources, 95 new, 45 with a likely optical identification. Each source carries a position, a 95 per cent confidence error radius, an offset from the [Anderson and van der Marel \(2010\)](#) cluster centre in units of the 155'' core radius, detected and exposure-corrected counts in the medium (0.5–4.5 keV), soft (0.5–1.5 keV) and hard (1.5–6.0 keV) bands, unabsorbed fluxes, the soft-to-hard colour $\log(\text{soft/hard})$, and an optical identification field whose null value defines the unidentified population.

2.2 The pulsar positions

Pulsar positions come from [Colom i Bernadich et al. \(2026\)](#) and are sorted into precision tiers before any matching is attempted, because a single match radius applied across a set with heterogeneous astrometry is the standard way to manufacture associations.

Tier 1, eight pulsars (A, B, C, D, E, G, H, K). Timing-derived positions at milliarcsecond scale. The radio positional uncertainty is set to zero, so the match radius is carried entirely by the X-ray error circle.

Tier 2, five pulsars (I, L, N, Q, S). Positions from orbital solutions with no formal uncertainty quoted, plus the SeeKAT localisation of S at $\pm 1.5''$. For the four without a stated error we adopt one unit in the last quoted digit as the uncertainty. That rule is ours and not the source paper’s, and it is why tier 2 is reported separately and not merged into a headline count.

Null, six pulsars (F, J, M, O, P, R). No position in [Colom i Bernadich et al. \(2026\)](#), attributed there to [Chen et al. \(2023\)](#). These six are absent from the cross-match entirely. Recovering their positions from the discovery paper would extend the probe set by six probes and is the cheapest improvement available to this analysis.

2.3 The published counterparts

[Zhao and Heinke \(2023\)](#) report X-ray counterparts for 11 of the 18 MSPs known at the time, of which seven carry a cross-reference to a [Henleywillis et al. \(2018\)](#) identifier: B→13d, C→23g, E→11f, G→24f, H→14c, K→21d, L→32d. Their tables were read from the published PDF page by page and the positions and luminosities re-verified against that read. Five of the eleven rows required correction of the luminosity uncertainty from a symmetric placeholder to the asymmetric interval the source paper quotes; central values were unaffected. We do not re-derive the luminosity-companion-mass correlation of that paper, which is their result and not ours.

2.4 Source 24f as a classification-supersession case

[Henleywillis et al. \(2018\)](#) classify source 24f as RGB/SGB-a, an anomalous red-giant-branch or sub-subgiant candidate. Later radio work identifies the same source as the X-ray counterpart of MSP G ([Zhao and Heinke, 2023](#); [Colom i Bernadich et al., 2026](#)), and after frame registration the timing position sits

0.071'' from it (Section 5.2). This is the one class assignment in the catalogue known to have been superseded by later radio work.

We carry 24f under its published class in every count reported in this paper and describe the supersession separately. Silently reclassifying one source because a later result is known would make the classification census a mixture of two epochs of knowledge, and the residual would then depend on which supersessions happened to be noticed. Reporting the case explicitly costs one paragraph and keeps the census reproducible from the published tables alone. It also sets the scale of the effect: one source in 233, in a catalogue whose identification work was done before any cluster MSP was known.

2.5 The control catalogue

The control is 47 Tucanae, from [Bhattacharya et al. \(2017\)](#), vendored from VizieR (J/MNRAS/472/3706) under the same byte-parse discipline. The record counts reproduce the source exactly: 370 sources within the half-mass radius, a 21-row MSP X-ray position-deviation table, and a 5-row table of new optical counterparts.

One trap in that paper deserves flagging because it would corrupt a control comparison built on the title alone. The title phrase “five new X-ray counterparts to millisecond radio pulsars” does not describe five new MSP counterparts. The five new optical counterparts of the paper’s table 5 are chromospherically active BY Dra and W UMa variable stars. The five MSPs named in the abstract are five pulsars given new spectral extractions. The MSP counterpart truth set for a control comparison is the 21-row deviation table, which covers effectively all known 47 Tuc MSPs, and that is the table used here.

Two asymmetries between control and target are relevant to the design and are not pipeline artifacts. The control’s MSP identification is close to complete, 21 counterparts against roughly 23 known pulsars, where ω Cen sits at 11 counterparts among 18 to 19 known. And the luminosity bands differ, 0.5–6 keV for the control against 0.5–10 keV for ω Cen, so no cross-cluster luminosity comparison is made anywhere in this paper.

