

Photometry of Magnetars in Spitzer Telescope Infrared Images

by

Darkhan Nurzhakyp

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Master of Science in Physics

at the

NAZARBAYEV UNIVERSITY

June 2026

© Nazarbayev University 2026. All rights reserved.

Author
Department of Physics
May 5, 2026

Certified by
Ernazar Abdikamalov
Associate Professor
Thesis Supervisor

Accepted by

Photometry of Magnetars in Spitzer Telescope Infrared Images

by

Darkhan Nurzhakyp

Submitted to the Department of Physics
on May 5, 2026, in partial fulfillment of the
requirements for the degree of
Master of Science in Physics

Abstract

We present a systematic mid-infrared photometric survey of 27 magnetars using archival Spitzer Space Telescope Infrared Array Camera (IRAC) mosaics. A custom Python pipeline combines forced aperture photometry at known X-ray positions with an ensemble PRF-injection method, in which synthetic point sources are injected at a 3×3 grid of positions in the real mosaic to empirically characterise the local confusion noise; the 5σ detection limit follows directly as five times the ensemble standard deviation. The pipeline is validated against the two magnetars with published Spitzer detections: the recovered $4.5\ \mu\text{m}$ flux of 1E 2259+586 ($4.86 \pm 0.23\ \mu\text{Jy}$) agrees with the literature value at the 1.4σ level, while the systematically lower fluxes for 4U 0142+61 are attributed to annular background over-subtraction in its heavily confused Galactic-plane field. Applied to the full sample, the 5σ limits span more than three orders of magnitude—from $\approx 10\ \mu\text{Jy}$ (1E 2259+586, $4.5\ \mu\text{m}$) to $\approx 34,700\ \mu\text{Jy}$ (SGR 1806–20)—confirming that confusion noise, not photon statistics, is the dominant sensitivity barrier for Spitzer in the Galactic plane. The spectral index for 4U 0142+61, the best-constrained detection, is $\alpha_{24} = -0.49 \pm 0.12$ between 4.5 and $8.0\ \mu\text{m}$, consistent with the wider-baseline JWST power-law slope $\alpha = -0.96$ measured by Hare et al. [2024]; the IRAC two-point measurement alone does not arbitrate between a non-thermal magnetospheric origin and a multi-temperature disk whose emission peak falls within the IRAC bands. Six targets with $4.5\ \mu\text{m}$ limits below $200\ \mu\text{Jy}$ (1E 2259+586, 4U 0142+61, SGR 1935+2154, SGR 2013+34, CXO J0100-7211, and SGR 0526-66) are identified as the highest-priority candidates for JWST MIRI follow-up.

Thesis Supervisor: Ernazar Abdikamalov

Title: Associate Professor

Acknowledgments

My deepest gratitude to Dr. Bruce Grossan, who taught me, mentored me, and shaped me as both a person and a scientist over many years. He was kind, supportive, and always motivating, and I carry forward everything I learned from him. May his name live on and his soul rest in peace.

I am grateful to my thesis supervisor Professor Ernazar Abdikamalov for his guidance and steady support throughout this work. I thank Professor Dana Alina, who first introduced me to astrophysics and has supported me from the very beginning of my path in the field.

I am grateful to everyone at the Ghalam Team for all their friendship and support.

I'm thankful to NASA/IPAC Infrared Science Archive for providing mosaic images and answering our questions.

Most importantly, I thank my family. Mom – you have supported and encouraged me throughout my life. Whatever I have achieved is because of your dedication and faith in me. Thank you for always being there.

Finally, I'm grateful to my close friends and my significant other who shared this long path with me.

Contents

1	Introduction	13
1.1	Magnetars and Their Energy Spectrum	13
1.2	Infrared Emission from Magnetars: Scientific Motivation	15
1.3	The Spitzer Space Telescope and IRAC	16
1.4	Photometry of Faint Point Sources	18
1.5	Thesis Overview	19
2	Methods	21
2.1	Photometry Procedure	21
2.2	Spitzer Data	26
2.2.1	Spitzer Infrared Data	26
2.2.2	Spitzer PRF	27
2.3	The Ensemble Photometry Procedure	29
2.4	5σ Limit Procedure	30
3	Results	35
3.1	Pipeline Validation: 4U 0142+61 and 1E 2259+586	35
3.2	Sensitivity Limits	37
4	Discussion	41
4.1	Determining the Spectral Slope	41
4.2	Comparison with Literature Values and Sources of Discrepancy	46
4.3	Constraints on the Undetected Magnetar Population	48

4.4	Limitations and Future Directions	49
5	Conclusion	51

List of Figures

- 2-1 Aperture photometry: pixel classification on a synthetic 17×17 pixel image. Blue shading marks aperture pixels (source flux), gold shading marks sky-annulus pixels (background estimate), and grey marks excluded pixels. Pixel values in counts are printed inside each cell. The equation $I = S - N_{\text{pix}} \cdot B$ is shown upper-left. Red star: source centre; arrows indicate the three defining radii r_{ap} , r_{in} , r_{out} 23
- 2-2 Spitzer IRAC $4.5 \mu\text{m}$ cutouts illustrating the contrast between easy and hard photometric environments. *(a) Easy field* — 4U 0142+61: the infrared counterpart is detected at $\text{SNR} = 36$ ($F_{4.5} = 23.3 \mu\text{Jy}$). The background is relatively smooth and uniform, making a single forced-aperture measurement reliable. *(b) Hard field* — SGR 1806–20: the field is dominated by bright, spatially structured Galactic-plane dust emission and an unresolved stellar population, with no point source visible at the magnetar position. The 5σ ensemble limit is $\approx 34,300 \mu\text{Jy}$, more than three orders of magnitude above the 4U 0142+61 detection — entirely because of background confusion. Red circle: photometric aperture ($r = 2.0$ native IRAC pixels $\approx 2.44''$). Yellow dashed circles: inner and outer sky annulus. Cyan crosshair: forced-photometry position from X-ray astrometry. Both cutouts are $60'' \times 60''$ and share the same angular scale. 25

2-3	Cryogenic IRAC Point Response Functions for all four channels at $100\times$ oversampling. <i>Left column</i> : 2D PRF on a logarithmic colour scale (inferno palette), showing the bright core, first Airy ring, and diffraction spikes out to $\pm 10''$. The cyan circle marks the photometric aperture ($r_{\text{ap}} = 2.0$ native IRAC pixels). <i>Right column</i> : azimuthally averaged radial profile (log scale, normalized to peak), with vertical lines indicating the aperture radius (solid red), inner annulus r_{in} (dashed gold), and outer annulus r_{out} (dotted gold). The PSF broadens systematically from channel 1 ($3.6\,\mu\text{m}$, $\text{FWHM} \approx 1.7''$) to channel 4 ($8.0\,\mu\text{m}$, $\text{FWHM} \approx 2.0''$), and the fraction of flux falling within r_{ap} decreases accordingly, motivating the channel-dependent aperture corrections listed in Table 2.1.	28
2-4	Spitzer IRAC ch1 ($3.6\,\mu\text{m}$) cutout of 4U 0142+61 showing the 3×3 ensemble injection grid. Each cyan aperture circle marks one synthetic PRF injection position; the red cross marks the catalogued X-ray position. Measuring the scatter of recovered fluxes across all 9 positions quantifies the local background confusion noise.	31
4-1	Infrared spectral energy distribution of 4U 0142+61. Red downward triangles mark our Spitzer IRAC ensemble 5σ upper limits (ch1, ch2, ch4). Blue squares show the Wang and Kaspi [2008] Spitzer detections; green diamonds show Hare et al. [2024] JWST NIRCam measurements. The purple dashed line is the Hare et al. [2024] best-fit power law ($\alpha = -0.96$ over $1.4\text{--}11\,\mu\text{m}$); the orange dash-dotted line shows a steady-state $\alpha = +1/3$ disk slope for reference. Both our IRAC measurements and the literature Spitzer values are consistent with the wide-baseline JWST power law; the IRAC two-point slope alone does not arbitrate between a non-thermal magnetospheric origin and a multi-temperature disk whose emission peak lies near $4\text{--}8\,\mu\text{m}$	43

List of Tables

2.1	IRAC photometric parameters.	26
3.1	4U 0142+61 one-shot photometry vs. literature. Channel 3 is omitted because no usable IRAC 5.8 μm mosaic covering the source position is available in the Spitzer Heritage Archive at the time of this analysis. .	36
3.2	1E 2259+586 one-shot photometry vs. literature	37
3.3	5σ detection limits (μJy) per magnetar per IRAC channel	38
4.1	Two-point spectral indices α for 4U 0142+61 ($f_\nu \propto \nu^\alpha$)	44

Chapter 1

Introduction

The broad outlines of stellar evolution are well understood: stars fuse hydrogen on the main sequence, exhaust their nuclear fuel, and complete their lives as white dwarfs, neutron stars, or black holes depending on their initial mass. Yet the extreme objects produced at the high-mass end of this sequence continue to challenge our understanding of dense matter and strong-field electrodynamics. Magnetars — a subclass of isolated neutron stars — are among the most energetically active compact objects known, and their infrared emission encodes unique information about the physical processes operating in their ultramagnetized environments. This thesis focuses on a fundamental observational problem: determining how faint a magnetar can be in the Spitzer Space Telescope infrared bands, and quantifying this limit rigorously across the known magnetar population.

1.1 Magnetars and Their Energy Spectrum

Magnetars are neutron stars powered with surface magnetic fields of order $B \sim 10^{14}$ – 10^{15} G, roughly two to three orders of magnitude stronger than the canonical $\sim 10^{12}$ G fields of ordinary rotation-powered pulsars [Duncan and Thompson, 1992, Thompson and Duncan, 1993]. The magnetar model was proposed by Duncan and Thompson [1992], who showed that a rapidly rotating proto-neutron star born with initial spin period $P_0 \lesssim 3$ ms could amplify its magnetic field to $\sim 10^{15}$ G through a convective

α - Ω dynamo operating during the first few seconds after core bounce. Unlike ordinary pulsars, whose luminosity is supplied by rotational kinetic energy, magnetar X-ray luminosities of 10^{33} – 10^{36} erg s $^{-1}$ greatly exceed their spin-down power, implying that the energy source is instead the slow ohmic dissipation and Hall drift of the intense internal magnetic field [Thompson and Duncan, 1995, 1996].

Observationally, magnetars appear as Anomalous X-ray Pulsars (AXPs) or Soft Gamma Repeaters (SGRs), both now understood to be manifestations of the same underlying class [Kaspi and Beloborodov, 2017, Mereghetti, 2008, Turolla et al., 2015, Enoto et al., 2019]. The Kaspi and Beloborodov [2017] review and the McGill Magnetar Catalog [Olausen and Kaspi, 2014] — the definitive observational census — document approximately 30 confirmed magnetars and candidates, each characterized by its spin period P , spin-down rate $\dot{P}$, and derived dipole field $B \approx 3.2 \times 10^{19} (P\dot{P})^{1/2}$ G. Several recent discoveries have extended the known population: Swift J1818.0–1607 is the fastest-rotating and youngest known magnetar [Esposito et al., 2020], and the detection of a luminous radio burst from SGR 1935+2154 [Bochenek et al., 2020, CHIME/FRB Collaboration et al., 2020] established the first direct link between magnetars and the cosmological fast radio burst population.

