Modeling our Milky Way

Astrostatistics | Big Data | Data Visualization

Catherine Zucker[1], Joshua Speagle[1], Edward Schlafly[2], Gregory Green[3], Douglas Finkbeiner[1], Alyssa Goodman[1,4], João Alves[4,5]

[1]Center for Astrophysics | Harvard & Smithsonian, [2]Lawrence Livermore National Laboratory, [3]Max Planck Institute for Astronomy, [4]Radcliffe Institute for Advanced Study, [5]University of Vienna

CENTER FOR ASTROPHYSICS
HARVARD & SMITHSONIAN

The last part, first.

The Radcliffe Wave

Each **red** dot marks a star-forming blob of gas whose distance from us has been accurately measured.

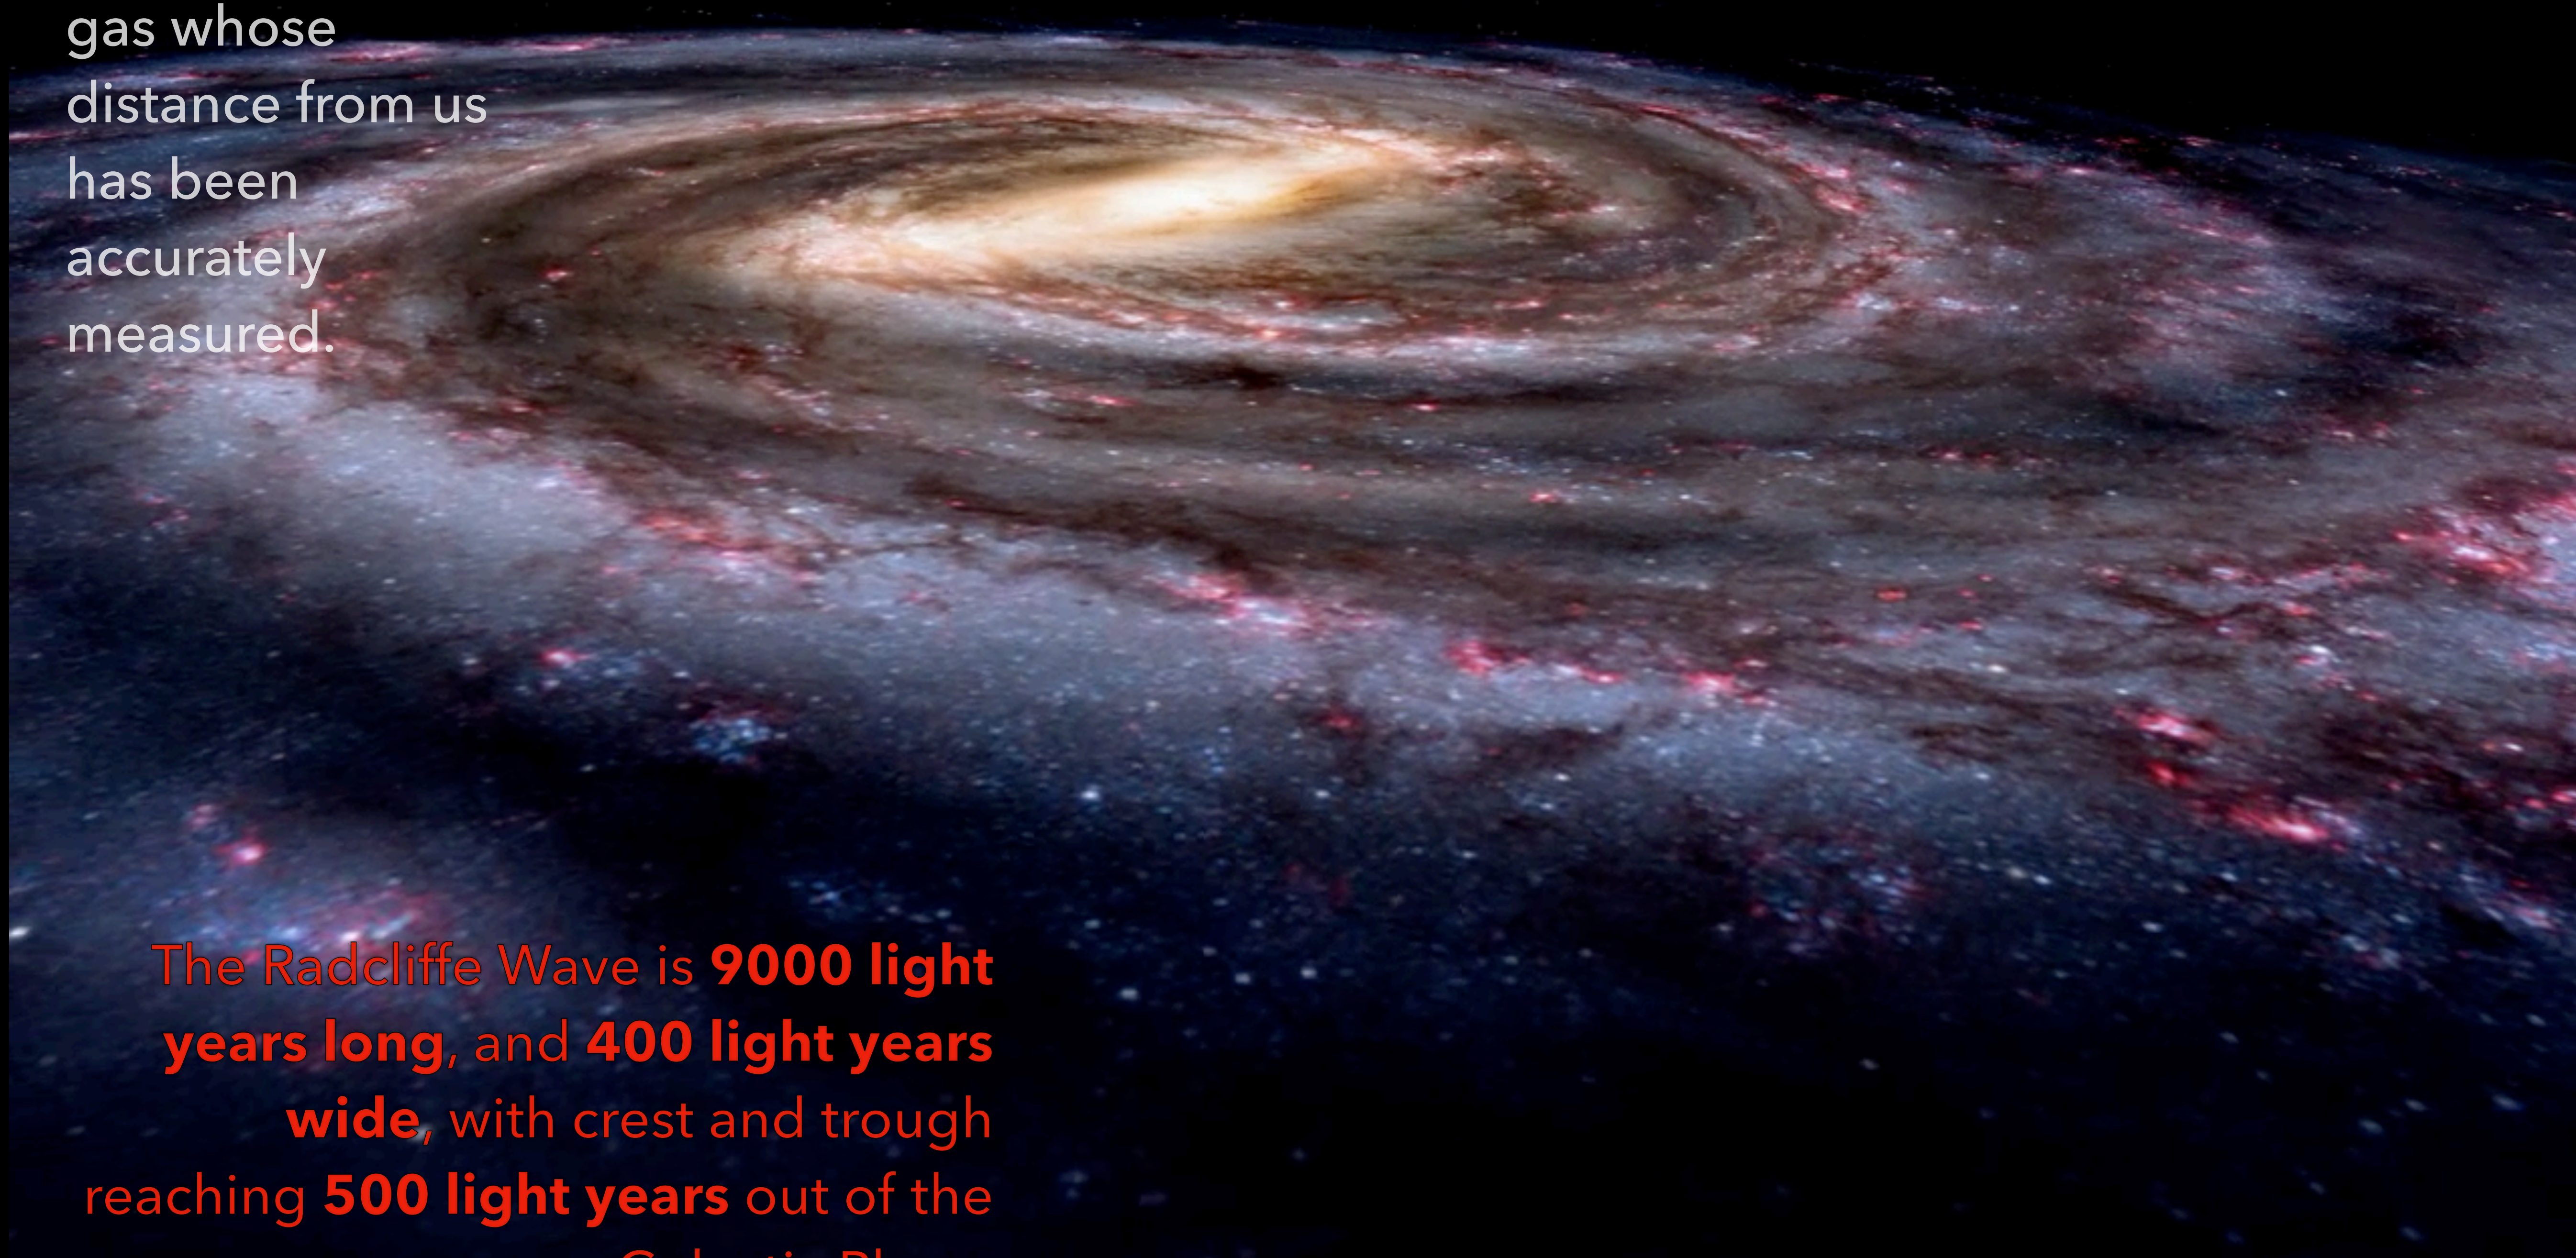

The Radcliffe Wave is **9000 light years long**, and **400 light years wide**, with crest and trough reaching **500 light years** out of the Galactic Plane.

Its gas mass is **more than three million times** the mass of the Sun.

The Milky Way
(Artist's Conception)

Real & Fake

Real Image of Actual
Spiral Galaxy

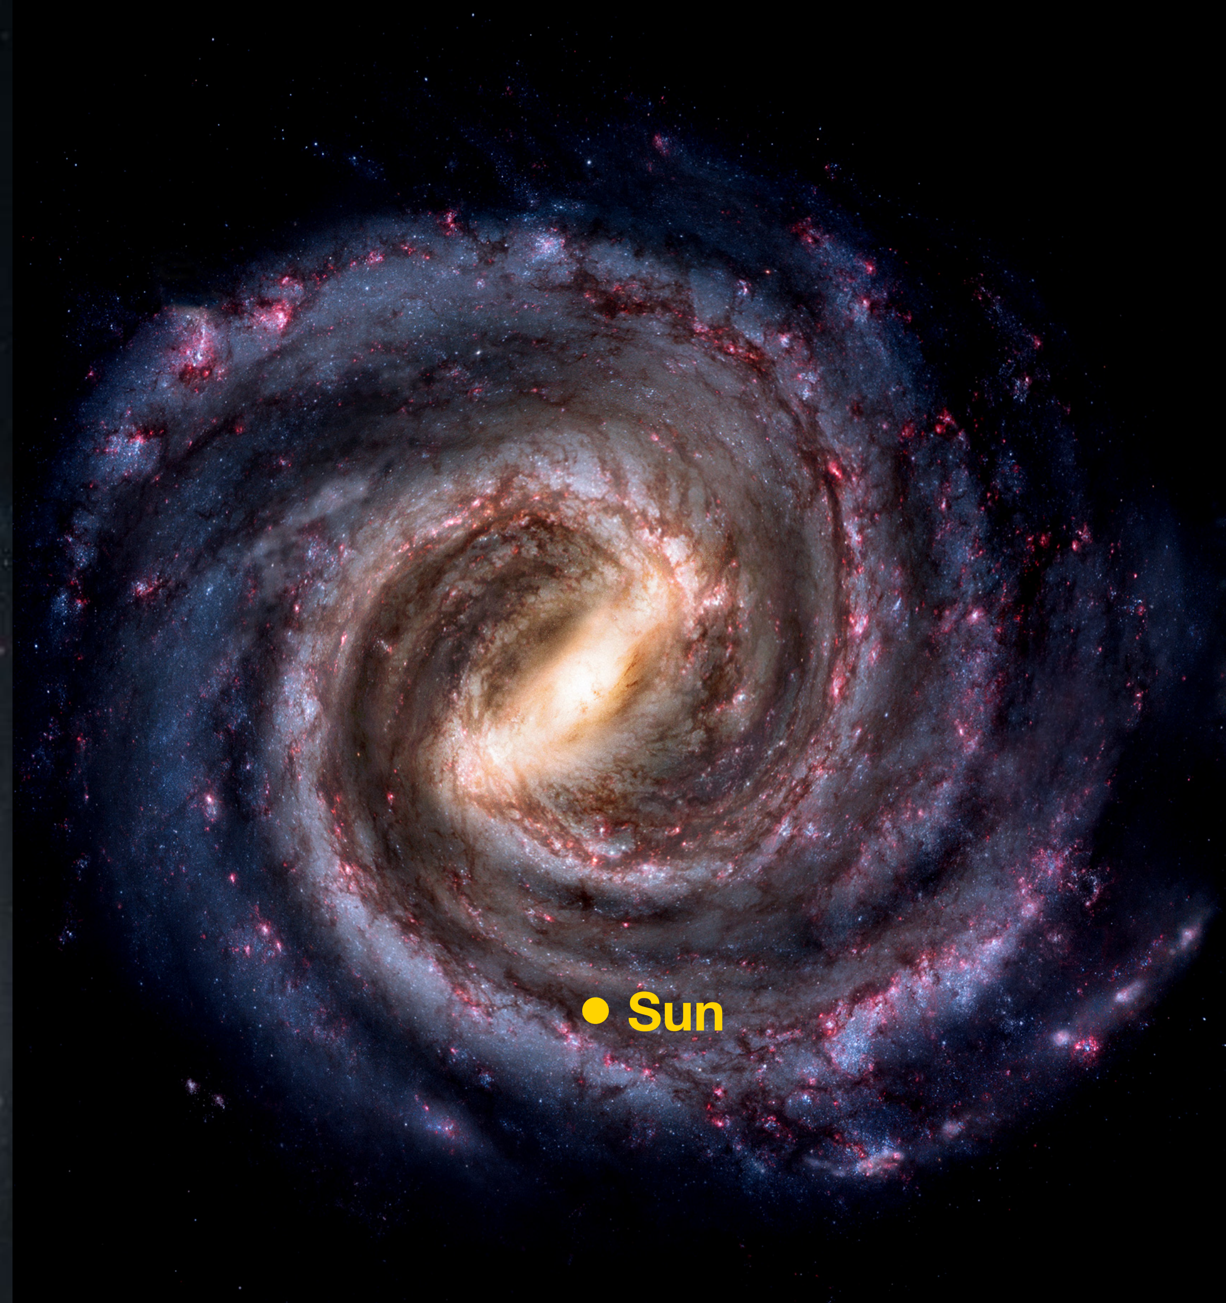

Cartoon Model of
our Milky Way

Stellar Nursery

3D Cartoon

Actual Image of the Orion Nebula

(in 2D, on the sky)

See Green+2014, 2015, 2018, 2019 for more details on stellar modeling

WARNING: schematic diagram, NOT to scale (credit A. Goodman, 2019)

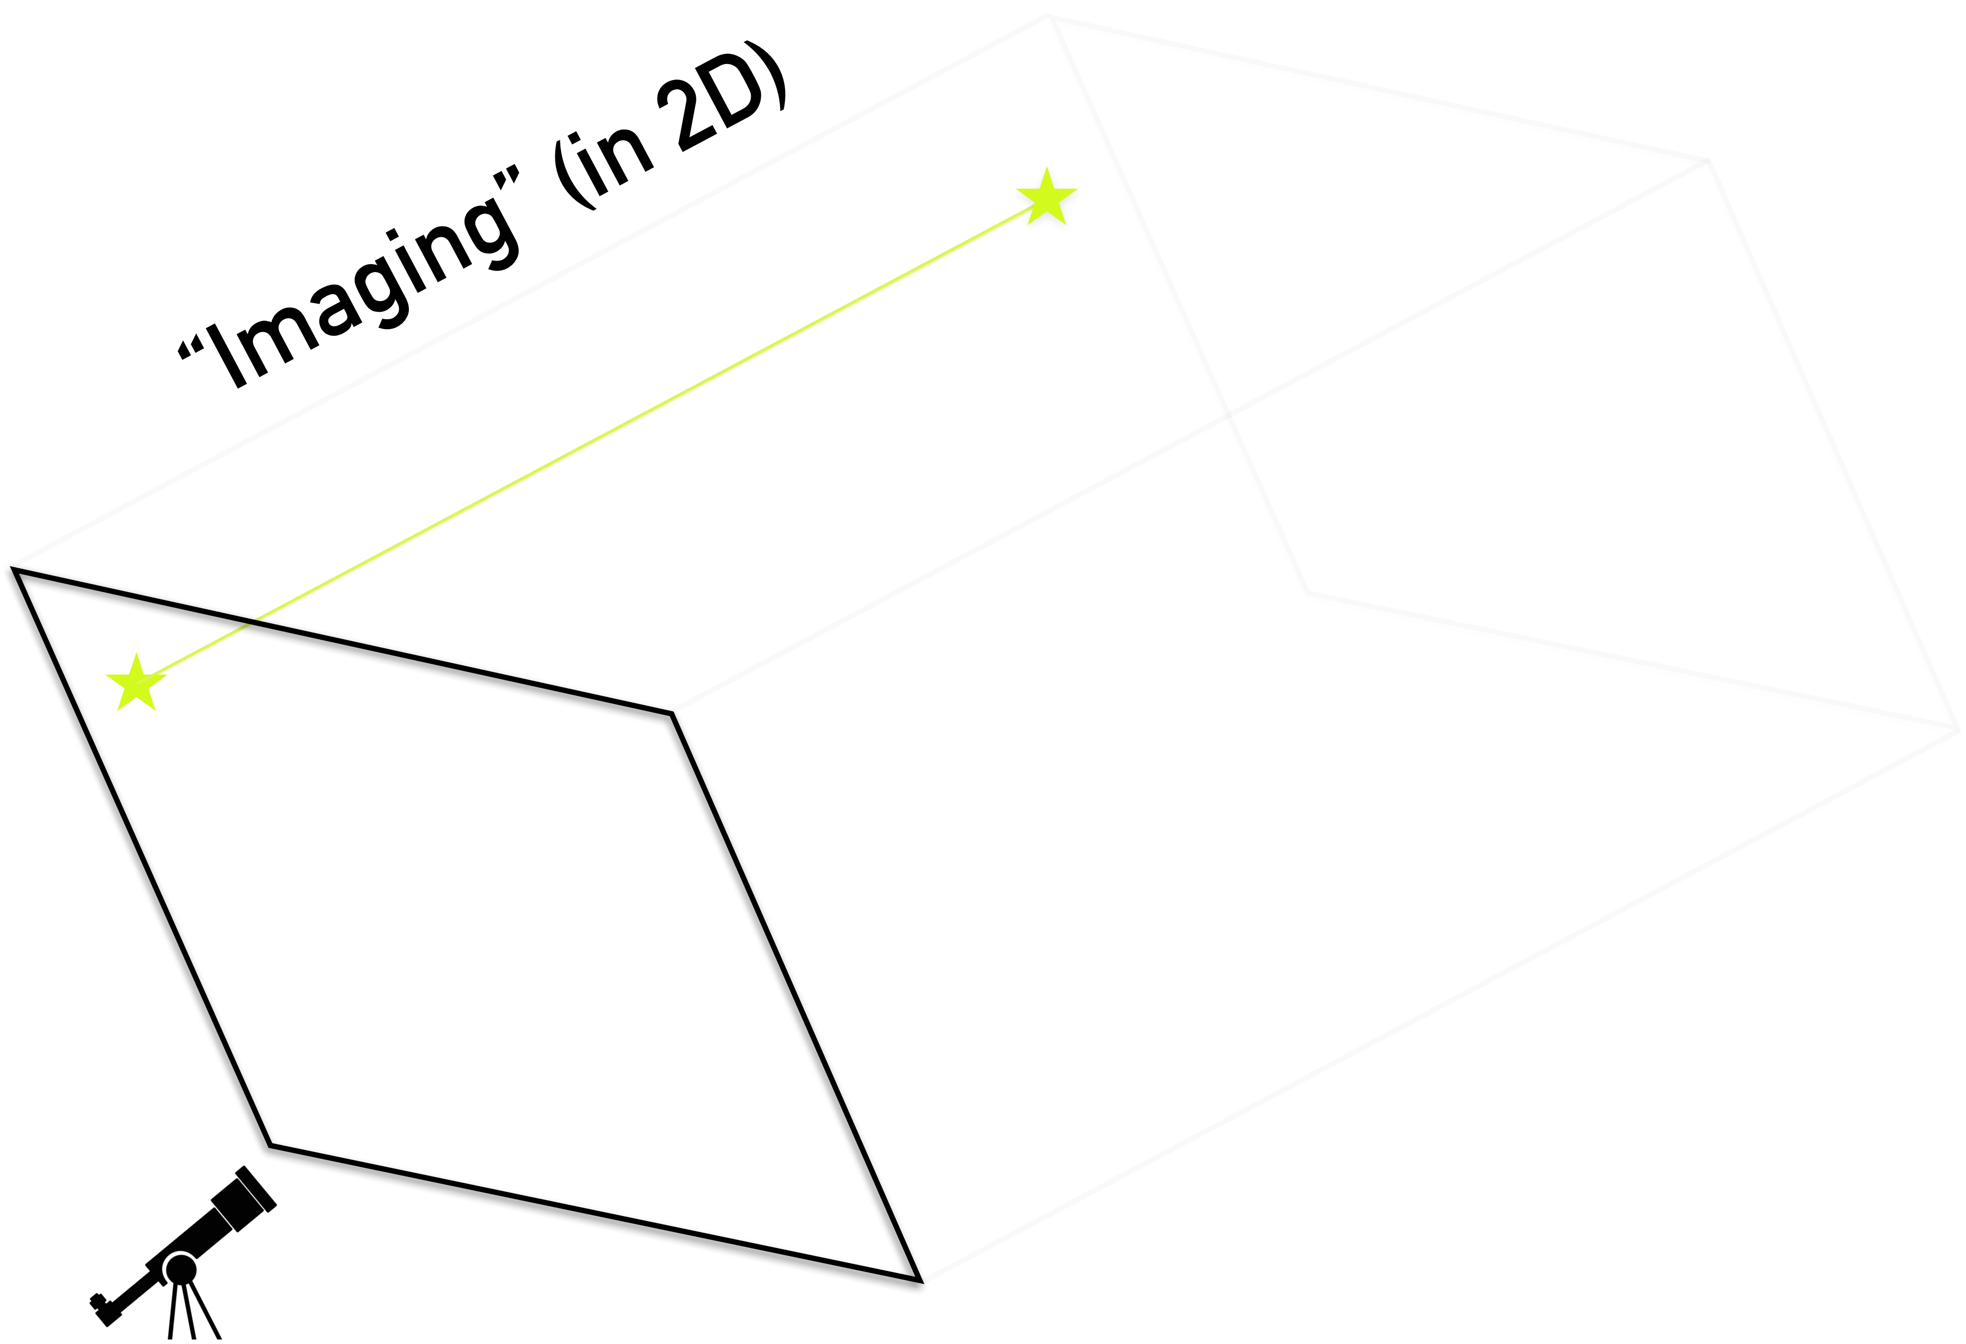

+

See Green+2014, 2015, 2018, 2019 for more details on stellar modeling

WARNING: schematic diagram, NOT to scale (credit A. Goodman, 2019)

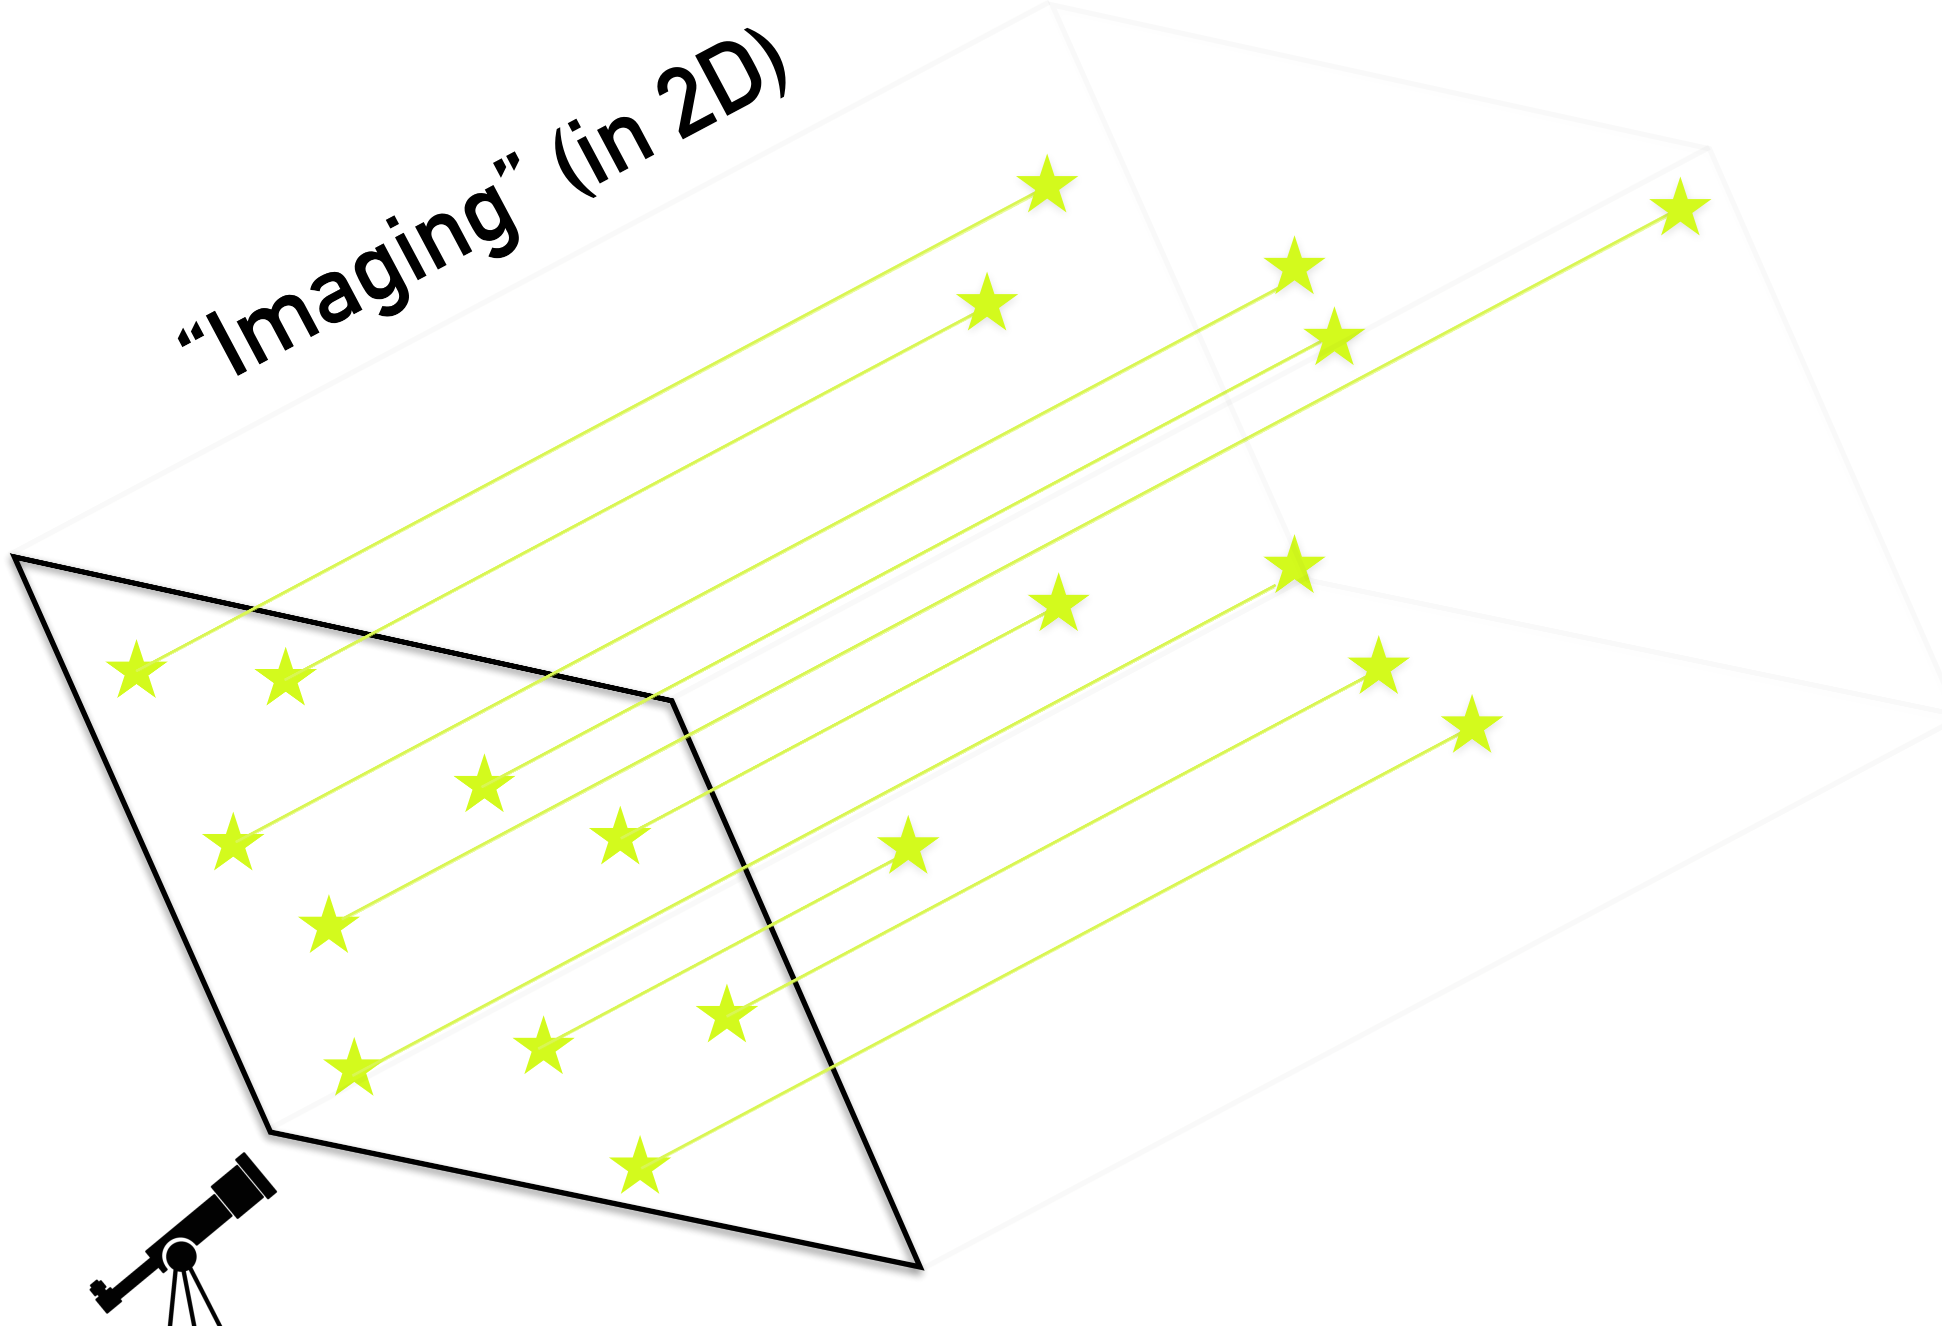

Pan-STARRS DECaPS 2MASS 2 MICRON ALL-SKY SURVEY + gaia

See Green+2014, 2015, 2018, 2019 for more details on stellar modeling

WARNING: schematic diagram, NOT to scale (credit A. Goodman, 2019)