3 Methods

3.1 Tiered cross-match

For each pulsar probe i and X-ray source j the match radius is

$$r_{ij} = \sqrt{\sigma_{\text{radio},i}^2 + \epsilon_{95,j}^2}, \quad (1)$$

where $\epsilon_{95,j}$ is the catalogue’s tabulated 95 per cent error circle radius and $\sigma_{\text{radio},i}$ is the radio positional uncertainty of the tier. A probe matches the nearest source inside its radius. Tiers are matched and reported separately throughout. The same code path, with no parameter changes, runs on the control.

3.2 Three chance-coincidence estimators

Every probe carries three independent estimates of the probability that its match is a coincidence, each computed with 2×10^4 trials at a fixed seed.

Azimuthal rotation. The probe position is rotated about the cluster centre by a uniform random angle, preserving its cluster-centric radius and therefore the local surface density, and the match is re-attempted.

Catalogue offset. The source catalogue is displaced by a random vector large compared with the match radius and small compared with the field, and the match is re-attempted at the probe’s true position.

Analytic Poisson. $P_{\text{chance}} = 1 - \exp(-\pi r_{ij}^2 \Sigma_{\text{local}})$, with Σ_{local} measured in an annulus at the probe’s cluster-centric radius.

The three answer slightly different questions. Rotation preserves the radial density profile and destroys azimuthal structure; offset preserves local clustering and destroys the absolute frame; the analytic estimate assumes a locally Poisson field. They are reported side by side for every probe and are never averaged or combined into a single number, because their disagreement is itself informative about which assumption the field violates.

3.3 Astrometric frame registration

Published associations provide anchors from which a bulk offset between the Chandra frame and the radio timing frame can be measured. We take the median of the anchor deviations, with the median absolute deviation as the spread, and apply the resulting shift to every probe position before the cross-match is run. Registration is reported both ways, before and after, throughout.

The limitation is structural and we state it rather than working around it. The anchors are published associations, so re-matching those same pairs after registration is not an independent test of them. The registration earns its place on the probes that are not anchors, where it shifts the search position by a measured amount, and on the resulting change in the tier-1 match count.

3.4 Membership as a probability

Membership is assigned from a two-component spatial model, a King-like cluster profile at the catalogue’s 155’’ core radius plus a uniform background, giving each source a probability P_{member} from its cluster-centric radius.

Fitting both components to the data internally fails, and the failure is reported as a diagnostic rather than buried. The internal fit is degenerate: it returns 198.33 members out of 200 fitted sources with a background of 32.7 per square degree, an answer wrong by roughly two orders of magnitude in the background and implying near-total membership. The cause is instrumental. Chandra sensitivity falls with off-axis angle, so the observed source density declines with radius for reasons unrelated to cluster membership, and the profile term absorbs the background’s own radial decline.

Membership probabilities in this paper therefore use external normalisations, varied explicitly rather than fitted: the number of cluster members is set to 40, 60 or 80, bracketing the $\sim 60 \pm 20$ estimate of [Henleywillis et al. \(2018\)](#), and the background surface density is set from the $\log N$ – $\log S$ relation of Section 3.5 at two flux limits. Every published membership statement in the catalogue is carried alongside the computed probability and never overwritten by it. Repairing the degeneracy properly requires the Chandra exposure map, which the vendored table does not carry.

3.5 Background accounting

Extragalactic background counts come from the soft-band (0.5–2.0 keV) cumulative number counts of [Moretti et al. \(2003\)](#),

$$N(> S) = N_S \frac{(2 \times 10^{-15})^{\alpha_1}}{S_0^{\alpha_1 - \alpha_2} S^{\alpha_2} + S^{\alpha_1}} \text{ deg}^{-2}, \quad (2)$$

Table 1: Control recovery on 47 Tucanae (Bhattacharya et al., 2017), 21 known MSP X-ray counterparts, 2×10^4 randomisation trials per probe at a fixed seed. Expected spurious counts are the number of matches each randomisation estimator returns over the full probe set.

Match radius scale	Recovered of 21	Expected spurious (rotation)	Expected spurious (offset)
1.0	18	0.39	0.10
1.5	20	0.86	0.23
2.5	21	2.39	0.64

with $\alpha_1 = 1.82$, $\alpha_2 = 0.60$, $S_0 = 1.48 \times 10^{-14}$ and $N_S = 6150$, fitted over 2.4×10^{-17} to 1×10^{-11} erg cm $^{-2}$ s $^{-1}$. The catalogue’s tabulated soft band is 0.5–2.0 keV, the same band, so no band correction is applied. As a transcription check the implementation reproduces that paper’s own integrated soft-band flux, 6.86×10^{-12} against the published 6.85×10^{-12} erg s $^{-1}$ cm $^{-2}$ deg $^{-2}$.