The broadband spectral energy distribution of magnetars spans many decades in frequency. In the X-ray band, quiescent spectra are characteristically described by a two-component model: a blackbody with temperature $kT \sim 0.3$ – 0.6 keV representing the thermally emitting neutron star surface plus a hard power-law tail ($\Gamma \sim 2$ – 4) from resonant cyclotron scattering in the twisted magnetosphere [Thompson et al., 2002, Rea et al., 2008]. At $E \gtrsim 10$ keV, a harder power-law component ($\Gamma \sim 1$) has been detected with INTEGRAL and NuSTAR in at least some magnetars (including 4U 0142+61), attributed to bulk Comptonization of thermal photons in the corona [den Hartog et al., 2008, Beloborodov and Xinyu, 2016]. In the radio band, transient emission has been detected from a handful of magnetars after outburst. The mid-infrared window, probed by Spitzer, is of special interest because it is sensitive to distinct physical components invisible in X-rays: a possible supernova fallback disk and the low-energy tail of magnetospheric non-thermal emission.

1.2 Infrared Emission from Magnetars: Scientific Motivation

The detection — or sensitive non-detection — of magnetar infrared counterparts addresses two competing theoretical frameworks for mid-infrared emission. The first is the **fallback disk model** [Chatterjee et al., 2000, Perna et al., 2000], in which supernova debris that failed to escape the gravitational well of the nascent neutron star forms a fossil accretion disk. Wang et al. [2006] first detected 4U 0142+61 with Spitzer IRAC at $4.5\ \mu\text{m}$ ($36.3 \pm 3.6\ \mu\text{Jy}$) and $8.0\ \mu\text{m}$ ($51.9 \pm 5.2\ \mu\text{Jy}$) and interpreted the rising mid-infrared spectral slope as a passive, X-ray-heated dust disk of mass ~ 10 Earth masses at temperature $\sim 920\ \text{K}$. Wang and Kaspi [2008] later refined these to $32.1 \pm 2.0\ \mu\text{Jy}$ and $59.8 \pm 8.5\ \mu\text{Jy}$ from four Spitzer epochs in February 2007. The second framework is the **magnetospheric model** [Thompson et al., 2002, Beloborodov and Thompson, 2007], where the hard power-law X-ray spectrum extends into the near-infrared and optical through resonant cyclotron scattering in a twisted, current-carrying magnetosphere. Ertan and Çalışkan [2006] developed actively irradiated disk models that could reproduce all optical and near-infrared observations of AXPs with a gaseous accretion disk rather than a passive dust disk [Ertan et al., 2009].

The physical distinction between these models has profound implications. A confirmed fallback disk would establish that the magnetar spin-down luminosity is partially fed by disk-star interaction rather than pure magnetic dipole radiation, directly constraining the post-supernova mass budget and neutron star spin evolution [Alpar, 2001, Benli and Ertan, 2016]. Correlated infrared/X-ray variability — as documented for 1E 2259+586 by Tam et al. [2004] after the 2002 outburst, where near-infrared and X-ray fluxes decayed in lockstep — favors magnetospheric coupling over a passive thermally reprocessing disk. The matter was further advanced by Hare et al. [2024], who obtained a JWST MIRI and NIRCам spectrum of 4U 0142+61 spanning $1.4\text{--}11\ \mu\text{m}$. Their best-fit single power law ($f_\nu \propto \nu^{-0.96}$, $A_V = 3.9 \pm 0.2$ mag) extends across nearly a decade in wavelength; reproducing it with a thermal-disk

component would require dust hot enough to dominate at $1.4\,\mu\text{m}$, deep in the Wien tail of any sublimation-temperature grain. Hare et al. therefore disfavour the simple passive-disk fit of Wang et al. [2006] and argue that the mid-infrared emission is dominated by a non-thermal magnetospheric component, although more elaborate disk-plus-power-law variants are not formally excluded.

Despite this progress, the mid-infrared properties of the vast majority of known magnetars remain unconstrained. Only two magnetars have confirmed Spitzer IRAC point-source detections: 4U 0142+61 [Wang et al., 2006, Wang and Kaspi, 2008] and 1E 2259+586 at $4.5\,\mu\text{m}$ with $6.3 \pm 1.0\,\mu\text{Jy}$ [Kaplan and Chakrabarty, 2009]. For 1E 1048.1–5937, Hulleman [2003] reported Spitzer 4.5 and $8.0\,\mu\text{m}$ upper limits excluding 4U 0142+61-like mid-infrared emission. SGR 1900+14 showed only an extended dust-free cavity ring at 16 and $24\,\mu\text{m}$ from the 1998 giant flare rather than a point source [Wachter et al., 2008]. Systematic mid-infrared upper limits across the full archival Spitzer dataset are therefore a critical contribution to the field, enabling population-level constraints on emission models that a single-source study cannot provide. The recent Chrimes et al. [2022] HST near-infrared survey identified six new magnetar counterpart candidates, further motivating comprehensive archival analysis at mid-infrared wavelengths.

1.3 The Spitzer Space Telescope and IRAC

The Spitzer Space Telescope was launched on 25 August 2003 into an Earth-trailing heliocentric orbit, placing it at a comfortable thermal distance from the Earth and minimizing the infrared background from our planet [Werner et al., 2004]. Its 85-cm primary mirror was cooled with a cryogenic liquid helium reservoir to $\sim 5.5\,\text{K}$, suppressing thermal self-emission and enabling unprecedented sensitivity at infrared wavelengths. The telescope carried three scientific instruments: the Infrared Array Camera (IRAC), the Infrared Spectrograph (IRS), and the Multiband Imaging Photometer for Spitzer (MIPS).

The Infrared Array Camera [Fazio et al., 2004] is the primary imaging instrument

relevant to this work. It employs a dichroic beamsplitter to enable simultaneous imaging in two channel pairs: channels 1 and 2 (3.6 and $4.5\ \mu\text{m}$, InSb detector arrays) and channels 3 and 4 (5.8 and $8.0\ \mu\text{m}$, Si:As IBC detector arrays). Each channel has a $5.2' \times 5.2'$ field of view with $\sim 1.22''$ native pixel scale and a 256×256 pixel array. The cryogenic mission delivered 5σ point-source sensitivities of approximately 6, 7, 36, and $54\ \mu\text{Jy}$ at channels 1–4 for 200-second integrations [Fazio et al., 2004, Hora et al., 2004]. Absolute photometric calibration was established by Reach et al. [2005] using A-type calibration stars near the north ecliptic pole, achieving **3% absolute accuracy** with long-term stability of 1.5% rms.

In May 2009, the cryogenic coolant was exhausted, ending the Cryogenic Mission after 5.7 years. The resulting temperature increase caused overwhelming dark current in the Si:As channels 3 ($5.8\ \mu\text{m}$) and 4 ($8.0\ \mu\text{m}$), rendering them non-functional. However, the InSb detectors in channels 1 ($3.6\ \mu\text{m}$) and 2 ($4.5\ \mu\text{m}$) retained their sensitivity at the elevated telescope temperature of $\sim 28.7\text{ K}$, enabling a **Warm Mission** that continued imaging operations until the mission was terminated on 30 January 2020 — a total mission duration of 6002 days. Hora et al. [2008] demonstrated that warm-mission photometric accuracy remained at ~ 1.2 – 1.3% in channels 1 and 2. All Spitzer data used in this work are retrieved from the **Spitzer Heritage Archive** at IRSA (Infrared Science Archive), which contains Basic Calibrated Data (BCD) and pipeline mosaics for the full mission [IRAC Instrument Team, 2021].

The IRAC mosaics of magnetar fields are characterized by two features that fundamentally affect photometric analysis. First, most known magnetars reside in the Galactic plane ($|b| \lesssim 5^\circ$), where IRAC images exhibit strong, spatially structured emission from interstellar dust heated by massive stars, molecular cloud surfaces (photodissociation regions), and unresolved stellar populations [Benjamin et al., 2003, Churchwell et al., 2009]. Kobulnicky et al. [2013] demonstrated that point-source completeness in Galactic-plane IRAC data depends primarily on this local background brightness rather than photon noise, with sensitivity losses of up to five magnitudes in bright photodissociation regions compared to low-background fields. Second, the IRAC PSF is critically sampled at the $\sim 1.22''$ native pixel scale, meaning that the

fraction of flux recorded in the central pixel depends sensitively on the sub-pixel position of a source — the pixel-phase effect — which must be accounted for through the Point Response Function (PRF) and aperture correction framework described in Chapter 2.

1.4 Photometry of Faint Point Sources

Aperture photometry measures the net flux of a source by summing pixel values within a circular aperture and subtracting a sky background estimate derived from a surrounding annulus. Standard implementations (e.g., DAOPHOT, SExtractor; Stetson 1987, Bertin and Arnouts 1996) treat the background as locally uniform and model the PSF from bright isolated stars in the same field. For the Spitzer pipeline, the recommended point-source extraction tool is MOPEX/APEX [Makovoz and Marleau, 2005], which performs PRF-fitting photometry, nonlinear matched filtering for faint source detection, and deblending for confused sources.

However, the standard MOPEX/APEX pipeline is not reliable for magnetar fields because it was optimized for sparse fields where isolated point sources can be identified by source-detection algorithms. In the bright, spatially varying backgrounds of Galactic-plane IRAC images, MOPEX source detection produces large numbers of spurious detections, systematic residuals from bright neighbors bias PRF fits for faint sources, and the formal APEX uncertainties do not reflect the true local noise [Kobulnicky et al., 2013]. Furthermore, magnetar positions are known precisely from X-ray astrometry (typically $\lesssim 1''$), making it unnecessary to rely on detection in the infrared image itself. Instead, **forced photometry** at the known X-ray position provides the optimal strategy: flux is measured at the specified coordinates regardless of whether a source is visible, yielding either a detection or an upper limit from the same measurement framework [Kashyap et al., 2010].

To derive accurate photometric uncertainties in complex backgrounds, this work employs an **ensemble photometry** approach in which synthetic sources (generated using the IRAC PRF) are injected at multiple positions near the target in the real

mosaic image. Aperture photometry on these injected sources measures the scatter induced by the actual background structure, providing a realistic local photometric noise estimate that subsumes confusion noise, background gradients, and PSF wing contamination from neighboring sources. The 5σ detection limit is then derived as five times this ensemble noise, following the statistical framework of Kraft et al. [1991] — with related developments for small-number Poisson statistics by Gehrels [1986] and Feldman and Cousins [1998] — and the practical methodology of Marengo et al. [2009] for IRAC faint-source limits.

1.5 Thesis Overview

The objective of this thesis is to estimate 5σ detection limits for magnetars observed by the Spitzer Space Telescope, using a custom ensemble photometry pipeline applied to archival IRAC mosaics. Chapter 2 (Methods) describes the full photometric pipeline: forced aperture photometry, PRF construction and injection, ensemble noise estimation, and 5σ limit derivation. Chapter 3 (Results) presents the pipeline validation on two detected sources (4U 0142+61 and 1E 2259+586) and the derived sensitivity limits for 27 magnetar targets across all four IRAC channels. Chapter 4 (Discussion) places these limits in the context of theoretical emission models and derives spectral slope constraints. Chapter 5 summarizes the main conclusions.

