Using the Amazon Compute Cloud to Analyze Large Data Sets

The Panchromatic Hubble Andromeda Treasury

Benjamin F. Williams
University of Washington

Ground Based Image

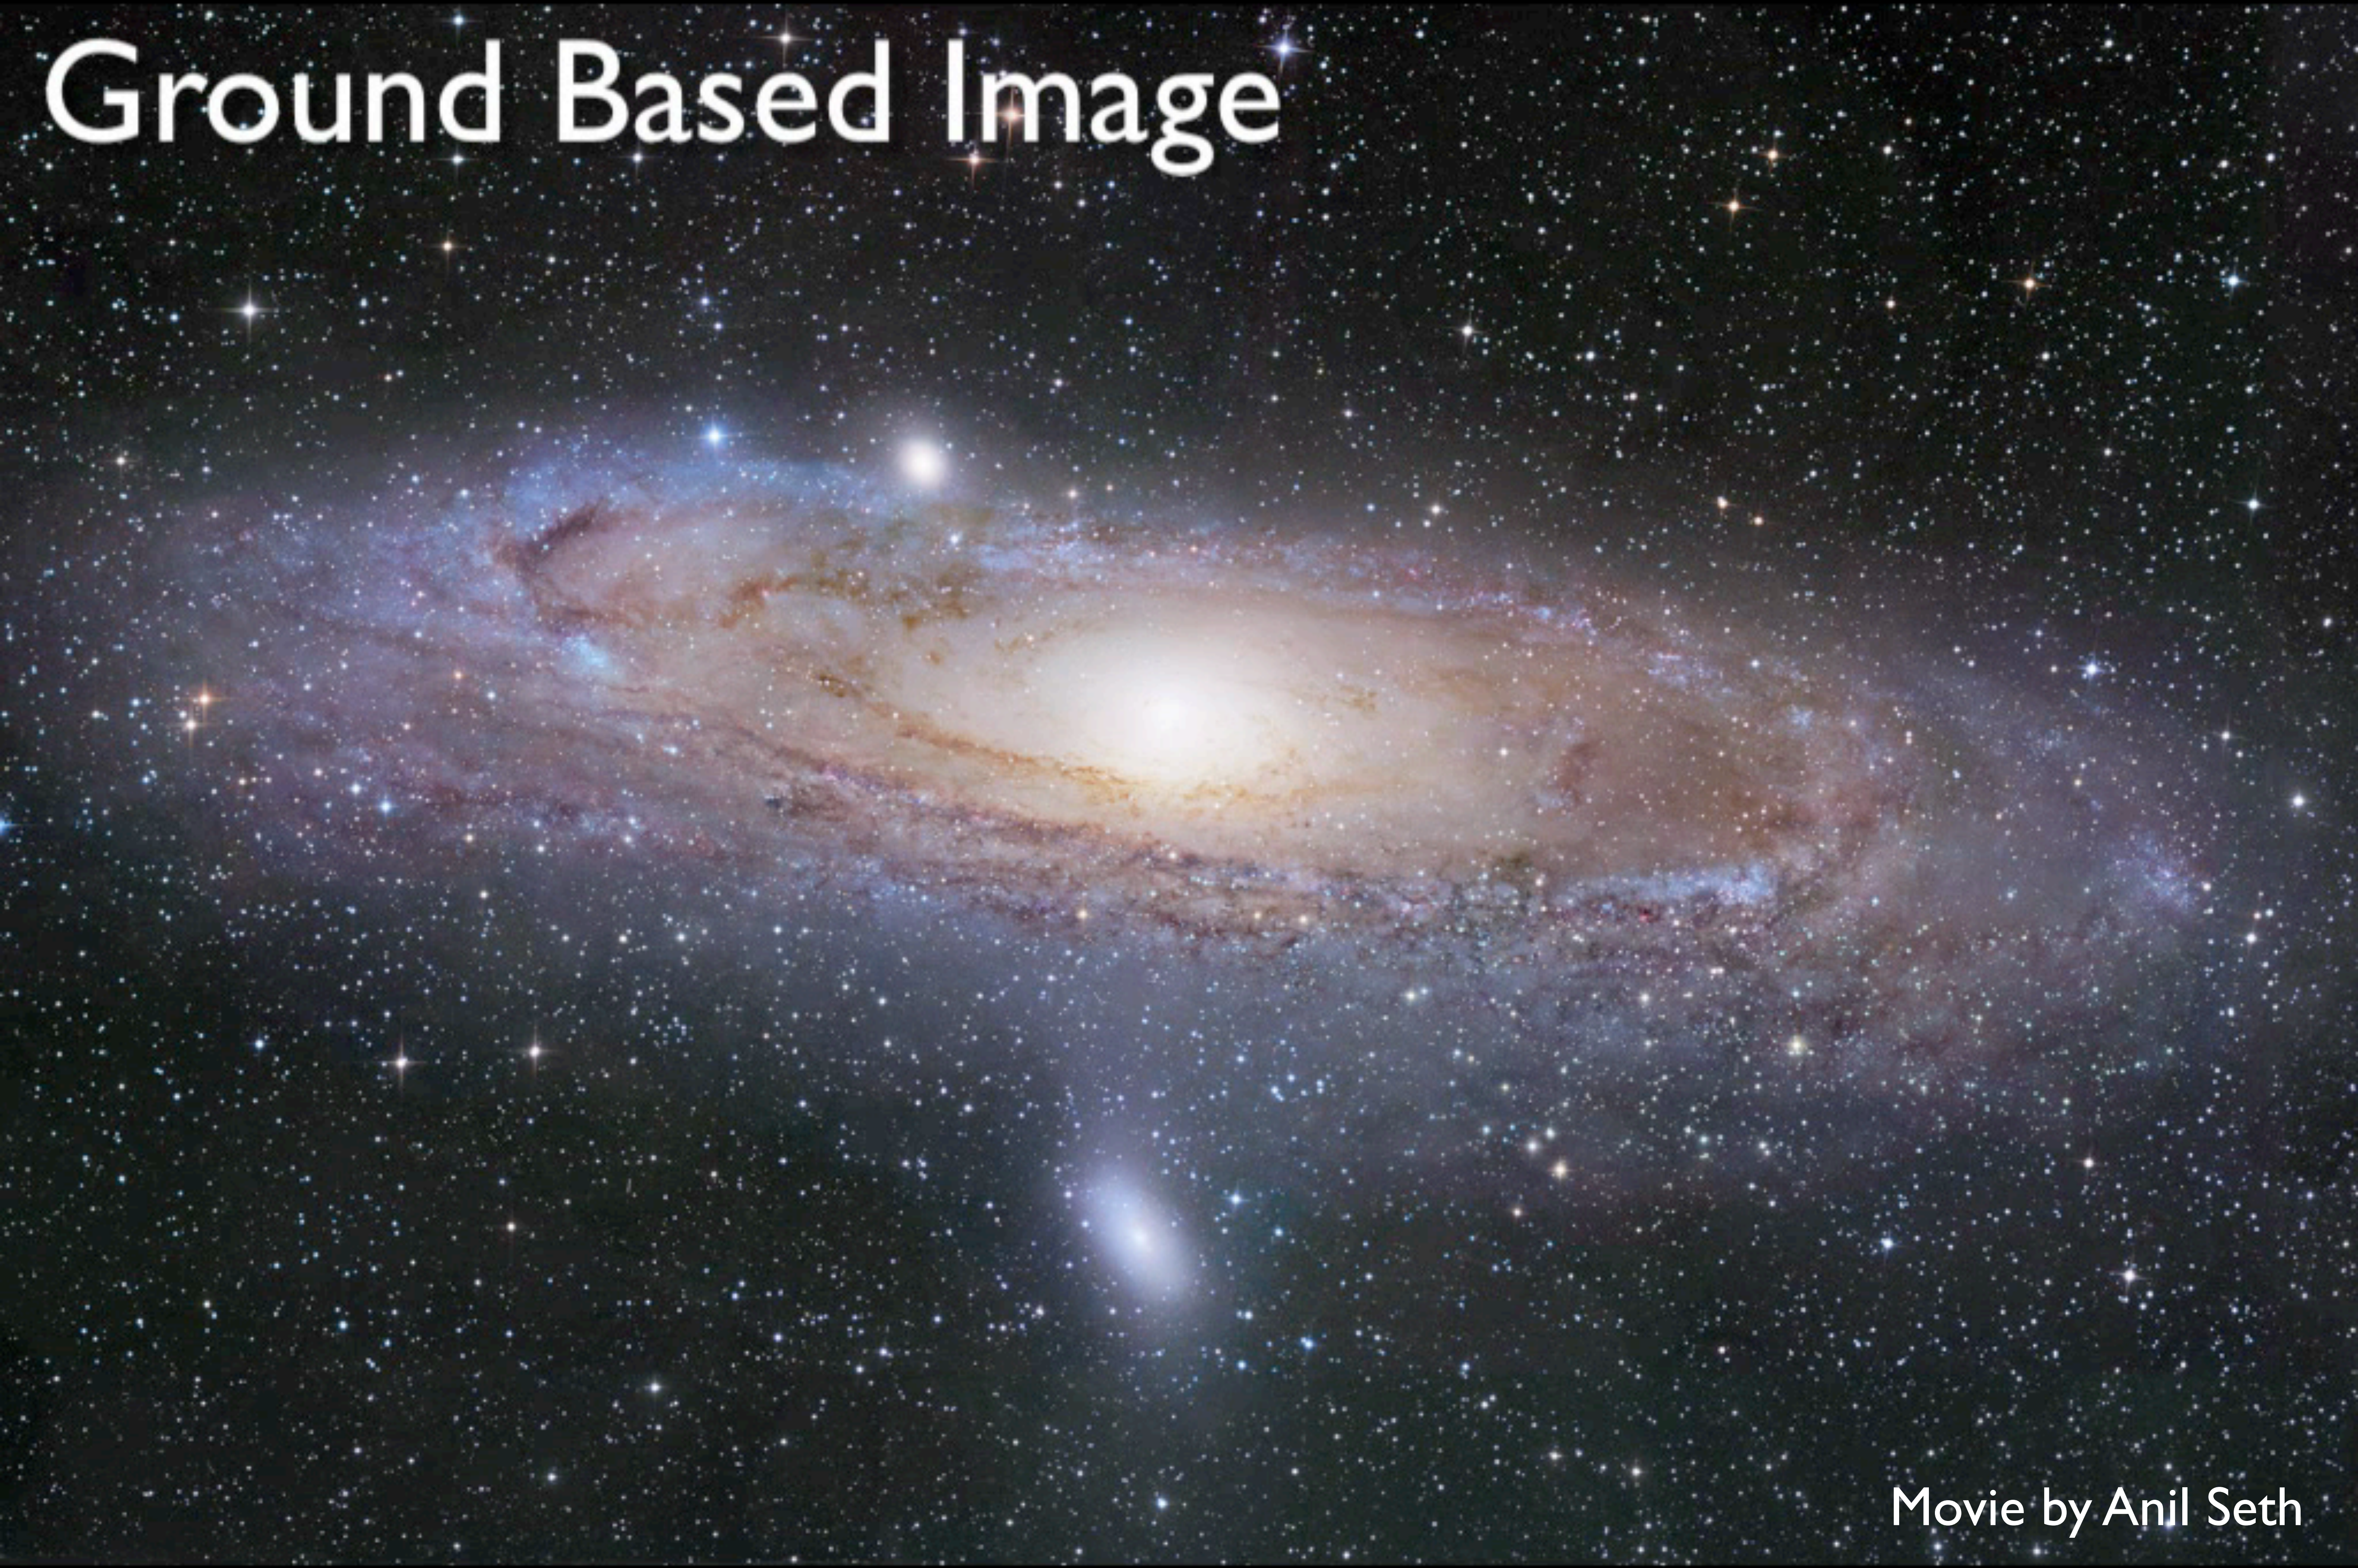