“Imaging” (in 2D)

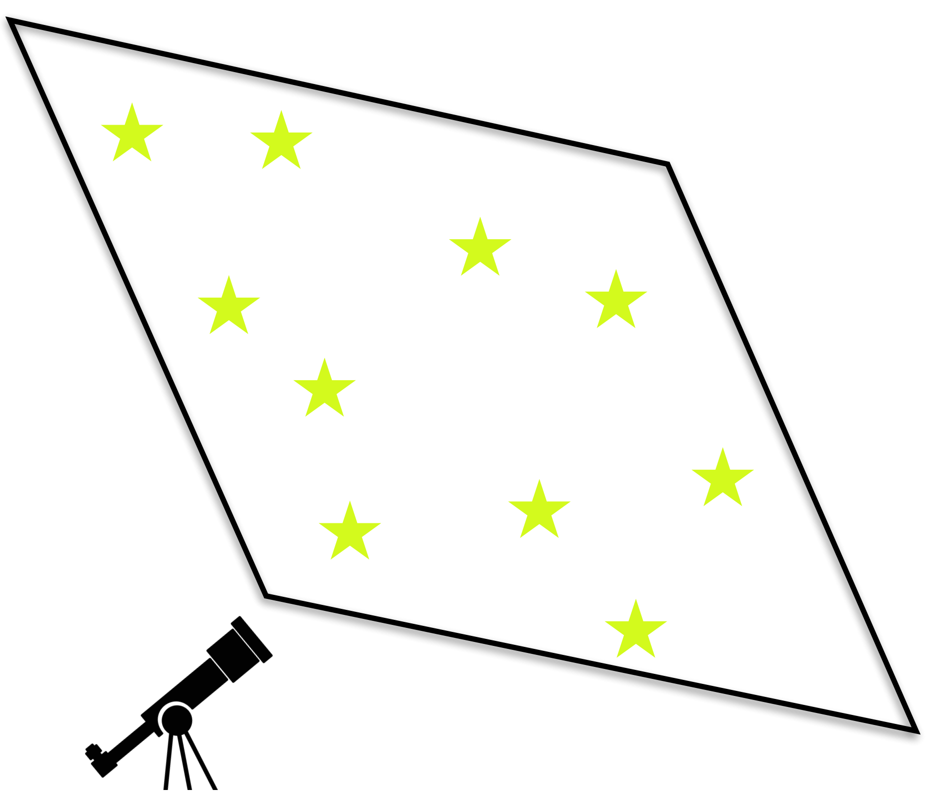

Pan-STARRS DECaPS 2MASS + gaia

See Green+2014, 2015, 2018, 2019 for more details on stellar modeling

WARNING: schematic diagram, NOT to scale (credit A. Goodman, 2019)

Extinction & Reddening, from Color Imaging

Can infer cloud's distance from dust's effects on stars.

Pan-STARRS DECaPS 2MASS 2 MICRON ALL-SKY SURVEY + gaia

Methodology

[From star colors to a 3D map of (at least part of) the Galaxy]

Part I.

Inferring Stellar Properties

How do we measure "star colors"? An introduction to photometry.

Photometry = set of "magnitudes" for each star:

$\mathbf{m} = \{m_u, m_g, m_r, m_i, m_z \ldots\}$

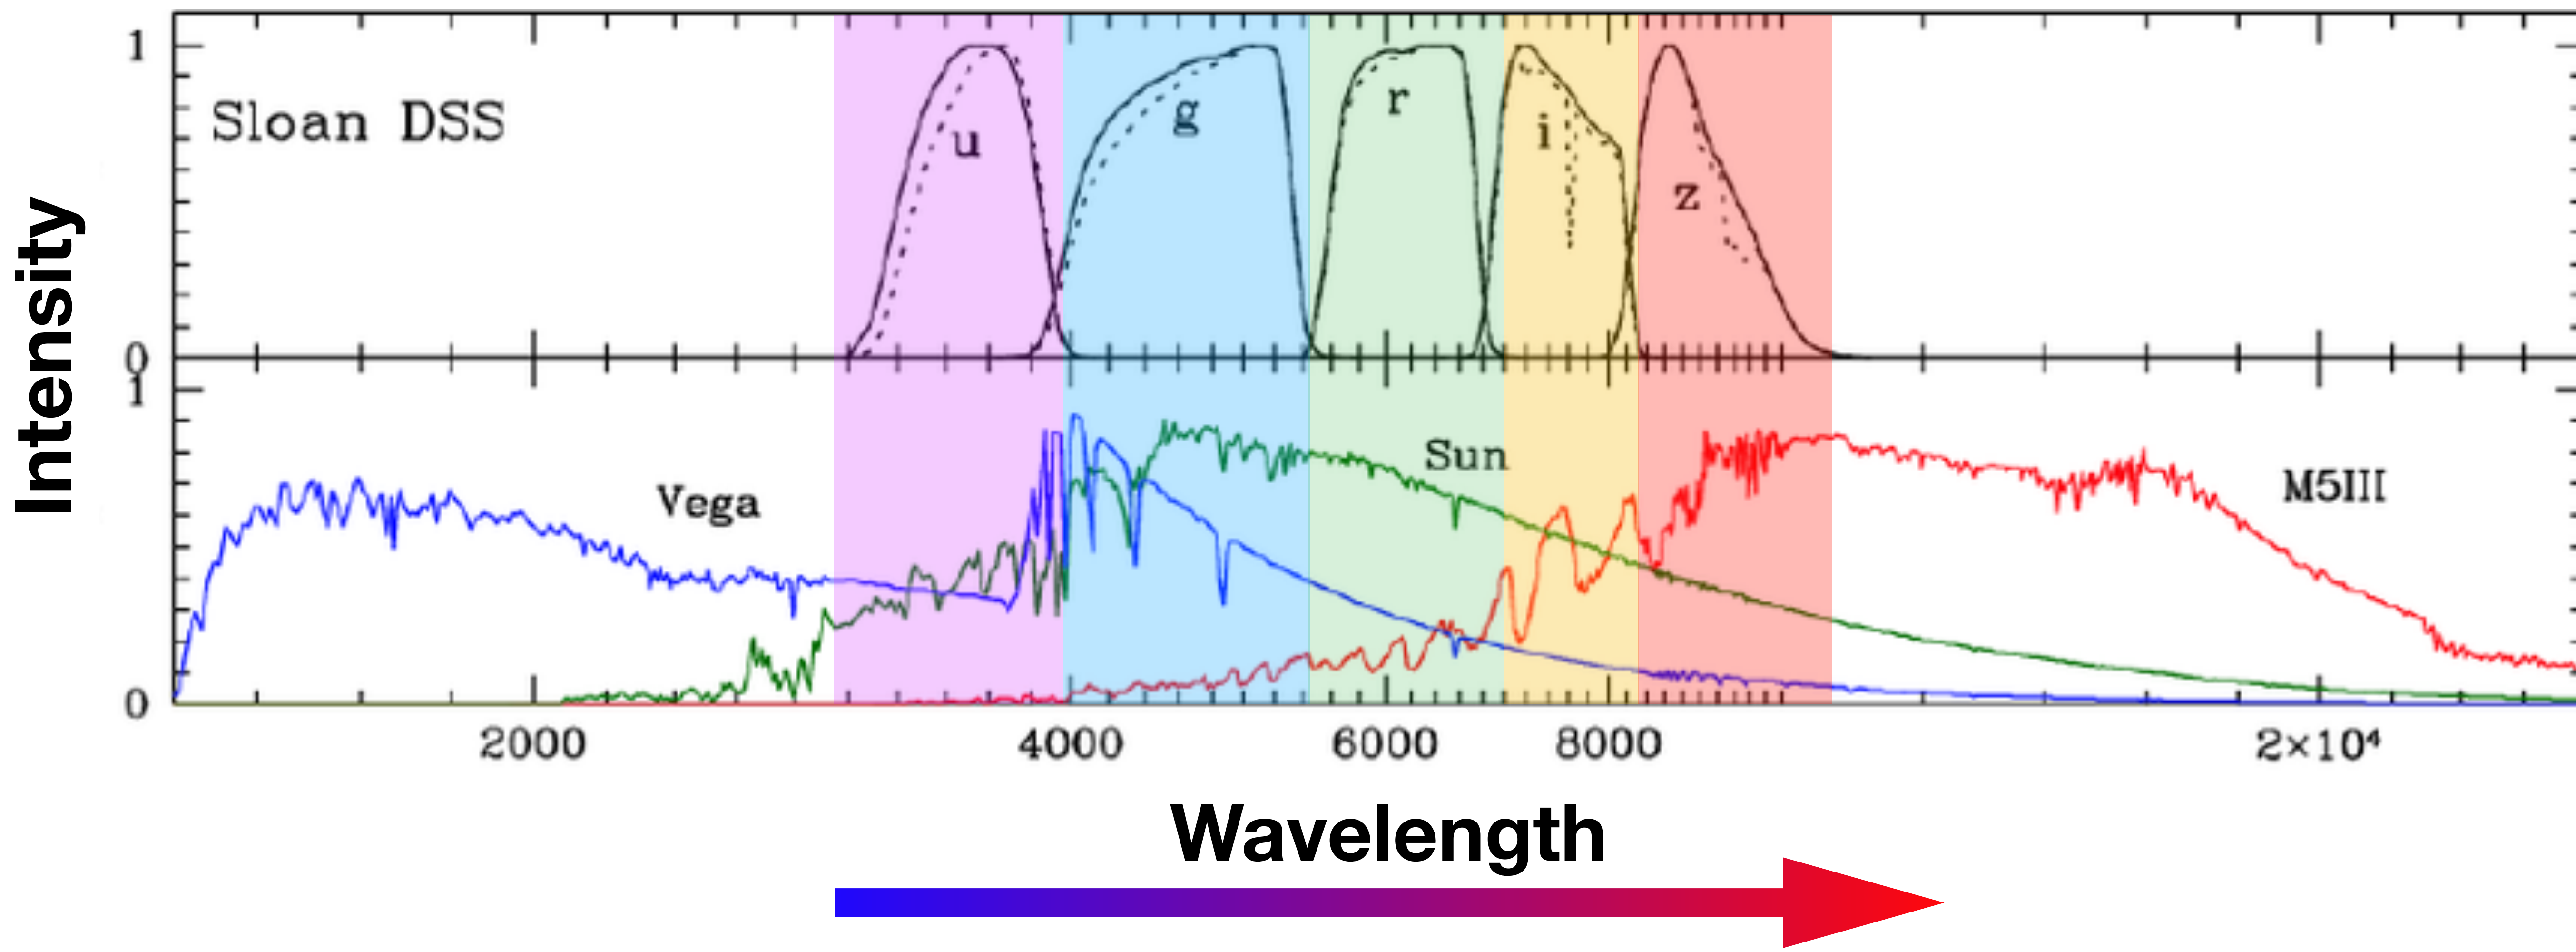

Photometry is available for BILLIONS of stars. Spectroscopy is available for millions.

Basic Idea

"Color-Color Space"

Green et al. (2014)

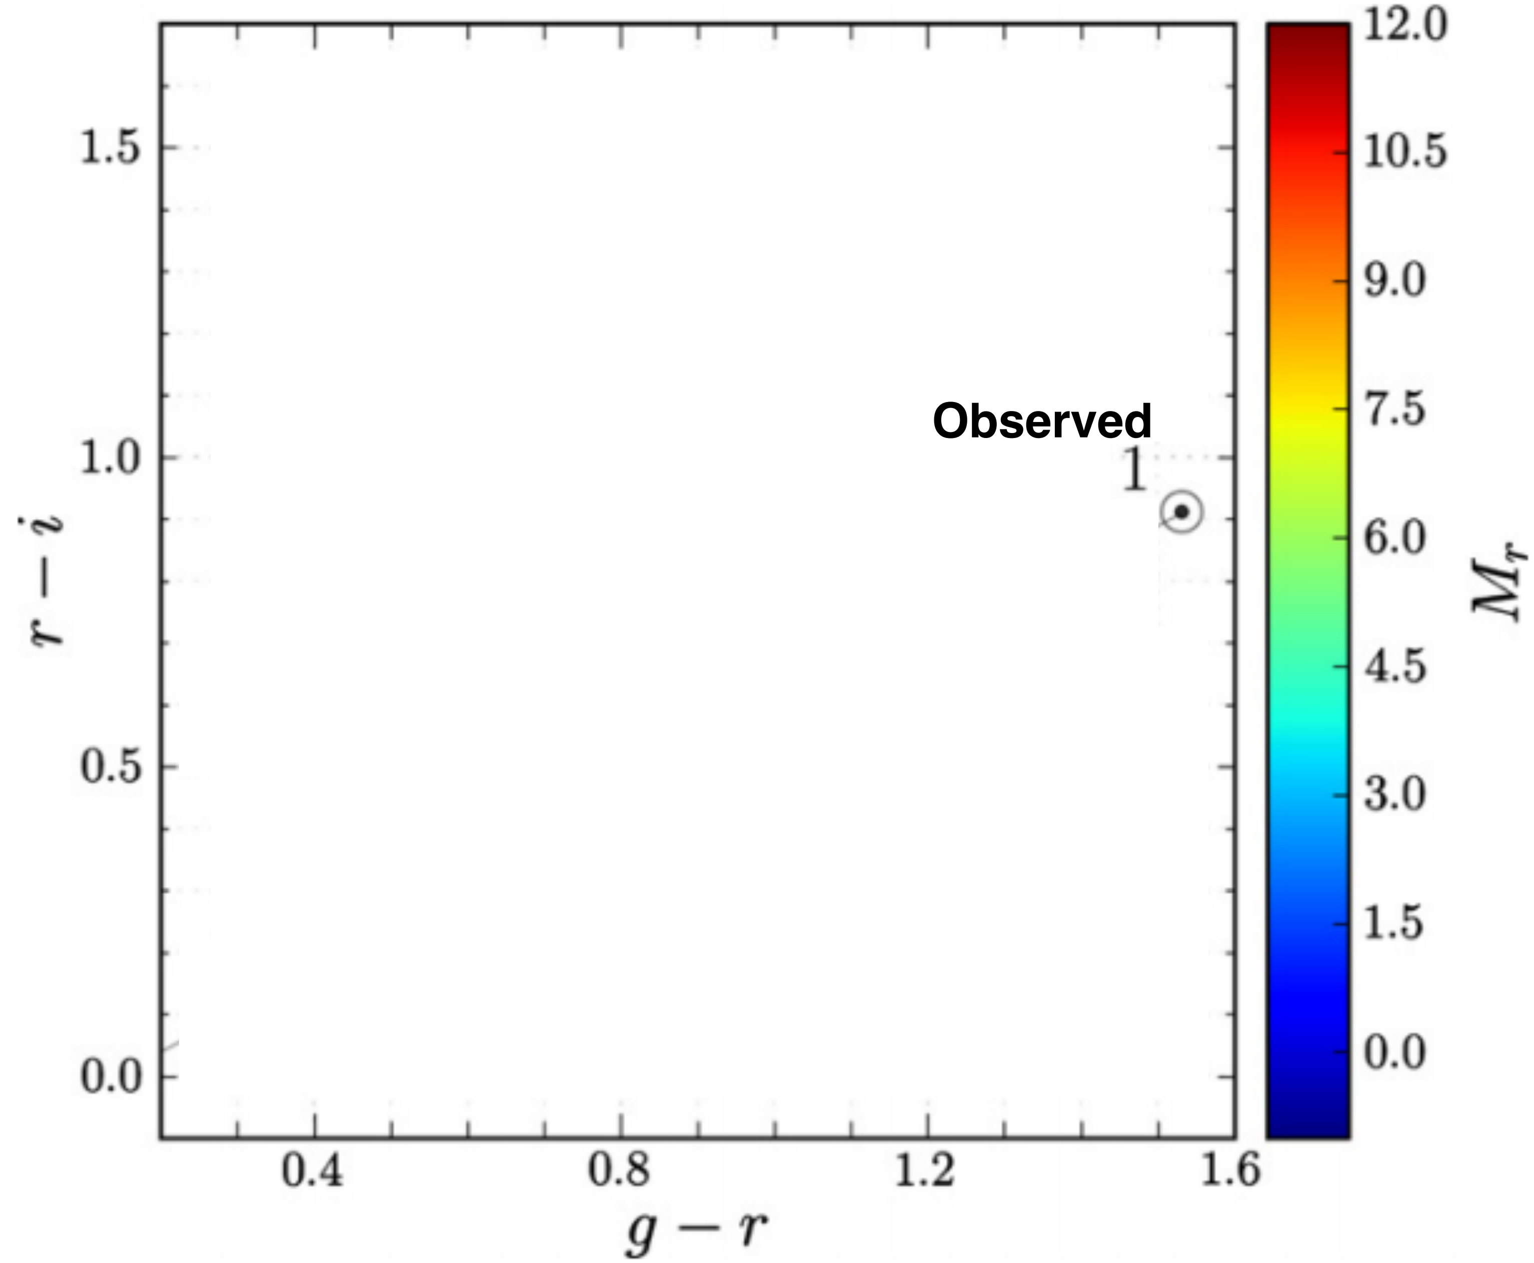

Increasing Mass

Basic Idea

"Color-Color Space"

Green et al. (2014)

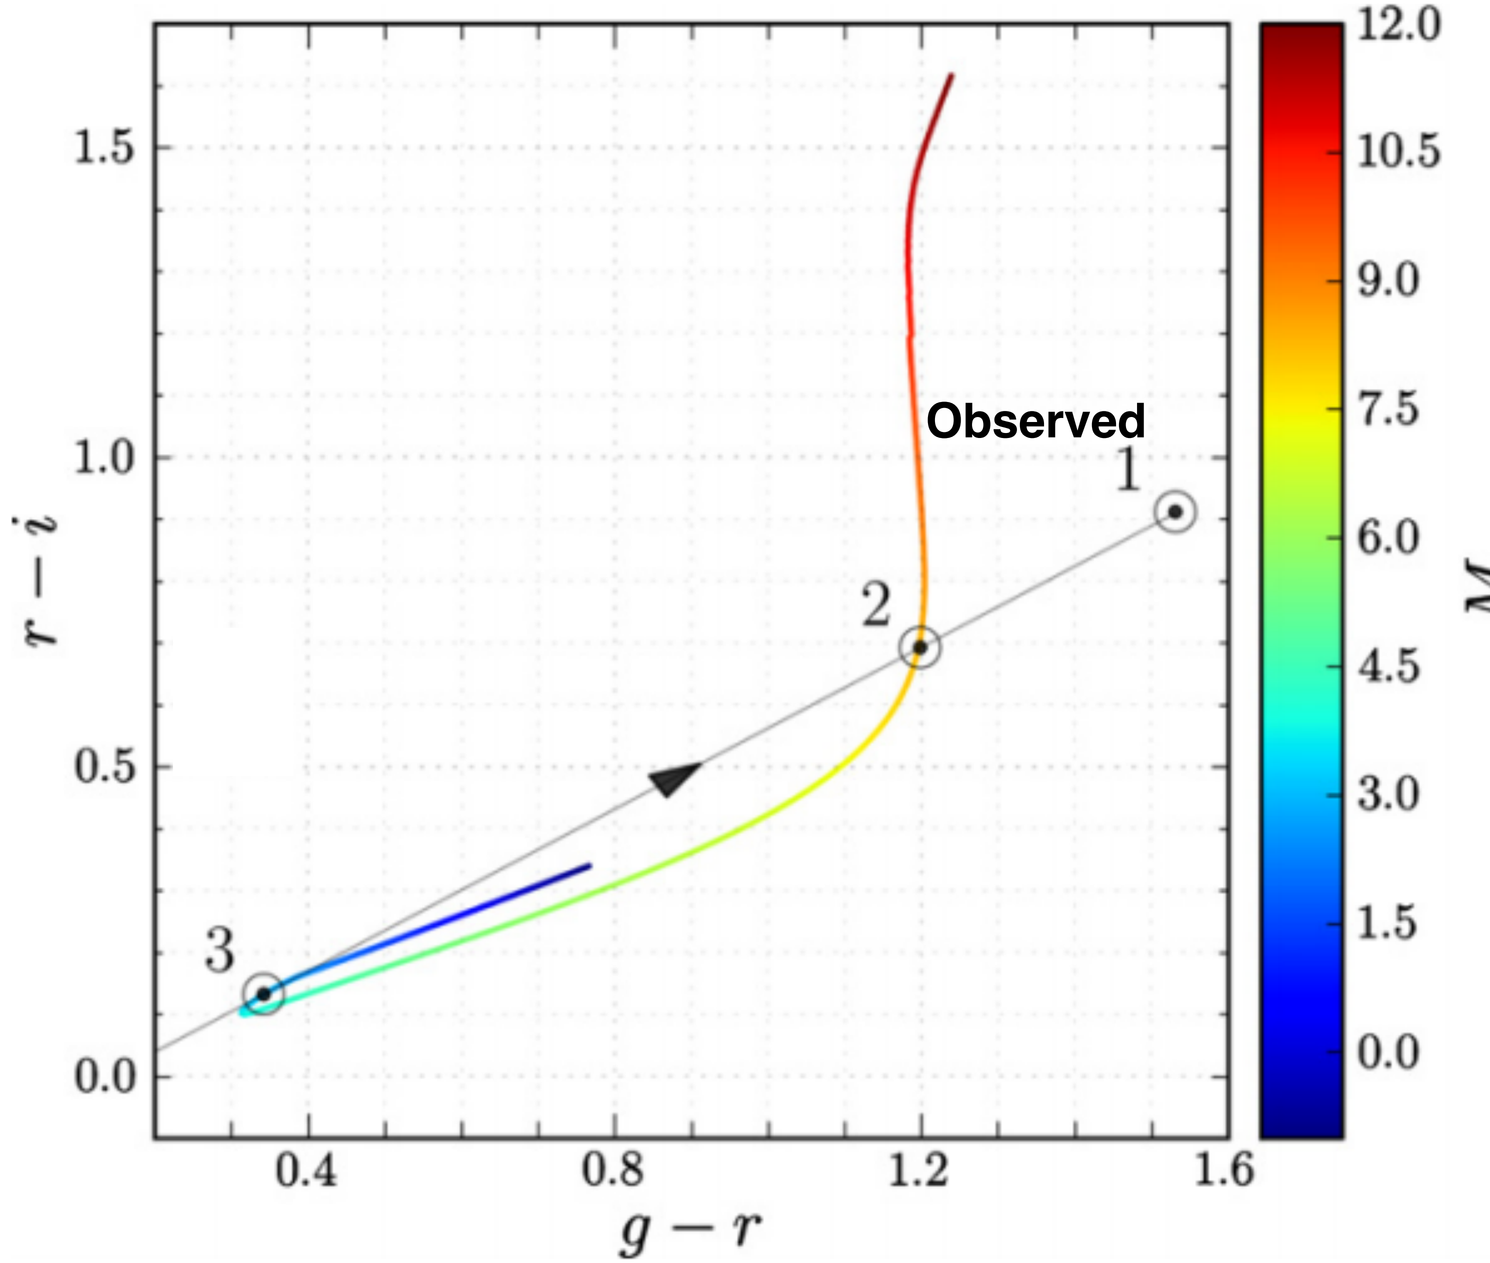

Basic Idea

"Color-Color Space"

Green et al. (2014)

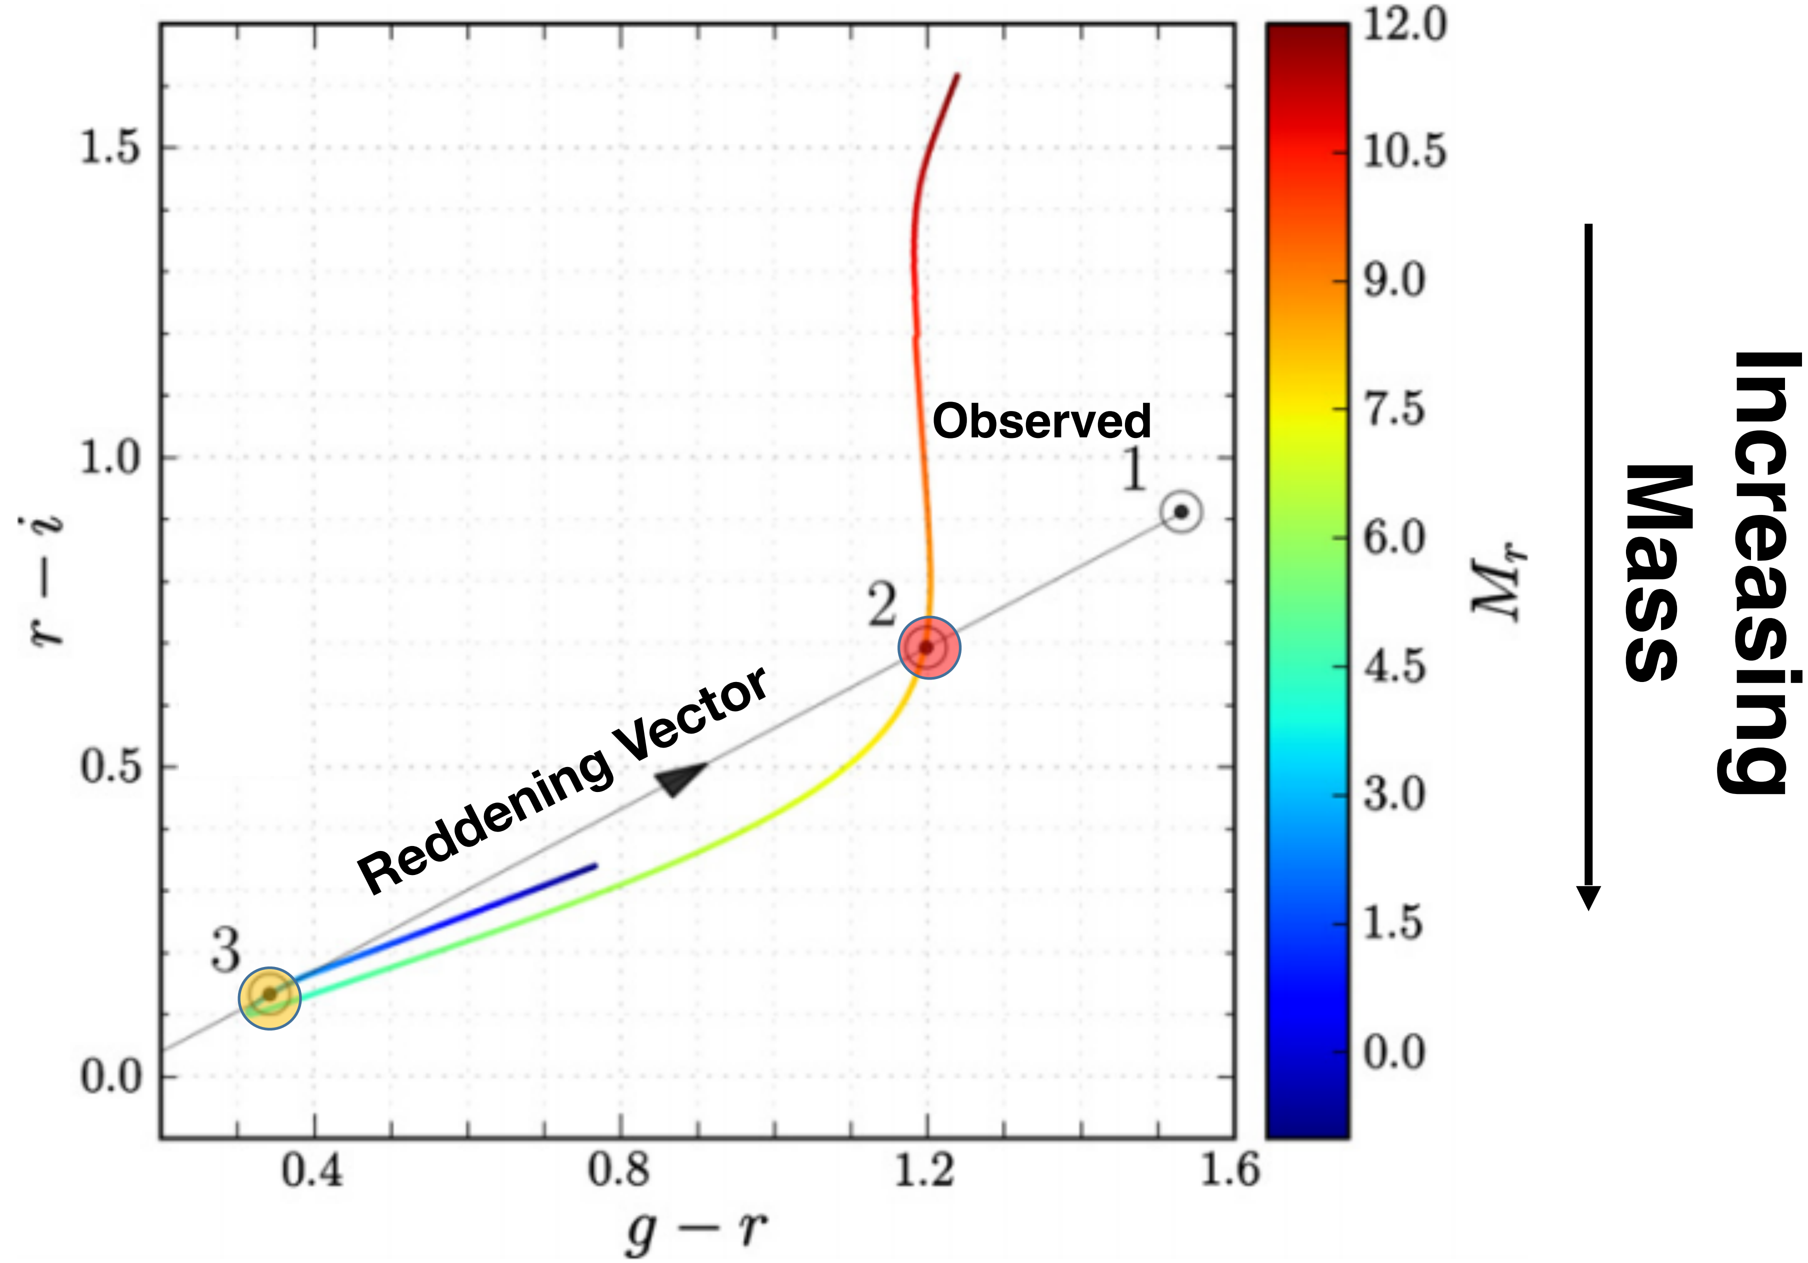

- **Low-mass solution**
  - Less reddening
  - Closer
- **High-mass solution**
  - More reddening
  - Further