Chapter 2

Methods

This chapter describes the photometric pipeline developed to measure or constrain the infrared flux of magnetars in archival Spitzer IRAC mosaics. Because most magnetars are undetected in the mid-infrared and lie in regions of complex, structured Galactic-plane emission, standard catalog-oriented photometry tools are inadequate. The pipeline instead combines forced aperture photometry at the known X-ray position, a PRF-injection simulation to characterize the local noise environment, and an ensemble statistical method to derive robust 5σ detection limits.

2.1 Photometry Procedure

Aperture photometry is the foundational technique used throughout this work. The precise implementation varies with observing mode: ground-based observations must contend with atmospheric seeing, which broadens the PSF and motivates large apertures, whereas space-based data with a stable, diffraction-limited PSF permit smaller apertures that minimise sky-noise contributions. In either case, the aperture radius and background-estimation algorithm must be chosen to maximise signal-to-noise while minimising systematic errors from boundary pixels and background gradients.

Aperture photometry estimates the net source flux I by summing pixel values inside a circular aperture of radius r_{ap} and subtracting the local sky background B estimated from an annulus at $[r_{\text{in}}, r_{\text{out}}]$:

$$I = S - N_{\text{pix}} \cdot B \quad (2.1)$$

Equation (2.1) states that the net source flux I is the total aperture sum S minus the scaled background $N_{\text{pix}} \cdot B$. The sky level B (counts per pixel) is estimated as the median of pixels in the surrounding annulus, making the result robust against low-level background offsets. As the aperture radius grows, $N_{\text{pix}} \propto r^2$ pixels are enclosed, so the background noise contribution $\sigma = \sqrt{N_{\text{pix}}} \sigma_{\text{sky}}$ grows as r — the fundamental trade-off in choosing r_{ap} .

In Figure 2-1, the circular aperture and annulus cut through the pixels. Boundary pixels are neither fully inside nor fully outside the radius. In practice, it must be decided how to treat the boundary pixels. For large apertures, the contribution of edge pixels is very small in contrast to the total number of enclosed pixels. In this case, ignoring the sub-pixel precision introduces negligible fractional error. However, for small apertures, it can introduce a significant bias in the measured flux and background estimation. From here, three methods can be distinguished: *exact*, *center*, and *subpixel*.

The *exact* method analytically computes the geometric overlap, giving the most accurate but computationally expensive result. The *center* method is a binary decision: if the pixel's center lies within the radius, the whole pixel is included. It is simple and efficient, yet for small apertures it introduces a fractional error. The *subpixel* method divides a pixel into N subpixels and uses the fraction inside the radius as a weight.

For the *exact* photometry method, Eq. (2.2) is substituted into Eq. (2.1). The total source flux S is calculated using a weighted sum of the pixel values (P_i) within the aperture, where w_i is the corresponding weight value for each pixel i :

$$S = \sum_i w_i P_i \quad (2.2)$$

The photometric signal-to-noise ratio (SNR) in the background-dominated regime is:




Figure 2-1: Aperture photometry: pixel classification on a synthetic 17×17 pixel image. Blue shading marks aperture pixels (source flux), gold shading marks sky-annulus pixels (background estimate), and grey marks excluded pixels. Pixel values in counts are printed inside each cell. The equation $I = S - N_{\text{pix}} \cdot B$ is shown upper-left. Red star: source centre; arrows indicate the three defining radii r_{ap} , r_{in} , r_{out} .

$$\text{SNR} = \frac{S_{\text{net}}}{\sigma}, \quad \sigma = \sqrt{N_{\text{pix}}} \sigma_{\text{sky}} \quad (2.3)$$

where N_{pix} is the number of pixels enclosed by the aperture and σ_{sky} is the pixel-to-pixel standard deviation of the background estimated in the annulus. This formulation differs from the full CCD equation

$$\text{SNR} = \frac{F_{\text{src}} t}{\sqrt{F_{\text{src}} t + N_{\text{pix}} \left(\sigma_{\text{sky}}^2 + F_{\text{sky}} t + \frac{R^2}{g^2} \right)}} \quad (2.4)$$

in that the read noise R and dark current contributions are negligible for Spitzer IRAC at the relevant flux levels, so σ_{sky} captures the dominant source of noise.

Motivation: Forced Photometry. Because most magnetars produce no visible point source in IRAC mosaics, photometry is performed at the known X-ray position regardless of whether a source can be seen by eye. This forced-photometry strategy yields either a flux measurement (for detected sources such as 4U 0142+61 and 1E 2259+586) or an upper limit (for undetected targets). It also ensures a uniform analysis framework across all targets. The contrast between an "easy" (low-background) and a "hard" (Galactic-plane confusion-limited) photometric environment is illustrated in Figure 2-2.

Motivation: Why Not MOPEX/APEX? The standard Spitzer pipeline tool MOPEX/APEX performs PRF-fitting photometry optimized for sparse fields where isolated point sources can be identified by a source-detection algorithm. In IRAC observations of Galactic-plane magnetar fields the background is dominated by extended dust emission and unresolved stellar crowding. Under these conditions APEX source detection is unreliable, systematic residuals from bright neighbors bias the fit, and the formal APEX uncertainties do not reflect the true local noise. Our approach bypasses source detection entirely and measures the local photometric noise empirically via the ensemble method described in Section 2.3.




Figure 2-2: Spitzer IRAC 4.5 μm cutouts illustrating the contrast between easy and hard photometric environments. (a) *Easy field* — 4U 0142+61: the infrared counterpart is detected at SNR = 36 ($F_{4.5} = 23.3 \mu\text{Jy}$). The background is relatively smooth and uniform, making a single forced-aperture measurement reliable. (b) *Hard field* — SGR 1806–20: the field is dominated by bright, spatially structured Galactic-plane dust emission and an unresolved stellar population, with no point source visible at the magnetar position. The 5σ ensemble limit is $\approx 34,300 \mu\text{Jy}$, more than three orders of magnitude above the 4U 0142+61 detection — entirely because of background confusion. Red circle: photometric aperture ($r = 2.0$ native IRAC pixels $\approx 2.44''$). Yellow dashed circles: inner and outer sky annulus. Cyan crosshair: forced-photometry position from X-ray astrometry. Both cutouts are $60'' \times 60''$ and share the same angular scale.

2.2 Spitzer Data

2.2.1 Spitzer Infrared Data

Spitzer IRAC mosaics used in this work are pipeline-processed Basic Calibrated Data (BCD) coadds retrieved from the Spitzer Heritage Archive [IRAC Instrument Team, 2021]. They are calibrated in surface brightness units of MJy sr^{-1} and drizzled to a common mosaic pixel scale of $0.6 \text{ arcsec pixel}^{-1}$.

Spitzer infrared data have a known issue whereby an observation can record a lower flux in the central pixel of a point source. This happens because the PSF of the telescope is undersampled: the full width at half maximum (FWHM) of the IRAC PSF is $\approx 1.7\text{--}2.0 \text{ arcsec}$ at $3.6\text{--}8.0 \mu\text{m}$, comparable to the native detector pixel size of $\approx 1.22 \text{ arcsec pixel}^{-1}$. Because the PSF is critically or sub-critically sampled, the fraction of flux recorded in the central pixel depends on the sub-pixel position of the source centroid within the pixel (pixel-phase effect). This sensitivity variation is folded into the Point Response Function (PRF) described in Section 2.2.2.

The unit conversion from mosaic pixel values (MJy sr^{-1}) to microjansky is:

$$F [\mu\text{Jy}] = C \times \text{APCOR}_{ch} \times \mathcal{M} \quad (2.5)$$

where $\mathcal{M} = \left(\frac{0.6\pi}{648000}\right)^2 \times 10^{12} = 8.47 \mu\text{Jy} (\text{MJy sr}^{-1} \text{ pixel})^{-1}$ is the solid angle of one 0.6-arcsec mosaic pixel expressed in the appropriate units, and APCOR_{ch} is the aperture correction for each IRAC channel (Table 2.1).

Table 2.1: IRAC photometric parameters.

Channel	Wavelength (μm)	Native pixel scale (arcsec/pix)	Aperture correction
1	3.6	1.221	1.125
2	4.5	1.213	1.120
3	5.8	1.222	1.135
4	8.0	1.220	1.221

Aperture corrections account for the fraction of total PSF flux that falls outside the chosen aperture radius ($r_{\text{ap}} = 2.0 \text{ native pixels} \approx 2.44 \text{ arcsec}$). Values are taken

from Table 4.7 of the IRAC Instrument Handbook (Warm Mission; IRAC Instrument Team 2021).

2.2.2 Spitzer PRF

The Point Response Function (PRF) is distinct from the telescope PSF. While the PSF describes the optical response of the 85-cm Spitzer primary mirror and IRAC optics, the PRF additionally folds in the intra-pixel sensitivity variation (pixel-phase effect) and the dithering and drizzling pattern used to construct the coadded mosaic. As a result, the PRF is the correct model to inject and fit when working with IRAC mosaics.

The cryogenic IRAC core PRFs provided by the IRAC team (via the APEX package) are delivered at $100\times$ oversampling relative to the native IRAC detector pixel scale. Core PRFs are generally not recommended for aperture photometry in crowded fields or regions with a strongly varying background. However, because our pipeline characterises the local noise empirically through the ensemble injection procedure rather than relying on PRF-fitting residuals, the core PRF is adequate for synthetic source injection. We use the cryogenic IRAC core PRFs at $100\times$ oversampling (Figure 2-3).

PRF processing pipeline. The raw oversampled PRF is processed in four steps before injection into a mosaic:

1. *Trim*: 50 pixels are removed from each border of the raw PRF array to eliminate the zero-padded edge region.
2. *Sub-pixel shift*: The PRF is shifted by the fractional pixel offset $(\delta x, \delta y) = (x - [x], y - [y])$ of the target coordinate within its mosaic pixel using array rolling. In the $100\times$ oversampled grid this corresponds to a shift of $(\delta x \times 100, \delta y \times 100)$ oversampled pixels, accurately capturing the pixel-phase effect for the specific target position.
3. *Congrid (rebin)*: The shifted PRF is downsampled from the $100\times$ oversampled




Figure 2-3: Cryogenic IRAC Point Response Functions for all four channels at 100× oversampling. *Left column:* 2D PRF on a logarithmic colour scale (inferno palette), showing the bright core, first Airy ring, and diffraction spikes out to $\pm 10''$. The cyan circle marks the photometric aperture ($r_{\text{ap}} = 2.0$ native IRAC pixels). *Right column:* azimuthally averaged radial profile (log scale, normalized to peak), with vertical lines indicating the aperture radius (solid red), inner annulus r_{in} (dashed gold), and outer annulus r_{out} (dotted gold). The PSF broadens systematically from channel 1 (3.6 μm , FWHM $\approx 1.7''$) to channel 4 (8.0 μm , FWHM $\approx 2.0''$), and the fraction of flux falling within r_{ap} decreases accordingly, motivating the channel-dependent aperture corrections listed in Table 2.1.

grid to the native mosaic pixel scale using linear interpolation. The resampling ratio is:

$$\eta = \frac{s_{\text{ch}}}{s_{\text{mosaic}} \times 100} \quad (2.6)$$

where s_{ch} is the native IRAC pixel scale for channel ch (Table 2.1) and $s_{\text{mosaic}} = 0.6 \text{ arcsec pixel}^{-1}$ is the drizzled mosaic pixel scale. After resampling the PRF is renormalized to unit sum, so it represents a source with total flux equal to the injection scale.