Hubble field of view is small

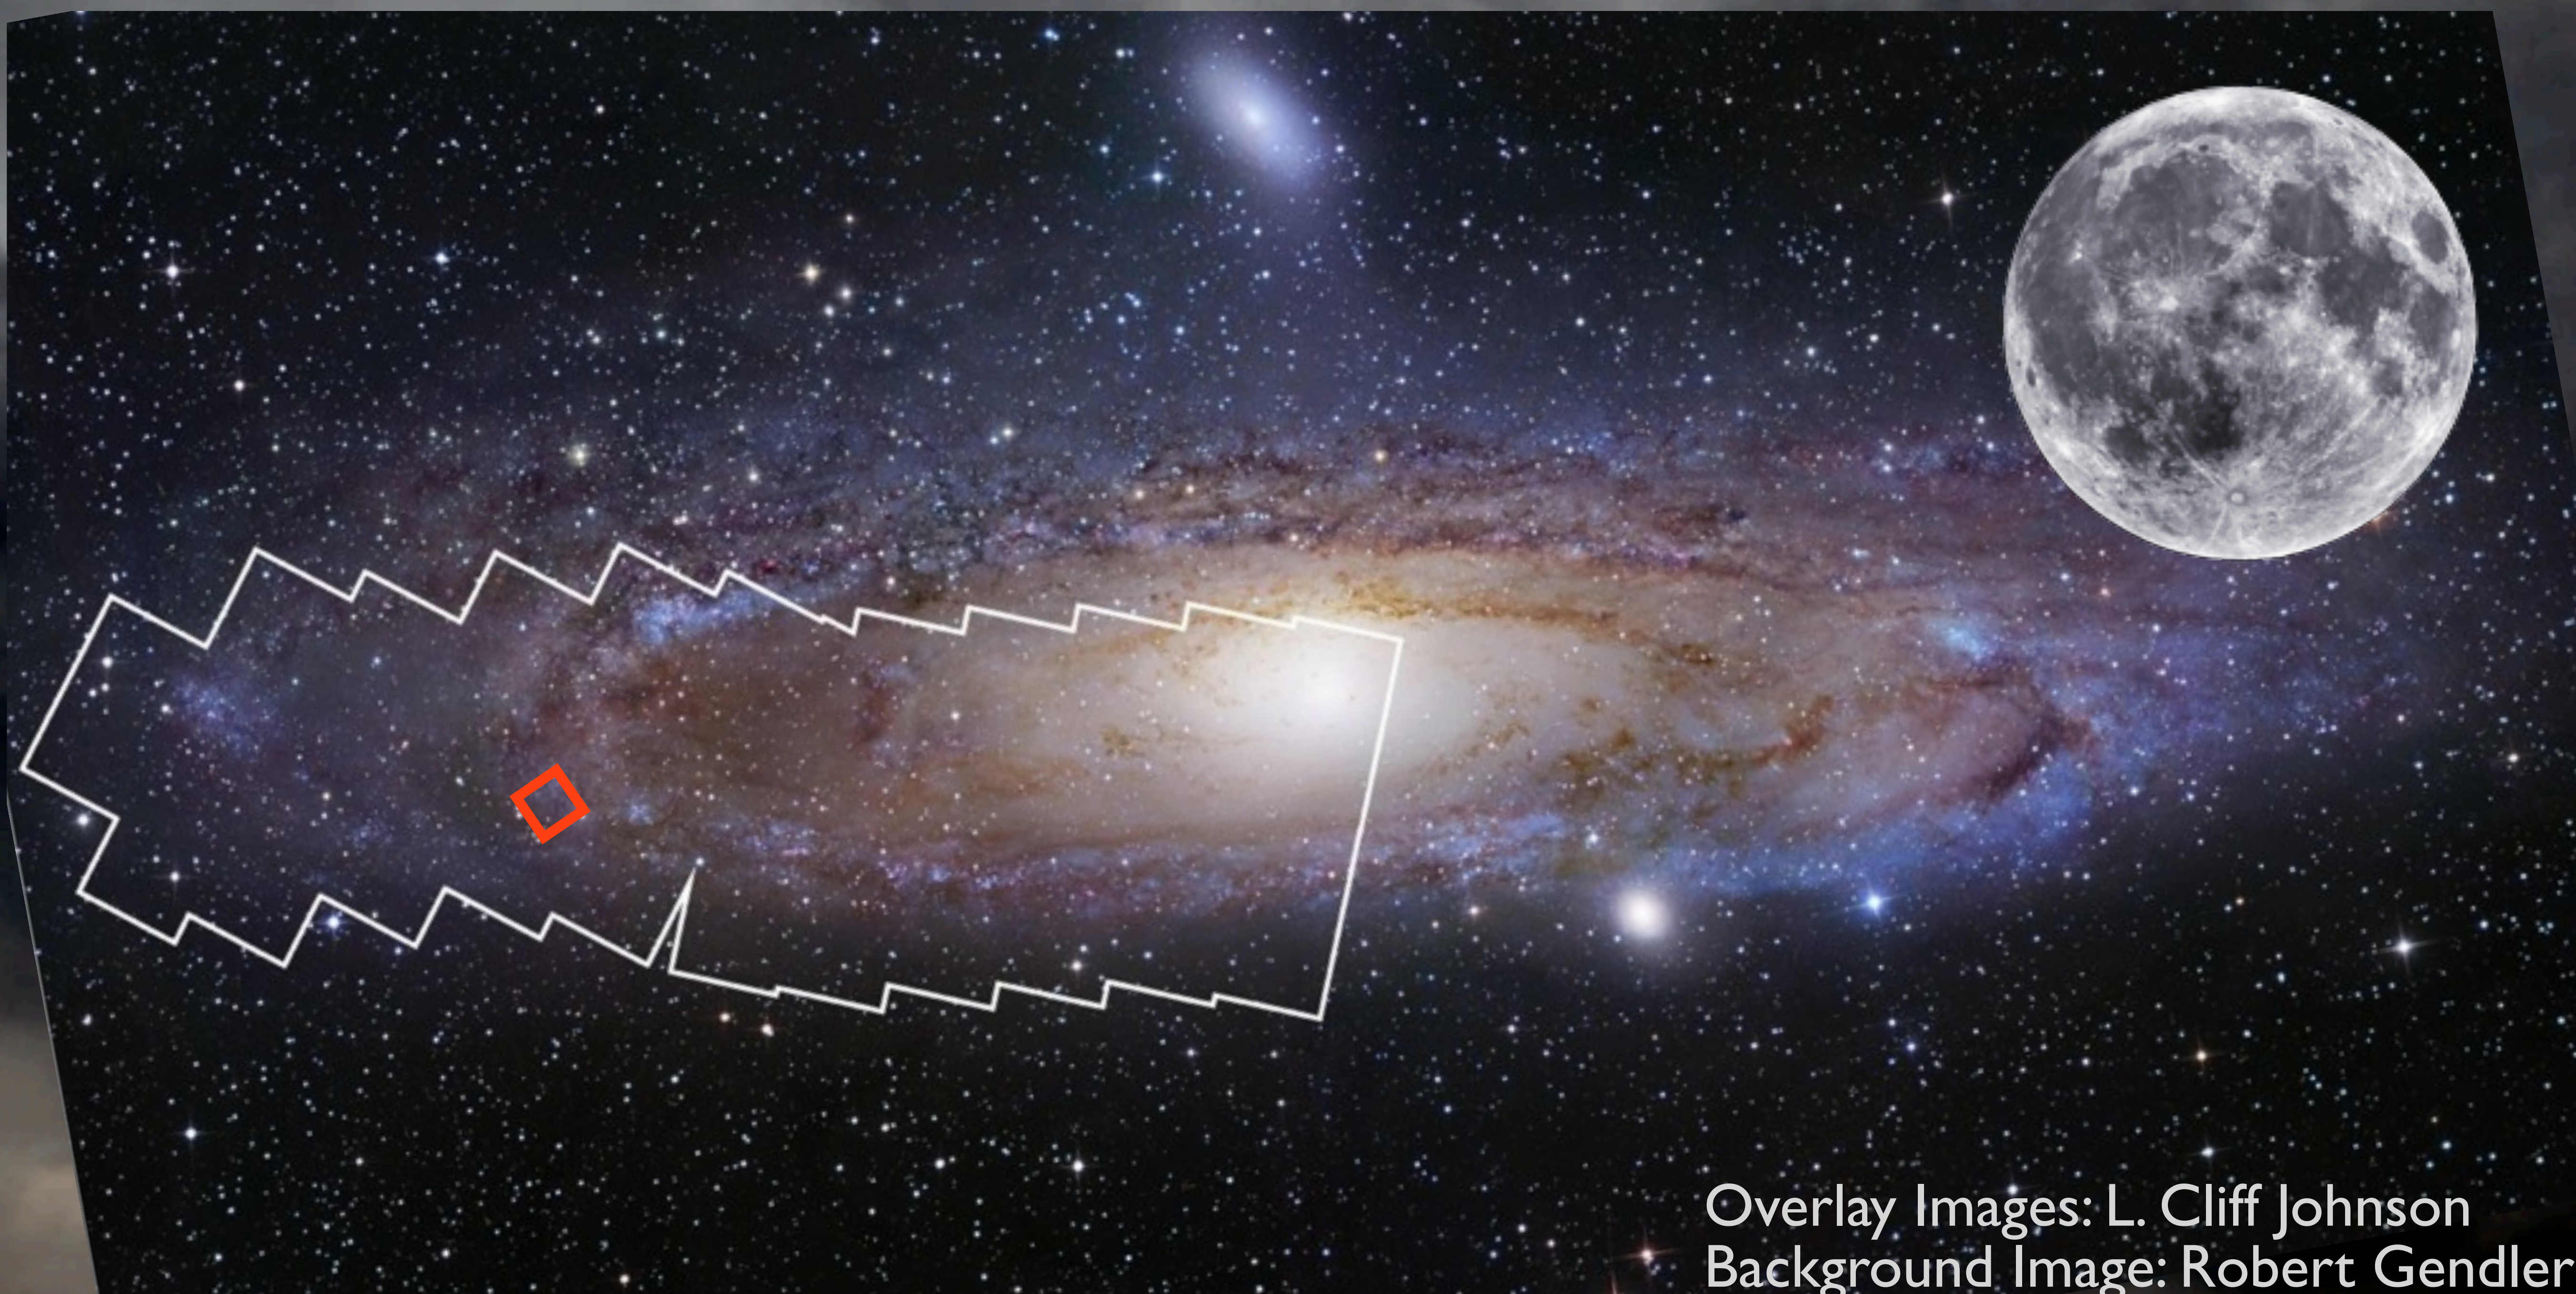