Basic Idea

- **Low-mass solution**
  - Less reddening
  - Closer
- **High-mass solution**
  - More reddening
  - Further

"Color-Color Space"

Green et al. (2014)

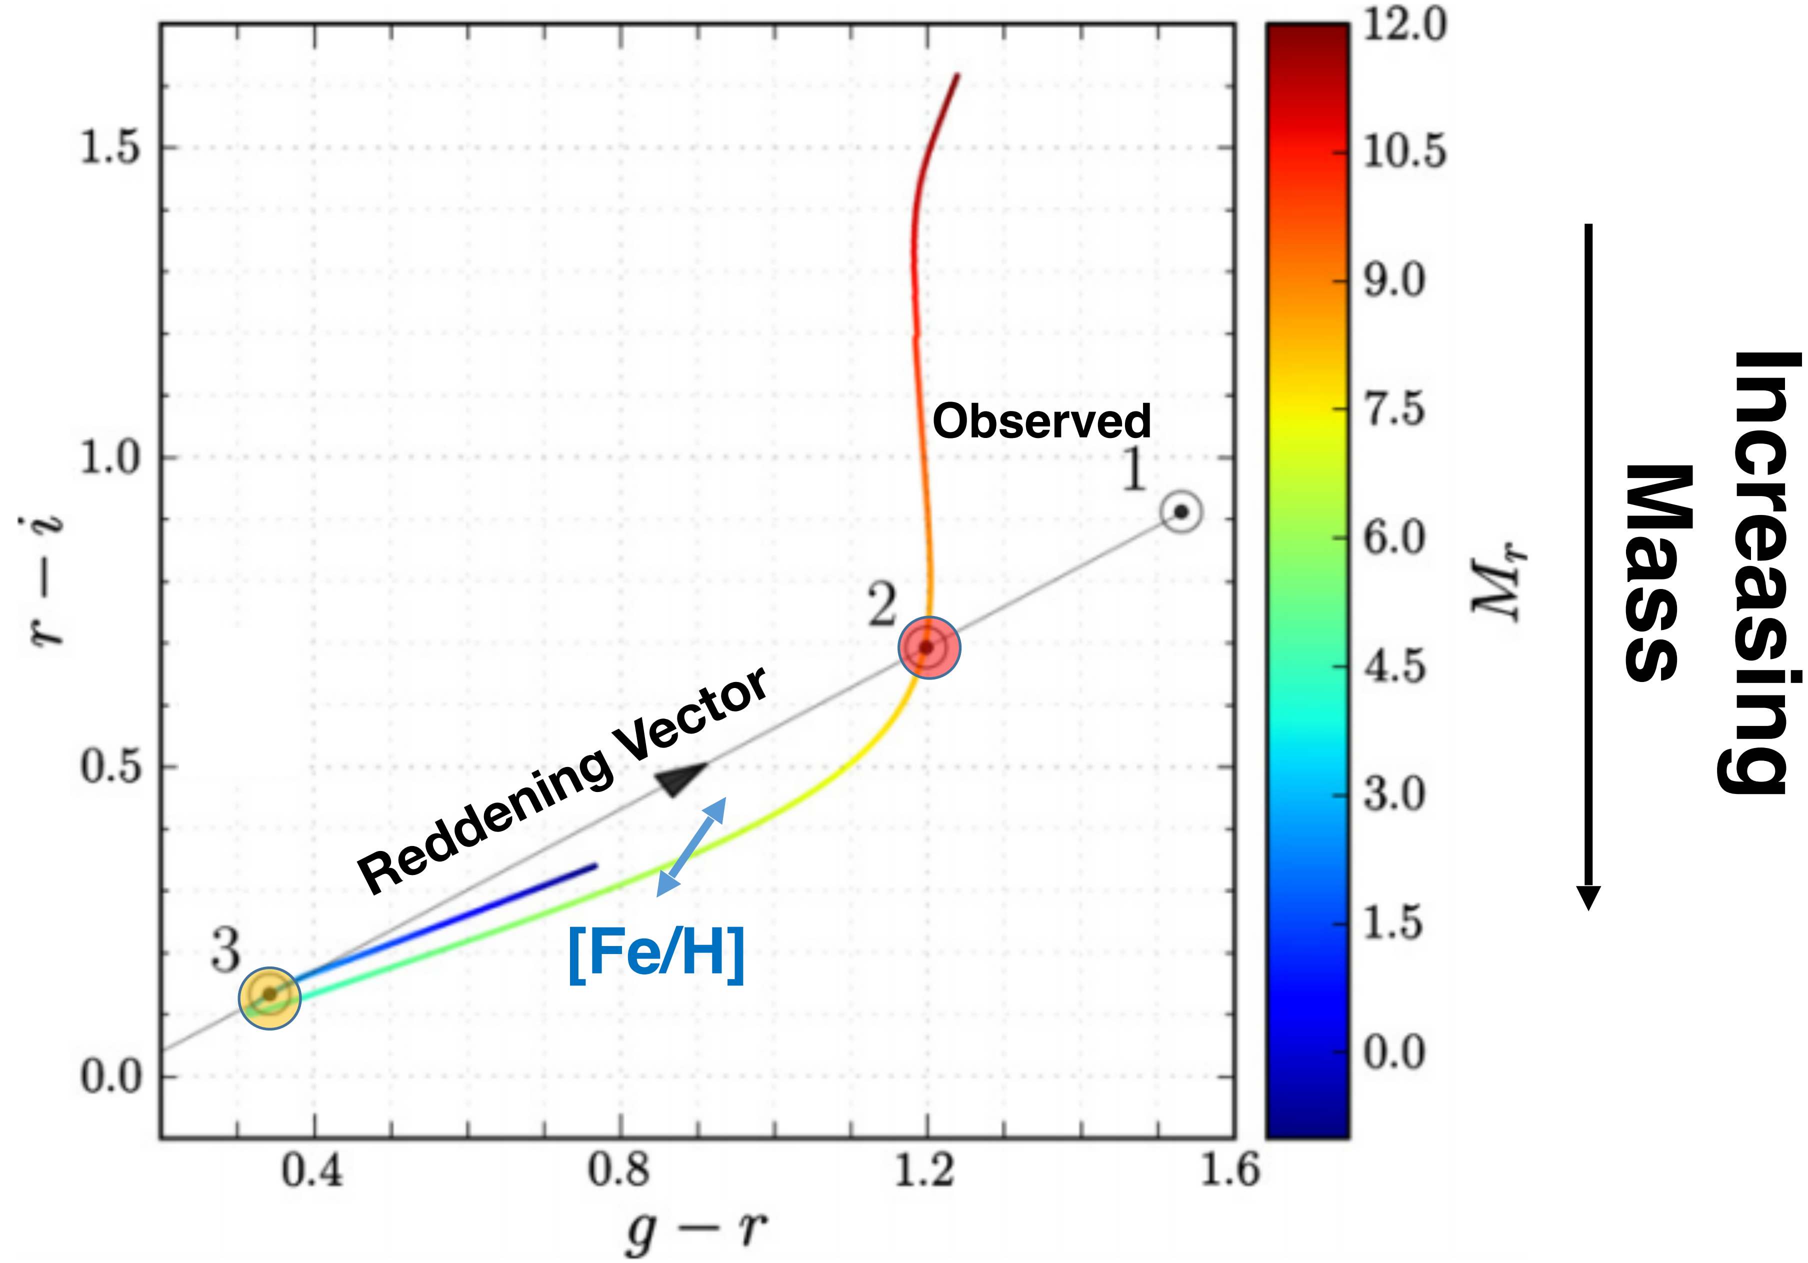

Basic Idea

"Color-Color Space"

Green et al. (2014)

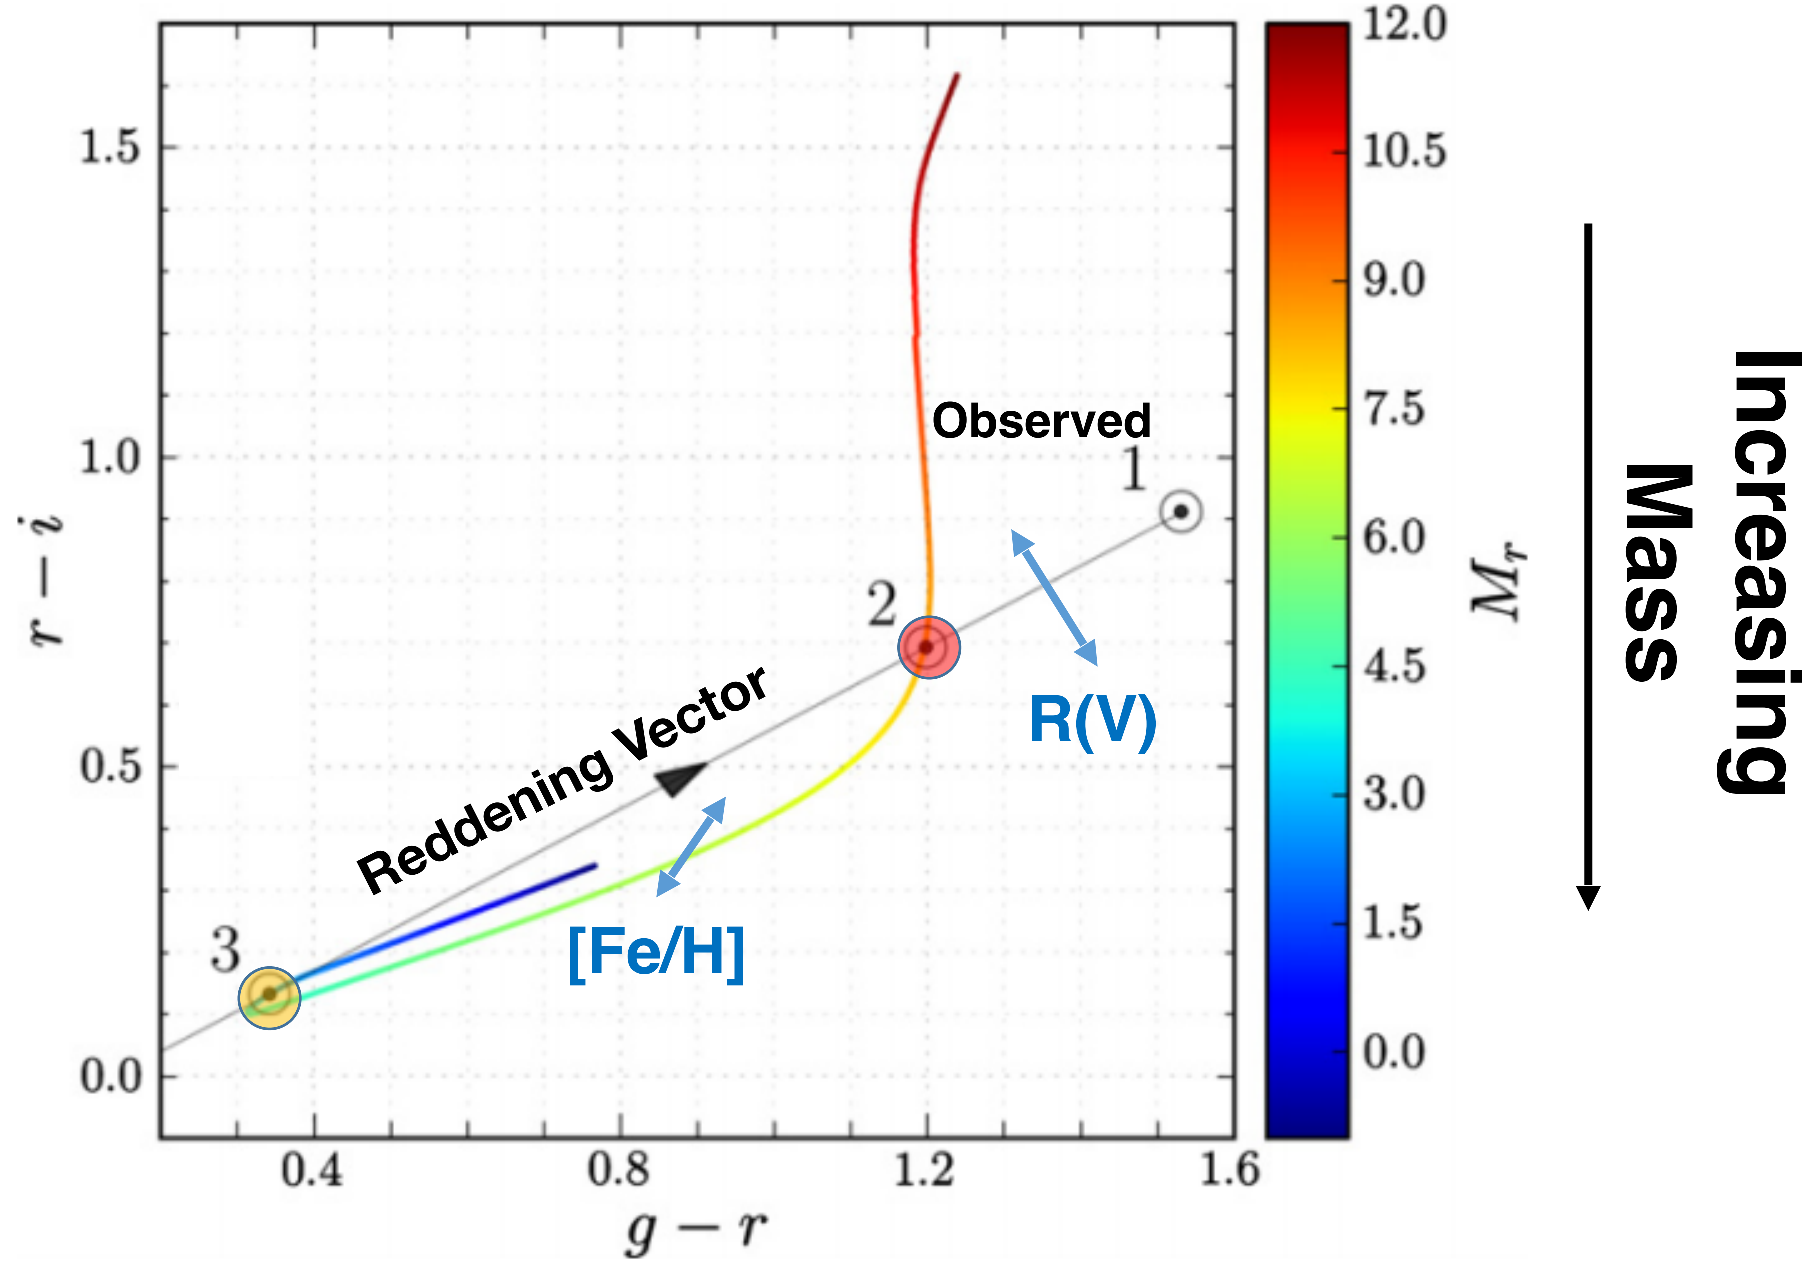

Low-mass solution
- Less reddening
- Closer

High-mass solution
- More reddening
- Further

Stellar Inference

- We model each star as having (predicted) observed magnitudes **m**.

- **m** is a function of:
  1. "type" of star (Mr, FeH)
  2. reddening from dust (A_V, R_V)
  3. distance (μ)

$$\mathbf{m} = \mathbf{m}_{\mathrm{int}}(M_r, [\mathrm{Fe/H}]) + A_V(\mathbf{R} + R_V\mathbf{R}') + \mu$$

- Five-parameter model

$$\mathbf{m}(\boldsymbol{\theta}) \equiv \mathbf{m}(M_r, [\mathrm{Fe/H}], A_V, R_V, \mu)$$

Stellar Inference

The posterior probability that observed magnitudes $\hat{\mathbf{m}}$ are consistent with our predicted photometry $\mathbf{m}(\boldsymbol{\theta})$ follows Bayes Rule:

Posterior

$$P(\boldsymbol{\theta}|\hat{\mathbf{m}}, \hat{\varpi}) \propto \mathcal{L}(\hat{\mathbf{m}}|\boldsymbol{\theta})\, \mathcal{L}(\hat{\varpi}|\mu)\, \pi(\boldsymbol{\theta})$$

$\boldsymbol{\theta}$: stellar type, reddening, distance

Stellar Inference

The posterior probability that observed magnitudes $\hat{\mathbf{m}}$ are consistent with our predicted photometry $\mathbf{m}(\boldsymbol{\theta})$ follows Bayes Rule:

Photometric Likelihood

$$P(\boldsymbol{\theta}|\hat{\mathbf{m}}, \hat{\varpi}) \propto \boxed{\mathcal{L}(\hat{\mathbf{m}}|\boldsymbol{\theta})}\, \mathcal{L}(\hat{\varpi}|\mu)\, \pi(\boldsymbol{\theta})$$

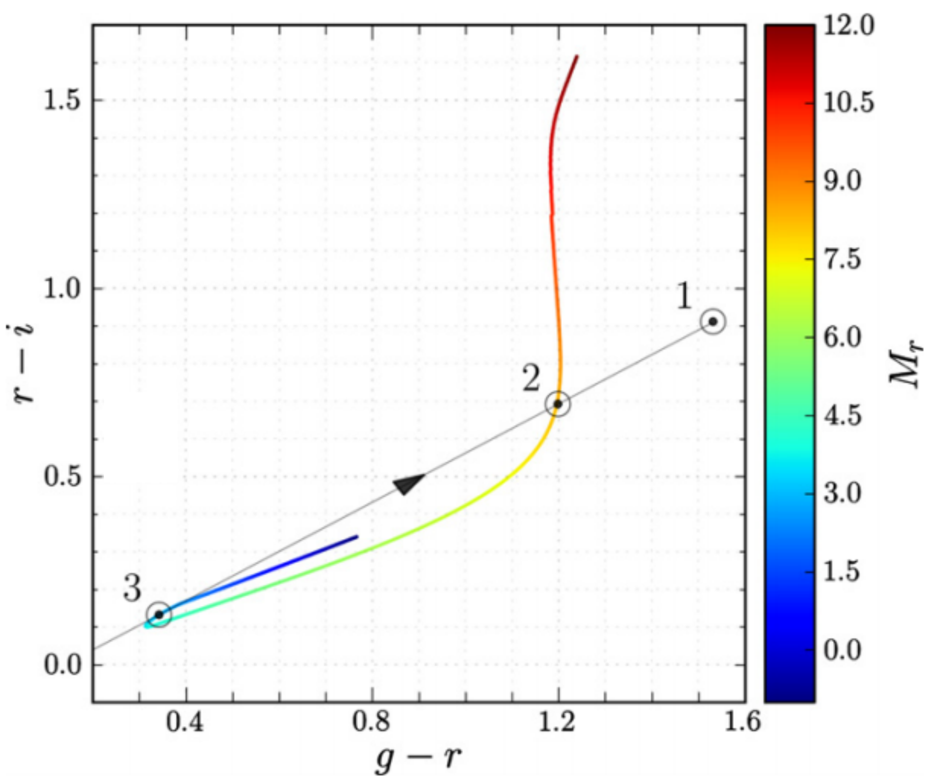

$$\mathcal{L}(\hat{\mathbf{m}}|\boldsymbol{\theta}) = \prod_b \frac{1}{\sqrt{2\pi}\sigma_b} \exp\left[-\frac{1}{2}\frac{(\mathbf{m}(\boldsymbol{\theta}) - \hat{\mathbf{m}})^2}{\sigma_b^2}\right]$$

Stellar Inference

The posterior probability that observed magnitudes $\hat{\mathbf{m}}$ are consistent with our predicted magnitudes $\mathbf{m}(\boldsymbol{\theta})$ follows Bayes Rule:

Priors

$$P(\boldsymbol{\theta}|\hat{\mathbf{m}}, \hat{\varpi}) \propto \mathcal{L}(\hat{\mathbf{m}}|\boldsymbol{\theta})\, \mathcal{L}(\hat{\varpi}|\mu)\, \pi(\boldsymbol{\theta})$$

3-D Galactic model

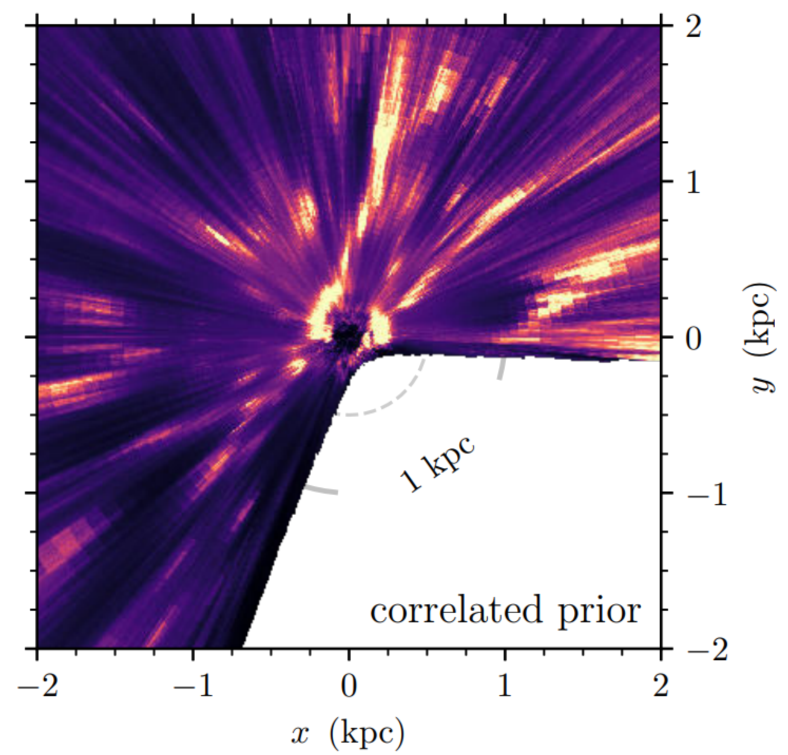

Stellar Inference

The posterior probability that observed magnitudes $\hat{\mathbf{m}}$ are consistent with our predicted photometry $\mathbf{m}(\boldsymbol{\theta})$ follows Bayes Rule:

Parallax Likelihood

$$P(\boldsymbol{\theta}|\hat{\mathbf{m}}, \hat{\varpi}) \propto \mathcal{L}(\hat{\mathbf{m}}|\boldsymbol{\theta}) \boxed{\mathcal{L}(\hat{\varpi}|\mu)} \pi(\boldsymbol{\theta})$$

Gaia DR2 Parallax Measurements

$$\mathcal{L}(\hat{\varpi}|\mu) = \frac{1}{\sqrt{2\pi}\sigma_{\varpi}} \exp\left[-\frac{1}{2}\frac{(\varpi(\mu) - \hat{\varpi})^2}{\sigma_{\varpi}^2}\right]$$

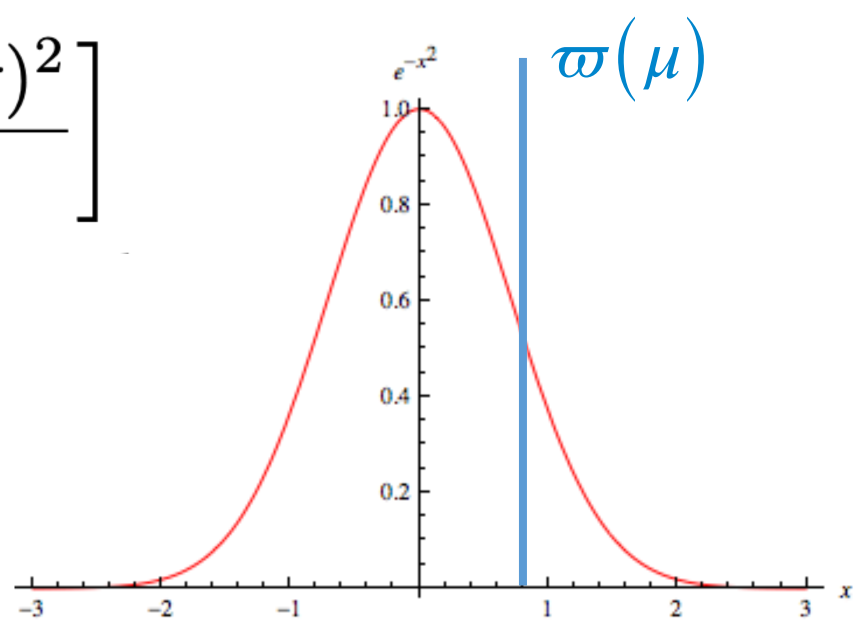

What is parallax?