4. *Scale and inject*: The normalized PRF is multiplied by the injection amplitude F_{inj} (in MJy sr^{-1}) and added to the mosaic at the target pixel coordinates using `place_PSF`.

2.3 The Ensemble Photometry Procedure

A single aperture measurement at the magnetar position yields a flux value but does not reliably quantify the photometric uncertainty in a spatially structured background. Background spatial gradients, unresolved neighboring sources, and detector residuals all introduce position-dependent noise that is not captured by the formal Poisson or background-scatter errors. The ensemble photometry method addresses this by empirically sampling the photometric noise distribution across many positions in the actual mosaic background near the target.

Grid placement. For each injection flux level F_{inj} (set by `FACTORS` $\times$ the estimated noise), the PRF is injected at $N = n_{\text{grid}}^2$ positions arranged on a regular $n_{\text{grid}} \times n_{\text{grid}}$ grid centered on the target. The default configuration uses $n_{\text{grid}} = 3$ (a $3 \times 3 = 9$ -position grid) with spacing $\Delta = 2.828$ mosaic pixels between adjacent positions; an example overlay is shown in Figure 2-4. At each grid position (x_i, y_i) the same procedure is applied: inject the PRF, run aperture photometry at the injection center, and record the recovered flux F_i^{meas} .

Aperture photometry parameters. Each measurement uses:

- Aperture radius: $r_{\text{ap}} = 2.0$ native IRAC pixels ≈ 2.44 arcsec.
- Background annulus: $r_{\text{in}} = 2.2$ native pixels, $r_{\text{out}} = 3.0$ native pixels.
- Background level: $B = \text{median}(\text{pixels in annulus})$.
- Net counts: $S_{\text{net}} = \sum_{\text{aperture}} P_i - N_{\text{pix}} \cdot B$.
- Photometric error in counts: $\sigma_{\text{phot}} = \sqrt{N_{\text{pix}}} \cdot \sigma_{\text{sky}}$.
- Flux: $F = S_{\text{net}} \cdot \text{APCOR}_{\text{ch}} \cdot \mathcal{M} [\mu\text{Jy}]$.

Ensemble standard deviation. The scatter of recovered fluxes at all N grid positions for a given injection level F_{inj} defines the ensemble standard deviation:

$$\sigma_{\text{ens}}(F_{\text{inj}}) = \text{std}\left(\{F_i^{\text{meas}} - F_{\text{inj}}\}_{i=1}^N\right) \quad (2.7)$$

This quantity measures how precisely a source of flux F_{inj} can be recovered in the real local background. Because it is derived from the actual image — including extended emission, crowding, and PSF wings from nearby sources — σ_{ens} is a more realistic noise estimate than any analytic formula applied to a single position.

The pipeline repeats this procedure at two injection flux levels, $f \in \{3.8, 7.0\} \times \sigma_{\text{est}}$, bracketing the anticipated 5σ detection threshold from below and above. If the background noise dominates (as expected for magnetar-faint flux levels), σ_{ens} should be approximately independent of F_{inj} , confirming the background-dominated regime.

2.4 5σ Limit Procedure

The 5σ detection limit $F_{5\sigma}$ is defined as the flux at which a point source would be recovered with signal-to-noise ratio $\text{SNR} = 5$ in the measured local background:

$$F_{5\sigma} = 5 \bar{\sigma}_{\text{ens}} \quad [\mu\text{Jy}] \quad (2.8)$$

where $\bar{\sigma}_{\text{ens}}$ is the mean ensemble standard deviation averaged over all injection flux levels:

4U 0142+61 — IRAC ch1 ($3.6\ \mu\text{m}$) · Ensemble 3×3 injection grid
Each panel: 28×28 mosaic-pixel cutout centred on one injection position




Figure 2-4: Spitzer IRAC ch1 ($3.6\ \mu\text{m}$) cutout of 4U 0142+61 showing the 3×3 ensemble injection grid. Each cyan aperture circle marks one synthetic PRF injection position; the red cross marks the catalogued X-ray position. Measuring the scatter of recovered fluxes across all 9 positions quantifies the local background confusion noise.

$$\bar{\sigma}_{\text{ens}} = \frac{1}{N_F} \sum_{k=1}^{N_F} \sigma_{\text{ens}}(F_{\text{inj},k}) \quad (2.9)$$

Derivation from ensemble data. The procedure is:

1. For each magnetar, IRAC channel, and injection flux level, compute $\sigma_{\text{ens}}(F_{\text{inj}})$ as the sample standard deviation (ddof=1) of the $N = 9$ recovered fluxes in the ensemble grid window.
2. Verify linearity: if the background dominates, σ_{ens} is approximately constant across the two FACTORS levels. A variation $> 15\%$ between levels triggers a warning and may indicate a partially source-dominated regime.
3. Average σ_{ens} over the two injection levels to obtain $\bar{\sigma}_{\text{ens}}$.
4. Compute $F_{5\sigma} = 5 \times \bar{\sigma}_{\text{ens}}$.

Background-dominated justification. For a source at the magnetar flux level ($F \lesssim 10 \mu\text{Jy}$), the Poisson noise from the source photons ($\propto \sqrt{F_{\text{src}}}$) is negligible compared to σ_{sky} in the IRAC background. The SNR reduces to:

$$\text{SNR} \approx \frac{F_{\text{src}}}{\bar{\sigma}_{\text{ens}}} \quad (2.10)$$

so $F_{5\sigma} = 5 \times \bar{\sigma}_{\text{ens}}$ is the appropriate threshold.

Validation strategy. The pipeline’s primary output is the 5σ detection limit $F_{5\sigma}$. Validating it requires two independent checks: a *sensitivity check*, that the ensemble noise correctly represents the field’s actual detection threshold; and a *flux-scale check*, that the one-shot forced photometry recovers the correct absolute flux when a real source is present. We carry out both on the two magnetars with published Spitzer detections (4U 0142+61: $32.1 \pm 2.0 \mu\text{Jy}$ at $4.5 \mu\text{m}$ and $59.8 \pm 8.5 \mu\text{Jy}$ at $8.0 \mu\text{m}$, Wang and Kaspi 2008; 1E 2259+586: $6.3 \pm 1.0 \mu\text{Jy}$ at $4.5 \mu\text{m}$, Kaplan and Chakrabarty 2009).

The **sensitivity check** compares $F_{5\sigma}$ derived at the source position to the published literature flux F_{lit} : $F_{\text{lit}} \gtrsim F_{5\sigma}$ confirms the pipeline could detect a real source at the

published level, $F_{\text{lit}} \approx F_{5\sigma}$ identifies a borderline-detectable case, and $F_{\text{lit}} \ll F_{5\sigma}$ would reveal a sensitivity overestimate. This is the directly relevant validation for the upper limits reported in Chapter 3.

The **flux-scale check** runs the pipeline in one-shot mode (no PRF injection; forced aperture photometry at the source position) and compares the recovered flux F_{rec} to F_{lit} via

$$\Delta m = m_{\text{recovered}} - m_{\text{lit}}. \quad (2.11)$$

In a clean low-background field $|\Delta m|$ should approach the IRAC absolute calibration accuracy of ~ 0.03 mag (3%; Reach et al. 2005). In confused fields, annulus-induced background bias may produce a larger $|\Delta m|$; in that case Δm bounds the absolute-scale offset as a known systematic but does not invalidate the ensemble noise estimate, which is calibrated against the empirical local scatter rather than against an absolute flux. Both checks are reported for the two validation targets in Section 3.1.

Chapter 3

Results

3.1 Pipeline Validation: 4U 0142+61 and 1E 2259+586

The two magnetars with published Spitzer IRAC detections — 4U 0142+61 [Wang et al., 2006, Wang and Kaspi, 2008] and 1E 2259+586 [Kaplan and Chakrabarty, 2009] — are used to validate the pipeline along the two axes introduced in Section 2.4: **sensitivity** (does the ensemble $F_{5\sigma}$ correctly represent the field’s actual detection threshold?) and **flux scale** (does one-shot forced photometry recover the literature flux?). The sensitivity check is the directly relevant validation for the upper limits tabulated in Section 3.2; the flux-scale check provides a complementary diagnostic that may fail in heavily confused fields without invalidating the limits.

4U 0142+61.

Sensitivity check. The ensemble 5σ limits at the source position are $F_{5\sigma} = 25.7, 33.7, 43.7 \mu\text{Jy}$ at channels 1, 2, and 4. The Wang and Kaspi [2008] fluxes ($32.1 \pm 2.0 \mu\text{Jy}$ at $4.5 \mu\text{m}$, $59.8 \pm 8.5 \mu\text{Jy}$ at $8.0 \mu\text{m}$) lie at $\approx 0.95 F_{5\sigma}$ at ch2 (a borderline-detectable case in the ensemble metric, with F_{lit} just below threshold) and $\approx 1.4 F_{5\sigma}$ at ch4 (well above threshold, $\sim 6.8\sigma$ in ensemble units). Both literature detections therefore sit in the expected location relative to the pipeline’s noise floor for real sources near the limit, validating the ensemble sensitivity estimate at this position.

Flux-scale check. Forced aperture photometry at the known X-ray position yields

significant one-shot detections at three IRAC channels. Table 3.1 gives the measured fluxes alongside the published values from Wang and Kaspi [2008].

Table 3.1: 4U 0142+61 one-shot photometry vs. literature. Channel 3 is omitted because no usable IRAC 5.8 μm mosaic covering the source position is available in the Spitzer Heritage Archive at the time of this analysis.

Channel	Wavelength (μm)	This work (μJy)	SNR	Literature (μJy)	Reference
1	3.6	10.15 ± 1.71	5.9	—	—
2	4.5	23.30 ± 0.65	36.1	32.1 ± 2.0	Wang and Kaspi [2008]
3	5.8	—	—	—	—
4	8.0	30.93 ± 1.93	16.0	59.8 ± 8.5	Wang and Kaspi [2008]

The source is detected at high significance at 3.6, 4.5, and 8.0 μm , confirming the infrared counterpart. Our 4.5 μm flux of $23.30 \pm 0.65 \mu\text{Jy}$ is $\approx 27\%$ below the Wang and Kaspi [2008] value of $32.1 \pm 2.0 \mu\text{Jy}$, and our 8.0 μm flux is $\approx 48\%$ below the published $59.8 \pm 8.5 \mu\text{Jy}$. Both discrepancies exceed the 3% IRAC absolute calibration accuracy and are likely attributable to differences in background estimation strategy and aperture definition: Wang et al. [2006] and Wang and Kaspi [2008] used point-spread-function fitting with MOPEX on individual BCD frames, while our method applies forced aperture photometry with median annulus subtraction on the coadded mosaic in a region of complex diffuse emission. Variability between epochs may also contribute, as 4U 0142+61 has shown mid-infrared variability in JWST data [Hare et al., 2024]. The flux-scale offset is bounded (~ 0.3 – 0.7 mag) and applies only to the absolute one-shot flux at this position; it does not propagate into the ensemble $F_{5\sigma}$, which is calibrated against the empirical local scatter and which the sensitivity check above shows to be self-consistent with the published detections.