Overlay Images: L. Cliff Johnson
Background Image: Robert Gendler

Coverage in 3 Cameras

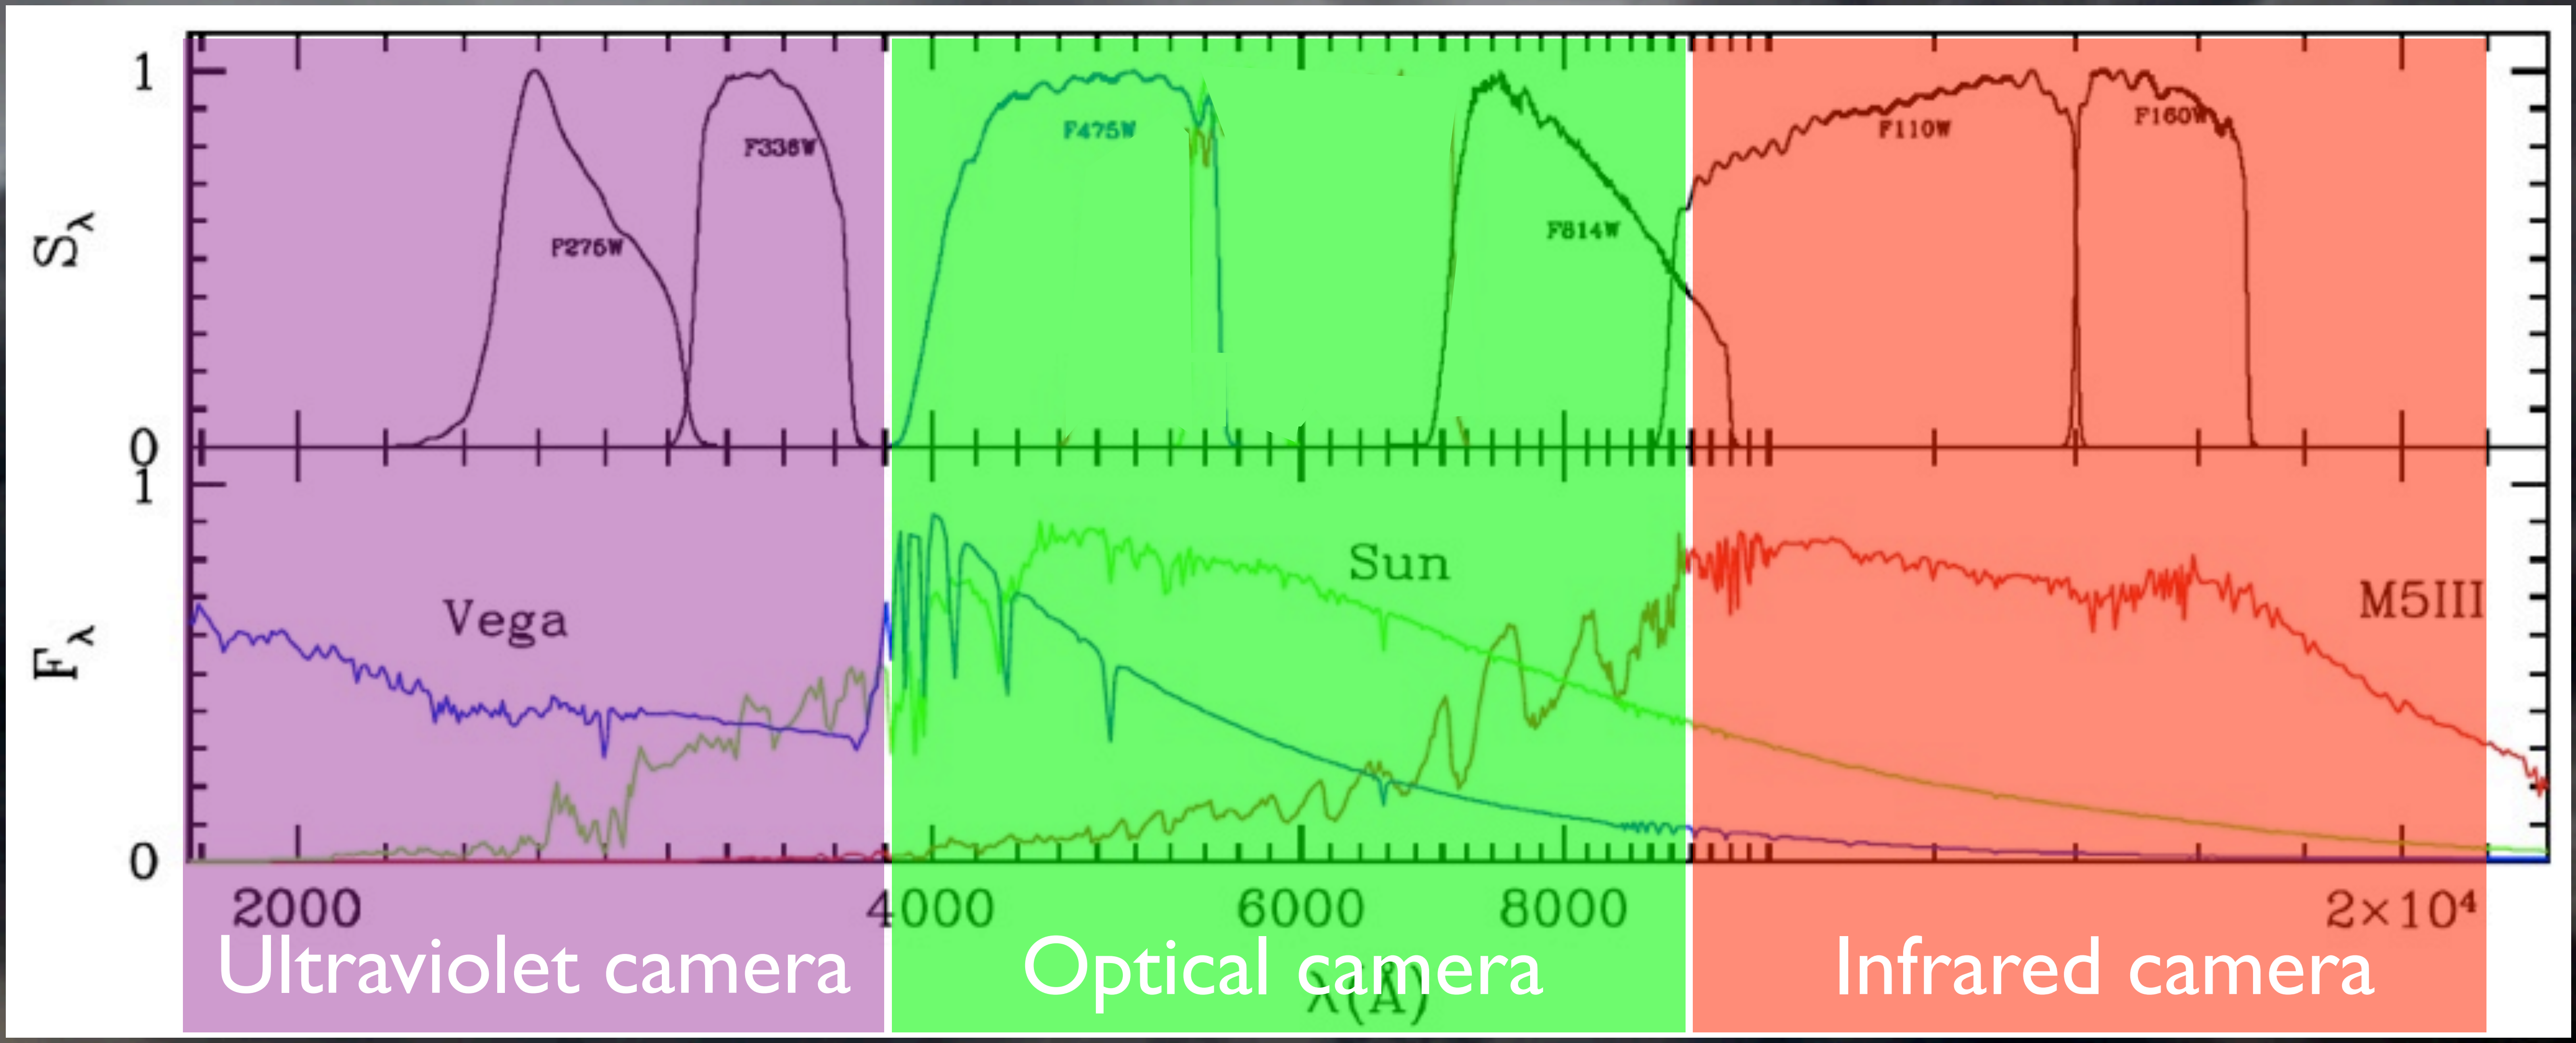

Characterizing