Stellar Inference

The posterior probability that observed magnitudes $\hat{\mathbf{m}}$ are consistent with our predicted photometry $\mathbf{m}(\boldsymbol{\theta})$ follows Bayes Rule:

Posterior		Photometric Likelihood	Parallax Likelihood	Priors
$P(\boldsymbol{\theta}\|\hat{\mathbf{m}}, \hat{\varpi})$	$\propto$	$\mathcal{L}(\hat{\mathbf{m}}\|\boldsymbol{\theta})$	$\mathcal{L}(\hat{\varpi}\|\mu)$	$\pi(\boldsymbol{\theta})$

$\boldsymbol{\Theta}$: stellar type, reddening, distance

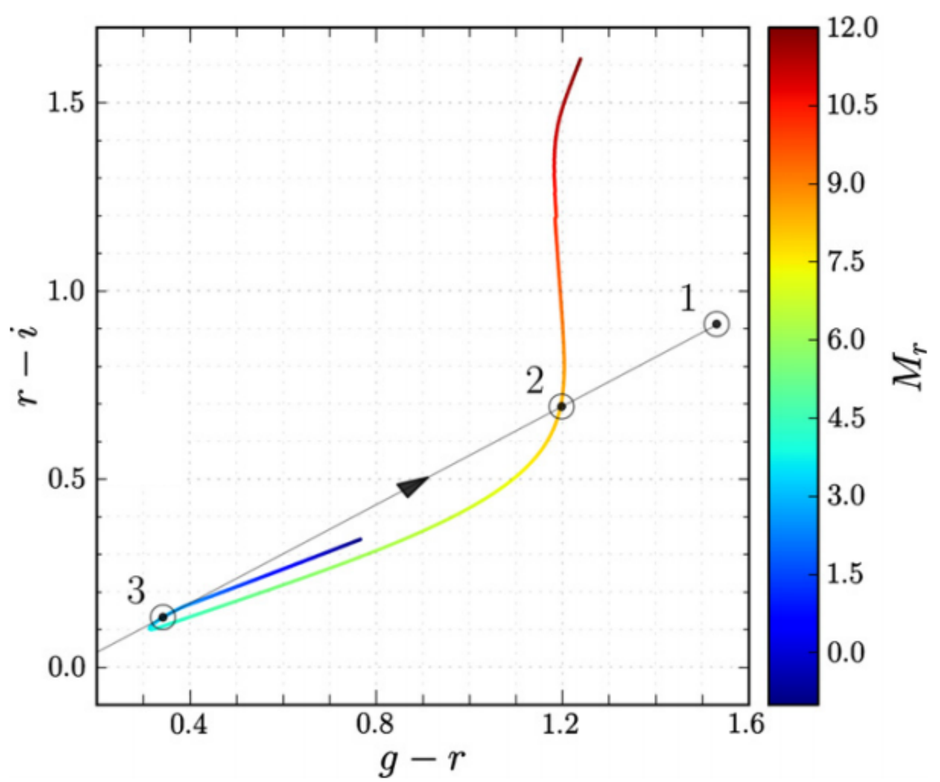

Gaia DR2 Parallax Measurements

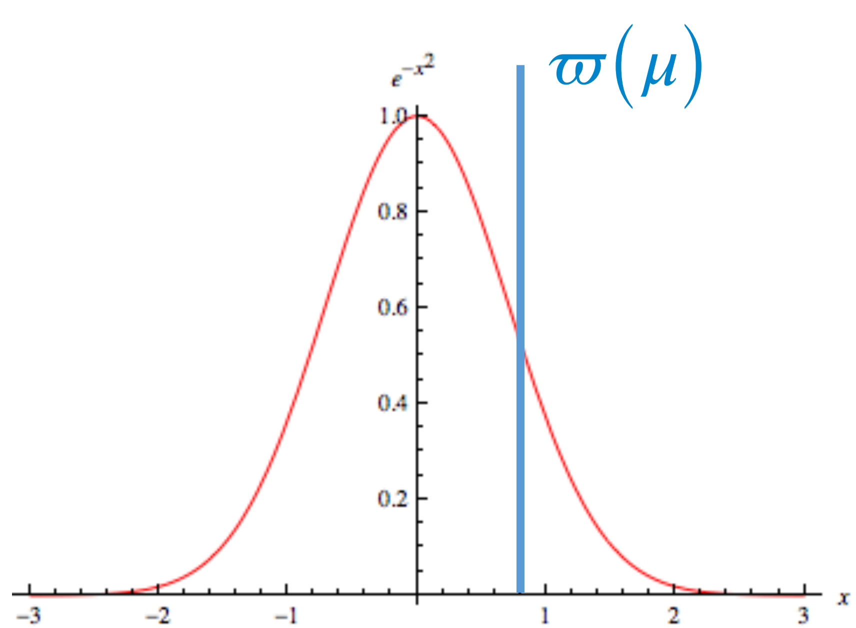

3-D Galactic model

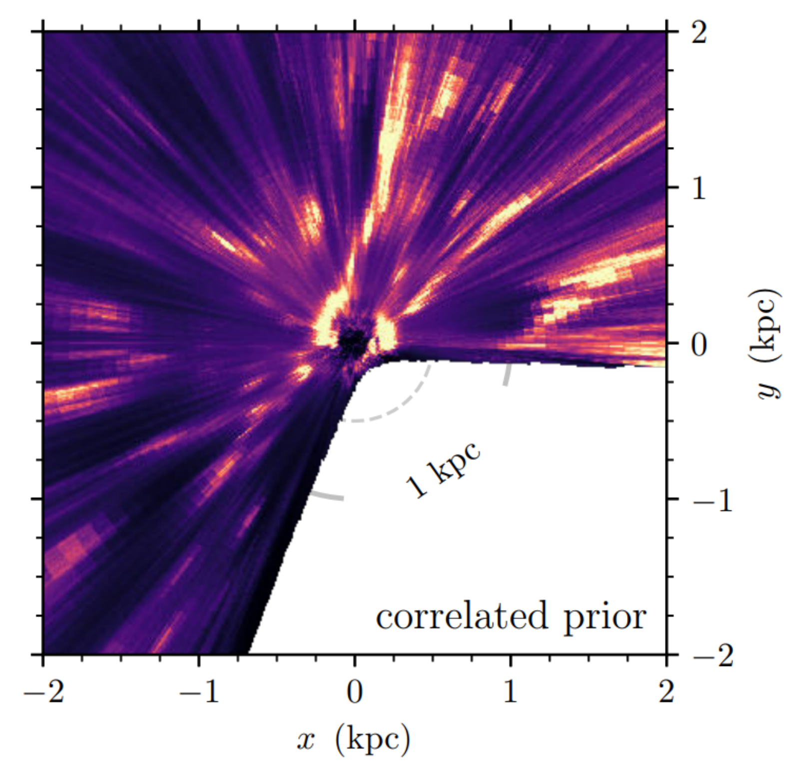

Stellar Inference

To derive stellar posteriors we adopt a grid-based approach over a sampling approach. Why?

- Multiple, widely separated solutions
- Posteriors have extended & complex degeneracies (need more samples and/or longer run times than MC methods)

Speagle et al. (2020a), in prep.

brutus

Et tu, Brute?

`brutus` is a Pure Python package whose core modules involve using "brute force" Bayesian inference to derive distances, reddenings, and stellar properties from photometry using a grid of stellar models.

The package is designed to be highly modular, with current modules including utilities for modeling individual stars, co-eval stellar associations, and stellar-based 3-D dust mapping.

Zucker & Speagle et al. (2019)

Per-Star Distance-Dust Posteriors

We marginalize over stellar "type" to get posteriors on distance + dust for individual stars

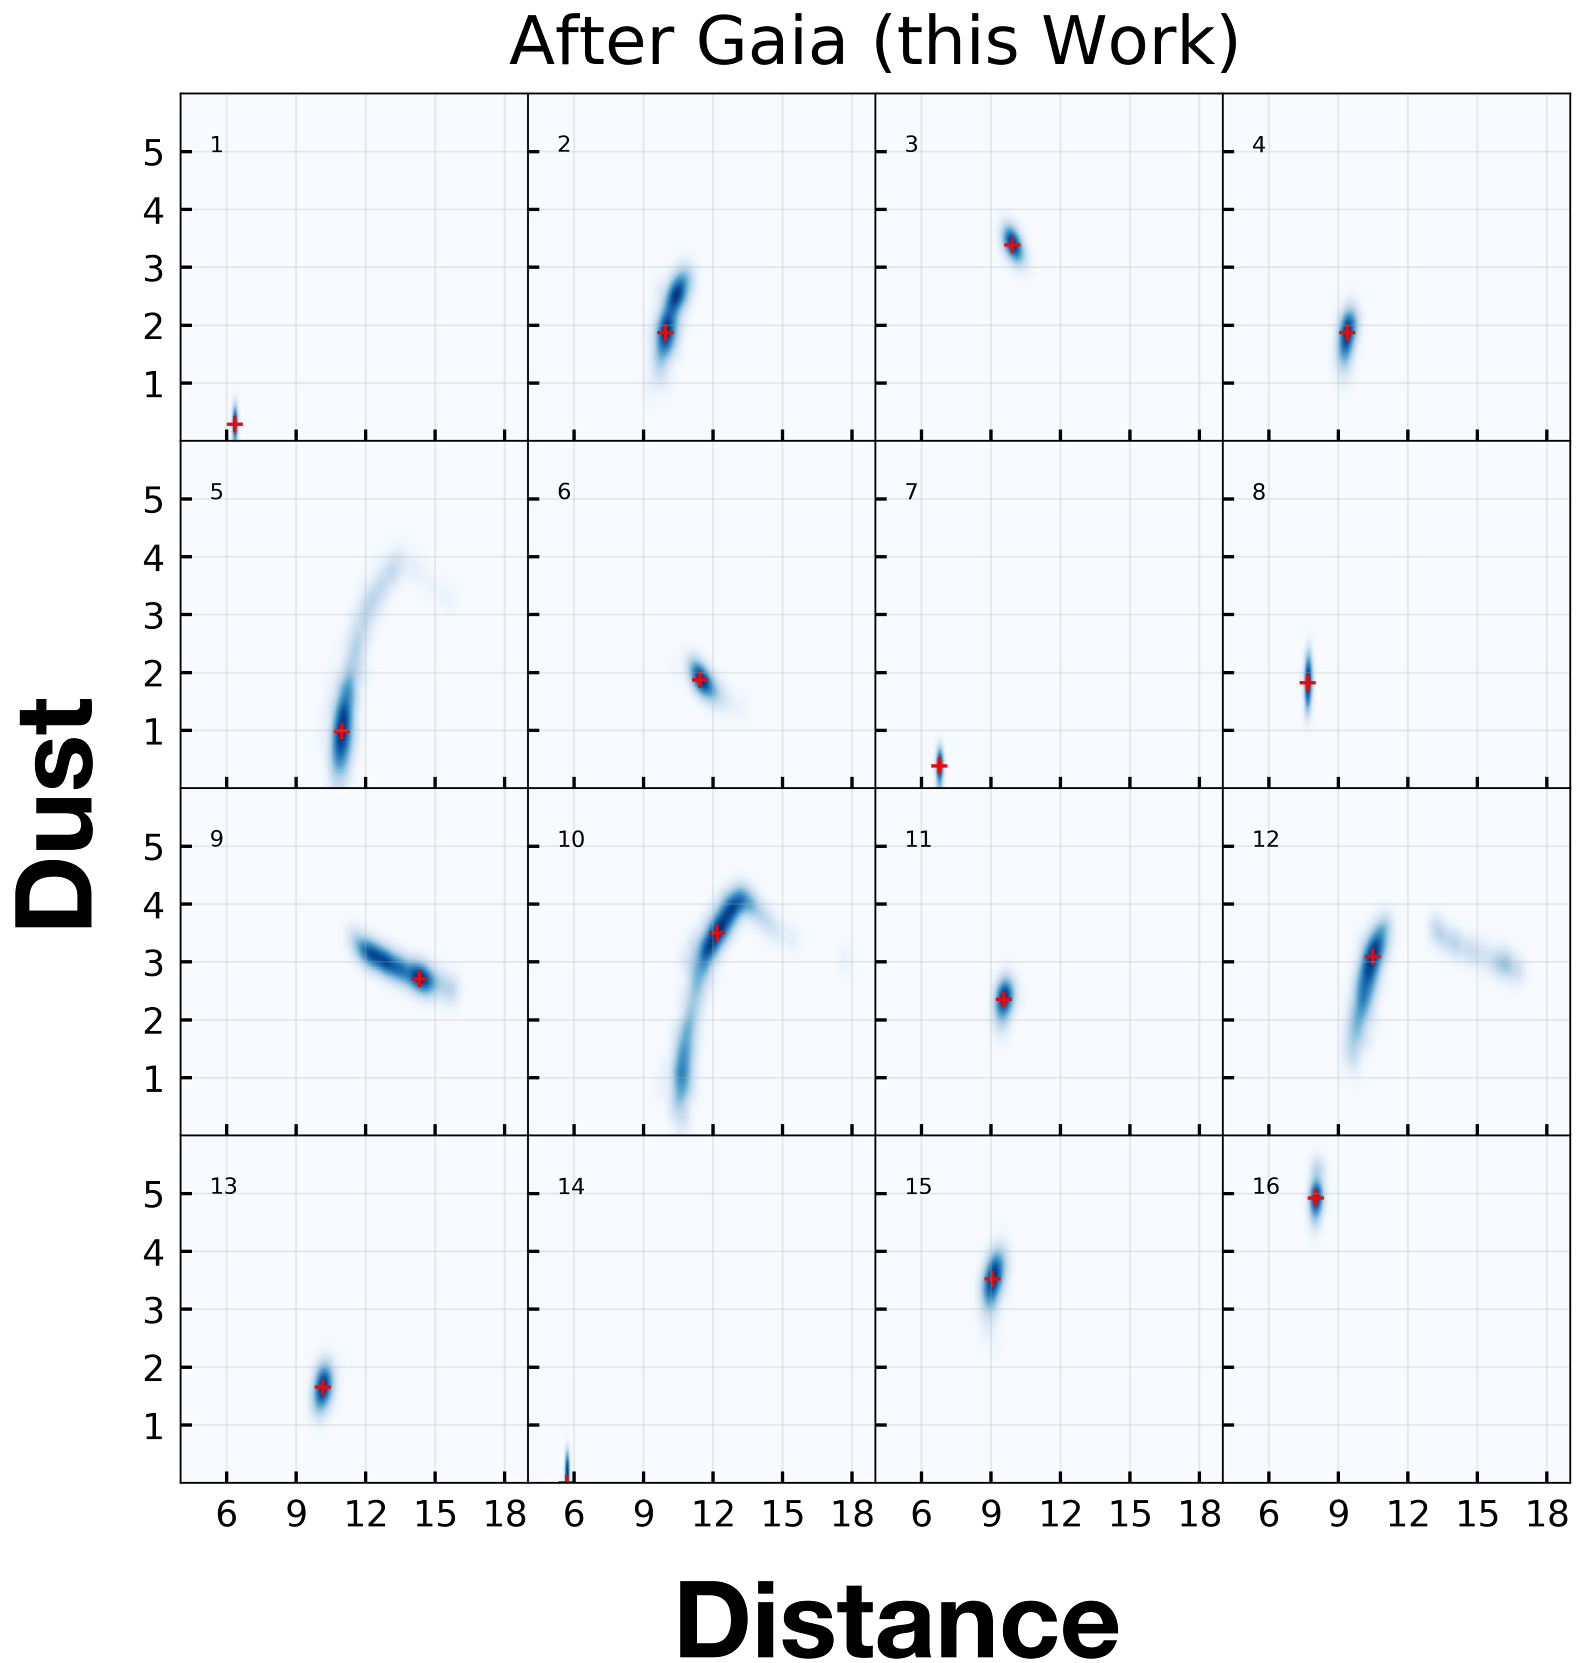

Part 2.

Inferring 3D Dust Cloud Distributions

Remember this Nebula?

We know where it lies on the sky…
But we need its distance.

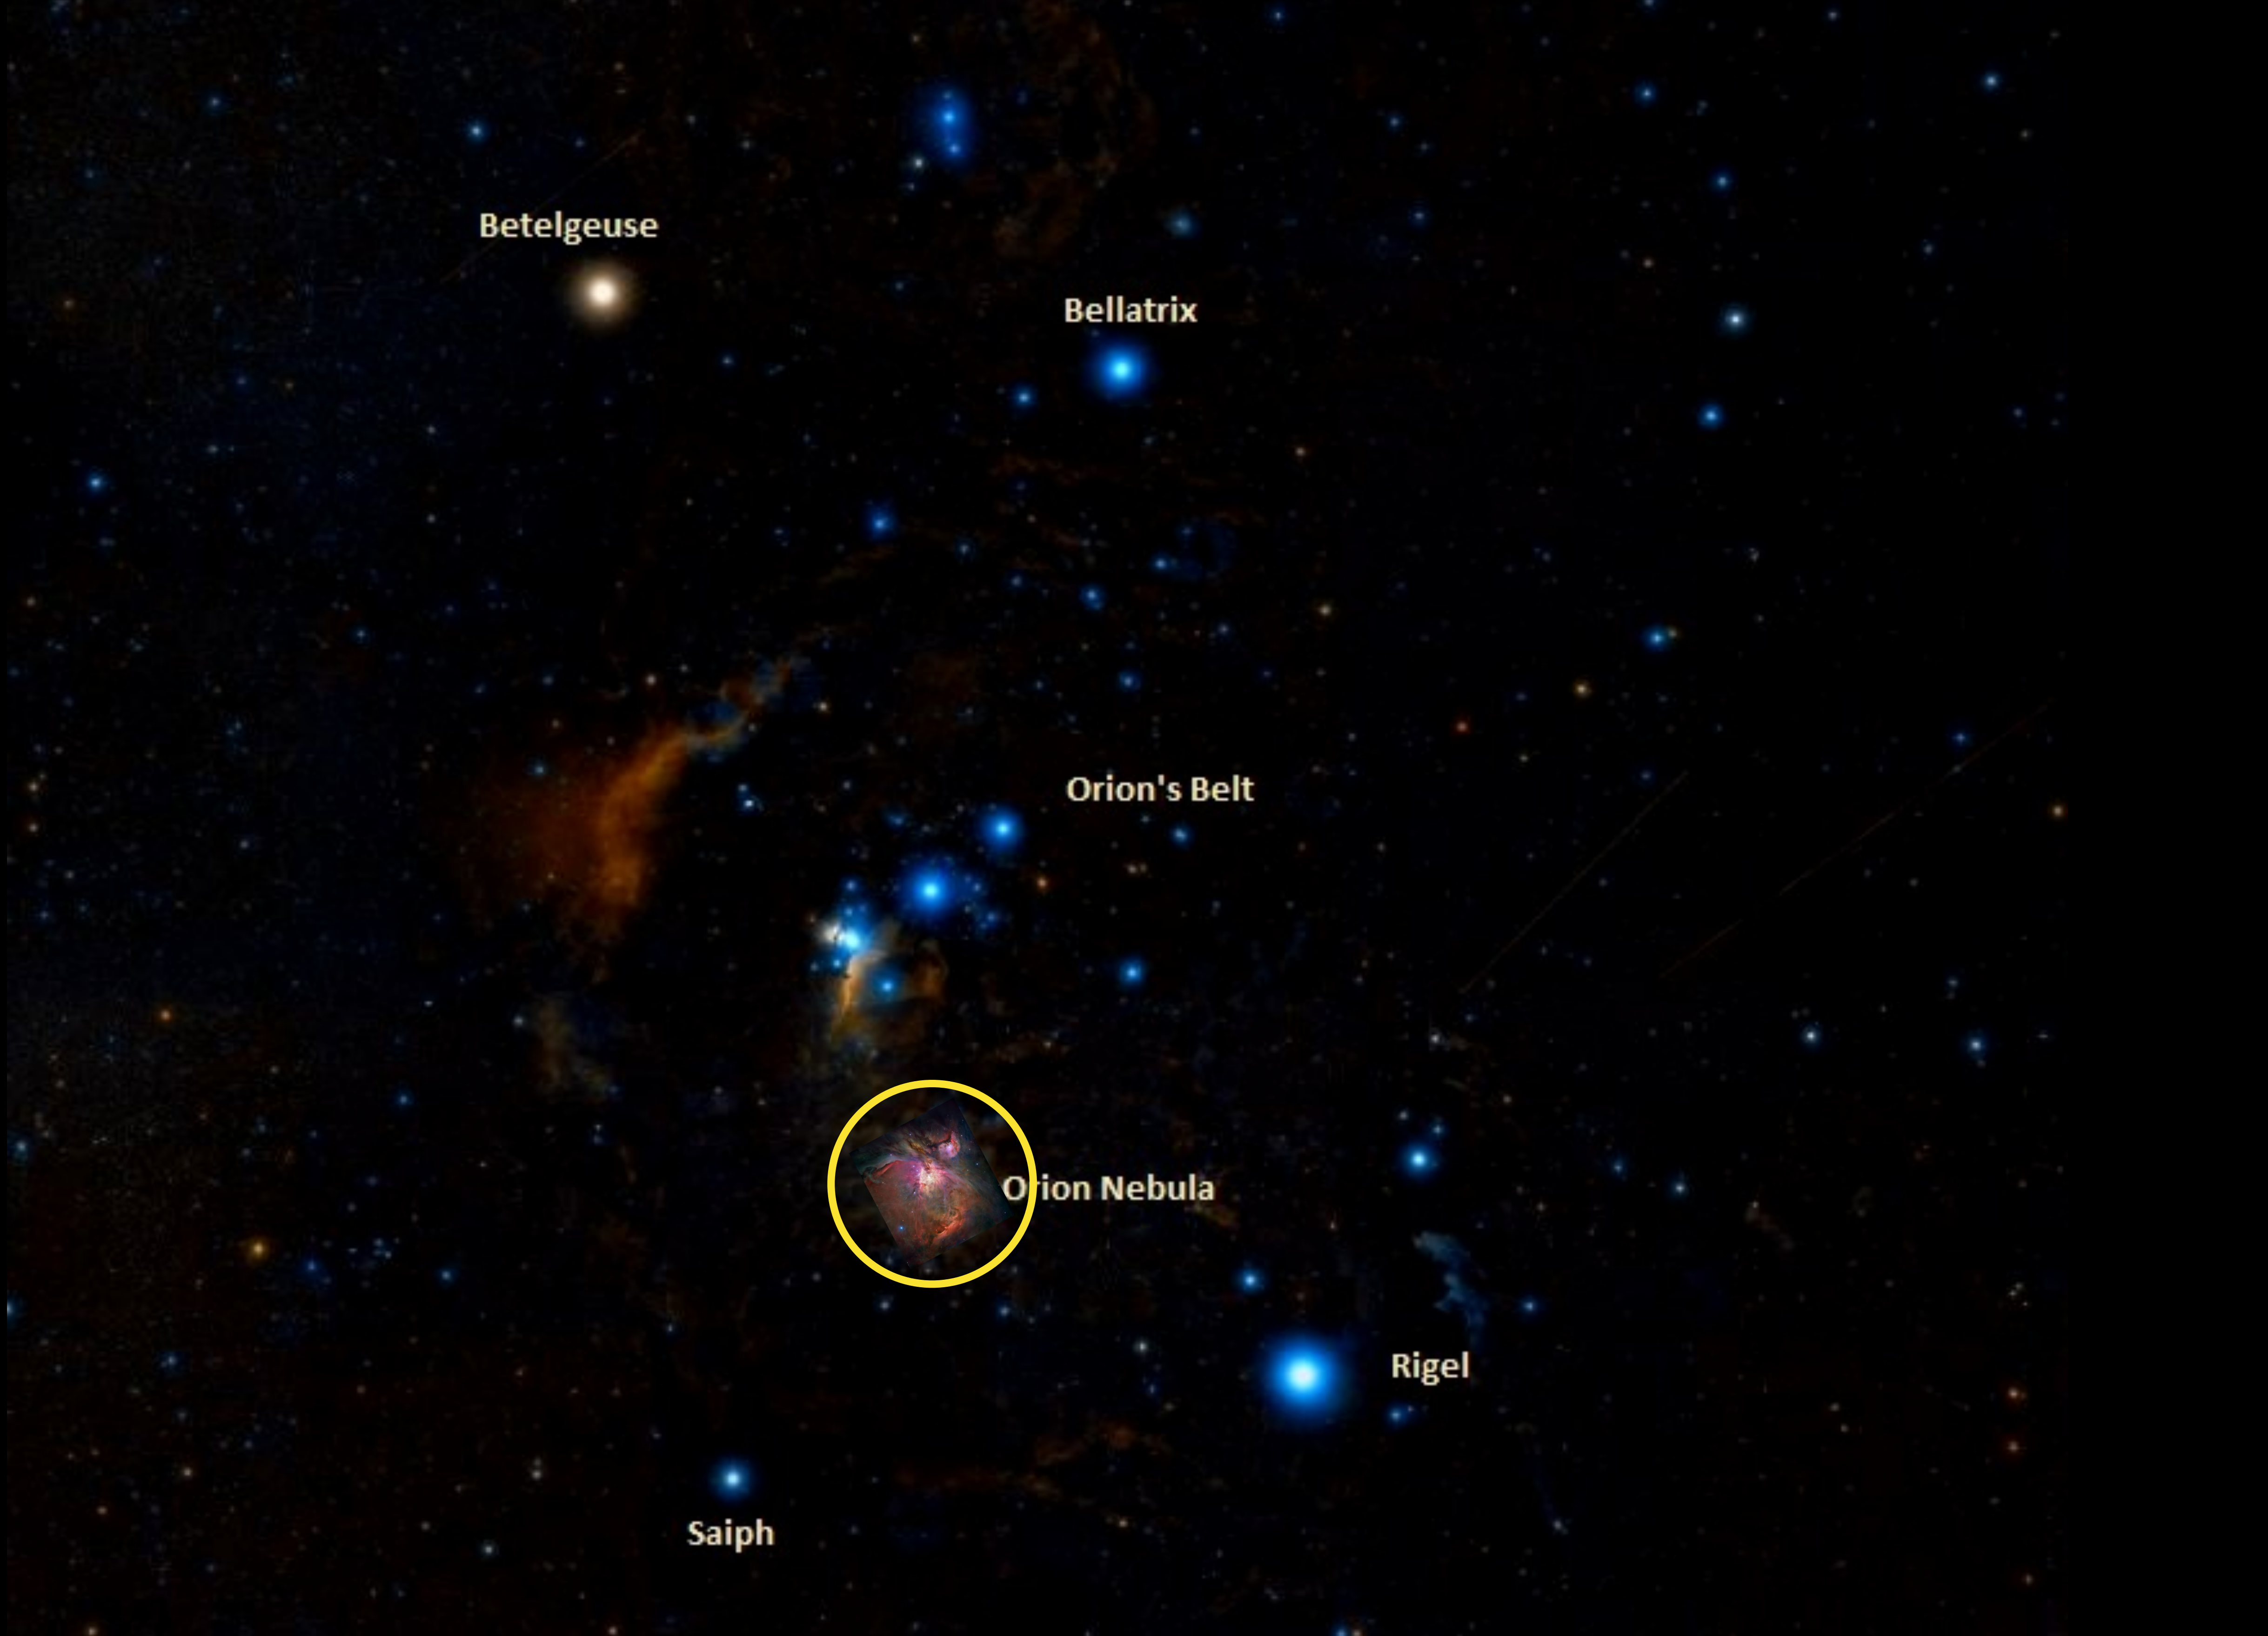

Estimating Cloud Distances

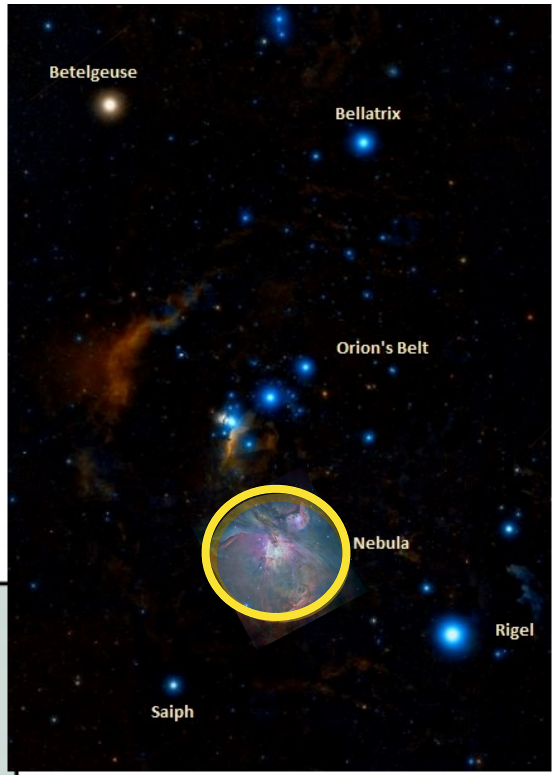

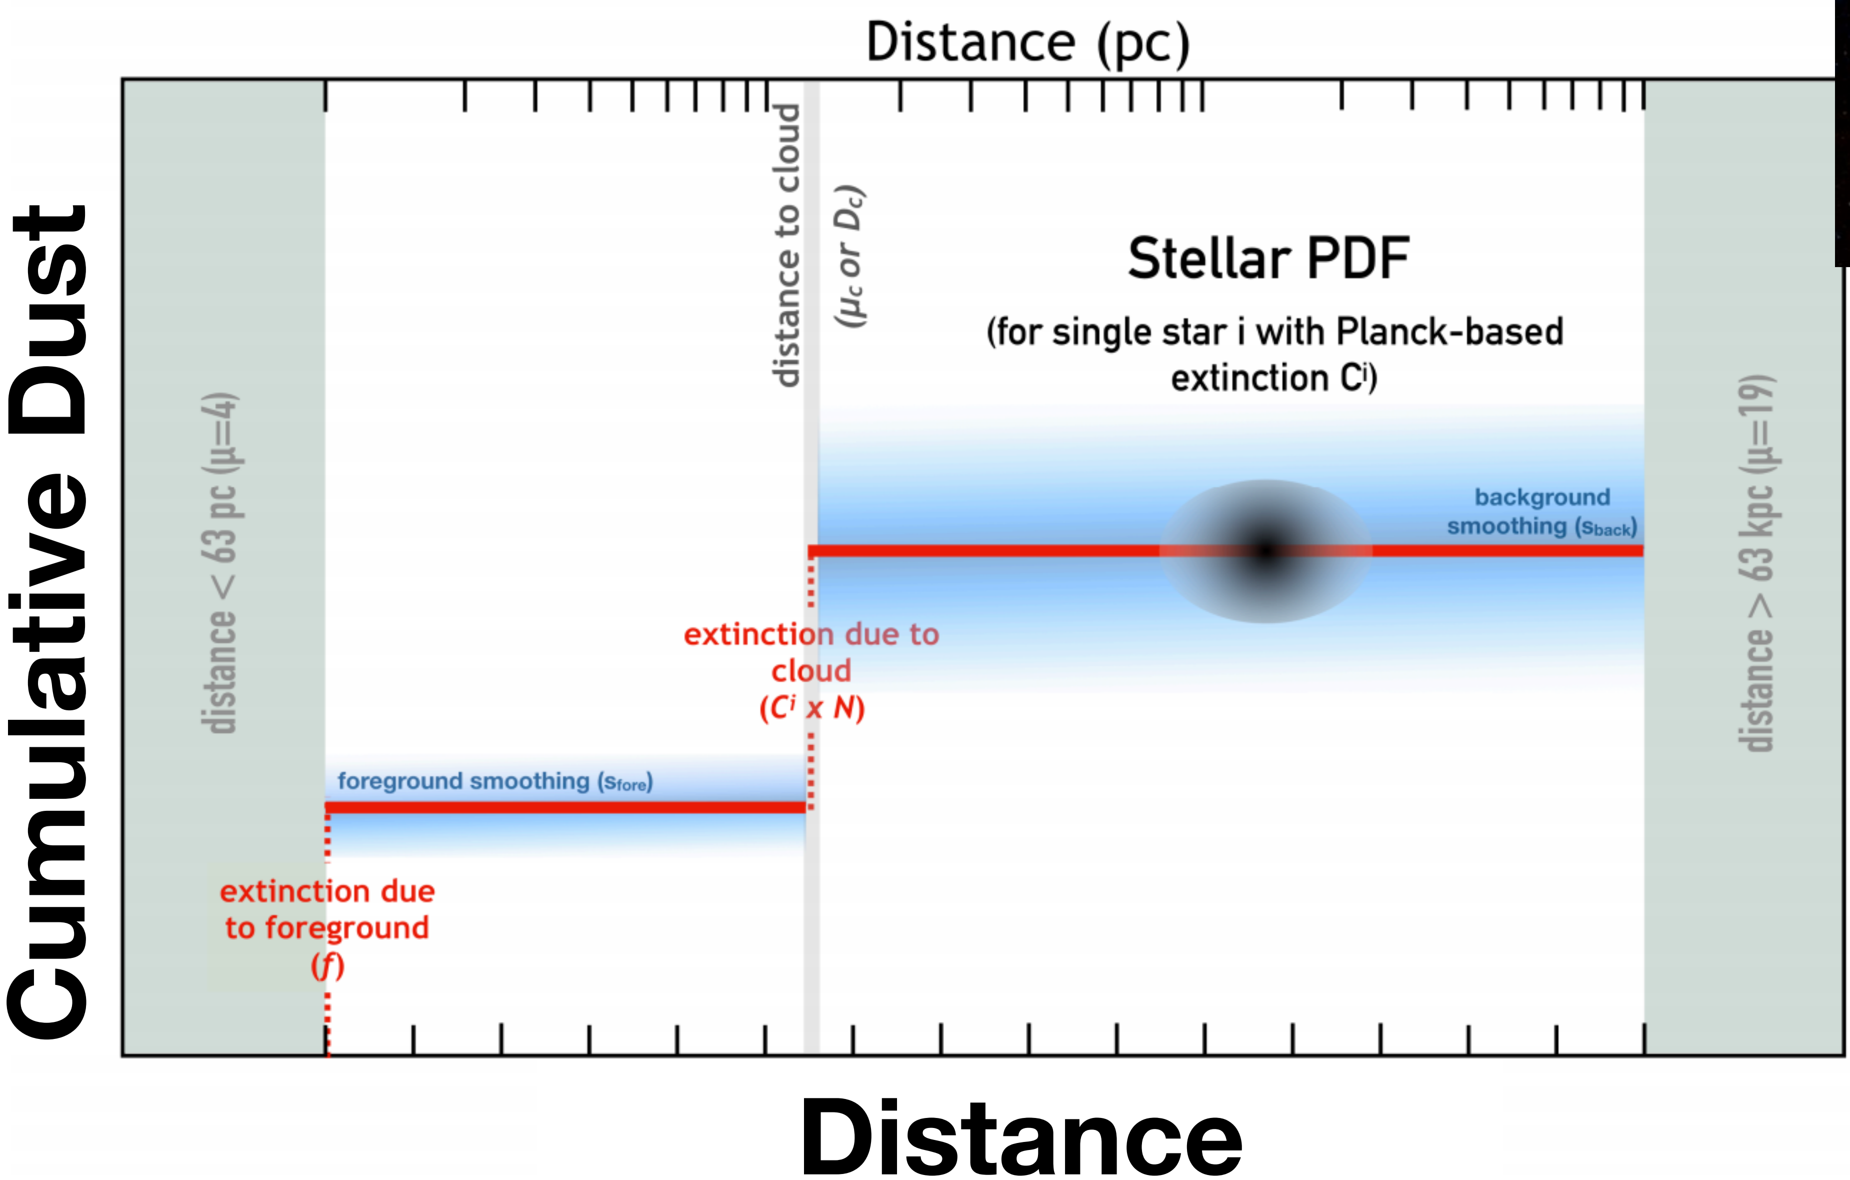

Zucker & Speagle et al. (2019)