1E 2259+586.

Sensitivity check. The ensemble 5σ limits at this position are 13.1, 10.8, 54.6, 37.9 μJy at channels 1–4. The Kaplan and Chakrabarty [2009] ch2 detection at $6.3 \pm 1.0 \mu\text{Jy}$ lies *below* our $F_{5\sigma, \text{ch2}} = 10.8 \mu\text{Jy}$: the published source sits below our pipeline’s 5σ threshold. This case directly demonstrates the central methodological point of the thesis. The formal one-shot uncertainty at this position is $\sigma_{\text{phot}} = 0.23 \mu\text{Jy}$ (per-pixel sky scatter in the annulus), which would imply a $\sim 21\sigma$ formal SNR for a $4.86 \mu\text{Jy}$

source. The ensemble noise at the same position is $\sigma_{\text{ens}} = 10.8/5 \approx 2.2 \mu\text{Jy}$ — nearly an order of magnitude larger — because confusion noise from background gradients and PSF wings of neighbouring sources is captured by the ensemble injection but not by the single-annulus sky estimate. The ensemble $F_{5\sigma}$ is therefore the appropriate detection threshold; formal SNR computed from σ_{phot} overstates the true significance.

Flux-scale check. Forced photometry at the position of 1E 2259+586 yields a highly significant one-shot detection at $4.5 \mu\text{m}$ and marginal signals at other channels (Table 3.2).

Table 3.2: 1E 2259+586 one-shot photometry vs. literature

Channel	Wavelength (μm)	This work (μJy)	SNR	Literature (μJy)	Reference
1	3.6	1.03 ± 0.93	1.1	—	—
2	4.5	4.86 ± 0.23	21.0	6.3 ± 1.0	Kaplan and Chakrabarty [2009]
3	5.8	1.42 ± 6.63	0.2	—	—
4	8.0	4.01 ± 1.63	2.5	—	—

The $4.5 \mu\text{m}$ flux of $4.86 \pm 0.23 \mu\text{Jy}$ lies 1.4σ below the Kaplan and Chakrabarty [2009] value of $6.3 \pm 1.0 \mu\text{Jy}$ in the combined uncertainties — a clean pass of the absolute-scale check in this lower-background field. The $3.6 \mu\text{m}$ and $5.8 \mu\text{m}$ measurements are consistent with non-detection ($\text{SNR} < 1.5$), and the $8.0 \mu\text{m}$ signal at $\text{SNR} = 2.5$ is marginal and cannot be claimed as a detection. The contrast between the unbiased flux-scale recovery here and the bounded ~ 0.3 – 0.7 mag offset for 4U 0142+61 supports the interpretation that the 4U 0142+61 mismatch is a background-driven systematic specific to its heavily confused field, not a calibration failure of the pipeline.

3.2 Sensitivity Limits

The 5σ detection limits derived from the ensemble photometry procedure are summarized in Table 3.3 for all 27 magnetar targets with available archival Spitzer IRAC data. Entries marked with "—" indicate channels for which no IRAC mosaic was available or for which the ensemble photometry failed the data-quality check (grid positions overlapping chip gaps or image edges).

The limits span more than three orders of magnitude, from $\approx 10 \mu\text{Jy}$ (1E 2259+586

Table 3.3: 5σ detection limits (μJy) per magnetar per IRAC channel

Magnetar	Ch 1 (3.6 μm)	Ch 2 (4.5 μm)	Ch 3 (5.8 μm)	Ch 4 (8.0 μm)
1E 1048.1-5937	1159.5	737.3	621.8	240.7
1E 2259+586	13.1	10.8	54.6	37.9
1RXS J1708-4009	2580.6	1490.0	1353.5	775.2
3XMM J1852+0033	—	1203.0	1586.7	1701.2
4U 0142+61	25.7	33.7	—	43.7
AX J1818-1559	9348.3	—	3737.9	975.7
AX J1845-0258	1071.2	869.7	1125.0	1062.6
CXO J0100-7211	224.5	140.4	96.5	157.5
CXO J1647-4552	4148.1	—	—	—
CXO J1714-3810	2376.9	1792.7	1458.9	2527.0
PSR J1622-4950	803.6	700.4	1856.0	—
PSR J1846-0258	1467.8	930.8	822.3	—
SGR 0526-66	125.3	169.4	304.3	333.0
SGR 0755-2933	—	43318.5	—	—
SGR 1627-41	—	2793.0	1491.2	3204.9
SGR 1745-2900	9352.4	—	—	—
SGR 1801-23	12546.3	8133.6	—	3736.3
SGR 1806-20	34730.0	34346.6	—	—
SGR 1808-20	2313.0	1472.2	1215.6	731.9
SGR 1833-0832	895.7	678.0	1807.3	2123.1
SGR 1900+14	—	6953.6	—	—
SGR 1935+2154	—	67.8	928.3	—
SGR 2013+34	89.5	82.0	—	—
Swift J1818.0-1607	5014.8	3493.3	3880.6	2252.4
Swift J1822.3-1606	13038.5	9415.6	4821.4	3289.2
Swift J1834.9-0846	2876.4	2277.6	2258.1	—
XTE J1810-197	370.3	—	849.6	1871.3

at $4.5\ \mu\text{m}$) to $\gtrsim 10^4\ \mu\text{Jy}$ (Galactic center region sources). This dynamic range directly reflects the variation in local background brightness across the IRAC mosaics. The most constraining limits are achieved for magnetars at high Galactic latitudes or in low-background environments: 1E 2259+586 ($F_{5\sigma} \approx 11\text{--}55\ \mu\text{Jy}$ across channels), 4U 0142+61 ($F_{5\sigma} \approx 26\text{--}44\ \mu\text{Jy}$), SGR 2013+34 ($F_{5\sigma} \approx 82\text{--}90\ \mu\text{Jy}$ at 3.6 and $4.5\ \mu\text{m}$), and SGR 0526-66 in the Large Magellanic Cloud ($F_{5\sigma} \approx 125\text{--}333\ \mu\text{Jy}$).

Conversely, the most elevated limits occur for magnetars projected near the Galactic center: SGR 1806-20 ($F_{5\sigma} \approx 34\,700$ and $34\,300\ \mu\text{Jy}$ at 3.6 and $4.5\ \mu\text{m}$) and SGR 1745-2900 ($F_{5\sigma} \approx 9350\ \mu\text{Jy}$ at $3.6\ \mu\text{m}$). In these crowded regions the diffuse stellar and dust emission so strongly dominates the background that the 5σ detection threshold rises to $\sim 30\text{--}35\ \text{mJy}$, roughly 10^3 times the flux of the brightest known Spitzer magnetar counterpart, rendering the Spitzer archive non-constraining for those targets. The large noise in these cases originates not from photon counting statistics but from the structured, spatially varying background that the ensemble method directly samples.

The 5.8 and $8.0\ \mu\text{m}$ channels (ch3 and ch4) consistently yield higher limits than 3.6 and $4.5\ \mu\text{m}$. This is consistent with the known IRAC sensitivity hierarchy: ch3 and ch4 have lower point-source sensitivity (instrumental noise floor ≈ 36 and $54\ \mu\text{Jy}$ at 5σ in nominal observing conditions; IRAC Instrument Handbook) and are more severely affected by thermal dust emission from the Galactic interstellar medium. For the majority of targets, the $4.5\ \mu\text{m}$ channel provides the most constraining limit, which is also the channel at which the only two confirmed Spitzer magnetar detections have been made.

Chapter 4

Discussion

The goal of this chapter is to place the photometric measurements of Chapter 3 into a physical context. Section 4.1 derives mid-infrared spectral slopes from the multi-band IRAC data and compares them to theoretical model predictions. Section 4.2 examines the discrepancies between our measurements and the literature for the two detected sources, attributing them to methodological differences and discussing their implications. Section 4.3 uses the ensemble 5σ limits to constrain the infrared emission of the undetected magnetar population. Section 4.4 addresses the limitations of the analysis and outlines directions for future work.

4.1 Determining the Spectral Slope

The mid-infrared spectral energy distribution of a magnetar encodes information about its emission mechanism. The two leading frameworks — a passive fallback dust disk and non-thermal magnetospheric radiation — predict different SED shapes, but as we discuss below, a two-point Spitzer slope is generally not sufficient on its own to discriminate between them; Spitzer photometry is most powerful when combined with wider-baseline near-infrared and JWST data.

Spectral index convention. We adopt the convention $f_\nu \propto \nu^\alpha$, where $\alpha > 0$ indicates flux rising with frequency (blue in f_ν space) and $\alpha < 0$ indicates flux falling with frequency (red in f_ν space). Note that Hare et al. [2024] use the equivalent

convention $f\nu \propto \nu^{-\alpha}$ and quote $\alpha = +0.96 \pm 0.02$; the underlying spectrum is the same, but the sign of the reported numerical value differs. Throughout this work all spectral indices are expressed in our $f\nu \propto \nu^{+\alpha}$ convention, so the Hare et al. value translates to $\alpha = -0.96$. The two-point spectral index between IRAC channels i and j is

$$\alpha_{ij} = \frac{\ln(f_j/f_i)}{\ln(\nu_j/\nu_i)} \quad (4.1)$$

where the IRAC band centres correspond to $\nu_1 = 8.33 \times 10^{13}$ Hz ($3.6 \mu\text{m}$), $\nu_2 = 6.67 \times 10^{13}$ Hz ($4.5 \mu\text{m}$), $\nu_3 = 5.17 \times 10^{13}$ Hz ($5.8 \mu\text{m}$), and $\nu_4 = 3.75 \times 10^{13}$ Hz ($8.0 \mu\text{m}$).

Model predictions. For a passive, X-ray-heated dust disk, the predicted spectral slope across the IRAC bands depends sensitively on where those bands sit relative to the disk’s emission peak. The best-fit Wang et al. [2006] model has an inner-edge temperature $T_{\text{in}} \sim 1200$ K and outer-edge $T_{\text{out}} \sim 700$ K; the corresponding f_ν peaks ($\lambda_{\text{peak}} = hc/(2.821 k_B T)$) are at ~ 4.3 and $\sim 7.3 \mu\text{m}$ respectively, so the $4.5 \mu\text{m}$ and $8.0 \mu\text{m}$ IRAC bands straddle the integrated disk peak rather than lying blueward of it. In this regime the integrated f_ν can rise with wavelength across IRAC — exactly the behaviour fit by Wang et al. themselves — so a negative two-point IRAC slope is not, by itself, in conflict with a passive disk. The asymptotic limits often quoted for disk emission ($\alpha \approx +1/3$ for a steady-state $T \propto r^{-3/4}$ disk in its mid-IR self-similar regime; $\alpha \approx 2 + \beta$ for an optically thin Rayleigh-Jeans tail with $\kappa_\nu \propto \nu^\beta$) apply only when both bands lie blueward or redward of the disk peak, which is not the case here. The active irradiated disk models of Ertan et al. [2009] likewise produce a range of mid-IR slopes depending on disk geometry and irradiation.