The published parameter uncertainties are used as a bracket rather than a confidence interval, built by moving the faint-end slope and the normalisation together to their extreme ends. They are quoted asymmetrically in the source and are used here only as symmetric widths, so no asymmetric error appears anywhere in this paper.

All area-normalised quantities refer to one circular aperture of radius 480'' inscribed in the ACIS-I field, 0.05585 deg 2 , containing 200 of the 233 catalogue sources. The 33 sources outside it are excluded from both sides of every comparison.

3.6 Control first

The control cluster is processed before the ω Cen residual is computed, and the residual is gated on the control passing. The gate criterion was fixed in advance: every unrecovered control counterpart must be a row the source catalogue itself records as deviating beyond the 95 per cent radius. A pipeline that misses counterparts the source table does not flag is failing on its own terms.

4 Control validation

Gate: passed. The pipeline recovers 18 of the 21 known 47 Tuc MSP X-ray counterparts at the 95 per cent match radius, 20 at 1.5 $\times$ that radius, and all 21 at 2.5 $\times$ (Table 1).

The three counterparts missed at the 95 per cent radius are Q, L and R, at separations of 0.34'', 0.63'' and 0.45''. These are the three rows of the source table whose normalised deviation D/X_{err} exceeds unity, at 1.04, 2.03 and 1.42. The pipeline reproduces the control catalogue’s own flag set. It does not merely recover a high fraction of counterparts; the identity of what it misses is predicted by the source table, which is a stronger statement than a recovery rate.

Two further checks come with the control. The median X-ray minus radio deviation over all 21 counterparts is 0.000'' in right ascension and $-0.070''$ in declination, a bulk offset consistent with zero, which is what an already-registered catalogue should show and which confirms that the radio positions reconstructed for this test carry no systematic error. Those radio positions are reconstructed as the X-ray position minus the tabulated deviation vector, and the reconstruction reproduces the table’s own total deviations to 0.01''.

The control result is an optimistic bound on match performance rather than an independent astrometric test, and we say so. Bhattacharya et al. (2017) registered their X-ray astrometric frame onto the radio frame

Table 2: Frame-registration anchors: deviations of the six published associations, and the residual separation after the median shift ($-0.081''$, $+0.411''$) is applied.

MSP	Chandra source	$\Delta\alpha$ (")	$\Delta\delta$ (")	residual after registration (")
B	13d	+0.062	+0.734	0.353
C	23g	+0.111	+0.142	0.331
E	11f	-0.072	+0.463	0.053
G	24f	-0.151	+0.400	0.071
H	14c	-0.216	-1.398	1.814
K	21d	-0.090	+0.422	0.014

using X-ray detections of 19 radio MSPs. A pipeline tested on a catalogue that was itself registered against the truth set is being asked an easier question than the one ω Cen poses, where the source performed no such registration. That is the argument for measuring the ω Cen frame offset directly rather than assuming the frames agree, which Section 5.1 does.

5 Results for Omega Centauri

5.1 Frame registration

Measured over the five associations that anchor the registration (B, C, E, G, K; the discrepant MSP H association is treated separately in Section 5.3), the Chandra frame sits at $\Delta\alpha = -0.072''$, $\Delta\delta = +0.422''$ relative to the radio frame, with median absolute deviations of $0.079''$ and $0.041''$. As a sensitivity check, including the discrepant sixth anchor (MSP H) shifts the median to $\Delta\alpha = -0.081''$, $\Delta\delta = +0.411''$ (MAD $0.103''$, $0.161''$; Table 2): both medians move by $\leq 0.011''$, inside either MAD, so no registered separation or match status in Table 3 changes at the $0.015''$ precision quoted throughout this paper, and no re-run of the tabulated cross-match is required.

Five of the six anchors fall between $+0.14''$ and $+0.73''$ in declination about a median of $+0.41''$. The sixth sits at $-1.40''$ and is treated in Section 5.3. A shift of this size is small compared with the field and large compared with the tier-1 match radii, which run from $0.36''$ to $0.49''$, so it changes the answer. Any ω Cen cross-match number quoted without it is quoting an unregistered frame.