Estimating Cloud Distances

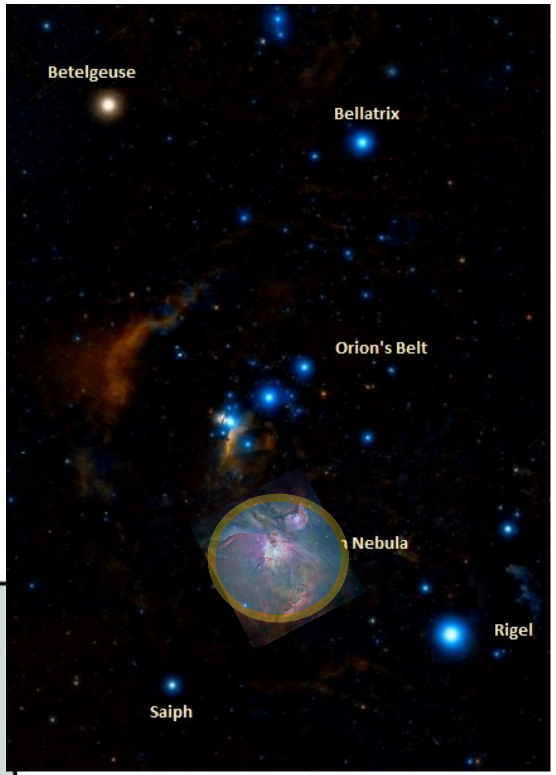

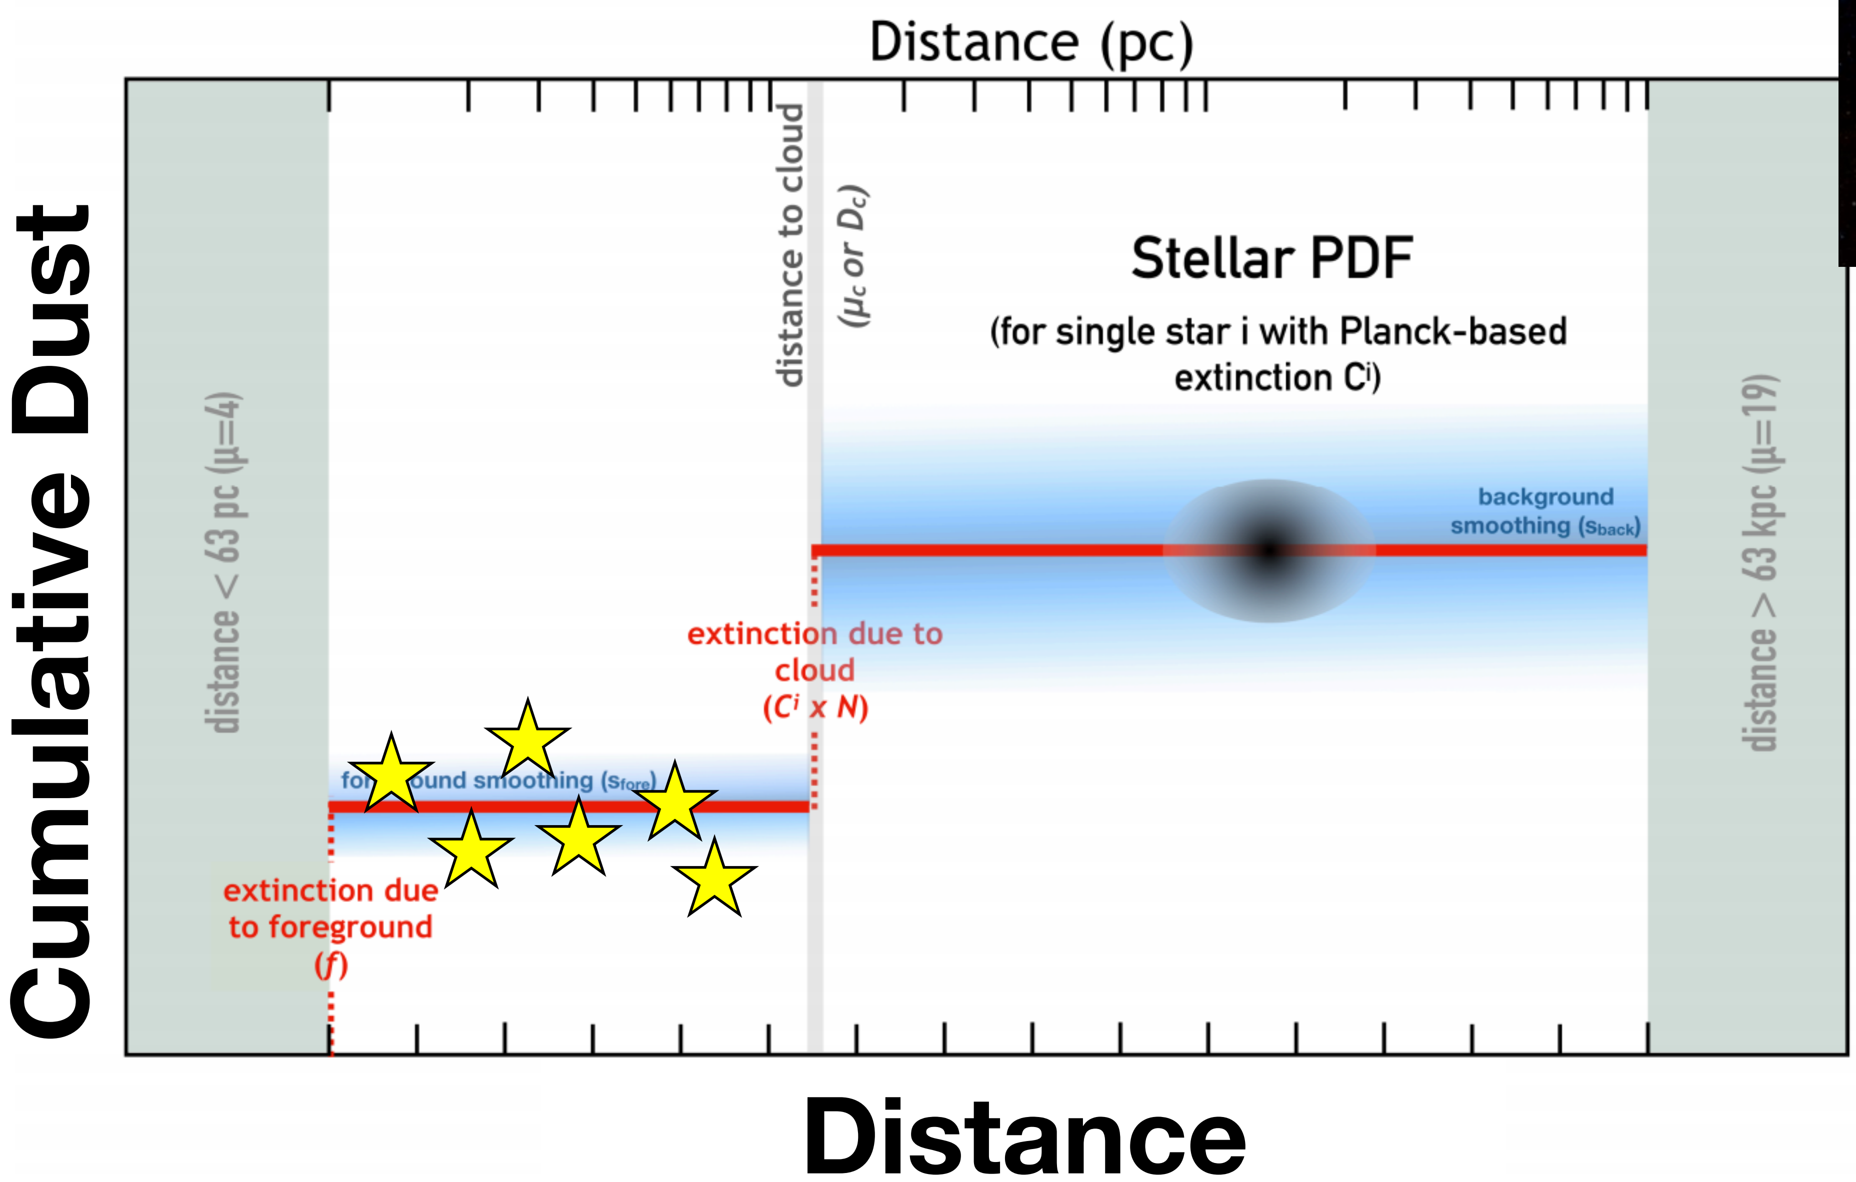

Zucker & Speagle et al. (2019)

Estimating Cloud Distances

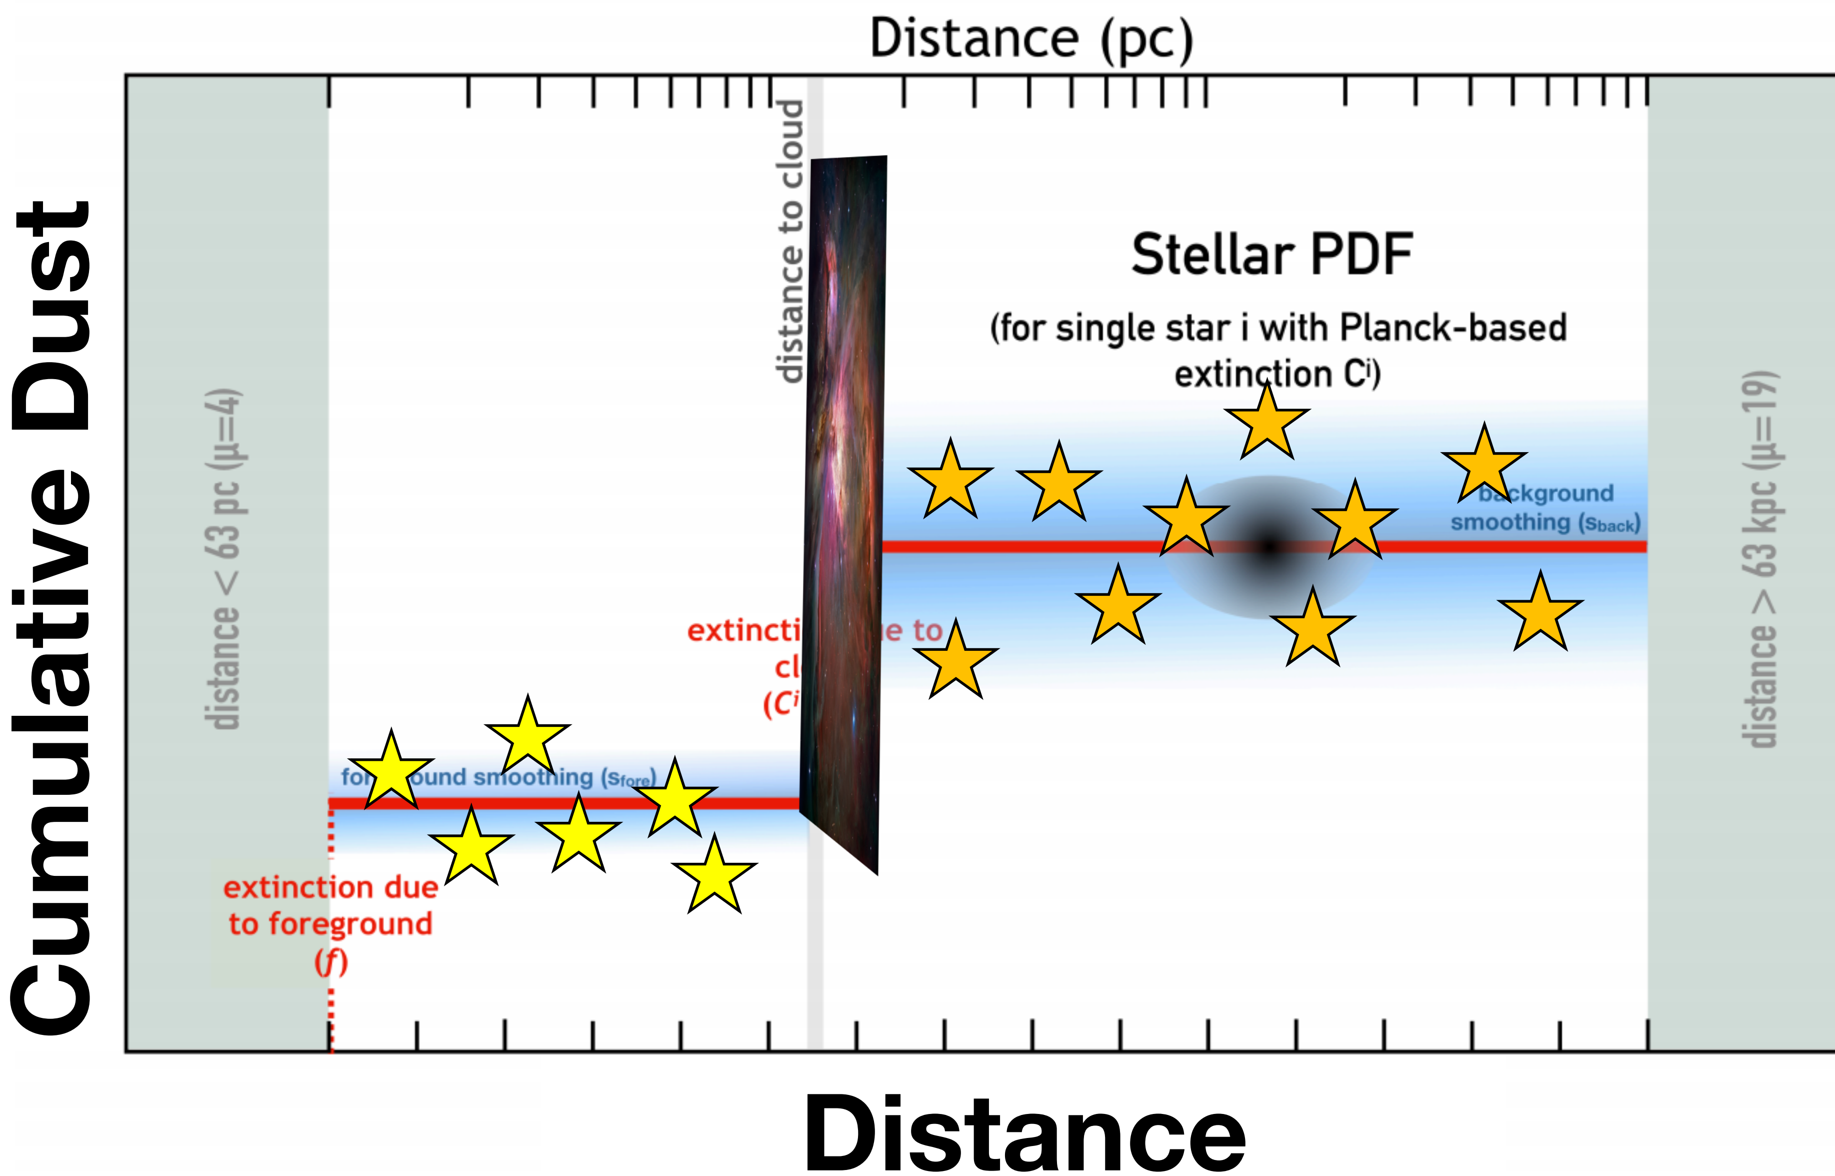

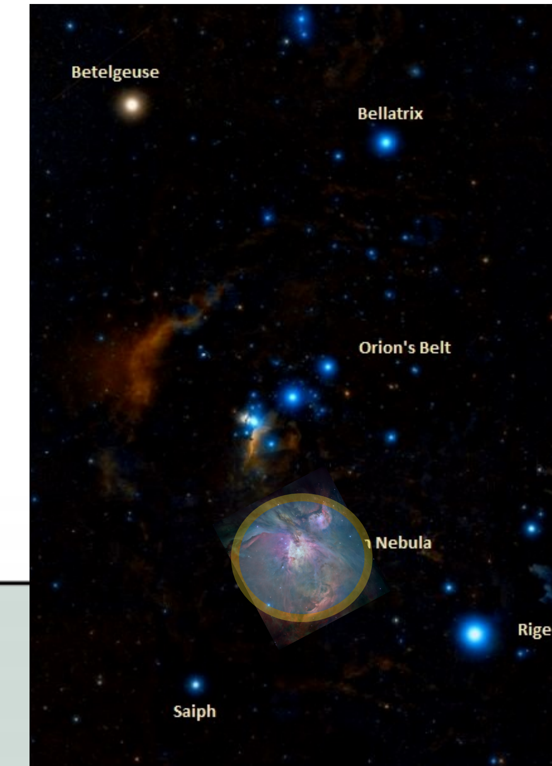

Zucker & Speagle et al. (2019)

Estimating Cloud Distances

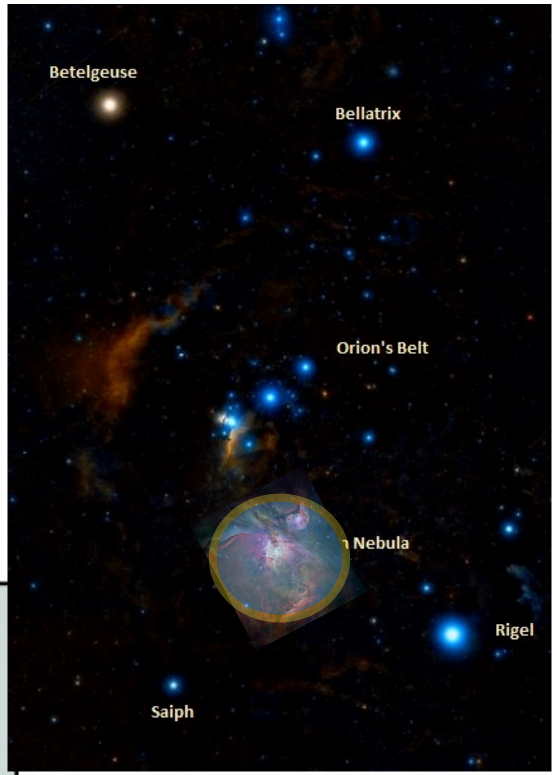

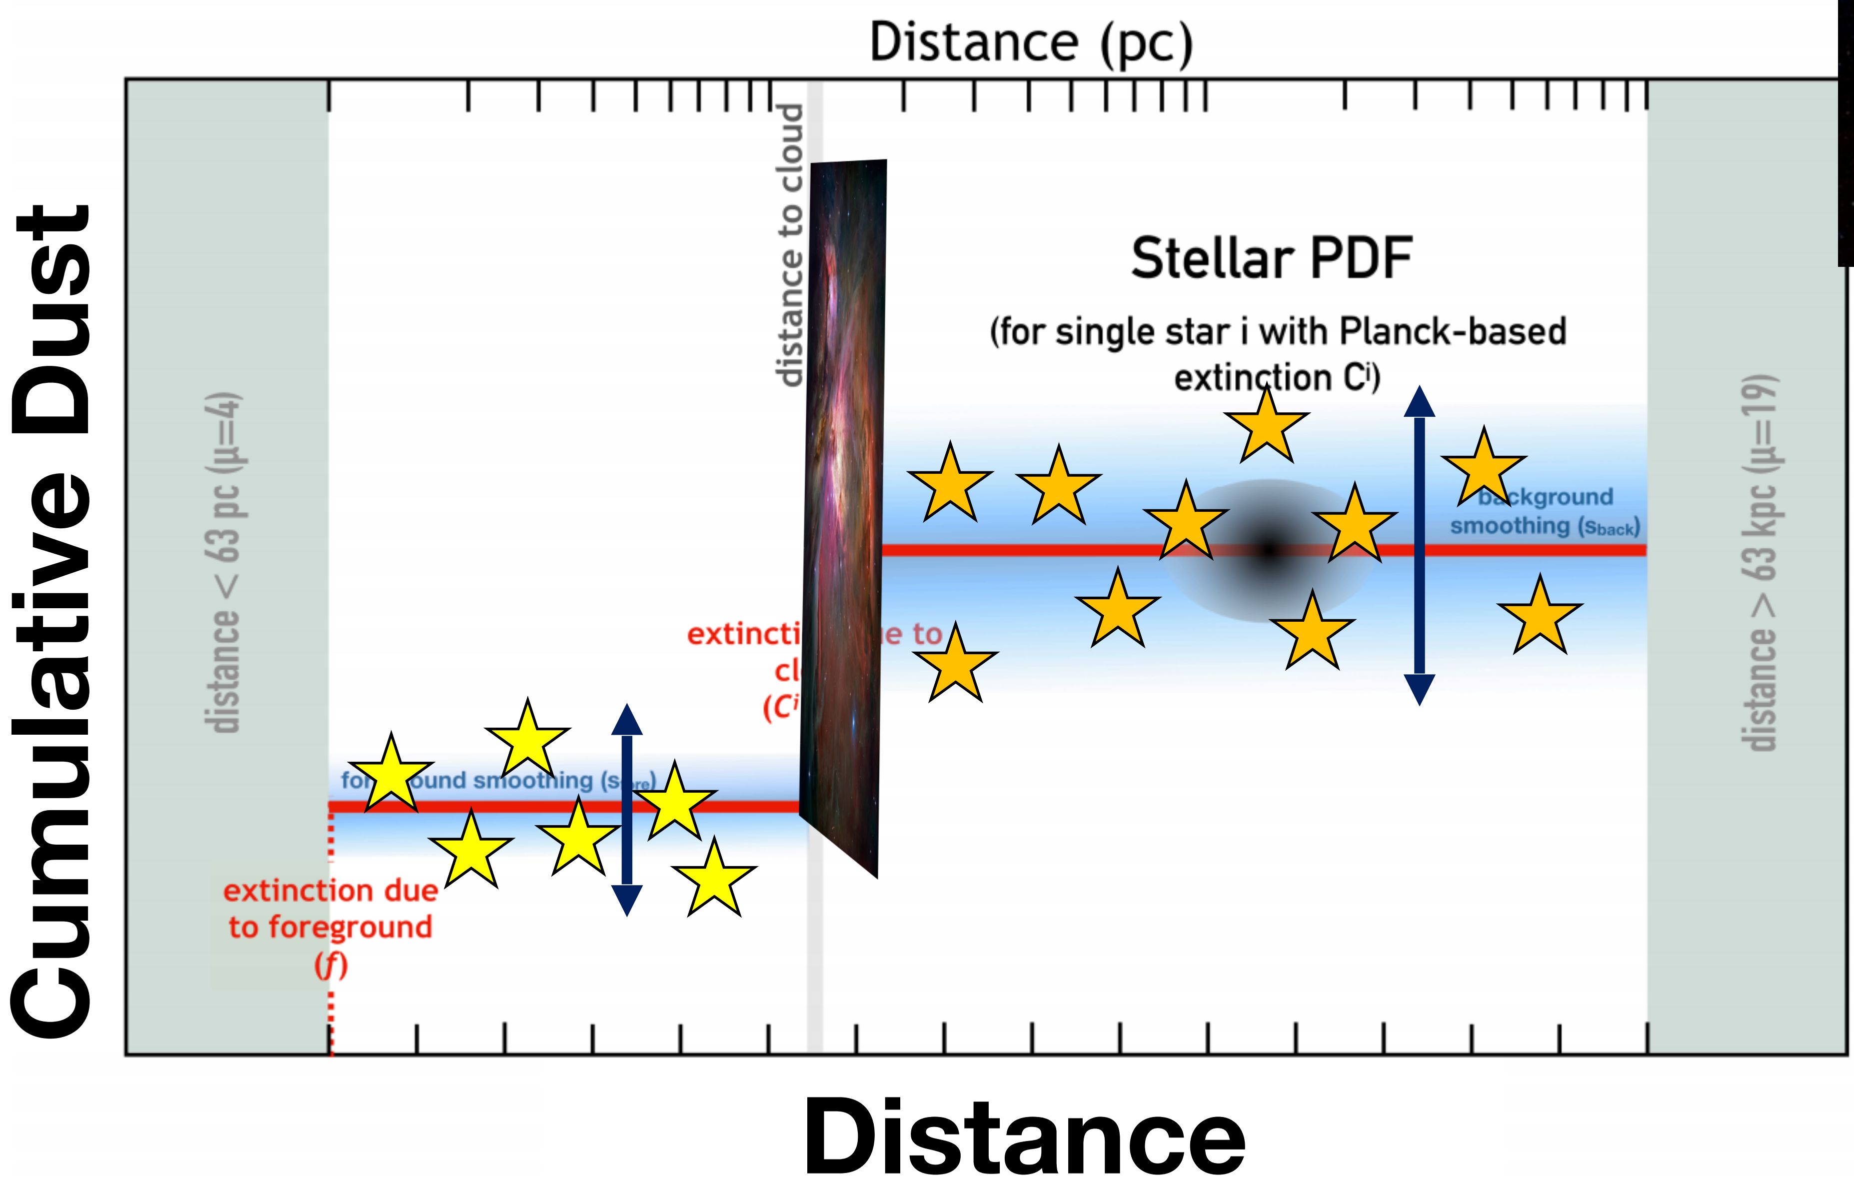

Zucker & Speagle et al. (2019)