For the magnetospheric model [Thompson et al., 2002, Beloborodov and Thompson, 2007], resonant cyclotron scattering in the twisted magnetosphere produces a non-thermal power law extending from X-ray energies down through the optical and infrared. Hare et al. [2024] measured $\alpha = -0.96$ for 4U 0142+61 with JWST over $1.4\text{--}11 \mu\text{m}$. The strongest constraint on a thermal-disk origin comes not from the slope

sign in the IRAC bands but from the wide wavelength baseline of this fit: a single power law extending to $1.4\,\mu\text{m}$ cannot be reproduced by any plausible-temperature dust component, since at $1.4\,\mu\text{m}$ even sublimation-temperature ($\sim 1200\,\text{K}$) grains lie deep in their Wien tail. Hare et al. therefore disfavour the simple passive-disk fit of Wang et al., while not formally excluding all disk variants (e.g., active or irradiated disks paired with an additional component).




Figure 4-1: Infrared spectral energy distribution of 4U 0142+61. Red downward triangles mark our Spitzer IRAC ensemble 5σ upper limits (ch1, ch2, ch4). Blue squares show the Wang and Kaspi [2008] Spitzer detections; green diamonds show Hare et al. [2024] JWST NIRCам measurements. The purple dashed line is the Hare et al. [2024] best-fit power law ($\alpha = -0.96$ over $1.4\text{--}11\,\mu\text{m}$); the orange dash-dotted line shows a steady-state $\alpha = +1/3$ disk slope for reference. Both our IRAC measurements and the literature Spitzer values are consistent with the wide-baseline JWST power law; the IRAC two-point slope alone does not arbitrate between a non-thermal magnetospheric origin and a multi-temperature disk whose emission peak lies near $4\text{--}8\,\mu\text{m}$.

Spectral slope of 4U 0142+61. Table 4.1 summarizes the two-point spectral indices derived from our one-shot photometry and from the Wang and Kaspi [2008]

literature values.

Table 4.1: Two-point spectral indices α for 4U 0142+61 ($f_\nu \propto \nu^\alpha$)

Channels	This work	Literature
Ch 1 $\rightarrow$ Ch 2	-3.72	—
Ch 2 $\rightarrow$ Ch 4	-0.49 ± 0.12	-1.09 (Wang & Kaspi 2008)
Ch 1 $\rightarrow$ Ch 4	-1.40	—
1.4–11 μm (JWST)	—	-0.96 (Hare+24)

The most physically reliable index is $\alpha_{24} = -0.49 \pm 0.12$, derived from the two highest-SNR detections (ch2: SNR = 36.1; ch4: SNR = 16.0). This negative slope is consistent with the wider-baseline JWST result of Hare et al. [2024] ($\alpha = -0.96$ over 1.4–11 μm) but, as noted above, a two-point IRAC slope alone cannot distinguish a non-thermal magnetospheric power law from a multi-temperature disk whose peak lies inside the IRAC bandpasses; both produce $\alpha < 0$ between 4.5 and 8.0 μm for the disk parameters of Wang et al. [2006]. The IRAC measurement is therefore most usefully read as an internally consistent point in the broader SED rather than as a model-discriminating diagnostic. The shallower-than-Hare value (-0.49 vs. -0.96) could reflect genuine spectral curvature between the IRAC and JWST bandpasses, the systematic ch4 under-recovery discussed in Section 4.2, or calibration differences between datasets. The very steep $\alpha_{12} = -3.72$ arises because our ch1 flux (10.15 μJy) is anomalously low compared to a smooth power law through channels 2 and 4; we discuss this in Section 4.2.

The ch2 $\rightarrow$ ch4 slope of -0.49 is significantly shallower than the literature slope of -1.09 derived from the Wang and Kaspi [2008] fluxes (32.1 and 59.8 μJy), which is itself close to Hare et al. [2024]. This difference ($\Delta\alpha \approx 0.6$) arises primarily from our lower ch4 flux (30.93 vs 59.8 μJy). All three Spitzer/JWST measurements agree on $\alpha < 0$ across IRAC; this is consistent with both a non-thermal magnetospheric origin and with a multi-temperature passive disk whose peak emission falls within the IRAC range, and the two scenarios are not separable from the IRAC two-point measurement alone.

Spectral slope of 1E 2259+586. For 1E 2259+586, only the ch2 detection (SNR = 21.0, $4.86 \pm 0.23 \mu\text{Jy}$) is statistically robust. The ch4 measurement ($4.01 \pm 1.63 \mu\text{Jy}$, SNR = 2.5) is marginal and subject to background systematics. Taking both values at face value yields $\alpha_{24} = +0.33$, suggesting a slightly rising spectral slope with frequency. However, given the low SNR of the ch4 measurement, this index is consistent within $\sim 1\sigma$ with a moderate disk slope ($\alpha \sim +1$) and within $\sim 1.8\sigma$ with the magnetospheric slope ($\alpha \sim -1$); given $\sigma_\alpha \approx 0.71$, the measurement is statistically consistent with either model. The ch2 upper limit from the 5σ ensemble ($F_{5\sigma, \text{ch2}} = 10.8 \mu\text{Jy}$) encloses the detected flux, confirming the detection is near the noise floor of the ensemble method for this field.

At ch1 the pipeline finds $1.03 \pm 0.93 \mu\text{Jy}$ (SNR = 1.1), consistent with non-detection. Extrapolating the Hare et al. [2024] magnetospheric slope of $\alpha = -0.96$ from the ch2 detection ($4.86 \mu\text{Jy}$ at $4.5 \mu\text{m}$) to $3.6 \mu\text{m}$ predicts $f_{\nu,1} = 4.86 \times (\nu_1/\nu_2)^{-0.96} = 4.86 \times (8.33/6.67)^{-0.96} \approx 3.9 \mu\text{Jy}$, which is well below our 5σ limit of $13.1 \mu\text{Jy}$ —consistent with non-detection under a magnetospheric spectral model. For 1E 2259+586 the Kaplan and Chakrabarty [2009] literature ch2 value of $6.3 \pm 1.0 \mu\text{Jy}$ would predict ch1 $\approx 5.1 \mu\text{Jy}$ under $\alpha = -0.96$, also below the $13.1 \mu\text{Jy}$ limit.

Spectral index upper limits for non-detected magnetars. For targets with only upper limits, the IRAC measurements constrain the mid-IR SED in the following way. If we assume the emission follows a power law with slope α , the ratio of flux upper limits between two channels bounds $|\alpha|$:

$$\alpha \geq \frac{\ln(F_{5\sigma,j}/F_{5\sigma,i})}{\ln(\nu_j/\nu_i)} \quad (4.2)$$

For example, SGR 2013+34 has $F_{5\sigma, \text{ch1}} = 89.5 \mu\text{Jy}$ and $F_{5\sigma, \text{ch2}} = 82.0 \mu\text{Jy}$, giving the bound $\alpha \geq \ln(82.0/89.5)/\ln(6.67/8.33) = -0.088/(-0.222) = 0.39$ *if the source flux at ch1 is at the 5σ detection threshold*. Under that assumption, a source with slope $\alpha < 0.39$ — including the standard steady-state disk slope ($\alpha \approx +1/3$) or any magnetospheric slope ($\alpha \approx -1$) — would have its ch2 flux exceed the $82.0 \mu\text{Jy}$ limit and would therefore be detected in ch2. We emphasise that this argument

constrains the joint (F_1, α) region only at the ch1 sensitivity boundary: a sufficiently faint magnetospheric source ($F_1 \ll F_{5\sigma, \text{ch1}}$) is fully consistent with non-detection in both channels. Steep disk slopes ($\alpha \gtrsim 0.4$) remain consistent with non-detection at any flux below the ch1 limit. For SGR 0526-66 in the Large Magellanic Cloud (at distance $d \approx 50$ kpc), the ch2 5σ limit of $169 \mu\text{Jy}$ corresponds to a luminosity upper limit of $\nu L_\nu \lesssim 4\pi d^2 \nu F_{5\sigma} \approx 3.4 \times 10^{34} \text{ erg s}^{-1}$ at $4.5 \mu\text{m}$. For comparison, the $4.5 \mu\text{m}$ luminosity of 4U 0142+61 at our measured flux is $\nu L_\nu \approx 4\pi d^2 \nu F \approx 2.4 \times 10^{31} \text{ erg s}^{-1}$ (for $d \approx 3.6$ kpc and $F = 23.3 \mu\text{Jy}$). The LMC limit is therefore roughly three orders of magnitude above the 4U 0142+61 detection level, placing any 4U 0142+61-like mid-infrared counterpart in SGR 0526-66 well below Spitzer’s sensitivity at that distance.

4.2 Comparison with Literature Values and Sources of Discrepancy

The measured fluxes for 4U 0142+61 are systematically lower than the Wang and Kaspi [2008] values at ch2 (23.30 vs $32.1 \mu\text{Jy}$; -27%) and ch4 (30.93 vs $59.8 \mu\text{Jy}$; -48%). In contrast, for 1E 2259+586 the ch2 flux of $4.86 \mu\text{Jy}$ agrees with Kaplan and Chakrabarty [2009] at the 1.4σ level ($6.3 \pm 1.0 \mu\text{Jy}$), well within the stated uncertainties. This contrast between the two sources points to a systematic effect that is stronger in the more background-confused 4U 0142+61 field.

Several factors contribute to the discrepancy.

Aperture photometry vs. PRF fitting. Wang et al. [2006] and Wang and Kaspi [2008] applied MOPEX PRF-fitting to individual Basic Calibrated Data frames before coadding, isolating the point-source contribution from the surrounding diffuse emission at native pixel scale. Our method applies circular aperture photometry with a local annular background subtraction to the coadded mosaic. In fields with strong extended emission gradients — as is the case for 4U 0142+61, which lies within an extended dust cloud — the median background estimated from an annulus at

$r_{\text{in}} = 2.2 - r_{\text{out}} = 3.0$ native pixels over-subtracts the local background beneath the source aperture if the background is rising toward the source. This over-subtraction is the most likely cause of our systematically lower fluxes, particularly at $8\,\mu\text{m}$ where the interstellar dust emission is strongest.

Aperture size and correction. Our aperture radius of $r_{\text{ap}} = 2.0$ native pixels ($\approx 2.44''$) is smaller than the $r_{\text{ap}} = 3$ pixels ($\approx 3.7''$) used in many published Spitzer analyses. Smaller apertures are more sensitive to centering errors and pick up a smaller fraction of the extended PRF wings, potentially introducing residual aperture correction errors even after applying the APCOR factors from the IRAC Instrument Handbook. At $8\,\mu\text{m}$, where the PRF is broader, this effect is larger, which is consistent with our larger fractional discrepancy in ch4.

Source variability. Hare et al. [2024] noted flux variability in 4U 0142+61 between the Spitzer epoch and their JWST observations. However, Wang and Kaspi [2008] found stable Spitzer fluxes across two observations separated by several months, and the variability documented by Hare et al. is at the factor-of-2 level in the MIRI $7.7\,\mu\text{m}$ band but not confirmed in the overlapping IRAC bandpasses. Source variability cannot be fully excluded but is unlikely to account for the entire 27–48% discrepancy between our measurements and the Wang values.

The agreement for 1E 2259+586 validates the pipeline under more favorable conditions: 1E 2259+586 lies in a lower-background field, making the annulus background estimate less prone to gradient over-subtraction. This supports the interpretation that the 4U 0142+61 discrepancy is a background estimation effect rather than a fundamental calibration failure.