5.2 Matches

Registration takes the tier-1 match count from three of eight to five of eight. Table 3 gives every tier-1 probe with its separation before and after registration and two of its three chance-coincidence probabilities; the third is discussed below.

The analytic Poisson estimator returns 0.8×10^{-3} to 1.2×10^{-3} for all eight probes, tracking the local surface density and showing no structure the other two do not also show. The five registered associations sit at separations of $0.014''$ to $0.353''$ with every estimator below 2.3×10^{-3} , and the expected number of spurious tier-1 matches over the whole probe set is 0.009 by rotation and 0.004 by offset. B's match is the marginal case of the five: its $0.353''$ separation sits close to its $0.36''$ match radius, so it is the association most sensitive to any further refinement of either position. The associations for B, C, E, G and K agree with the published identifications for the same pulsars, an expected outcome: all five are registration anchors, so the informative content is in the separations and chance rates, not the identities.

Registration is what makes two of these associations tight. Zhao and Heinke (2023) quote offsets of $2.65''$ for pulsar G and $2.42''$ for pulsar K between their radio and X-ray positions. With the timing

Table 3: Tier-1 cross-match, eight timing-position probes. Separations in arcseconds; P_{chance} from independent estimators at 2×10^4 trials, reported separately and never combined. The match radius is $\sqrt{\sigma_{\text{radio}}^2 + \epsilon_{95}^2}$ with $\sigma_{\text{radio}} = 0$ for this tier.

Pulsar	Match	Sep. before registration	Sep. after registration	Match radius	P_{chance} (rotation)	P_{chance} (offset)
A	–	23.468	23.309	–	$< 5 \times 10^{-5}$	0.0003
B	13d	0.736	0.353	0.36	0.0022	0.0004
C	23g	0.180	0.331	0.46	0.0006	0.0008
D	–	19.915	19.596	–	0.0010	0.0004
E	11f	0.469	0.053	0.49	0.0022	0.0001
G	24f	0.428	0.071	0.45	0.0013	0.0006
H	–	1.415	1.814	–	$< 5 \times 10^{-5}$	0.0004
K	21d	0.431	0.014	0.38	0.0011	0.0006

positions of [Colom i Bernadich et al. \(2026\)](#) and the measured frame shift applied, the same two associations sit at $0.071''$ and $0.014''$.

Tier 2 returns no matches from five probes, at separations of $4.1''$ to $24.1''$, against an expected spurious count of 0.261 by rotation. Pulsar L, which [Colom i Bernadich et al. \(2026\)](#) note as positionally consistent with source 32d at roughly 3σ , sits $4.06''$ from its nearest catalogue source and outside the match radius our uncertainty rule produces. That rule is ours, not the source paper’s, so the non-match is reported as a consequence of a stated convention rather than as evidence against their association.

Six pulsars (F, J, M, O, P, R) carry no position in [Colom i Bernadich et al. \(2026\)](#) and are absent from the cross-match entirely. Every match count in this paper is conditional on that omission.

5.3 The MSP H association is discrepant

[Zhao and Heinke \(2023\)](#) quote an offset of $0.73''$ between MSP H and their X-ray source h1, and cross-reference h1 to [Henleywillis et al. \(2018\)](#) source 14c. The timing position of [Colom i Bernadich et al. \(2026\)](#) sits $1.415''$ from the catalogue position of 14c, and $1.814''$ from it after frame registration. The deviation runs in the opposite declination sense to the other five anchors, $-1.398''$ against a $+0.42''$ median.

Two explanations are available and we do not choose between them. Either 14c is not the [Henleywillis et al. \(2018\)](#) counterpart of source h1, in which case the cross-reference between the two catalogues is wrong for this row while both positions may be correct, or one of the two published positions is in error. Both are checkable, and the check is specific: the [Zhao and Heinke \(2023\)](#) source list needs comparing against the [Henleywillis et al. \(2018\)](#) catalogue position by position. We have not done that, and it is not a question the vendored tables settle on their own.

The consequence for this paper is bounded. The primary frame registration (Section 5.1) already excludes H; including it as a sensitivity check moves both medians by $\leq 0.011''$, well inside the six-anchor MAD of $0.103''/0.161''$, and changes no registered separation or match status at the $0.015''$ precision quoted elsewhere in this paper. The discrepancy matters as a data-quality finding about the counterpart chain rather than because it moves any number here.