Estimating Cloud Distances

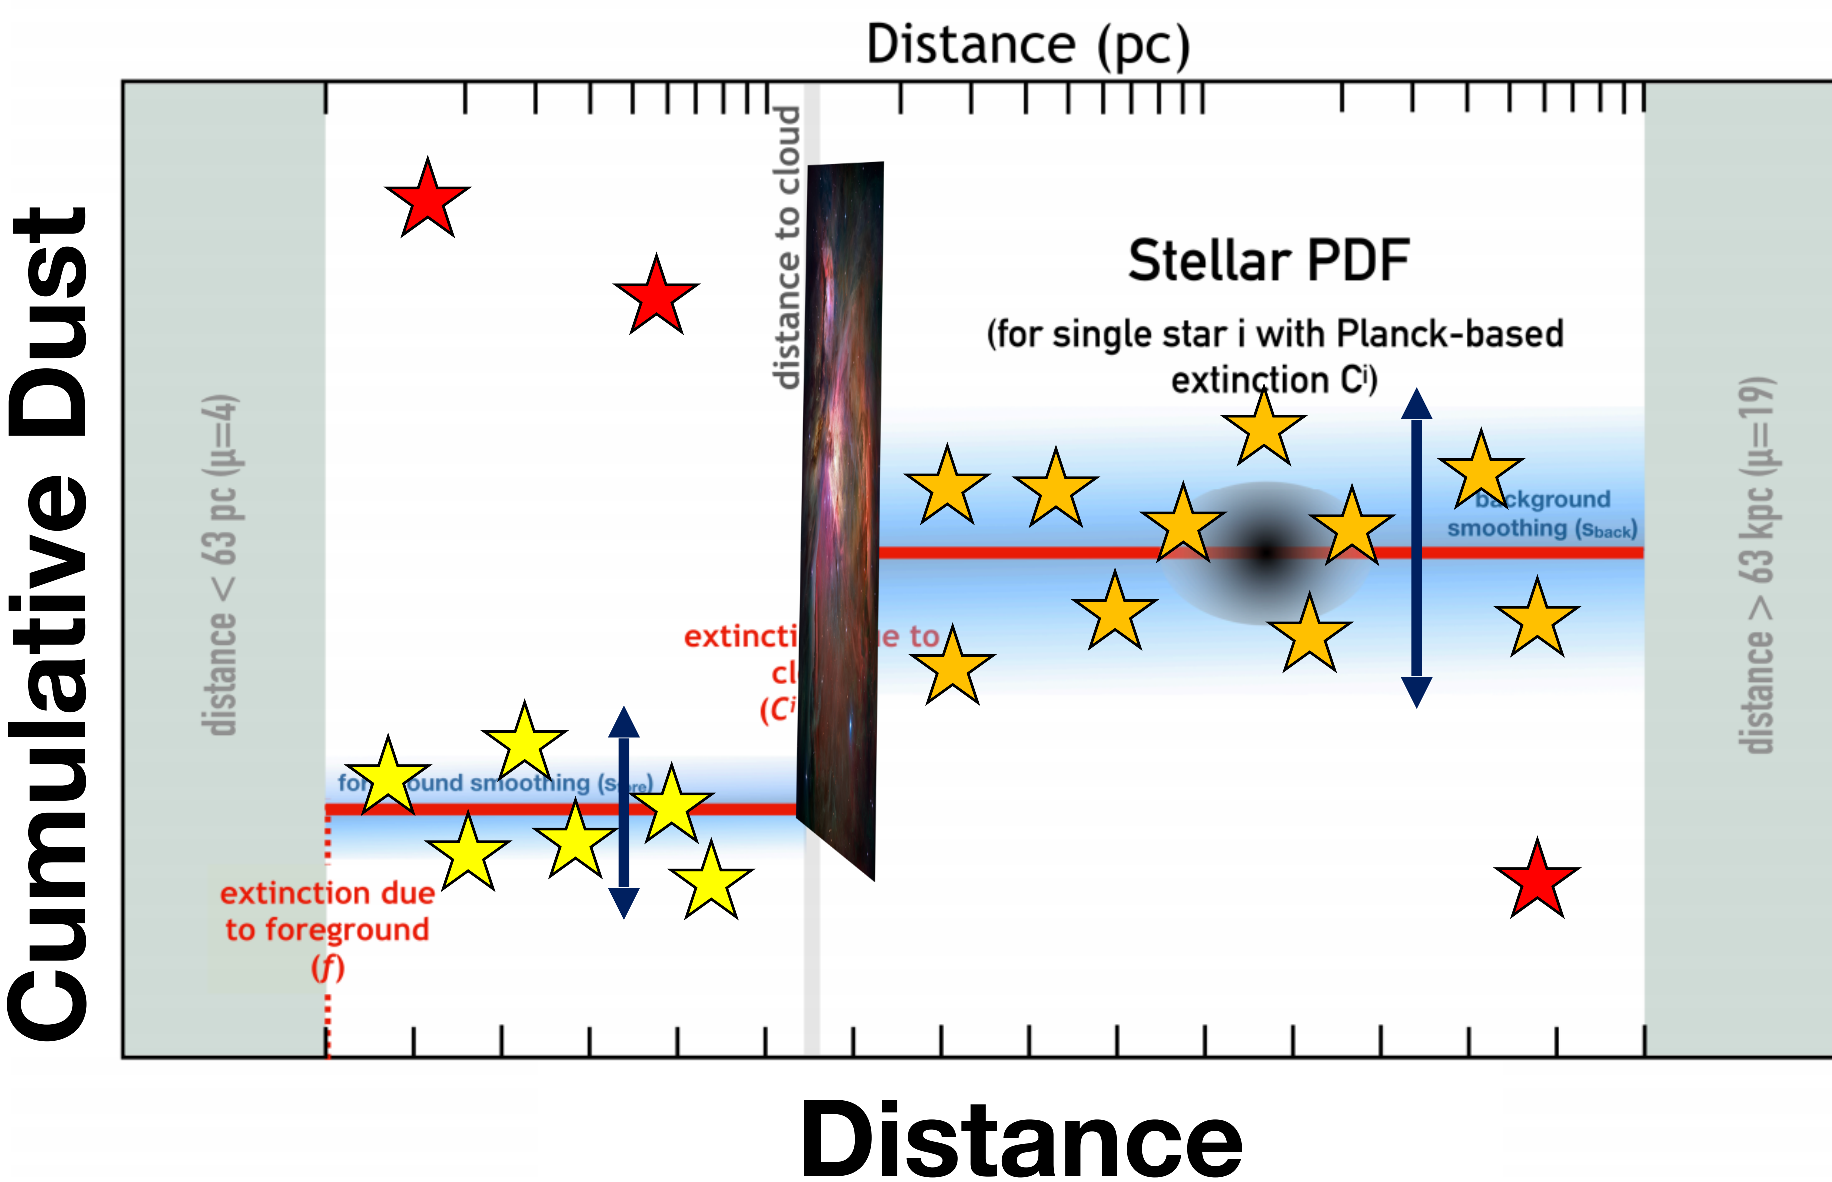

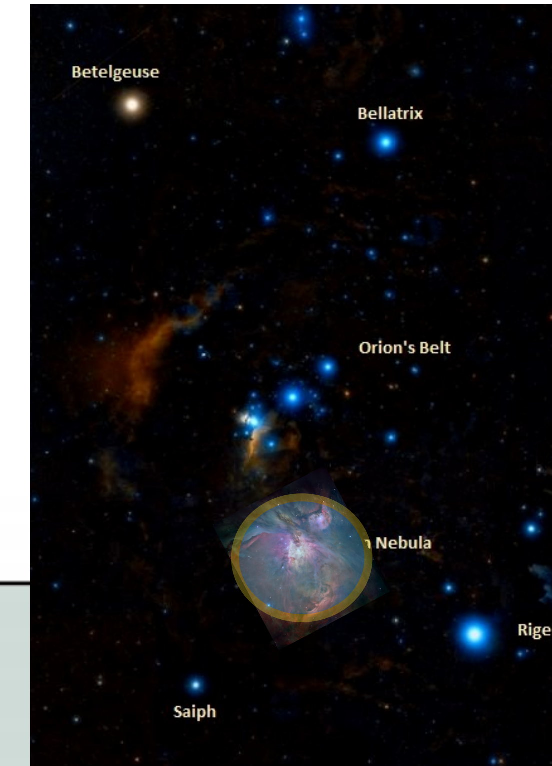

Zucker & Speagle et al. (2019)

Formalism

Model Parameters

α = {d_{cloud}, bunch of nuisance parameters}

Line-of-sight dust model **Per-Star Posterior**

$$\mathcal{L}(\boldsymbol{\alpha}|\hat{\mathbf{m}}_i) = \int \boxed{P(\boldsymbol{\alpha}|\mu, A_V)}\,\boxed{P(\mu, A_V|\hat{\mathbf{m}}_i)}\, d\mu\, dA_V$$

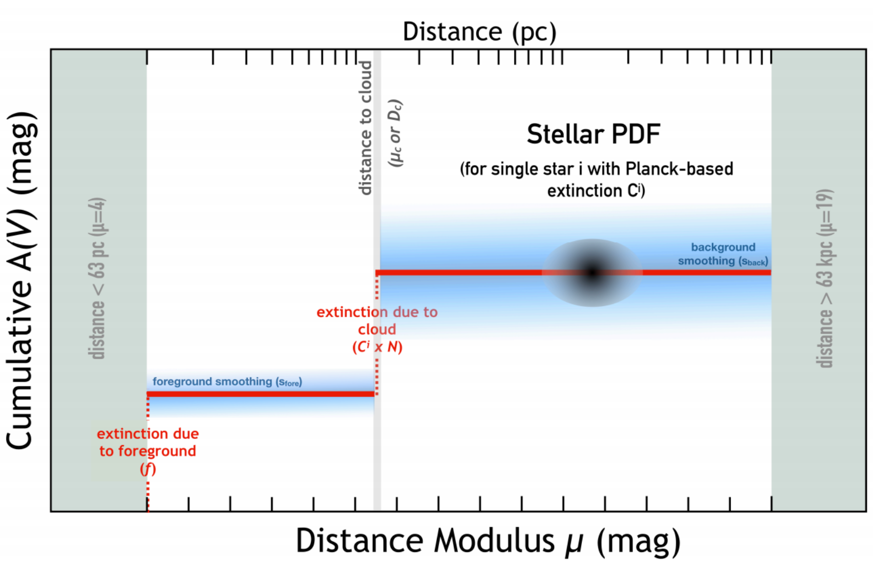

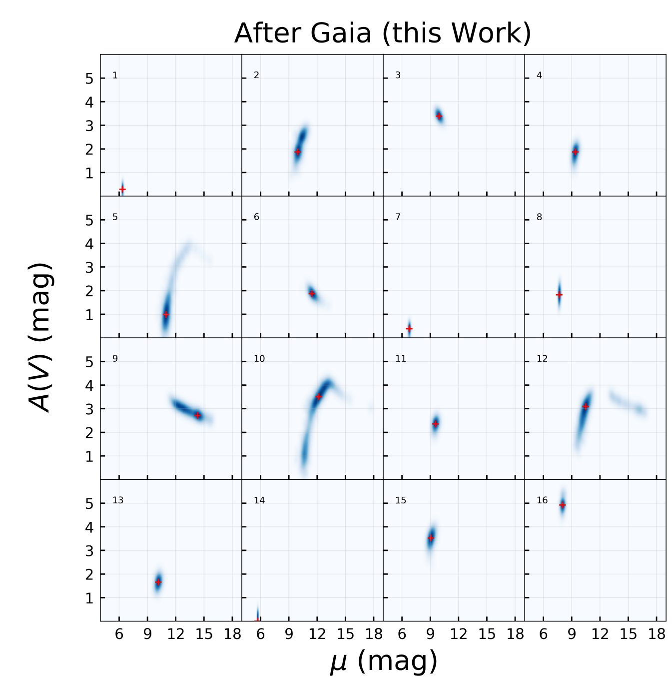

Formalism

We sample from our six parameter model (cloud distance + 5 nuisance parameters) using **dynamic nested sampler** dynesty

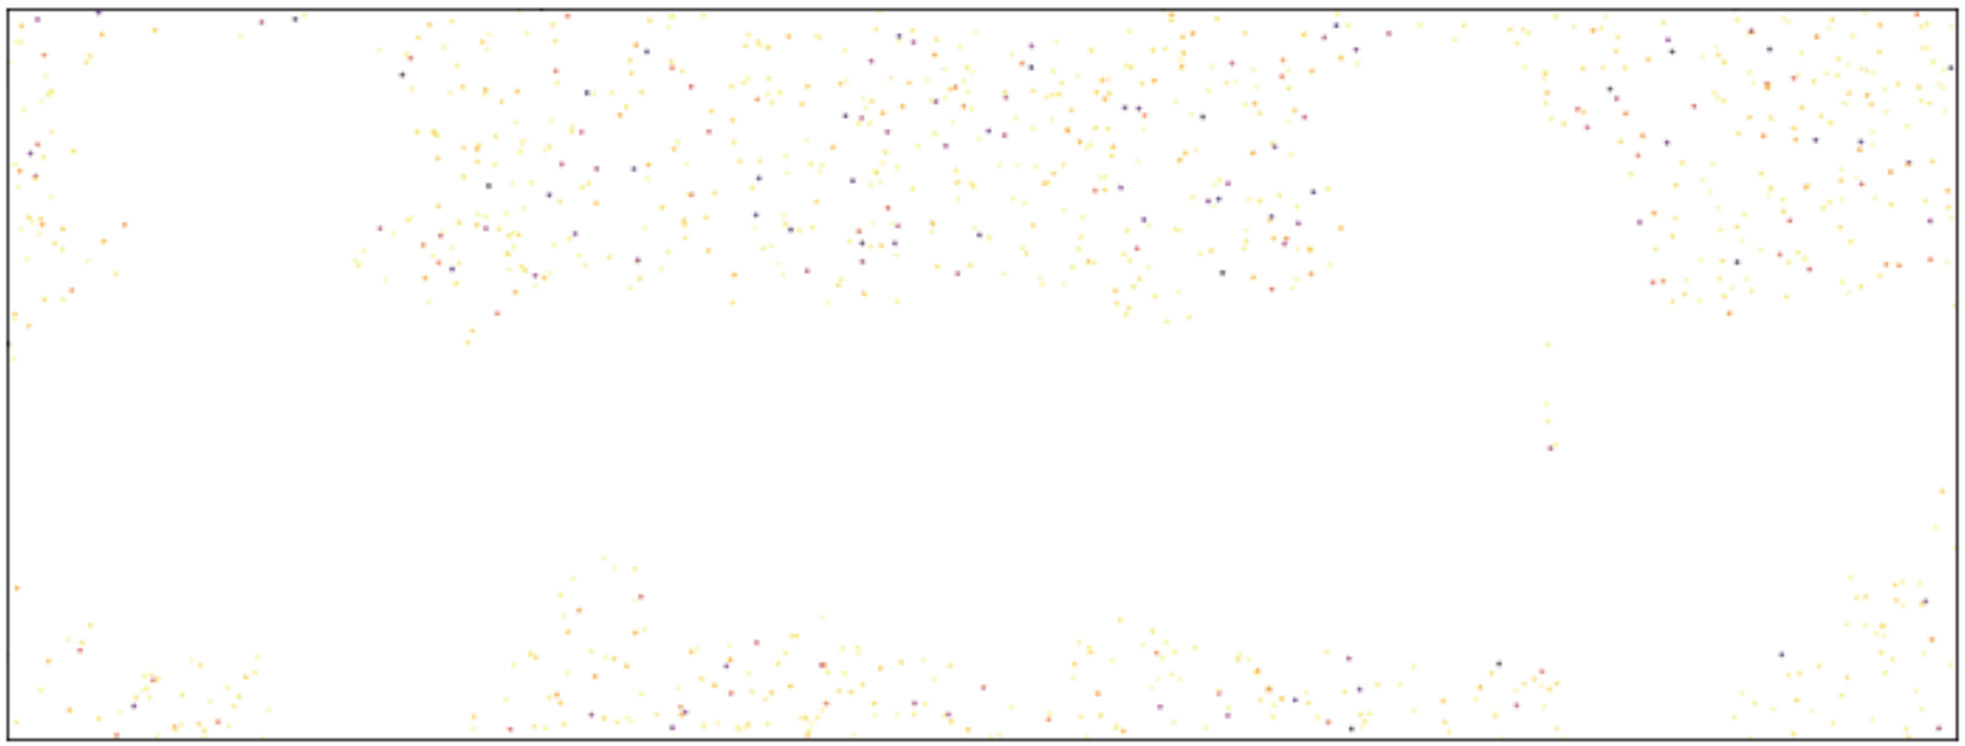

Three main advantages:

1. Can characterize complex uncertainties in real-time (Ferozet al. 2009).
2. Allocates samples more efficiently (Higson et al. 2017b).
3. Possesses well-motivated stopping criteria (Skilling 2006; Speagle 2020)

Speagle 2020

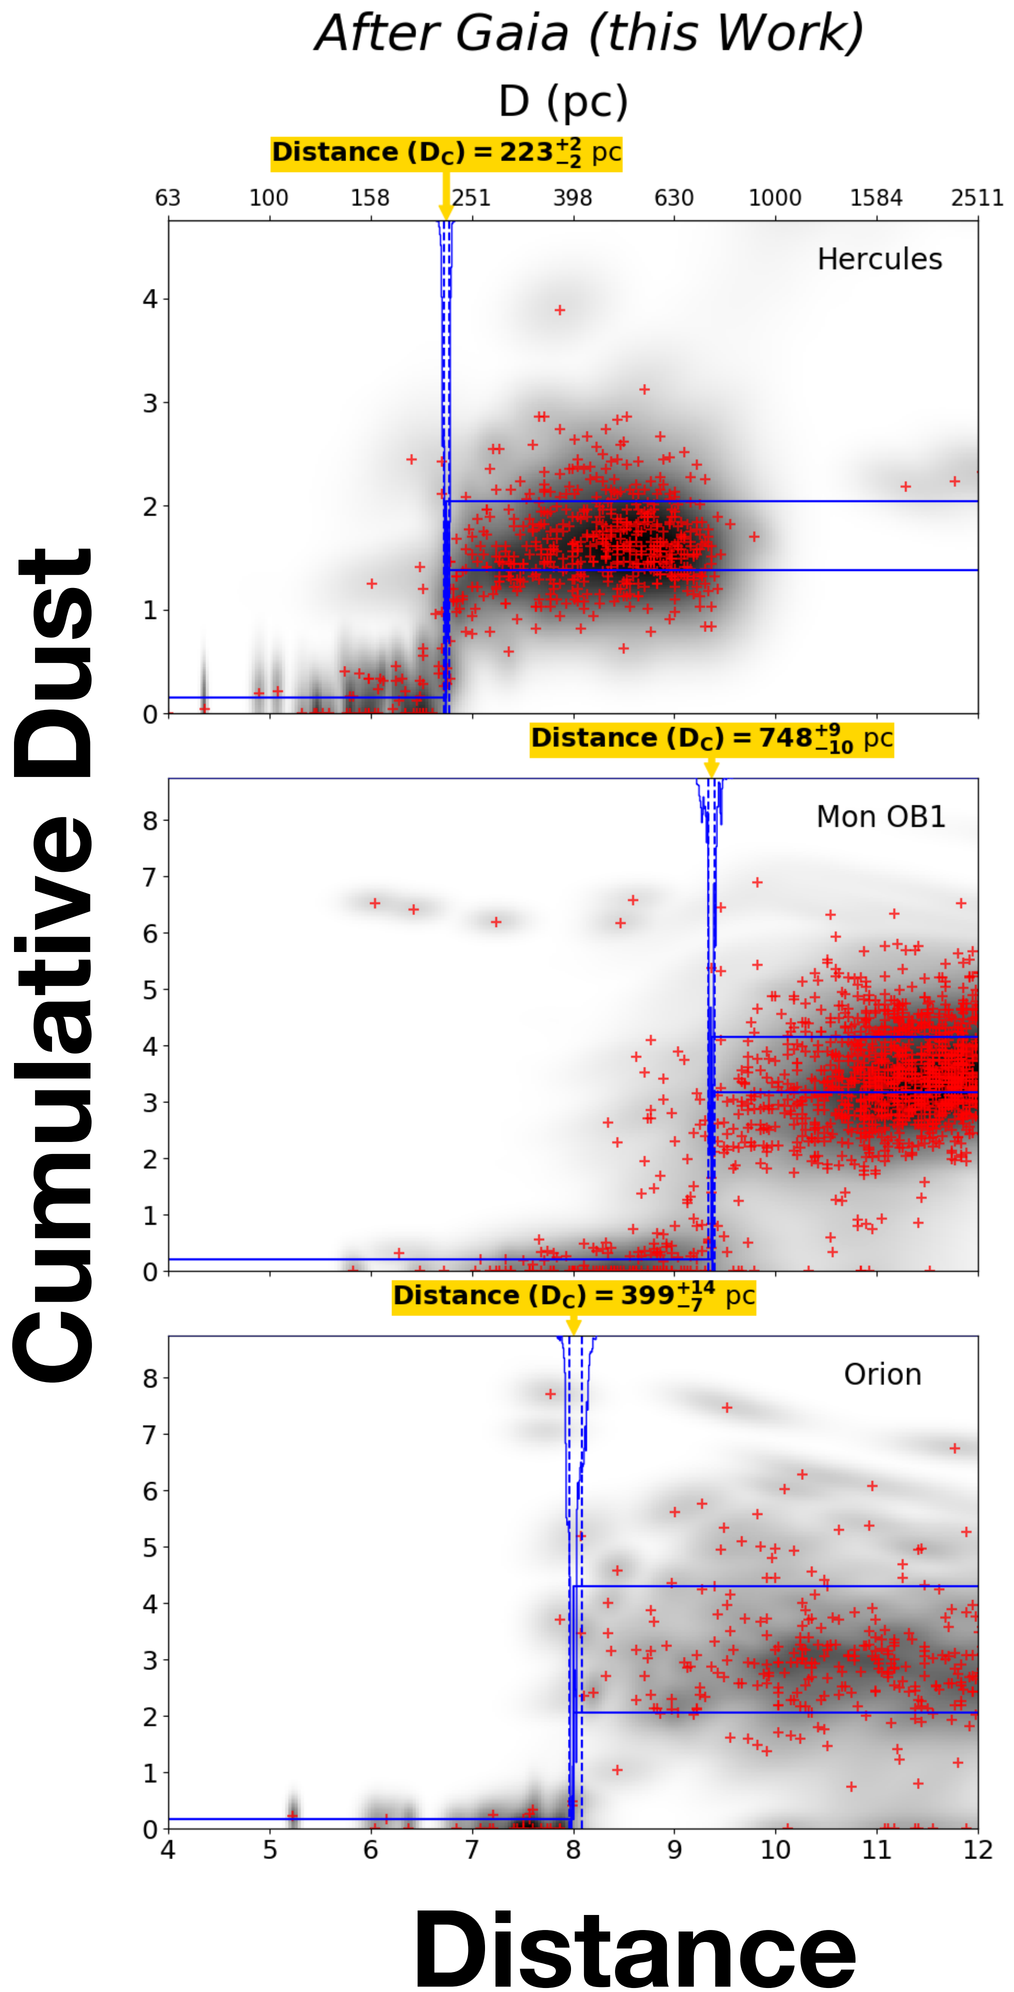

After Gaia (this Work)
D (pc)
Distance (D$_C$) = 223^{+2}_{-2} pc
Hercules
Distance (D$_C$) = 748^{+9}_{-10} pc
Mon OB1
Distance (D$_C$) = 399^{+14}_{-7} pc
Orion
Cumulative Dust
Distance

Results

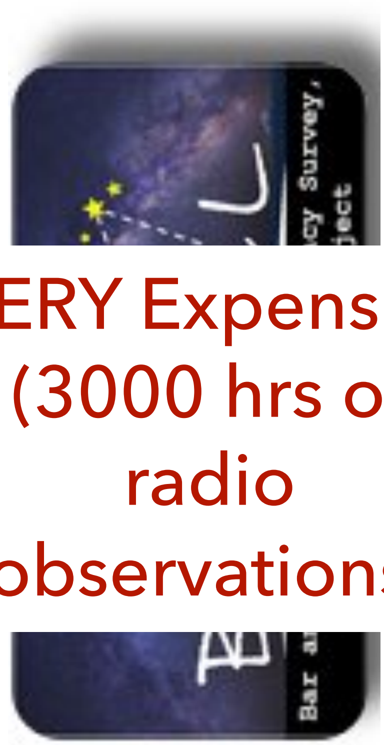

VERY Expensive
(3000 hrs of
radio
observations)

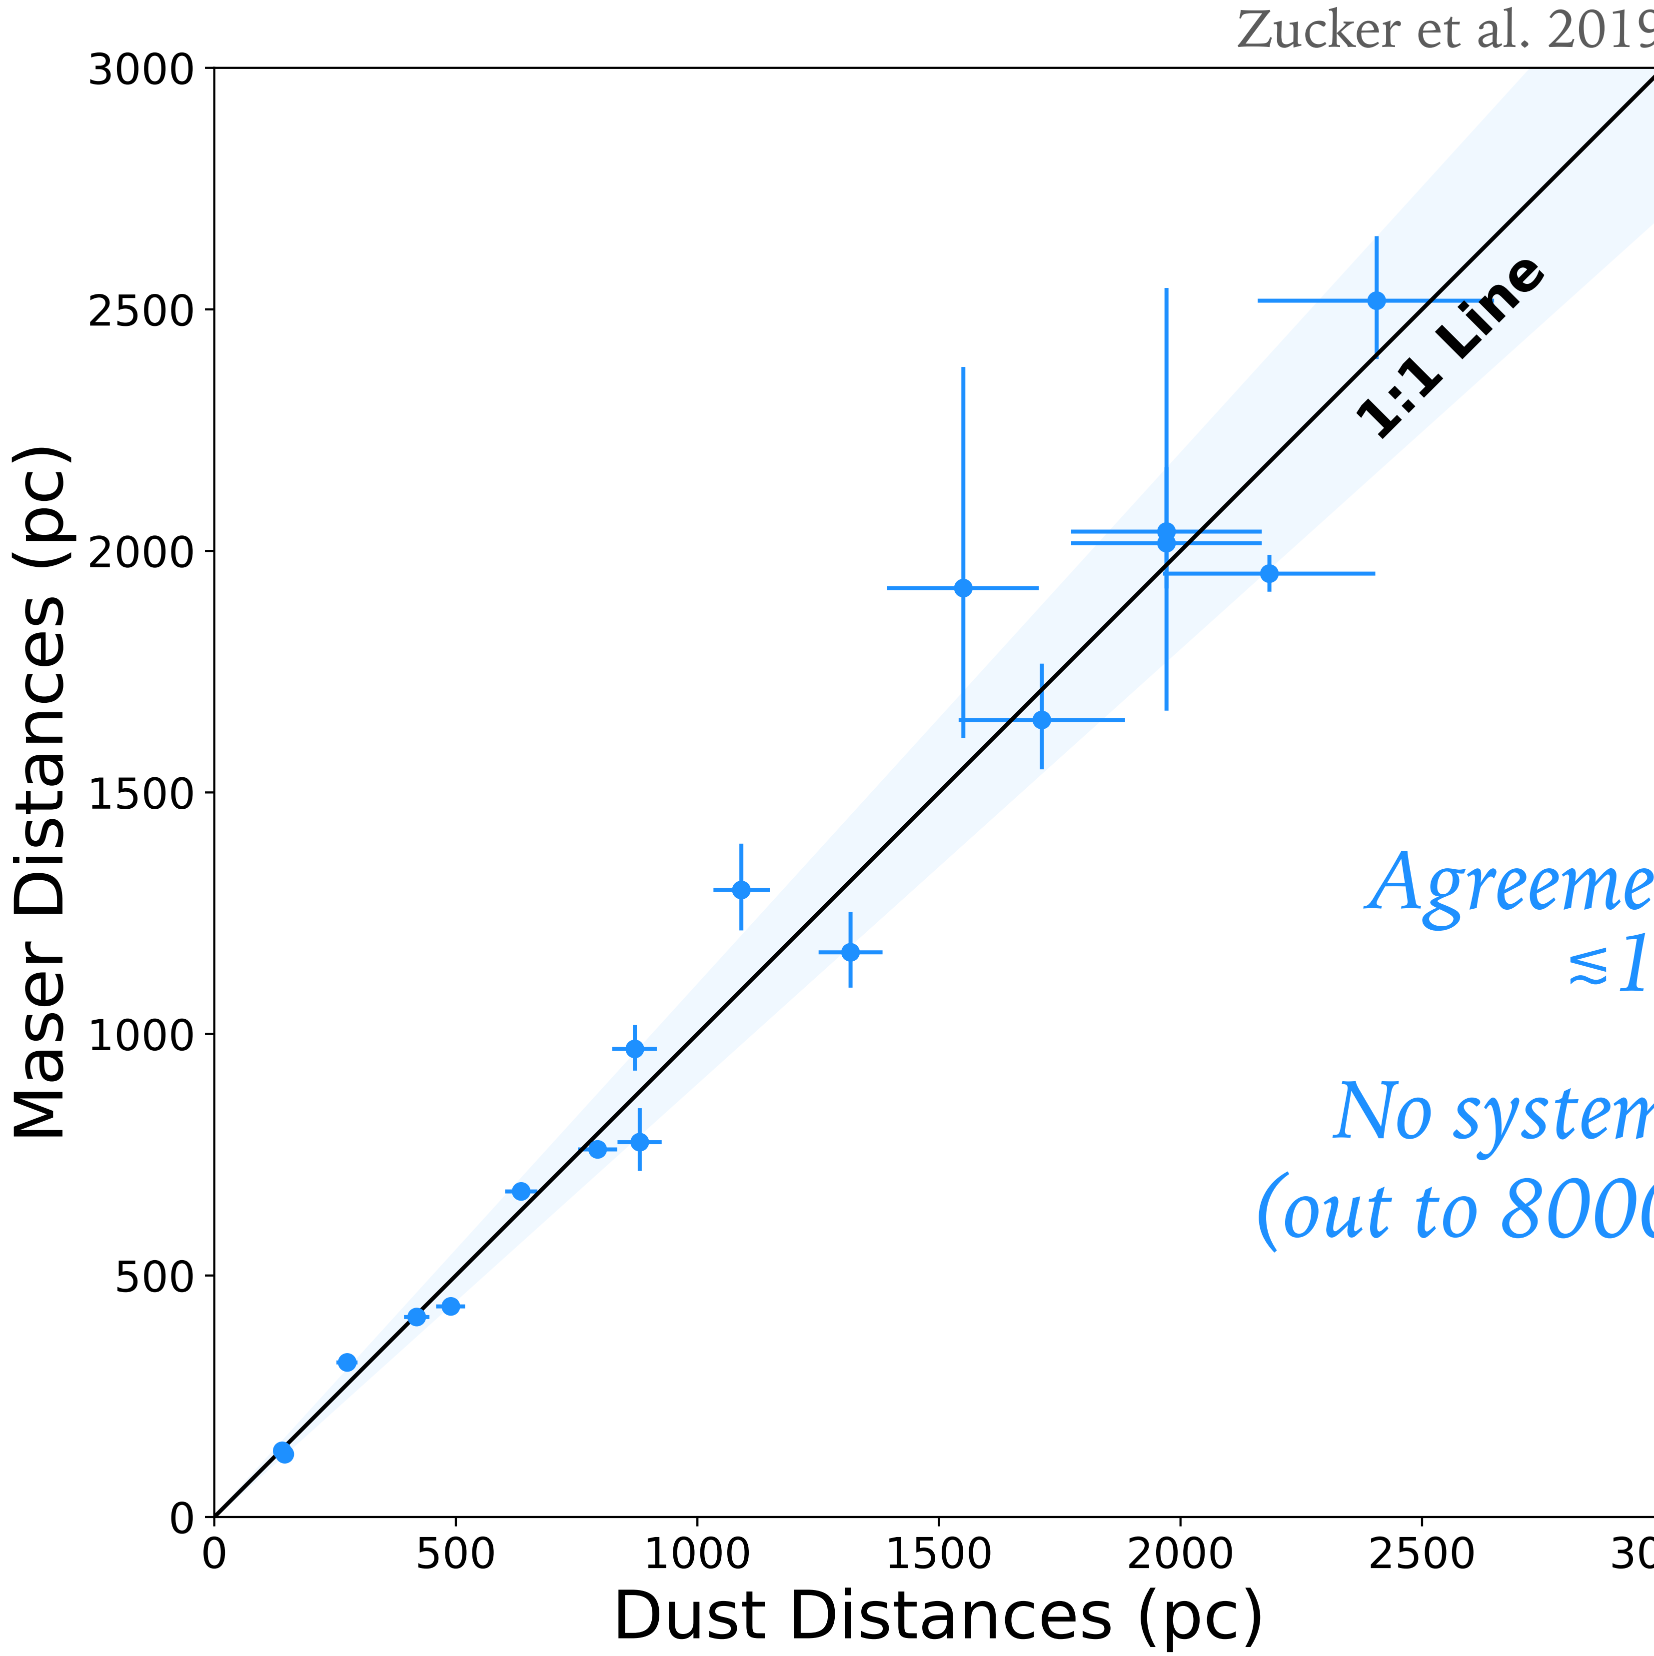

Agreement within
≲10%

No systematic offset
(out to 8000 light years)