Implications for ensemble limit validation. Recall from Section 3.1 that the 5σ ensemble limits for 4U 0142+61 are $F_{5\sigma,\text{ch2}} = 33.7\,\mu\text{Jy}$ and $F_{5\sigma,\text{ch4}} = 43.7\,\mu\text{Jy}$. The ch2 literature value ($32\,\mu\text{Jy}$) falls just below the ch2 ensemble limit, and the ch4 literature value ($60\,\mu\text{Jy}$) lies above the ch4 limit, confirming that sources at the Wang+06 flux level are detectable by our pipeline in both channels. This confirms that the ensemble method correctly characterizes the photometric noise level even when the absolute flux from the one-shot measurement is offset by background systematics. The

ensemble procedure is therefore reliable for deriving upper limits even in moderately confused fields.

4.3 Constraints on the Undetected Magnetar Population

Of the 27 targets analysed, 25 yield no statistically significant one-shot detection ($\text{SNR} < 3$) in any IRAC channel. The ensemble 5σ limits can be translated directly into physical constraints as follows.

Luminosity upper limits. For a target at distance d , the ch2 ($4.5\,\mu\text{m}$) 5σ limit gives a luminosity constraint:

$$\nu L_\nu|_{4.5\,\mu\text{m}} < 4\pi d^2 \nu_2 F_{5\sigma,\text{ch2}} \quad (4.3)$$

For the best-constrained non-detected target, SGR 2013+34, $F_{5\sigma,\text{ch2}} = 82\,\mu\text{Jy}$. If SGR 2013+34 is at a distance comparable to other magnetars in the Galactic disk ($d \sim 5\text{--}10\,\text{kpc}$), this corresponds to $\nu L_\nu < 2\text{--}7 \times 10^{32}\,\text{erg s}^{-1}$ at $4.5\,\mu\text{m}$, roughly one order of magnitude above the detected $4.5\,\mu\text{m}$ luminosity of 4U 0142+61 ($\nu L_\nu \approx 2.4 \times 10^{31}\,\text{erg s}^{-1}$ at $d \approx 3.6\,\text{kpc}$). Distances are poorly constrained for most magnetars; adopting the dispersion-measure-based distance estimates from the McGill Catalog [Olausen and Kaspi, 2014], the majority of upper limits translate to $\nu L_\nu \lesssim 10^{32}\text{--}10^{33}\,\text{erg s}^{-1}$ in the best-constrained channels at typical Galactic distances, roughly one to two orders of magnitude above the confirmed $4.5\,\mu\text{m}$ luminosity of 4U 0142+61.

Model exclusions. The active irradiated disk models of Ertan et al. [2009] predict mid-IR fluxes that scale with the X-ray irradiation luminosity L_X . For magnetars with $L_X \sim 10^{33}\text{--}10^{35}\,\text{erg s}^{-1}$ and assuming the same irradiation efficiency as in the 4U 0142+61 model, predicted $4.5\,\mu\text{m}$ fluxes are of order $\sim 1\text{--}100\,\mu\text{Jy}$ at $5\,\text{kpc}$ [Ertan et al., 2009]. Our best ensemble limits ($\sim 80\text{--}130\,\mu\text{Jy}$ for SGR 2013+34 and SGR 0526-66) are not yet deep enough to exclude this range for typical magnetar X-ray luminosities, but they do exclude very bright disk scenarios with $F_{4.5} \gtrsim 100\,\mu\text{Jy}$ that

would be expected for a disk as luminous as 4U 0142+61 scaled to the same distance.

Impact of background confusion. The factor of > 1000 spread in 5σ limits across the sample (from $\sim 11 \mu\text{Jy}$ for 1E 2259+586 to $\sim 34,700 \mu\text{Jy}$ for SGR 1806-20) is driven entirely by local background brightness, confirming the finding of Kobulnicky et al. [2013] that confusion noise — not photon noise — limits Spitzer sensitivity in the Galactic plane. The full range of ensemble 5σ limits is tabulated in Table 3.3. For the ~ 8 targets with limits exceeding $\sim 10^3 \mu\text{Jy}$ (SGR 1806-20, SGR 1745-2900, SGR 1801-23, Swift J1818, Swift J1822, AX J1818-1559, SGR 1627-41, CXO J1647-4552), the Spitzer archival data provide no meaningful constraint, and JWST with its superior angular resolution and sensitivity is the appropriate instrument to attempt mid-IR detection.

4.4 Limitations and Future Directions

Systematic uncertainties in aperture photometry. The ensemble method robustly estimates the noise induced by the real background, but it does not fully account for systematic over-subtraction of background beneath faint sources in strongly rising backgrounds. A PRF-fitting approach applied to individual BCD frames — as employed by Wang et al. [2006] — would be less affected by large-scale background gradients. Future work should apply PRF-fitting to the same archival mosaics and compare the results to the aperture photometry presented here; any systematic difference would quantify the background over-subtraction bias.

PRF model accuracy. The cryogenic APEX PRFs used here represent the average PRF over the full cryogenic focal plane, whereas the true PRF has a documented position-dependent variation of $\sim 10\%$ across the array [IRAC Instrument Team, 2021]. For individual targets whose mosaic positions are near the edges of the IRAC array, this introduces a systematic uncertainty in the injected PRF shape and hence in the ensemble noise estimate. This effect is expected to be sub-dominant compared to background confusion for most targets.

Multi-wavelength SED constraints. The spectral indices derived in Section 4.1

are based on at most three IRAC bands and span only a factor of ~ 2 in frequency ($3.6\text{--}8\,\mu\text{m}$). A physically complete SED model requires coverage from the near-infrared (J , H , K bands, $\sim 1\text{--}2\,\mu\text{m}$) through the mid-infrared to the far-infrared. For 4U 0142+61, combining the Hulleman et al. [2000, 2004] optical/NIR measurements, our IRAC photometry, and the Hare et al. [2024] JWST spectrum would provide a 13-band SED from 0.6 to $11\,\mu\text{m}$ — sufficient to constrain both the spectral slope and any spectral break between the near-IR and mid-IR. For non-detected targets, the most impactful improvement would come from adding deep ground-based or HST near-IR photometry at K -band to constrain the slope shortward of the IRAC bands, as pursued by Chrimes et al. [2022] for the HST sample.

JWST follow-up. For the most physically interesting targets with constraining Spitzer limits ($F_{5\sigma} \lesssim 200\,\mu\text{Jy}$ at $4.5\,\mu\text{m}$: 1E 2259+586, 4U 0142+61, SGR 1935+2154, SGR 2013+34, CXO J0100-7211, SGR 0526-66), JWST MIRI observations would provide $\sim 10\times$ deeper sensitivity at comparable wavelengths with $\sim 7\times$ better angular resolution [Hare et al., 2024]. The MIRI F560W and F770W filters overlap the IRAC ch3 ($5.8\,\mu\text{m}$) and ch4 ($8.0\,\mu\text{m}$) bands, respectively, and would enable direct spectral slope measurements at the sub- μJy level, sufficient to distinguish the passive disk, active disk, and magnetospheric emission models across the full magnetar population.

Chapter 5

Conclusion

This thesis set out to determine how faint a magnetar can be in the Spitzer Space Telescope infrared bands, and to quantify that limit rigorously and uniformly across the known magnetar population. To that end, a custom Python photometric pipeline was developed that combines forced aperture photometry at known X-ray positions with an ensemble PRF-injection method for empirical noise estimation, yielding 5σ detection limits that faithfully reflect the true local photometric noise rather than idealized Poisson statistics.

The pipeline was validated along two independent axes against the two magnetars with published Spitzer IRAC detections: a sensitivity check ($F_{5\sigma}$ vs. literature flux) and an absolute-flux-scale check (one-shot photometry vs. literature flux). The sensitivity check is the directly relevant validation for the upper limits that constitute the main scientific output, and it succeeds at both targets: for 4U 0142+61 the literature $4.5\ \mu\text{m}$ flux ($32.1\ \mu\text{Jy}$) and $8.0\ \mu\text{m}$ flux ($59.8\ \mu\text{Jy}$) bracket the ensemble thresholds (33.7 and $43.7\ \mu\text{Jy}$) in the expected way, while for 1E 2259+586 the literature ch2 detection at $6.3\ \mu\text{Jy}$ sits just below the ensemble threshold of $10.8\ \mu\text{Jy}$ — directly demonstrating that confusion noise (captured empirically by the ensemble) dominates over the formal per-pixel sky scatter (which would have implied $\sim 20\sigma$ formal SNR at this position). The absolute-flux-scale check passes for 1E 2259+586 (1.4σ agreement) and shows a bounded $\sim 27\text{--}48\%$ offset for 4U 0142+61 attributable to annular background over-subtraction in its heavily confused Galactic-plane field; that offset bounds the

absolute scale at this position but does not propagate into the ensemble $F_{5\sigma}$.

Applying the pipeline to 27 magnetar targets with archival Spitzer IRAC mosaics, the ensemble method yields 5σ flux limits spanning more than three orders of magnitude: from $\approx 10 \mu\text{Jy}$ for 1E 2259+586 at $4.5 \mu\text{m}$ to $\approx 34,700 \mu\text{Jy}$ for SGR 1806-20 at $3.6 \mu\text{m}$. This extraordinary dynamic range is driven entirely by local background brightness, confirming that confusion noise — not photon counting statistics — is the limiting factor for Spitzer sensitivity in the Galactic plane. For the approximately eight targets with limits exceeding $\sim 10^3 \mu\text{Jy}$ (including SGR 1806-20, SGR 1745-2900, and the Galactic center sources), the archival Spitzer data provide no meaningful physical constraint on mid-infrared emission.

For 4U 0142+61, whose multi-band photometry is statistically robust, the two-point IRAC spectral indices are negative ($\alpha < 0$). The most reliable index, $\alpha_{24} = -0.49 \pm 0.12$ between 4.5 and $8.0 \mu\text{m}$, is consistent with the wider-baseline JWST result of Hare et al. [2024] ($\alpha = -0.96$ over $1.4\text{--}11 \mu\text{m}$). On its own, however, the IRAC two-point measurement does not arbitrate between a non-thermal magnetospheric power law and a multi-temperature passive disk whose f_ν emission peak falls within the IRAC bandpasses (the Wang et al. 2006 disk model has $T_{\text{in}} \sim 1200 \text{ K}$ and $T_{\text{out}} \sim 700 \text{ K}$, with peaks near $4\text{--}7 \mu\text{m}$); both classes of model produce $\alpha < 0$ between ch2 and ch4. The argument against a thermal-disk dominance therefore rests on the wide wavelength baseline of the Hare et al. [2024] fit rather than on the IRAC slope alone. The shallower IRAC value relative to JWST is plausibly explained by genuine spectral curvature between the bandpasses, by stronger background over-subtraction at $8.0 \mu\text{m}$ reducing our recovered f_4/f_2 ratio, or by source variability between epochs.