Figure 1: Tier-1 separations before and after the $(-0.081'', +0.411'')$ frame registration, with the match radius of each registered association marked. Registration moves B, E, G and K below their radii and moves H further away. A and D have no catalogue source within $19''$ under either frame. Values from `analysis/results.omega.cen.json`.

Table 4: Published class census over all 233 catalogue sources (Henleywillis et al., 2018), membership as published. Source 24f is counted here under its published class (Section 2.4).

Class	Published membership	n
unidentified	no published statement	188
CV / CV candidate	member	18
AGN / AGN candidate	non-member	9
anomalous RGB / sub-subgiant	member	5
variable star	member	4
variable star	non-member	4
foreground star	non-member	3
quiescent LMXB	member	1
CH star (symbiotic candidate)	member	1

5.4 Classification census

Published classes are taken verbatim from Henleywillis et al. (2018) and are not re-derived; the parenthesis convention of that catalogue marks non-members. The parse reproduces the paper’s own abstract composition (Table 4): 18 cataclysmic variables and candidates, one quiescent low-mass X-ray binary, four member variable stars, five anomalous-RGB or sub-subgiant stars, and one CH star, against 45 identifications in total.

Membership probabilities under the external normalisations of Section 3.4 are given in Table 5 and plotted in Figure 2. The spread across assumptions is the result. Summed membership over the aperture runs from 53.1 to 86.0 sources as the assumed member count moves from 40 to 80 and the flux limit moves between its two settings, and the number of individual sources crossing $P_{\text{member}} > 0.5$ runs from 0 to 67 over the same range. A binary membership label extracted from any single row of that table would be an artifact of the normalisation chosen.

Table 5: Membership sensitivity. Rows are the six combinations of assumed cluster-member count and background flux limit. Sums run over the 200 sources inside the 480'' aperture.

Assumed members	Flux limit	$N(P_{\text{member}} > 0.5)$	$\sum P_{\text{member}}$	$\sum P_{\text{member}}$ (unidentified)
40	core limiting	0	53.1	41.2
40	5th percentile	6	55.9	43.3
60	core limiting	27	69.5	54.0
60	5th percentile	38	72.7	56.5
80	core limiting	55	82.5	64.3
80	5th percentile	67	86.0	67.0




Figure 2: Membership ladder. Summed membership probability over the 480'' aperture, and over its unidentified subset, against the assumed external member normalisation at two background flux limits. The count of sources individually exceeding $P_{\text{member}} = 0.5$ (dotted) moves from 0 to 55 over the same range, which is why membership is reported as a summed probability under stated assumptions and not as a label. Values from analysis/results_omega_cen.json.

Table 6: Expected extragalactic background over the 480'' aperture at four candidate flux limits, from [Moretti et al. \(2003\)](#). Brackets from the published parameter uncertainties used as symmetric widths. The row marked † lies below the relation’s fitted flux range.

Flux limit	S (erg cm ⁻² s ⁻¹)	N per deg ²	Expected in field	Bracket
core limiting, scaled to soft	1.24×10^{-16}	2827	157.9	99.0–232.5
faintest tabulated soft†	1.00×10^{-17}	12852	717.8	417.2–1140.5
fifth percentile soft (adopted)	1.40×10^{-16}	2630	146.9	92.4–215.5
median soft	9.90×10^{-16}	787	43.9	29.4–60.7

5.5 Background accounting

Table 6 gives the expected background over the 0.05585 deg² aperture at four candidate flux limits. The limit adopted for the residual is the fifth-percentile soft flux of the catalogue itself, 1.4×10^{-16} erg cm⁻² s⁻¹, giving 146.9 expected background sources with a bracket of 92.4 to 215.5. The core limiting flux scaled to the soft band gives a nearly identical 157.9, and the two lie well inside each other’s brackets. The faintest tabulated flux lies below the fitted range of the [Moretti et al. \(2003\)](#) relation and is shown only to mark where extrapolation begins.

5.6 The residual census

Inside the 480'' aperture the catalogue holds 200 sources, of which 41 carry a published identification (27 published members, 14 published non-members) and 159 do not. Eleven sources are confirmed MSP X-ray counterparts. Subtracting the adopted background expectation from the unidentified count gives

$$N_{\text{residual}} = 159 - 146.9 (92.4\text{--}215.5) = 12.1 (-56.5, +66.6). \quad (3)$$

The residual is background-dominated. Its central value is 8 per cent of the quantity subtracted and its bracket spans zero by a wide margin on both sides. The conservative reading is that the unidentified population inside this aperture is consistent with the extragalactic background alone, and that no excess is required by these data at this flux limit. Nothing in the rest of this paper should be read as claiming otherwise. What the residual does supply is a scale: an identification programme working through this field is working against a background that supplies of order 150 sources, and any cluster population it hopes to extract is a minority of the unidentified list.