Uses “cheap” publicly
available photometry.
Hours of CPU time

See also Reid+2014, 2016,
Brunthaler+2011, Loinard+2013
Ortiz-Leon+2017a,b, Galli+2018 for
maser/compact radio source references

DISTRIBUTION OF LOCAL CLOUDS

Zucker et al. 2020
Green et al. 2019

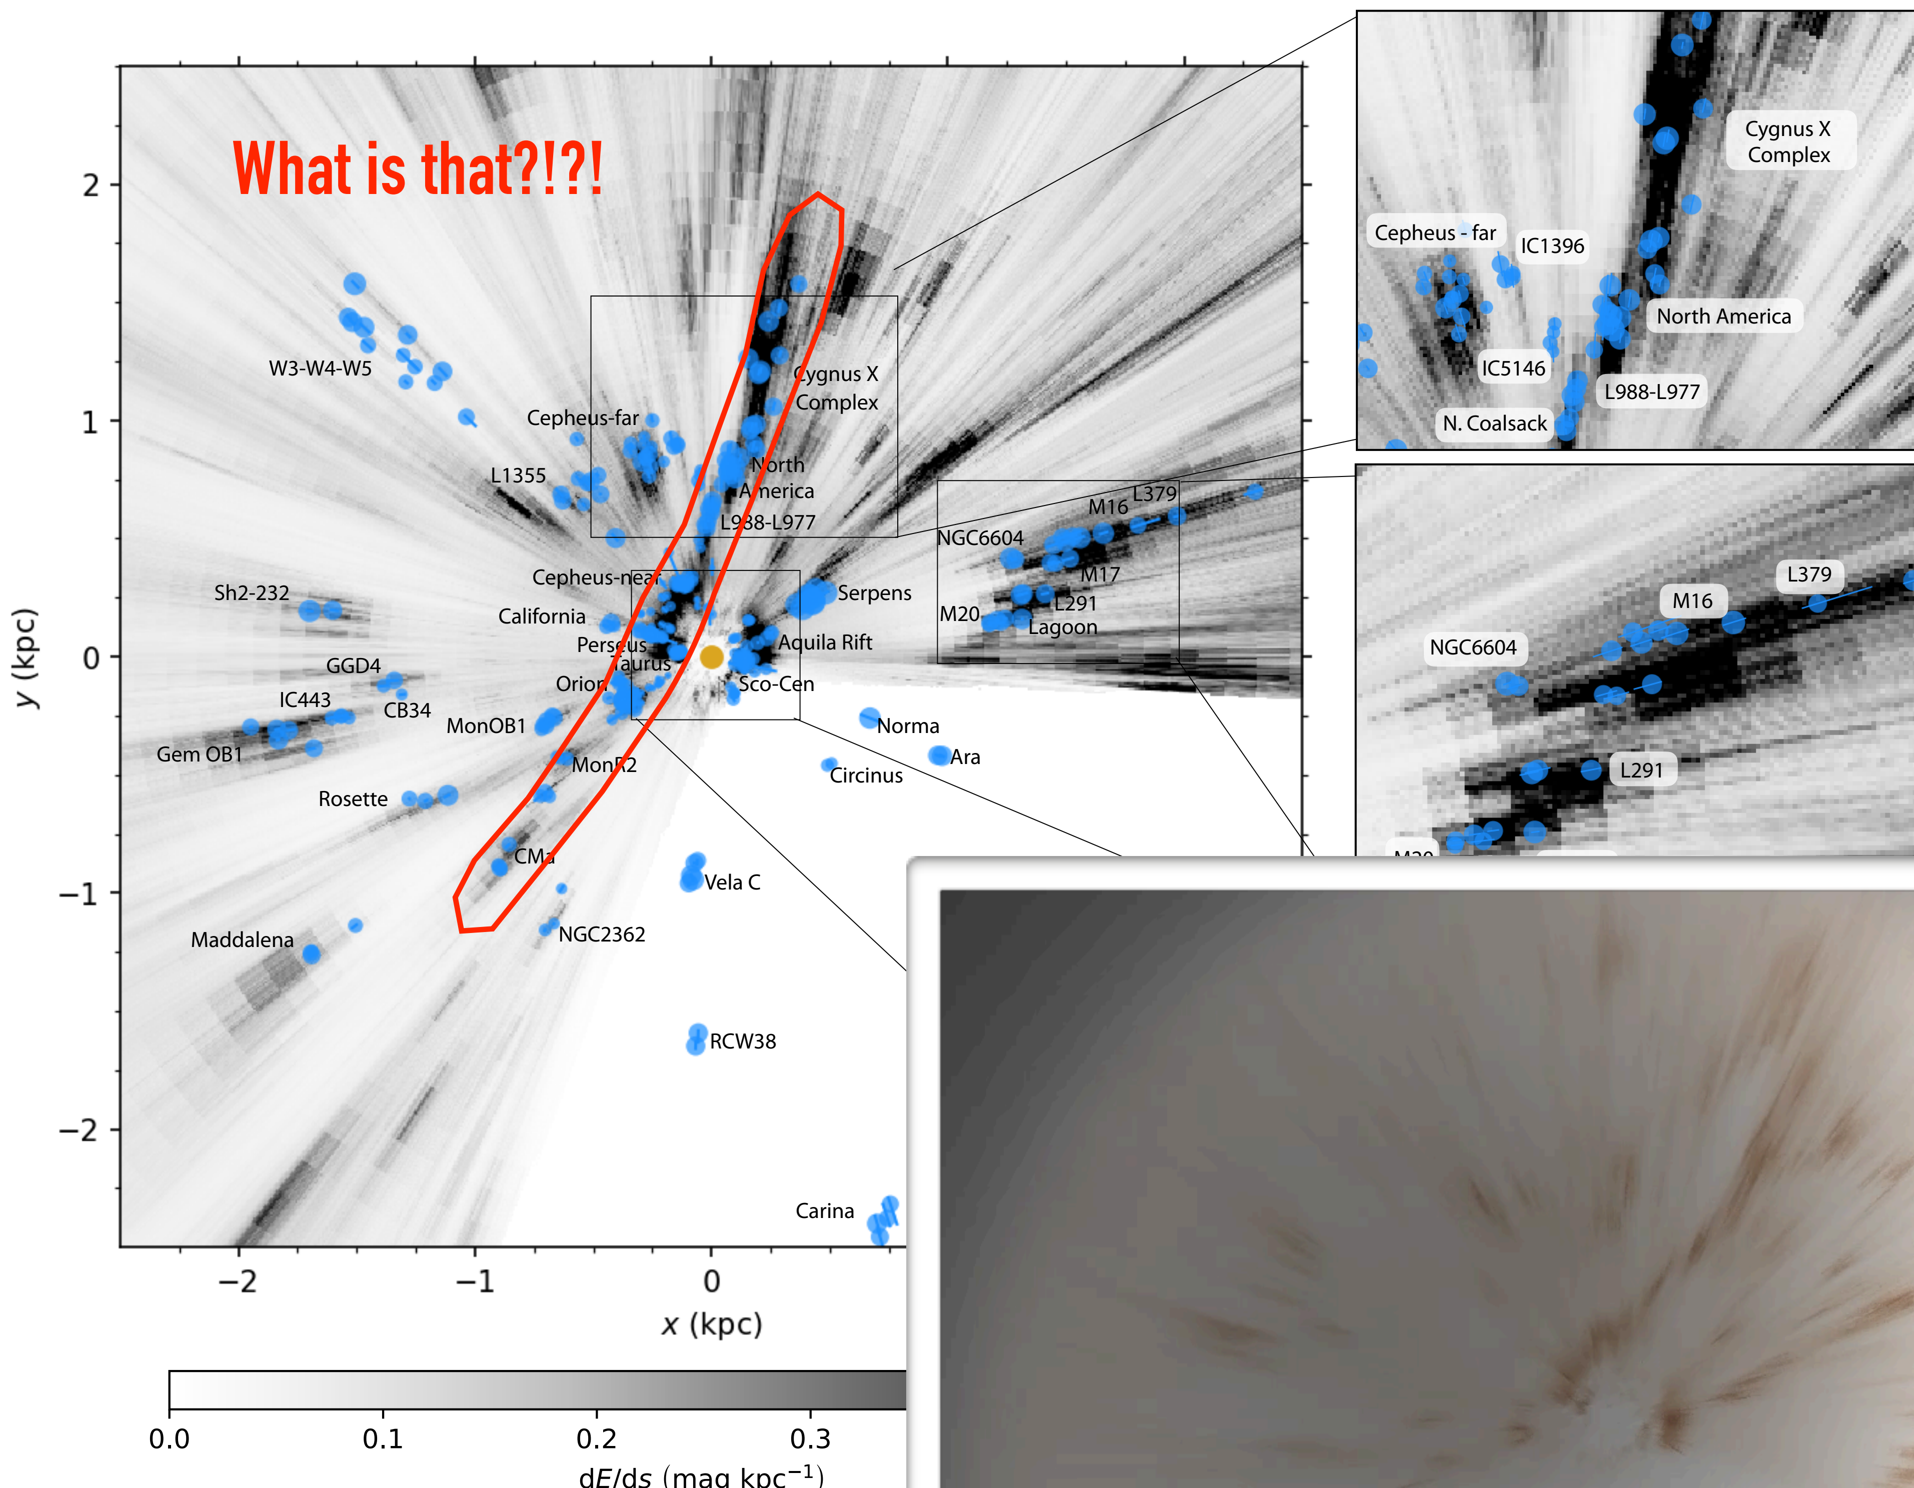

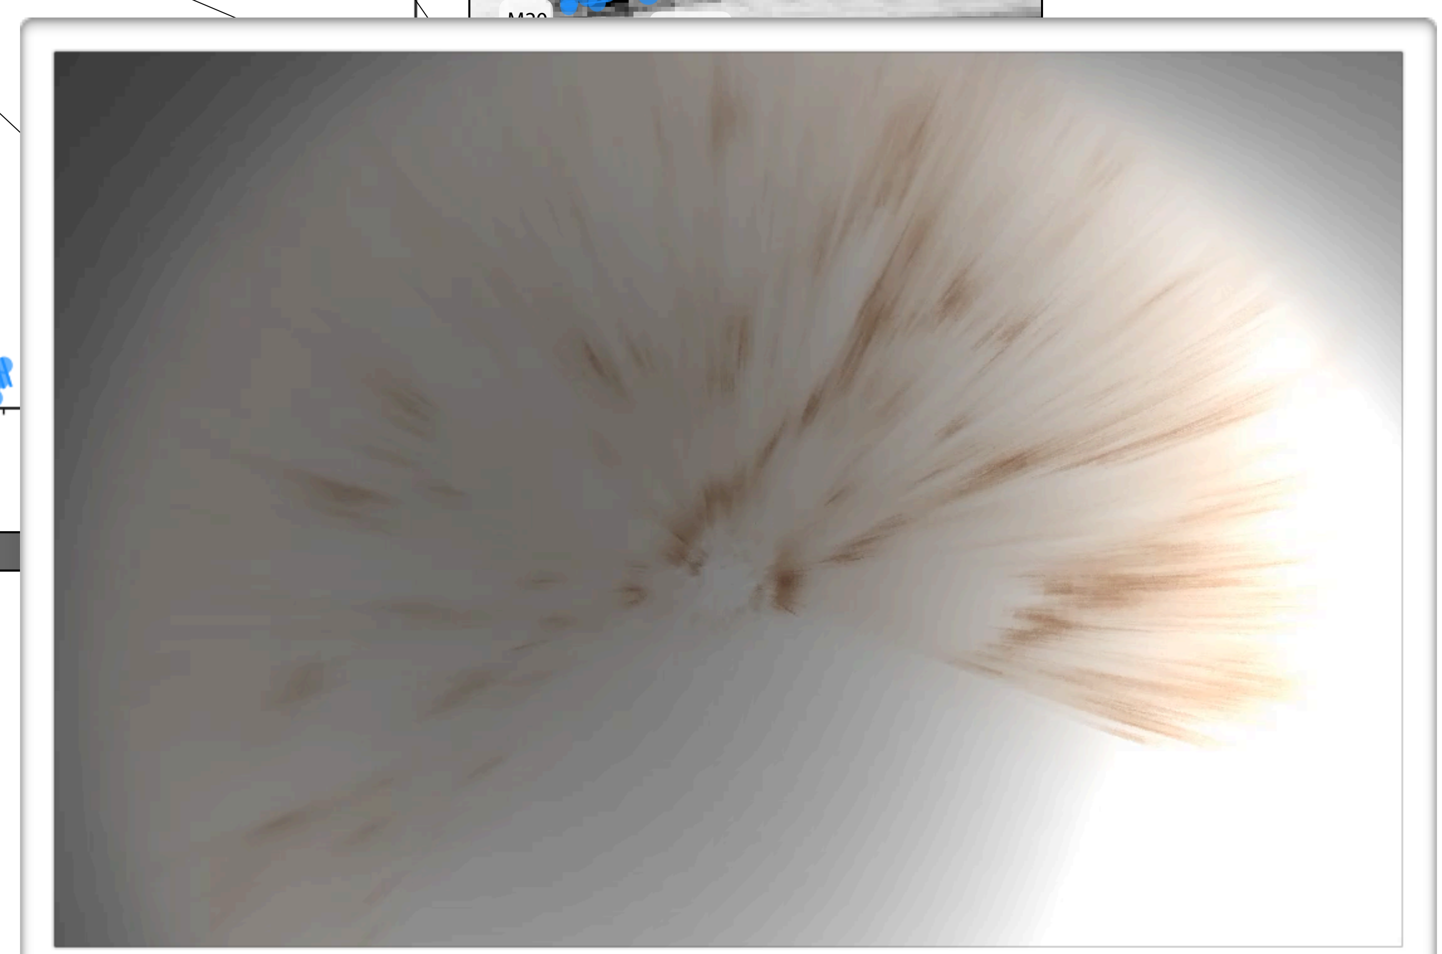

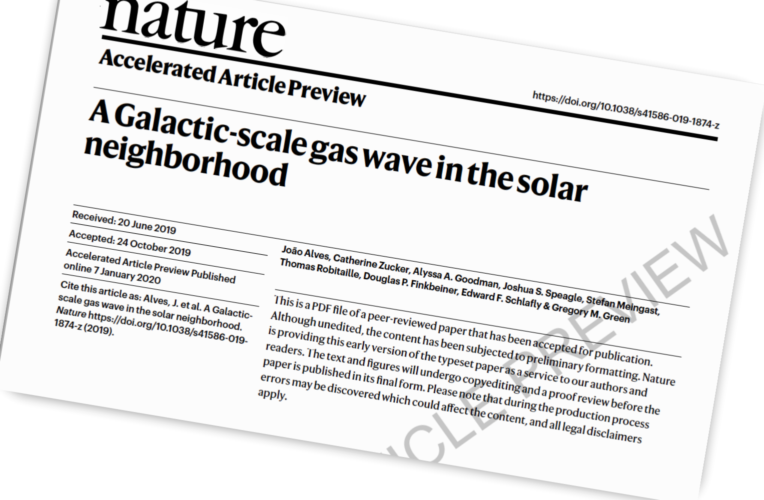
nature
Accelerated Article Preview
https://doi.org/10.1038/s41586-019-1874-z
A Galactic-scale gas wave in the solar neighborhood
Received: 20 June 2019
Accepted: 24 October 2019
Accelerated Article Preview Published online 7 January 2020
Cite this article as: Alves, J. et al. A Galactic-scale gas wave in the solar neighborhood. Nature https://doi.org/10.1038/s41586-019-1874-z (2019).
João Alves, Catherine Zucker, Alyssa A. Goodman, Joshua S. Speagle, Stefan Meingast, Thomas Robitaille, Douglas P. Finkbeiner, Edward F. Schlafly & Gregory M. Green
This is a PDF file of a peer-reviewed paper that has been accepted for publication. Although unedited, the content has been subjected to preliminary formatting. Nature is providing this early version of the typeset paper as a service to our authors and readers. The text and figures will undergo copyediting and a proof review before the paper is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers apply.
ACCELERATED ARTICLE PREVIEW

The "Radcliffe" Wave

tinyurl.com/radwave

Alves+2020

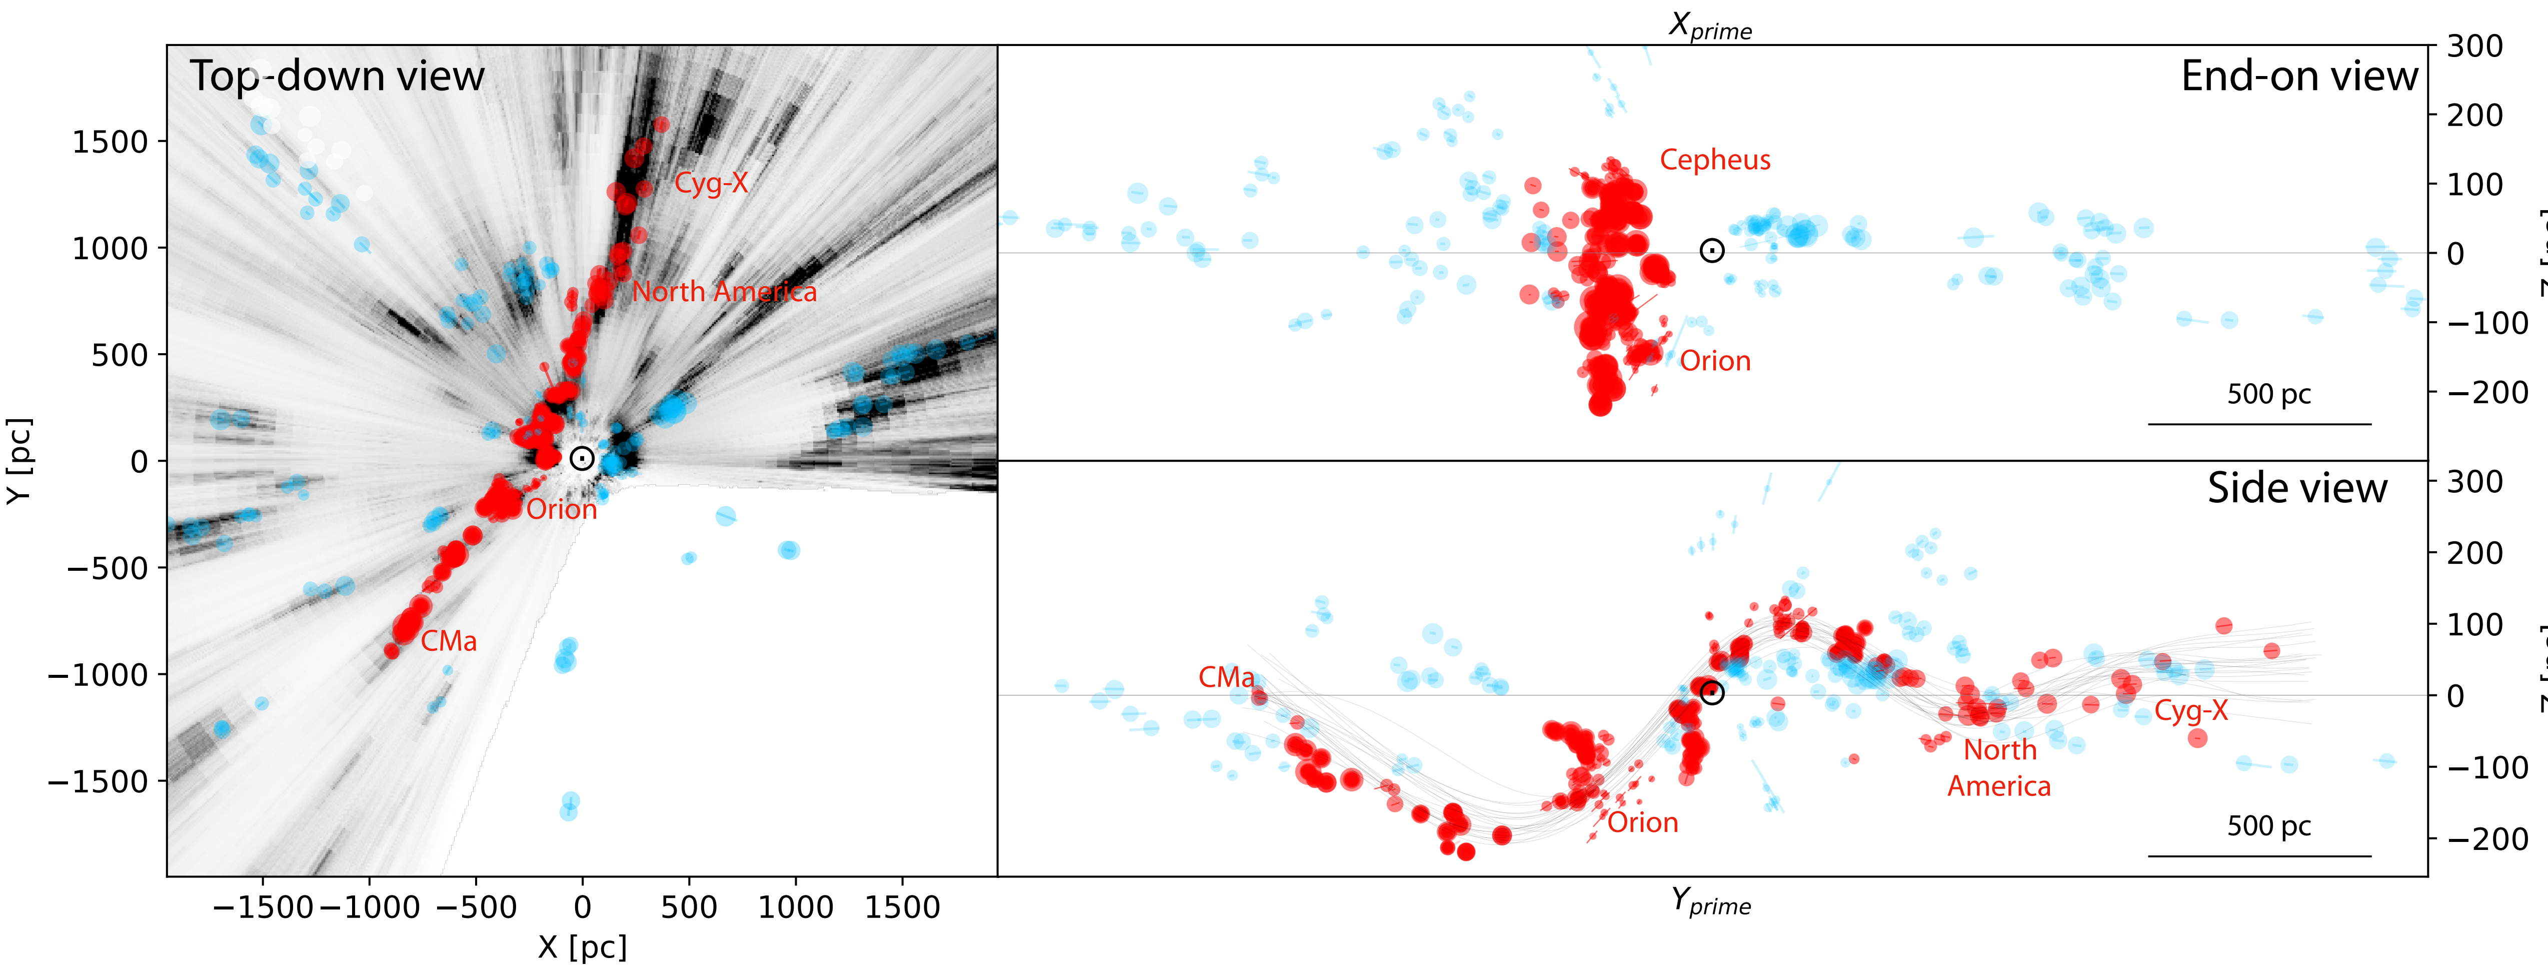

João Alves, Catherine Zucker, Alyssa Goodman, Joshua Speagle,
Stefan Meingast, Thomas Robitaille, Douglas Finkbeiner, Edward F. Schlafly,
and Gregory Green 2020, ***Nature***

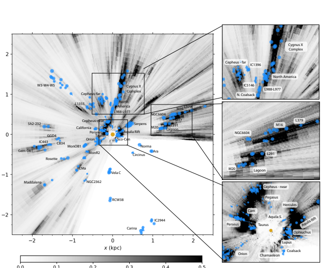

Alves et al. 2020

SCHEMATIC CARTOON(!)

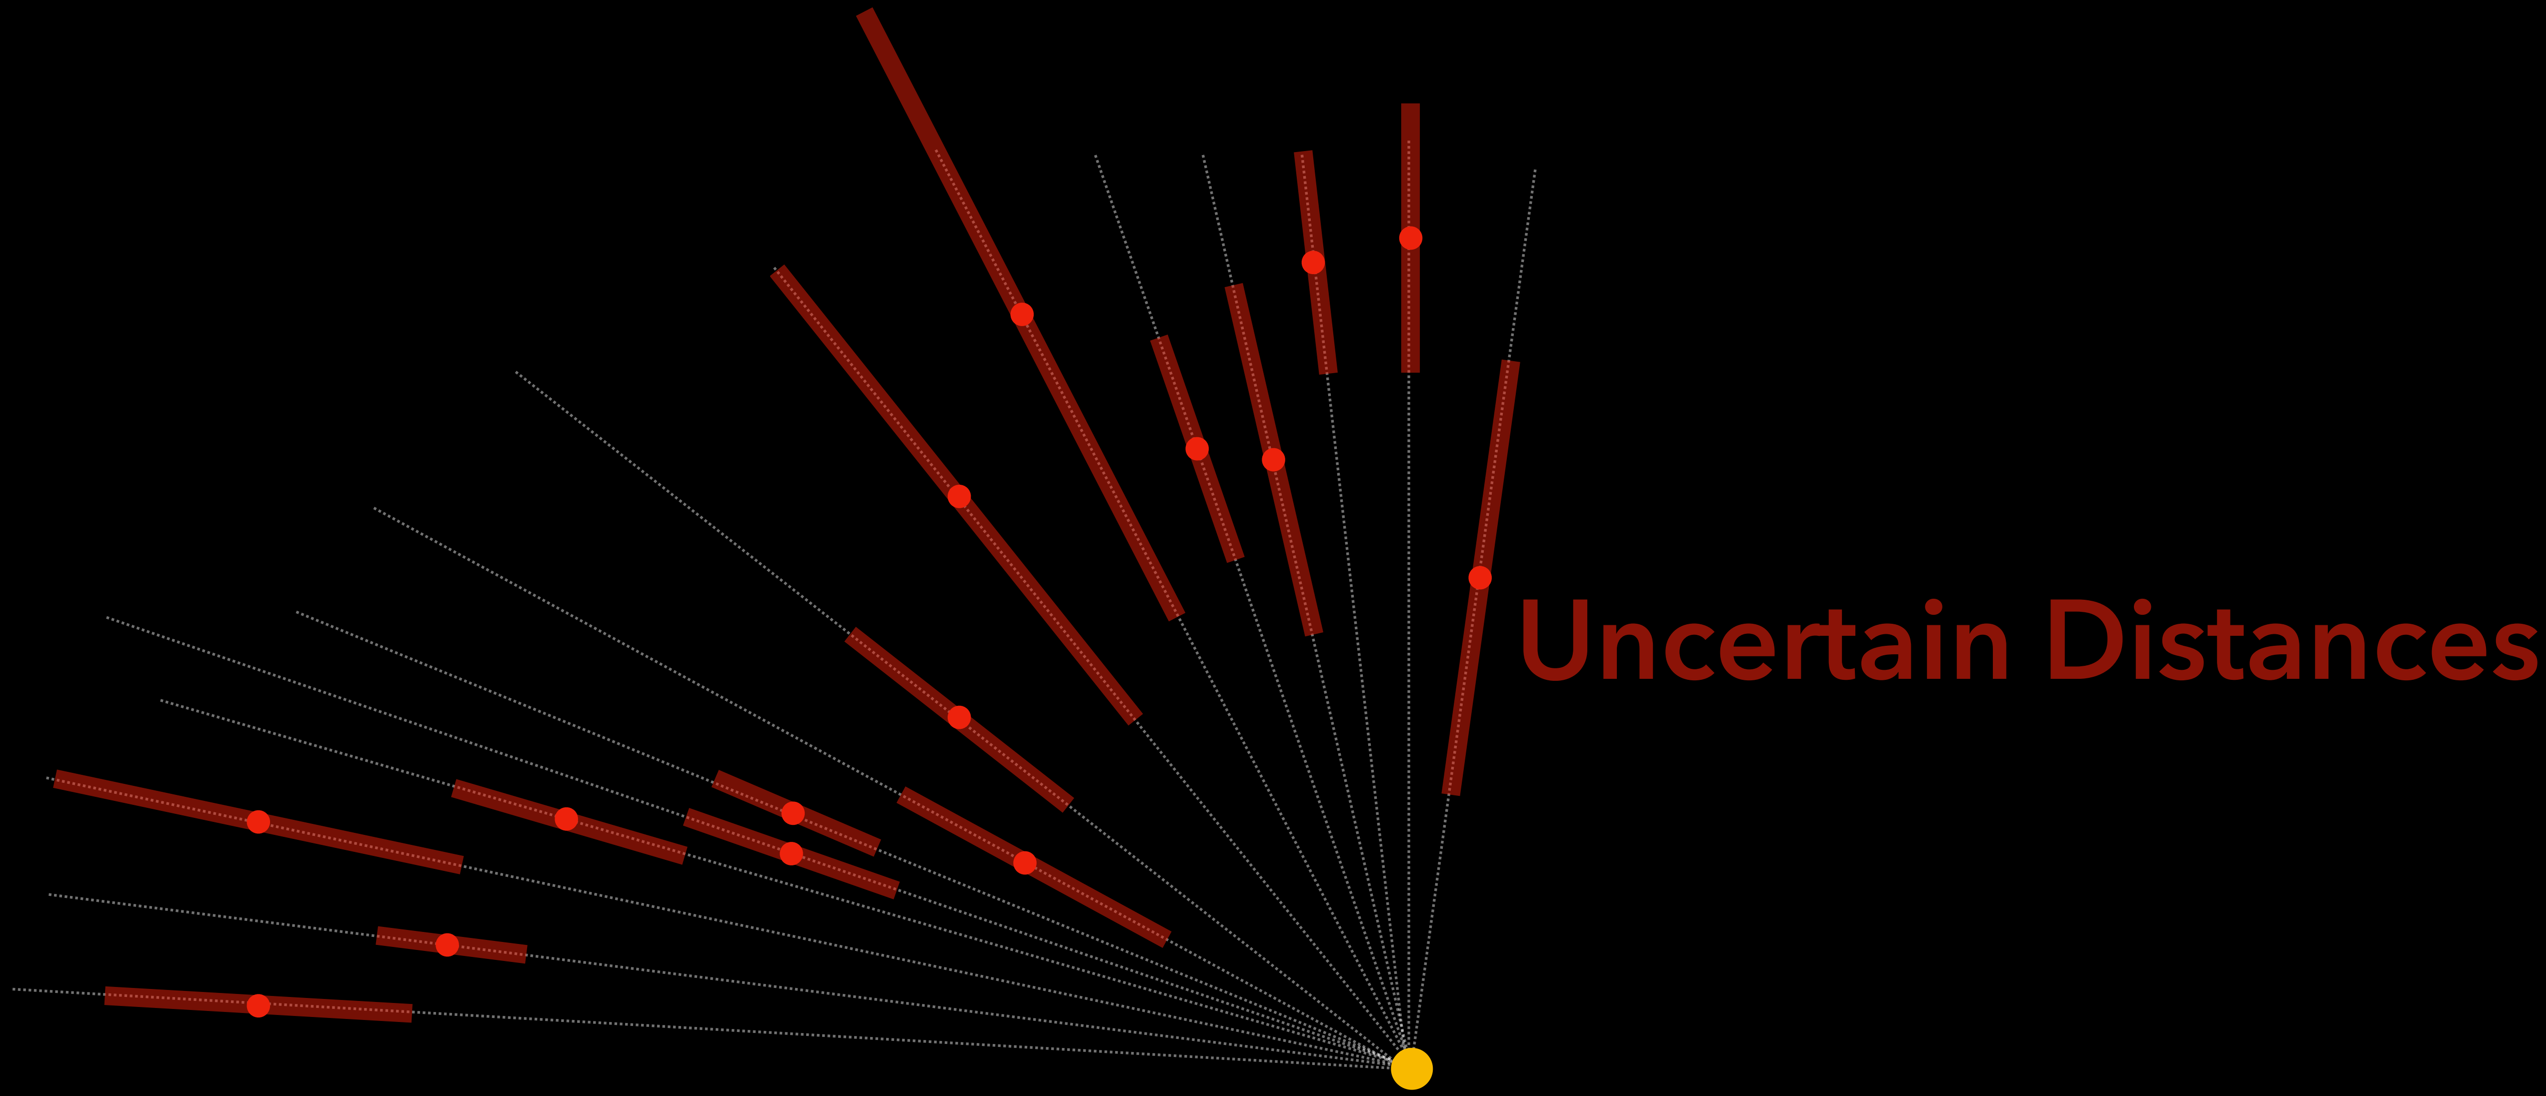

Distances estimates **BEFORE** 3D dust mapping & Gaia (~30%)

SCHEMATIC CARTOON(!)