For the non-detected magnetar population, the best ensemble limits ($\sim 80\text{--}130 \mu\text{Jy}$ for SGR 2013+34 and SGR 0526-66) are not yet deep enough to exclude the full range of theoretical disk and magnetospheric models at typical magnetar distances, but they do exclude bright disk scenarios comparable to 4U 0142+61 scaled to similar distances, and they constrain mid-IR luminosities to $\nu L_\nu \lesssim 10^{32}\text{--}10^{33} \text{ erg s}^{-1}$ at typical Galactic distances for the best-constrained targets. For sources in low-background environments (1E 2259+586, 4U 0142+61, SGR 1935+2154, SGR 2013+34, CXO

J0100-7211, SGR 0526-66), JWST MIRI follow-up at the sub- μ Jy level with $\sim 7\times$ better angular resolution would provide the sensitivity needed to detect or conclusively rule out disk emission in these targets and to extend spectral slope measurements across the full IRAC-to-MIRI wavelength range.

The ensemble photometry pipeline developed for this thesis is implemented in the open-source Python package `IRACMagLim` and is directly applicable to any Spitzer IRAC mosaic. Its design — forced photometry at a known position, PRF injection at a grid of nearby positions, and noise estimation from the scatter of recovered fluxes — is general and can be adapted for archival analysis with other infrared observatories. The 27-target 5σ limit catalog presented in Table 3.3 constitutes the most comprehensive mid-infrared sensitivity survey of the known magnetar population to date and provides a quantitative baseline for planning and interpreting future JWST and ground-based infrared observations of these extreme objects.

Bibliography

- M. Ali Alpar. On young neutron stars as propellers and accretors with conventional magnetic fields. *ApJ*, 554:1245–1254, 2001. doi: 10.1086/321393.
- Andrei M. Beloborodov and Christopher Thompson. Corona of magnetars. *ApJ*, 657: 967–993, 2007. doi: 10.1086/508917.
- Andrei M. Beloborodov and L. Xinyu. Magnetar heating. *ApJ*, 833:261, 2016. doi: 10.3847/1538-4357/833/2/261.
- Robert A. Benjamin et al. GLIMPSE: I. an SIRTf legacy project to map the inner galaxy. *PASP*, 115:953–964, 2003. doi: 10.1086/376696.
- O. Benli and Ümit Ertan. Long-term evolution of anomalous x-ray pulsars and soft gamma-ray repeaters. *MNRAS*, 457:4114—4122, 2016. doi: 10.1093/mnras/stw235.
- Emmanuel Bertin and Stephane Arnouts. SExtractor: Software for source extraction. *A&AS*, 117:393–404, 1996. doi: 10.1051/aas:1996164.
- Christopher D. Bochenek et al. A fast radio burst associated with a galactic magnetar. *Nature*, 587:59–62, 2020. doi: 10.1038/s41586-020-2872-x.
- Pranab Chatterjee, Lars Hernquist, and Ramesh Narayan. An accretion model for anomalous x-ray pulsars. *ApJ*, 534:373–379, 2000. doi: 10.1086/308748.
- CHIME/FRB Collaboration et al. A bright millisecond-duration radio transient from a galactic magnetar. *Nature*, 587:54–58, 2020. doi: 10.1038/s41586-020-2863-y.
- Ashley A. Chrimes et al. New candidates for magnetar counterparts from a deep search with the hubble space telescope. *MNRAS*, 512:6093–6121, 2022. doi: 10.1093/mnras/stac792.
- Ed Churchwell et al. The spitzer/glimpse surveys: A new view of the milky way. *PASP*, 121:213–230, 2009. doi: 10.1086/597811.
- P. R. den Hartog, L. Kuiper, W. Hermsen, V. M. Kaspi, R. Dib, J. Knödlseider, and F. P. Gavriil. Detailed high-energy characteristics of AXP 4U 0142+61. multi-year observations with INTEGRAL, RXTE, XMM-Newton, and ASCA. *A&A*, 489: 245–261, 2008. doi: 10.1051/0004-6361:200809390.

- Robert C. Duncan and Christopher Thompson. Formation of very strongly magnetized neutron stars: Implications for gamma-ray bursts. *ApJL*, 392:L9–L13, 1992. doi: 10.1086/186413.
- Teruaki Enoto, Shota Kisaka, and Shinpei Shibata. Observational diversity of magnetized neutron stars. *Reports on Progress in Physics*, 82:106901, 2019. doi: 10.1088/1361-6633/ab3def.
- Ümit Ertan and Şirin Çalışkan. Optical and infrared emission from the anomalous x-ray pulsars and soft gamma-ray repeaters. *ApJL*, 649:L87–L90, 2006. doi: 10.1086/508347.
- Ümit Ertan et al. On the evolution of anomalous x-ray pulsars and soft gamma ray repeaters with fallback disks. *ApJ*, 702:1309, 2009. doi: 10.1088/0004-637X/702/2/1309.
- Paolo Esposito et al. A very young radio-loud magnetar. *ApJL*, 896:L30, 2020. doi: 10.3847/2041-8213/ab9742.
- Giovanni G. Fazio et al. The infrared array camera (irac) for the spitzer space telescope. *ApJS*, 154:10–17, 2004. doi: 10.1086/422843.
- Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Physical Review D*, 57:3873–3889, 1998. doi: 10.1103/PhysRevD.57.3873.
- Neil Gehrels. Confidence limits for small numbers of events in astrophysical data. *ApJ*, 303:336–346, 1986. doi: 10.1086/164079.
- Jeremy Hare et al. Probing the spectrum of the magnetar 4u 0142+61 with jwst. *ApJ*, 972:176, 2024. doi: 0.3847/1538-4357/ad5ce5.
- J. L. Hora et al. Photometry using the infrared array camera on the spitzer space telescope. *AJ*, 120:1233, 2008. doi: 10.1086/593217.
- Joseph L. Hora et al. In-flight performance and calibration of the infrared array camera for the spitzer space telescope. *Proc. SPIE*, 5487:77–92, 2004. doi: 10.1117/12.551380.
- F. Hulleman. Optical and near-infrared observations of anomalous x-ray pulsars. *arXiv:astro-ph*, 2003. doi: 10.48550/arXiv.astro-ph/0111561.
- Frank Hulleman, Marten H. van Kerkwijk, and Shrinivas R. Kulkarni. An optical counterpart to the anomalous x-ray pulsar 4u 0142+61. *Nature*, 408:689–692, 2000. doi: 10.1038/35047110.
- Frank Hulleman, Marten H. van Kerkwijk, and Shrinivas R. Kulkarni. The anomalous x-ray pulsar 4U 0142+61: Variability in the infrared and a spectral break in the optical. *A&A*, 416:1037–1045, 2004. doi: 10.1051/0004-6361:20034501.

- IRAC Instrument Team. IRAC Instrument Handbook. IRSA/IPAC, Caltech, 2021. Version 2.3. Available at irsa.ipac.caltech.edu.
- David L. Kaplan and Deepto Chakrabarty. A mid-infrared counterpart to the magnetar 1e 2259+586. *ApJ*, 700:149–154, 2009. doi: 10.1088/0004-637X/700/1/149.
- Vinay L. Kashyap, David A. van Dyk, Alanna Connors, et al. On computing upper limits to source intensities. *ApJ*, 719:900–914, 2010. doi: 10.1088/0004-637X/719/1/900.
- Victoria M. Kaspi and Andrei M. Beloborodov. Magnetars. *ARA&A*, 55:261–301, 2017. doi: 10.1146/annurev-astro-081915-023329.
- Henry A. Kobulnicky, Brian L. Babler, Michael J. Alexander, Marilyn R. Meade, Barbara A. Whitney, and Edward Churchwell. Effects of diffuse background emission and source crowding on photometric completeness in Spitzer IRAC surveys: The GLIMPSE catalogs and archives, 2013.
- Ralph P. Kraft, David N. Burrows, and John A. Nousek. Determination of confidence limits for simple poisson processes. *ApJ*, 374:344–355, 1991. doi: 10.1086/170124.
- David Makovoz and Francine R. Marleau. Point-source extraction with mopex. *PASP*, 117:1113–1128, 2005. doi: 10.1086/432977.
- Massimo Marengo et al. Irac point source photometry guide. Spitzer Science Center Technical Note, 2009. Spitzer Heritage Archive.
- Sandro Mereghetti. The strong-field regime of neutron stars: Soft gamma-ray repeaters and anomalous x-ray pulsars. *A&AR*, 15:225–287, 2008. doi: 10.1007/s00159-007-0011-z.
- Steve A. Olausen and Victoria M. Kaspi. The mcgill magnetar catalog. *ApJS*, 212:6, 2014. doi: 10.1088/0067-0049/212/1/6.
- Rosalba Perna, Lars Hernquist, and Ramesh Narayan. Emission spectra of fallback disks around young neutron stars. *ApJ*, 541:344–350, 2000. doi: 10.1086/309422.
- Nanda Rea et al. Resonant cyclotron scattering in magnetars’ emission. *ApJ*, 686:1245–1260, 2008. doi: 10.1086/591310.
- William T. Reach et al. Absolute calibration of the infrared array camera on the spitzer space telescope. *PASP*, 117:978–990, 2005. doi: 10.1086/432670.
- Peter B. Stetson. DAOPHOT: A computer program for crowded-field stellar photometry. *PASP*, 99:191–222, 1987. doi: 10.1086/131977.
- Curtis R. Tam et al. Correlated infrared and X-ray flux changes following the 2002 june outburst of the anomalous X-ray pulsar 1e 2259+586. *ApJL*, 617:L53–L56, 2004. doi: 10.1086/426963.

- Christopher Thompson and Robert C. Duncan. Neutron star dynamos and the origins of pulsar magnetism. *ApJ*, 408:194–217, 1993. doi: 10.1086/172580.
- Christopher Thompson and Robert C. Duncan. The soft gamma repeaters as very strongly magnetized neutron stars. i. radiative mechanism for outbursts. *MNRAS*, 275:255–300, 1995. doi: 10.1093/mnras/275.2.255.
- Christopher Thompson and Robert C. Duncan. The soft gamma repeaters as very strongly magnetized neutron stars. ii. quiescent neutrino, x-ray, and alfvén wave emission. *ApJ*, 473:322–342, 1996. doi: 10.1086/178147.
- Christopher Thompson, Maxim Lyutikov, and Shrinivas R. Kulkarni. Electrodynamics of magnetars: Implications for the persistent x-ray emission and spin-down of the soft gamma repeaters and anomalous x-ray pulsars. *ApJ*, 574:332–355, 2002. doi: 10.1086/340586.
- Roberto Turolla, Silvia Zane, and Anna L. Watts. Magnetars: The physics behind observations. a review. *Reports on Progress in Physics*, 78:116901, 2015. doi: 10.1088/0034-4885/78/11/116901.
- Stefanie Wachter et al. An infrared ring around the magnetar sgr 1900+14. *Nature*, 453:626–628, 2008. doi: 10.1038/nature06987.
- Zhongxiang Wang and Victoria M. Kaspi. Spitzer observations of anomalous x-ray pulsar 4u 0142+61. *ApJ*, 675:695–697, 2008. doi: 10.1086/526311.
- Zhongxiang Wang, Deepto Chakrabarty, and David L. Kaplan. A debris disk around an isolated young neutron star. *Nature*, 440:772–775, 2006. doi: 10.1038/nature04669.
- Michael W. Werner et al. The spitzer space telescope mission. *ApJS*, 154:1–9, 2004. doi: 10.1086/422992.