Three accounting conventions change the number and are stated here. Published non-member AGN identifications sit in the identified column, so the subtraction removes a predicted background from a population that already excludes background sources someone has taken the trouble to identify, which biases the residual low. The aperture excludes 33 catalogue sources from both sides, so this is not a statement about the full 233. And the eleven confirmed MSP X-ray counterparts are retained in the unidentified column, since the catalogue’s own optical-identification field predates the pulsar timing set; removing them would lower the residual by that count, well inside the quoted bracket.

5.7 The MSP-like region

The confirmed counterparts define a region in the colour-flux plane from this cluster’s own data. Four of them ($n = 4$: 23g, 11f, 24f, 21d) carry both a tabulated colour and a flux and anchor it; a fifth recovered counterpart lacks a tabulated colour and is dropped. The primary construction is a bounding box on those same four anchors: counting unidentified sources inside it gives 39, of which 14.85 survive membership

Table 7: Population predictions from [Yin et al. \(2024\)](#), carried as a bracket. The final two columns apply the X-ray detection fraction measured in this cluster under its two denominators: 0.611 (11 of the 18 MSPs [Zhao and Heinke \(2023\)](#) searched) and 0.579 (11 of the 19 known today, conservative).

Model	Predicted total	Undiscovered	X-ray-detectable (0.611)	X-ray-detectable (0.579)
escape velocity	324	305	186.4	176.6
encounter rate	6	0	0.0	0.0

weighting at the adopted normalisation. A 2σ ellipse on the same $n = 4$ anchors, reported as a secondary and wider construction, gives 53 (19.44 weighted). Eleven unidentified sources carry neither a colour nor a flux and cannot be placed in the plane at all. Neither region screens for background AGN, so the 39 and 53 counts are upper bounds on MSP-like membership rather than membership estimates.

This is not the “more than 30 MSP-like” count of [Henleywillis et al. \(2018\)](#). That criterion compares against MSPs in other clusters and is not tabulated in the source paper, so it cannot be reproduced from the machine-readable table. The region used here is defined from four counterparts in this cluster, a different construction with a different systematic. The two numbers are not interchangeable, and the numerical proximity of 39 to “more than 30” is a coincidence of construction rather than a confirmation.

6 The population disagreement

Two published estimates of the total pulsar population of ω Cen are available and they disagree by a factor of 54. [Yin et al. \(2024\)](#) report an empirical-Bayesian escape-velocity scaling relation that predicts 324 pulsars at this cluster’s escape velocity of 62.2 km s^{-1} . The encounter-rate estimate they tabulate alongside it gives 6.

Against 19 known pulsars those totals imply 305 and 0 undiscovered. The X-ray detection fraction measured in this cluster depends on which denominator is used: [Zhao and Heinke \(2023\)](#) found 11 counterparts among the 18 MSPs known and searched at the time of their work (0.611), while 19 pulsars are known today and the 19th was never searched, not a non-detection, so the conservative variant against today’s full count is 11 of 19 (0.579). We carry both rather than choosing: applying 0.611 gives 186.4 and applying 0.579 gives 176.6 X-ray-detectable undiscovered pulsars from the escape-velocity model; the encounter-rate model gives 0.0 under either fraction, since it predicts zero undiscovered pulsars (Table 7).

The escape-velocity number exceeds the residual of 12.1 by more than an order of magnitude, and exceeds the full unidentified count of 159 once the background is not subtracted. That bounds something. It bounds the extrapolation, or the assumed detection fraction, or the background subtraction, and this analysis does not decide which. Three reasons keep each candidate live.

The extrapolation. The [Yin et al. \(2024\)](#) relation was fitted on FAST-visible clusters. ω Cen at $\delta = -47^\circ$ is not FAST-visible, so applying the relation here reaches outside the sample it was fitted on, in a cluster whose escape velocity places it at the top of the predicted range.

The detection fraction. The measured 0.611 (or 0.579 against today’s full known count) is the fraction of searched pulsars with a Chandra counterpart. Applying it to the undiscovered population assumes those pulsars share the X-ray luminosity distribution of the discovered ones. Radio-faint or large-radius pulsars need not, and the discovered set is selected in ways that correlate with X-ray brightness, so either fraction is more plausibly an upper bound than an unbiased estimate.