Distances estimates **AFTER** 3D dust mapping & Gaia (~5%)

Modeling the Radcliffe Wave

- Model the Radcliffe Wave as a quadratic function in (X,Y, Z) space with respect to three “anchor points”: (x_0,y_0,z_0), (x_1,y_1,z_1), (x_2,y_2,z_2)

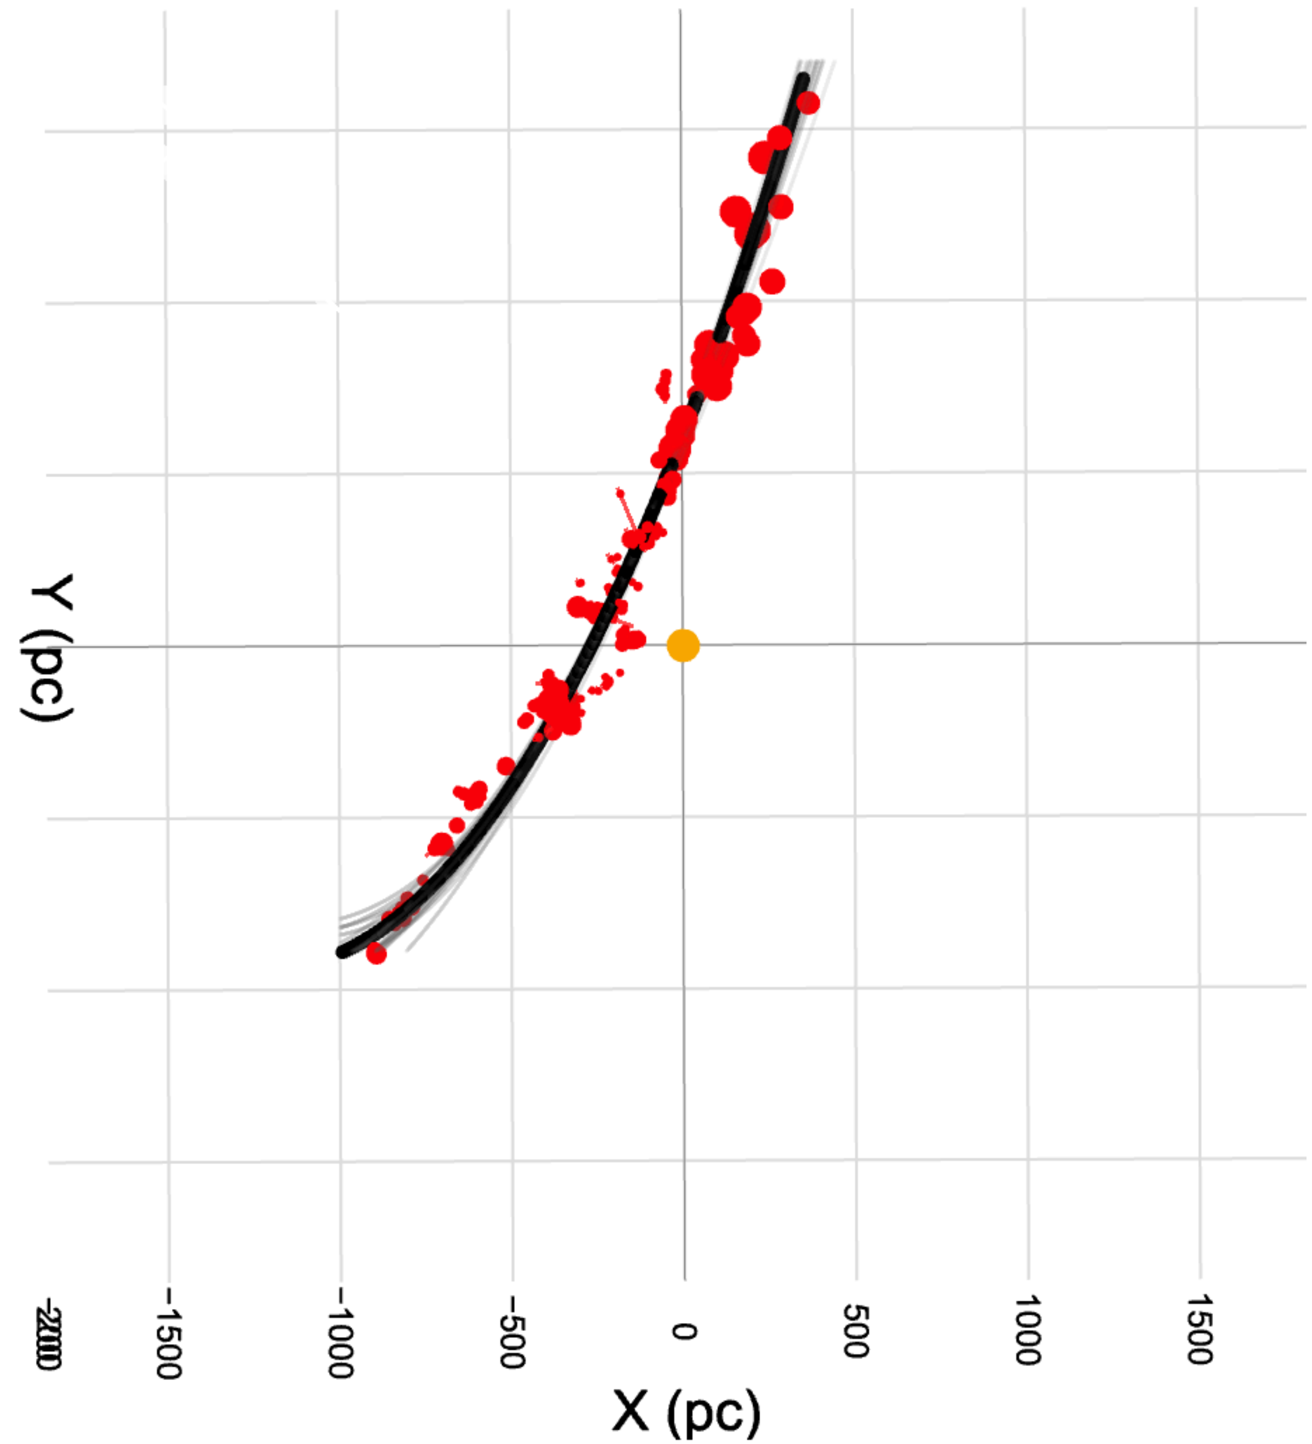

Modeling the Radcliffe Wave

- Undulating behavior parameterized as a damped sinusoidal function with decaying period and amplitude

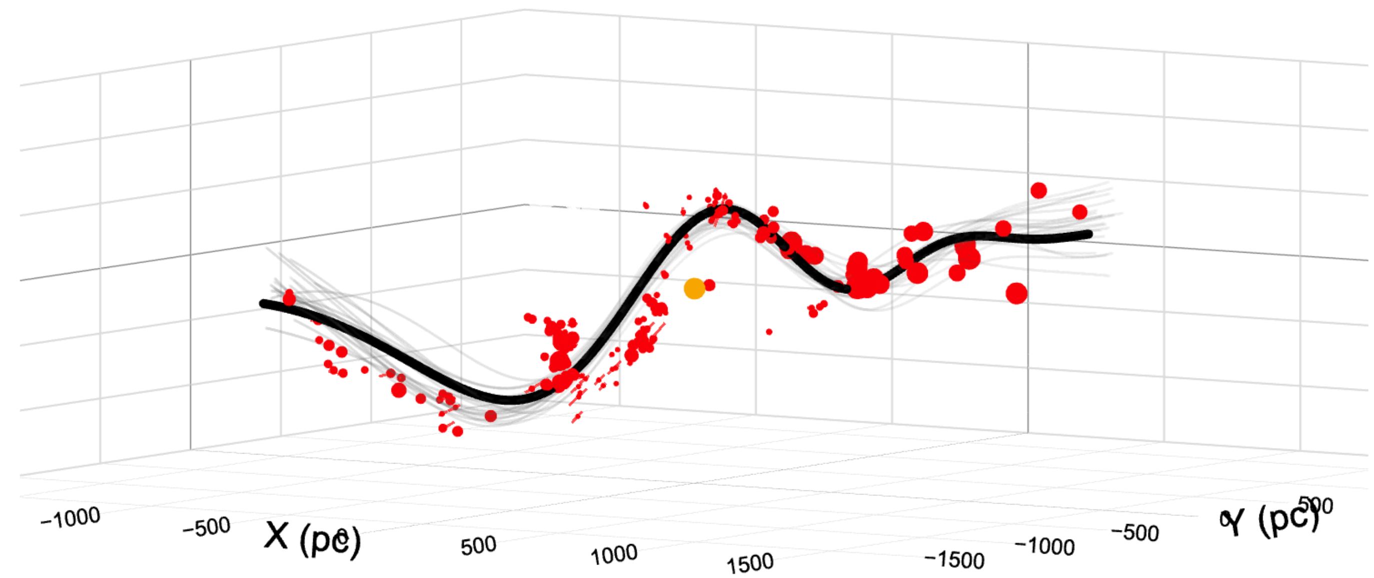

$$d(t) = ||(x,y,z)(t) - (x_0,y_0,z_0)|| = \sqrt{(x-x_0)^2+(y-y_0)^2+(z-z_0)^2}$$

$$\Delta z(t) = A \times \exp\left[-\delta\left(\frac{d(t)}{\text{kpc}}\right)^2\right] \times \sin\left[\left(\frac{2\pi d(t)}{P}\right)\left(1+\frac{d(t)/d_{\max}}{\gamma}\right)+\phi\right]$$

Euclidean distance from start of wave, parameterized by “t" ($d(t)$)

Amplitude (A)

Rate of decay of amplitude (δ)

Period (P)

Rate of decay of period (γ)

Max distance between start & end of wave ($d_{\max}$)

Phase (ϕ)

Modeling the Radcliffe Wave

- Distance of each cloud ***d_{cloud}*** (red points) relative to our model is assumed to be normally distributed with some unknown scatter ***σ***:

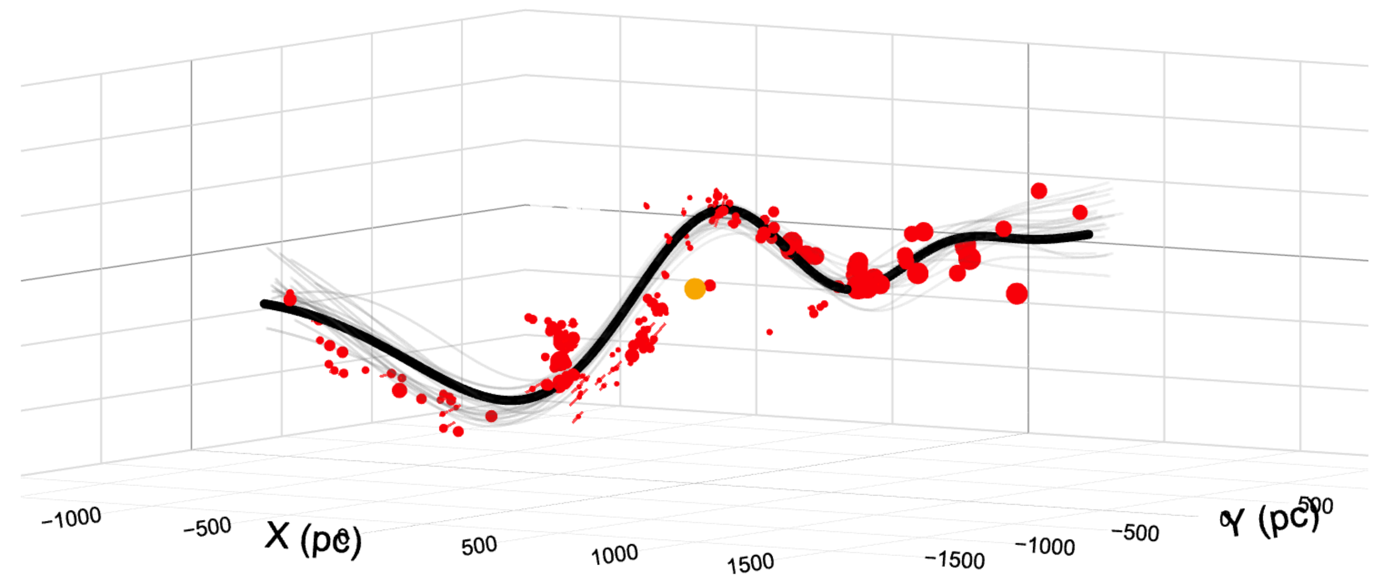

$$d_{\text{cloud}} = \min_t \left(||(x_{\text{cloud}}, y_{\text{cloud}}, z_{\text{cloud}}) - (x_{\text{wave}}, y_{\text{wave}}, z_{\text{wave}})(t)|| \right)$$

- We account for structure "off" the Wave by fitting a mixture model.
    - Some fraction ***f*** of clouds unassociated with Wave are distributed quasi-uniformly in a large volume around the sun
    - Remaining ***1-f*** of clouds associated with Wave

Modeling the Radcliffe Wave

- Likelihood of a realization of our 16-parameter 3D model is given by:

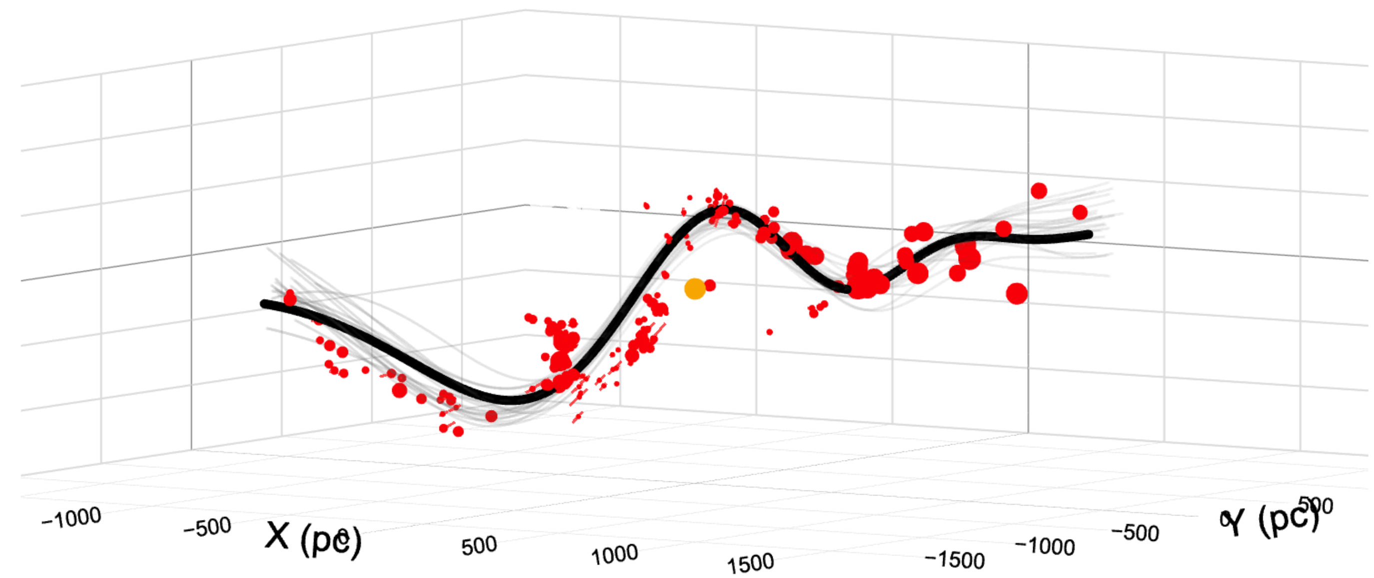

$$\mathcal{L}(\theta) = \prod_{i=1}^{n} \left[(1-f)\mathcal{L}_{\text{cloud},i}(\theta) + f\mathcal{L}_{\text{unif},i}(\theta)\right] \quad (3)$$

where

$$\mathcal{L}_{\text{cloud},i}(\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2}\frac{d^2_{\text{cloud},i}}{\sigma^2}\right], \quad \mathcal{L}_{\text{unif},i}(\theta) = 10^{-7}$$

- Generate samples from posterior using the nested sampling code dynesty (Speagle 2020)

Modeling the Radcliffe Wave

- Using our samples, we associate particular clouds with the Wave by computing the mean odds ratio averaged over the posterior:

$$\langle R_i \rangle = \int \frac{(1-f)\mathcal{L}_{\mathrm{cloud},i}(\theta)}{f\mathcal{L}_{\mathrm{unif},i}(\theta)} \mathcal{P}(\theta) d\theta$$

  - We classify all clouds with $\langle R_i \rangle > 1$ as being part of the Wave

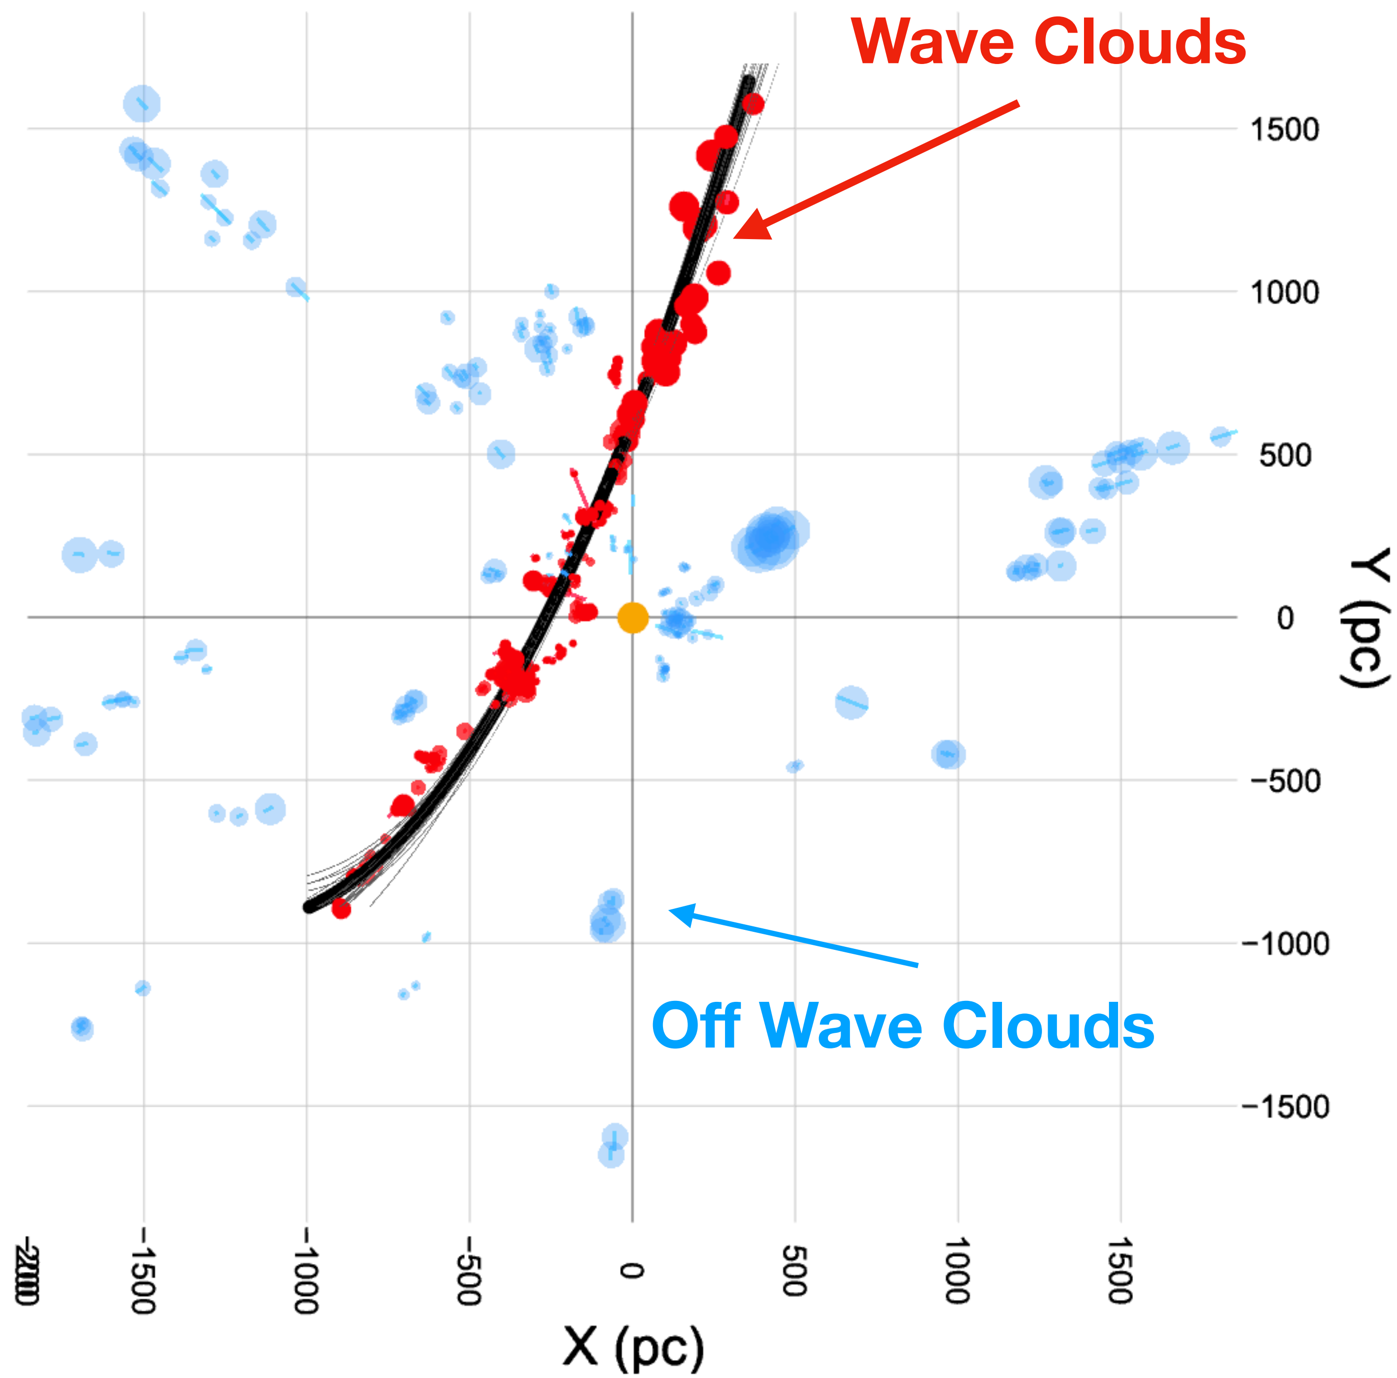

The Future…

- Update our stellar modeling pipeline (switch from "empirical" models to "theoretical" models) to so we can see through more dust at father distances

- Only using a small fraction of the available photometry (~ 1 billion out of ~ 5 billion stars). Incorporate more data, at more wavelengths!

- Star and dust modeling is currently decoupled. These properties should be jointly estimated in the context of a hierarchical model
    - (block-)Gibbs schemes?
    - Importance resampling?

- Your ideas here…

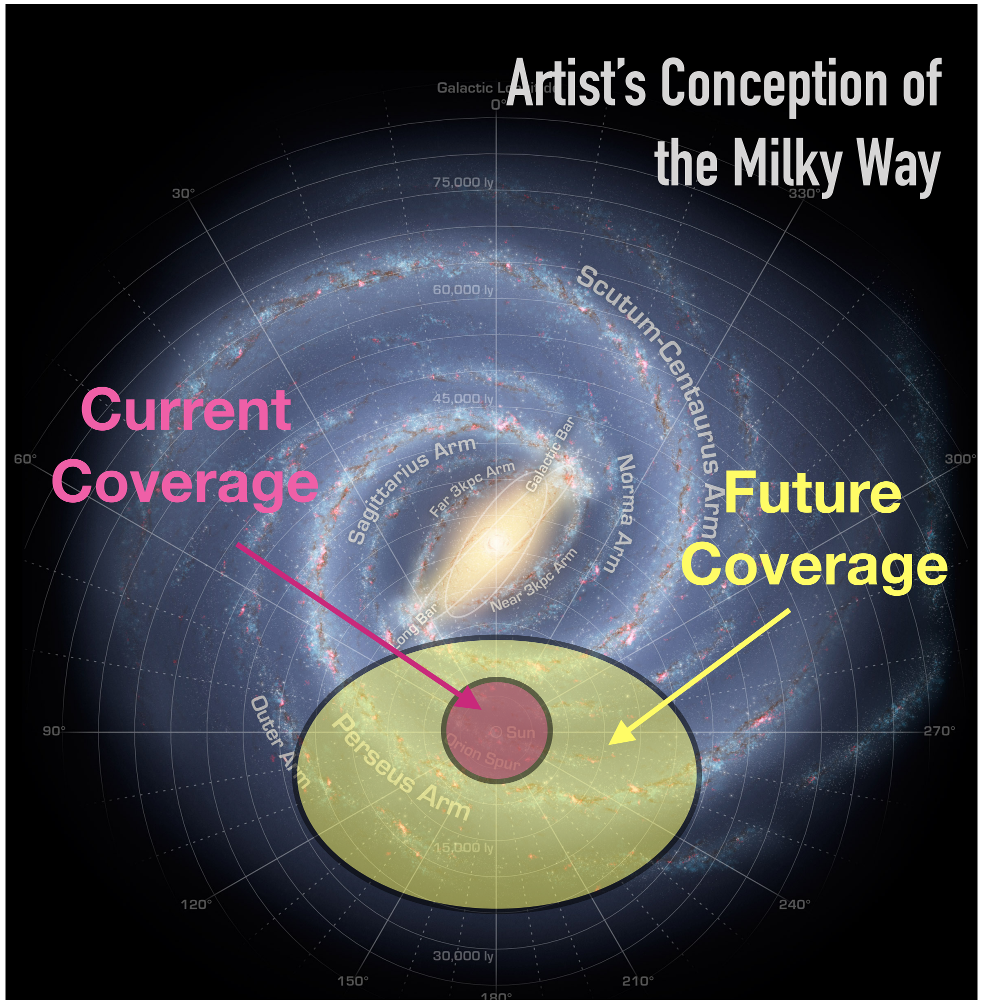

Thanks! Any (more?) questions?