The background subtraction. Eddington bias inflates the faint fluxes on which the background limit is computed, which makes the expected background a lower bound and the residual an upper bound. Correcting the residual downward would widen the gap rather than close it.

We carry the disagreement as a disagreement. Adjudicating it requires the identification work this census is meant to enable, not a preference asserted here.

7 Discussion

7.1 What would identify the residual

The residual is an input to identification work, and the useful output of this paper is a specific list of what that work requires. Each item below is a concrete, currently blocked step, ordered by cost.

The six missing pulsar positions. F, J, M, O, P and R carry no position in [Colom i Bernadich et al. \(2026\)](#). Recovering them from [Chen et al. \(2023\)](#) would extend the probe set by six probes and is the cheapest available improvement to every number in Section 5.

The counterpart-chain check. The MSP H discrepancy of Section 5.3 is resolvable by comparing the [Zhao and Heinke \(2023\)](#) source list against the [Henleywillis et al. \(2018\)](#) catalogue position by position. It requires no new observation.

Exposure-map membership. The degeneracy of Section 3.4 is instrumental and is fixed by the Chandra exposure map, which turns a spatial fit that cannot separate the cluster profile from the sensitivity profile into one that can. Membership would then be fitted rather than normalised externally, and the six-row sensitivity table of Table 5 would collapse to one row with an error bar.

Variability with epoch-resolved data. No variability analysis appears in this paper. Both the 2000 and 2012 Chandra epochs exist and the catalogue tabulates a flux ratio between them, but a single ratio does not support a statistically valid variability test at these count rates, and epoch-resolved event data are not staged here. This is the clearest gap between what this paper does and what a full identification programme needs, since variability is what separates quiescent binaries from background AGN when colour does not.

Deeper radio positions. Sub-arcsecond positions for the tier-2 pulsars would move five probes out of a tier whose uncertainty rule is a convention of ours rather than a measurement.

7.2 The census as an anomaly screen

A complete accounting of a cluster’s X-ray sources doubles as a screen for sources fitting no known class, a standard byproduct of catalogue work worth one sentence of acknowledgement in a series that elsewhere considers non-astrophysical hypotheses ([Swanson, 2026](#)). This census does not perform that screen. The residual computed here is background-dominated, consistent with zero, and defined by the absence of an optical identification rather than by any positive property; treating it as a population of anomalies would require the identification work Section 7.1 lists as outstanding, and until that work is done the residual carries no information about the nature of its members.

8 Conclusion

We have cross-matched the 233-source Chandra catalogue of [Henleywillis et al. \(2018\)](#) against the 19-pulsar timing set of [Colom i Bernadich et al. \(2026\)](#), validated the pipeline on 47 Tucanae first, and reported what is left.

The control recovers 18 of 21 known counterparts at the 95 per cent match radius and misses the three rows the control catalogue itself flags. The ω Cen Chandra frame carries a measured $+0.41''$ declination offset against the radio frame; registering it takes tier-1 matches from three of eight to five of eight, and

the five registered associations sit between $0.01''$ and $0.35''$ with chance-coincidence probabilities below 2.3×10^{-3} under three independent estimators. The published association of MSP H with source 14c does not survive registration, sitting at $1.81''$ in the wrong declination sense, which we record as a data-quality finding requiring a position-by-position check of the counterpart chain.

The residual census inside a $480''$ aperture is 12.1 unidentified sources ($-56.5, +66.6$) after subtracting a background expectation of 146.9 (92.4–215.5), which is to say it is consistent with the background alone. Between 39 and 53 unidentified sources fall in a colour-flux region defined by this cluster’s own confirmed counterparts, a quantity constructed differently from the catalogue’s own “more than 30” and not interchangeable with it. The two published population predictions differ by a factor of 54, and the larger over-predicts the residual by more than an order of magnitude, which bounds the extrapolation, the detection fraction, or the background subtraction without deciding among them.

Data availability

The analysis pipeline (`crossmatch.py`, `classify.py`, `residual.py`, `run_g.py`), its fixed seed, its two results files and the mechanical analysis report accompany this paper under `paper/g/analysis/`. The vendored source catalogues, their checksums and their provenance record are under `paper/g/data/`. Both source catalogues are public at VizieR (J/MNRAS/479/2834 and J/MNRAS/472/3706).

A Caveat checklist

The following are stated as limitations of this analysis and are not corrected for.

1. **Eddington bias.** Fluxes near the detection limit are biased high, so a background count computed above a limit derived from those same fluxes is a lower bound on contamination. No deboosting is possible without the detection-significance distribution, which the vendored table does not carry.
2. **Completeness varies with off-axis angle** across the ACIS-I field. A single flux limit is applied over the whole area, which overestimates the effective depth at large radius and therefore overestimates the background count there while underestimating the true source count.
3. **No variability analysis** is performed. Epoch-resolved data are not staged and the single tabulated 2012/2000 flux ratio does not support a Poisson-honest test. Explicitly out of scope, and the largest departure from what a full identification programme requires.
4. **The residual is an input to future identification work.** It is not a detection and it is not evidence for any hypothesis.
5. **Chance-coincidence rates** are reported per probe by three independent estimators and are never collapsed into one number.
6. **Membership** is reported as a probability under stated external normalisations, never as a binary label, and every published membership statement is carried alongside rather than overwritten.
7. **The tier-2 positional uncertainty rule** is a convention of this analysis (one unit in the last quoted digit), not a value stated by the source paper.
8. **The escape-velocity population relation** was fitted on FAST-visible clusters and ω Cen is not FAST-visible, so its prediction here is an extrapolation.

9. **The background-relation parameter uncertainties** are published asymmetrically and are used here only as symmetric bracket widths; no asymmetric error is quoted.
10. **The internal membership fit is degenerate** and is reported as a diagnostic only; membership uses external normalisations throughout.
11. **The control is an optimistic bound.** The control catalogue was itself astrometrically registered against its MSP truth set, so control recovery bounds match performance from above rather than testing astrometry independently.
12. **Six pulsars are absent** from the cross-match for want of a published position, so every match count is conditional on a 13-probe subset of the 19-pulsar timing set.

References

- Anderson, J. and van der Marel, R. P. (2010). New limits on an intermediate-mass black hole in Omega Centauri. I. Hubble Space Telescope photometry and proper motions. *The Astrophysical Journal*, 710:1032–1062.
- Bhattacharya, S., Heinke, C. O., Chugunov, A. I., Freire, P. C. C., Ridolfi, A., and Bogdanov, S. (2017). Chandra studies of the globular cluster 47 Tucanae: A deeper x-ray source catalogue, five new x-ray counterparts to millisecond radio pulsars, and new constraints to r-mode instability window. *Monthly Notices of the Royal Astronomical Society*, 472:3706–3721.
- Chen, W., Freire, P. C. C., Ridolfi, A., Barr, E. D., Stappers, B., Kramer, M., et al. (2023). MeerKAT discovery of 13 new pulsars in Omega Centauri. *Monthly Notices of the Royal Astronomical Society*, 520:3847–3856.
- Colomo Bernadich, M., Dai, S., Abbate, F., Kerr, M., Bachetti, M., Buchner, S., Johnston, S., Burgay, M., Possenti, A., Ridolfi, A., Freire, P. C. C., Kramer, M., Stappers, B., et al. (2026). A joint MeerKAT and Parkes view of Omega Centauri: New TRAPUM searches and pulsar timing. *arXiv e-prints*, page arXiv:2603.21845. Under review.
- Henleywillis, S., Cool, A. M., Haggard, D., Heinke, C., Callanan, P., and Zhao, Y. (2018). A deep x-ray survey of the globular cluster Omega Centauri. *Monthly Notices of the Royal Astronomical Society*, 479:2834–2852.
- Moretti, A., Campana, S., Lazzati, D., and Tagliaferri, G. (2003). The resolved fraction of the cosmic x-ray background. *The Astrophysical Journal*, 588:696–703.
- Swanson, T. (2026). Engineered intermediate-mass black hole systems: Infrastructure constraints, observable residue, and a multi-messenger adjudication framework. Preprint; omegacentauri.me.
- Yin, D., Zhang, L., Qian, L., Pan, Z., Li, D., et al. (2024). An empirical-Bayesian estimate of the pulsar populations of globular clusters. *The Astrophysical Journal Letters*, 969:L7.
- Zhao, J. and Heinke, C. O. (2023). A Chandra x-ray study of millisecond pulsars in the globular cluster Omega Centauri: A correlation between spider pulsar companion mass and x-ray luminosity. *Monthly Notices of the Royal Astronomical Society*.