all star types requires ultraviolet optical and infrared measurements

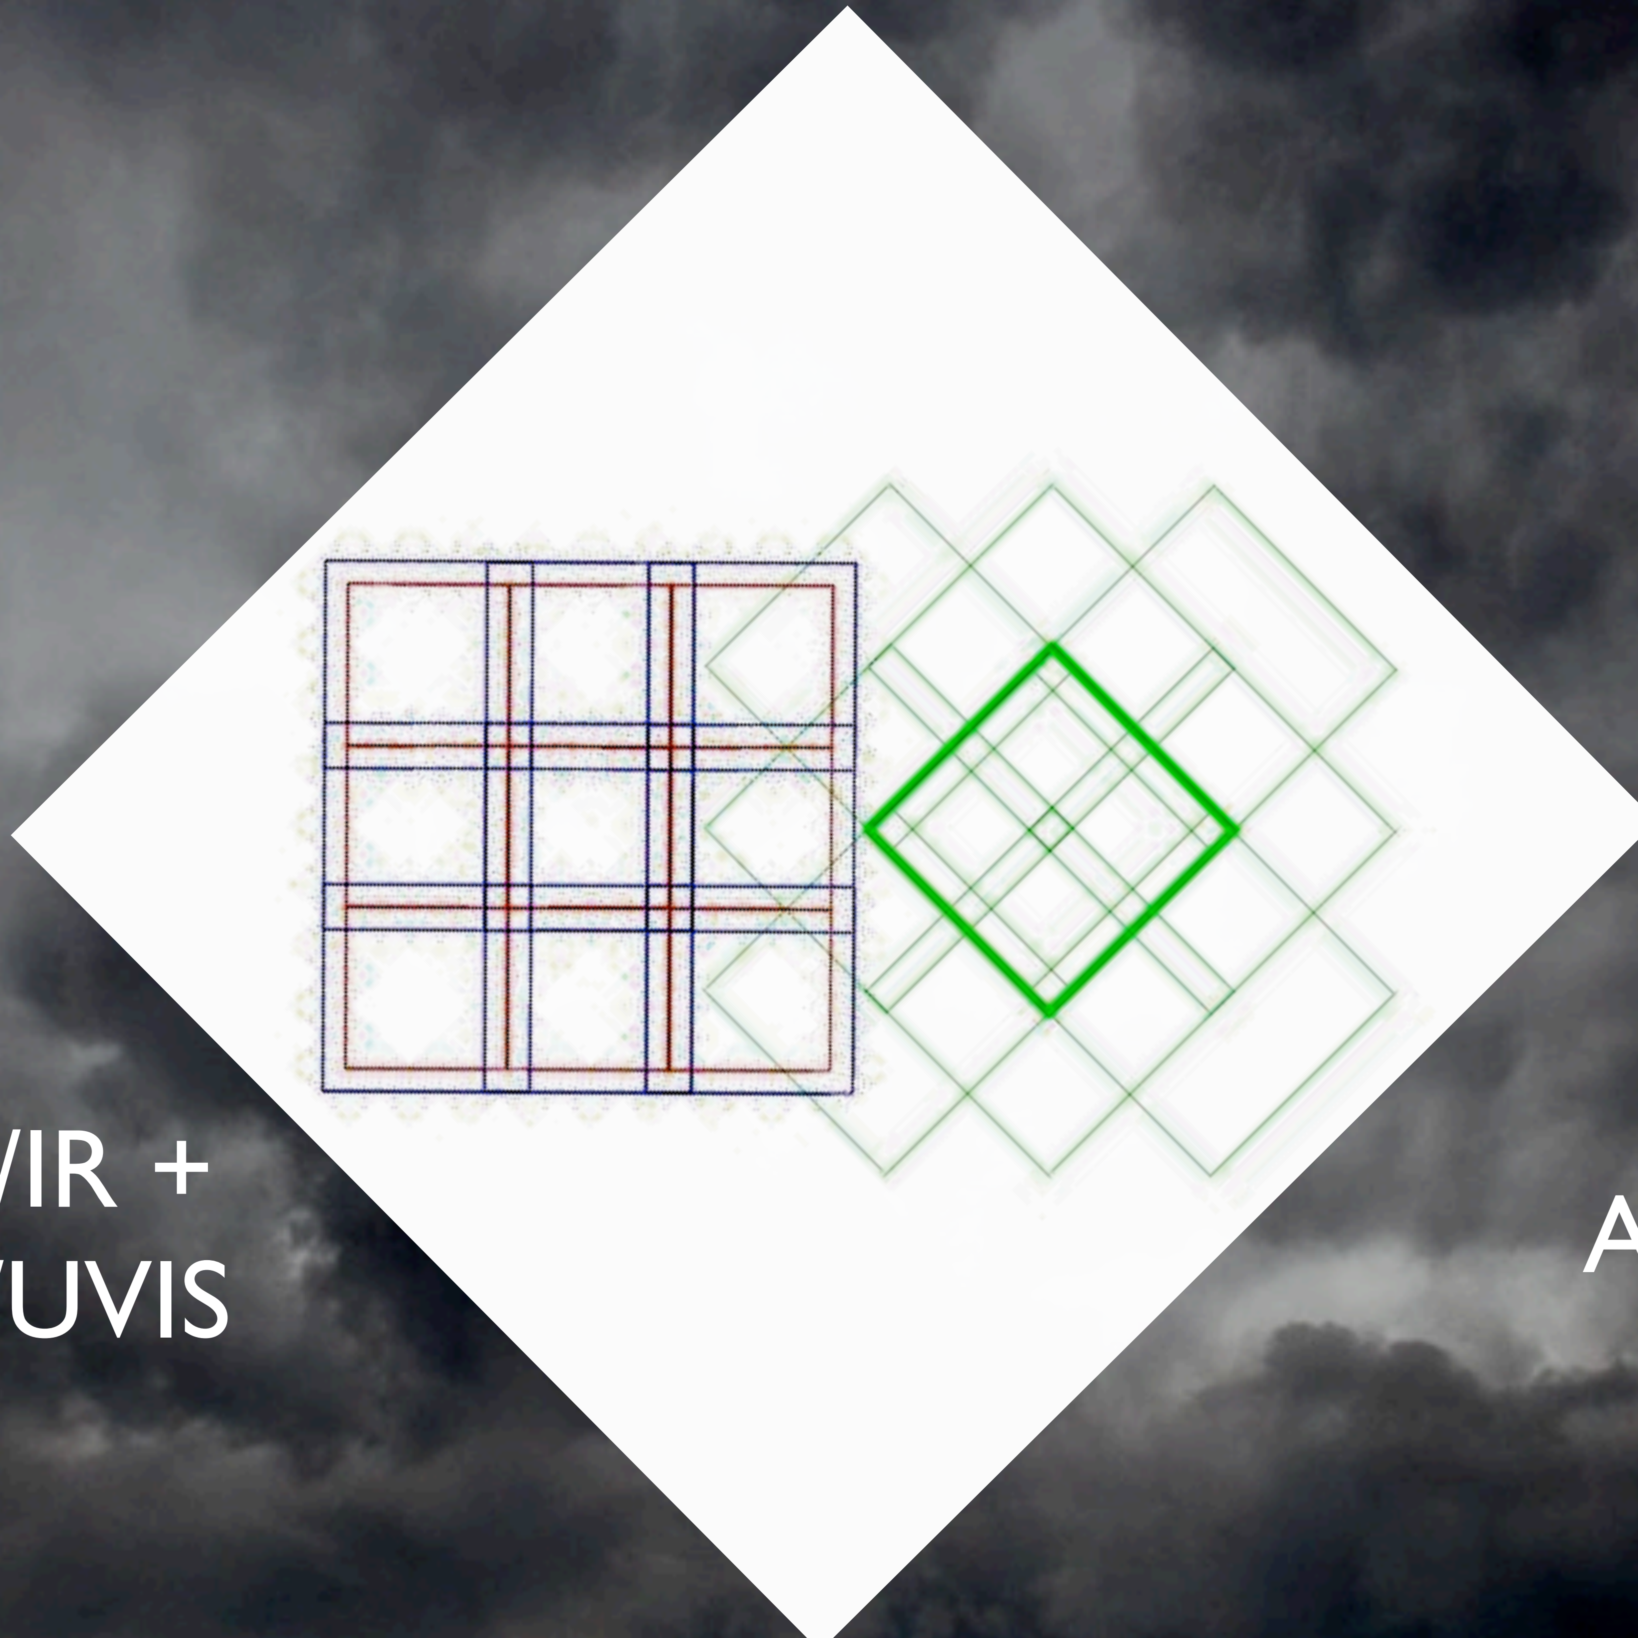

WFC3/IR +
WFC3/UVIS

ACS

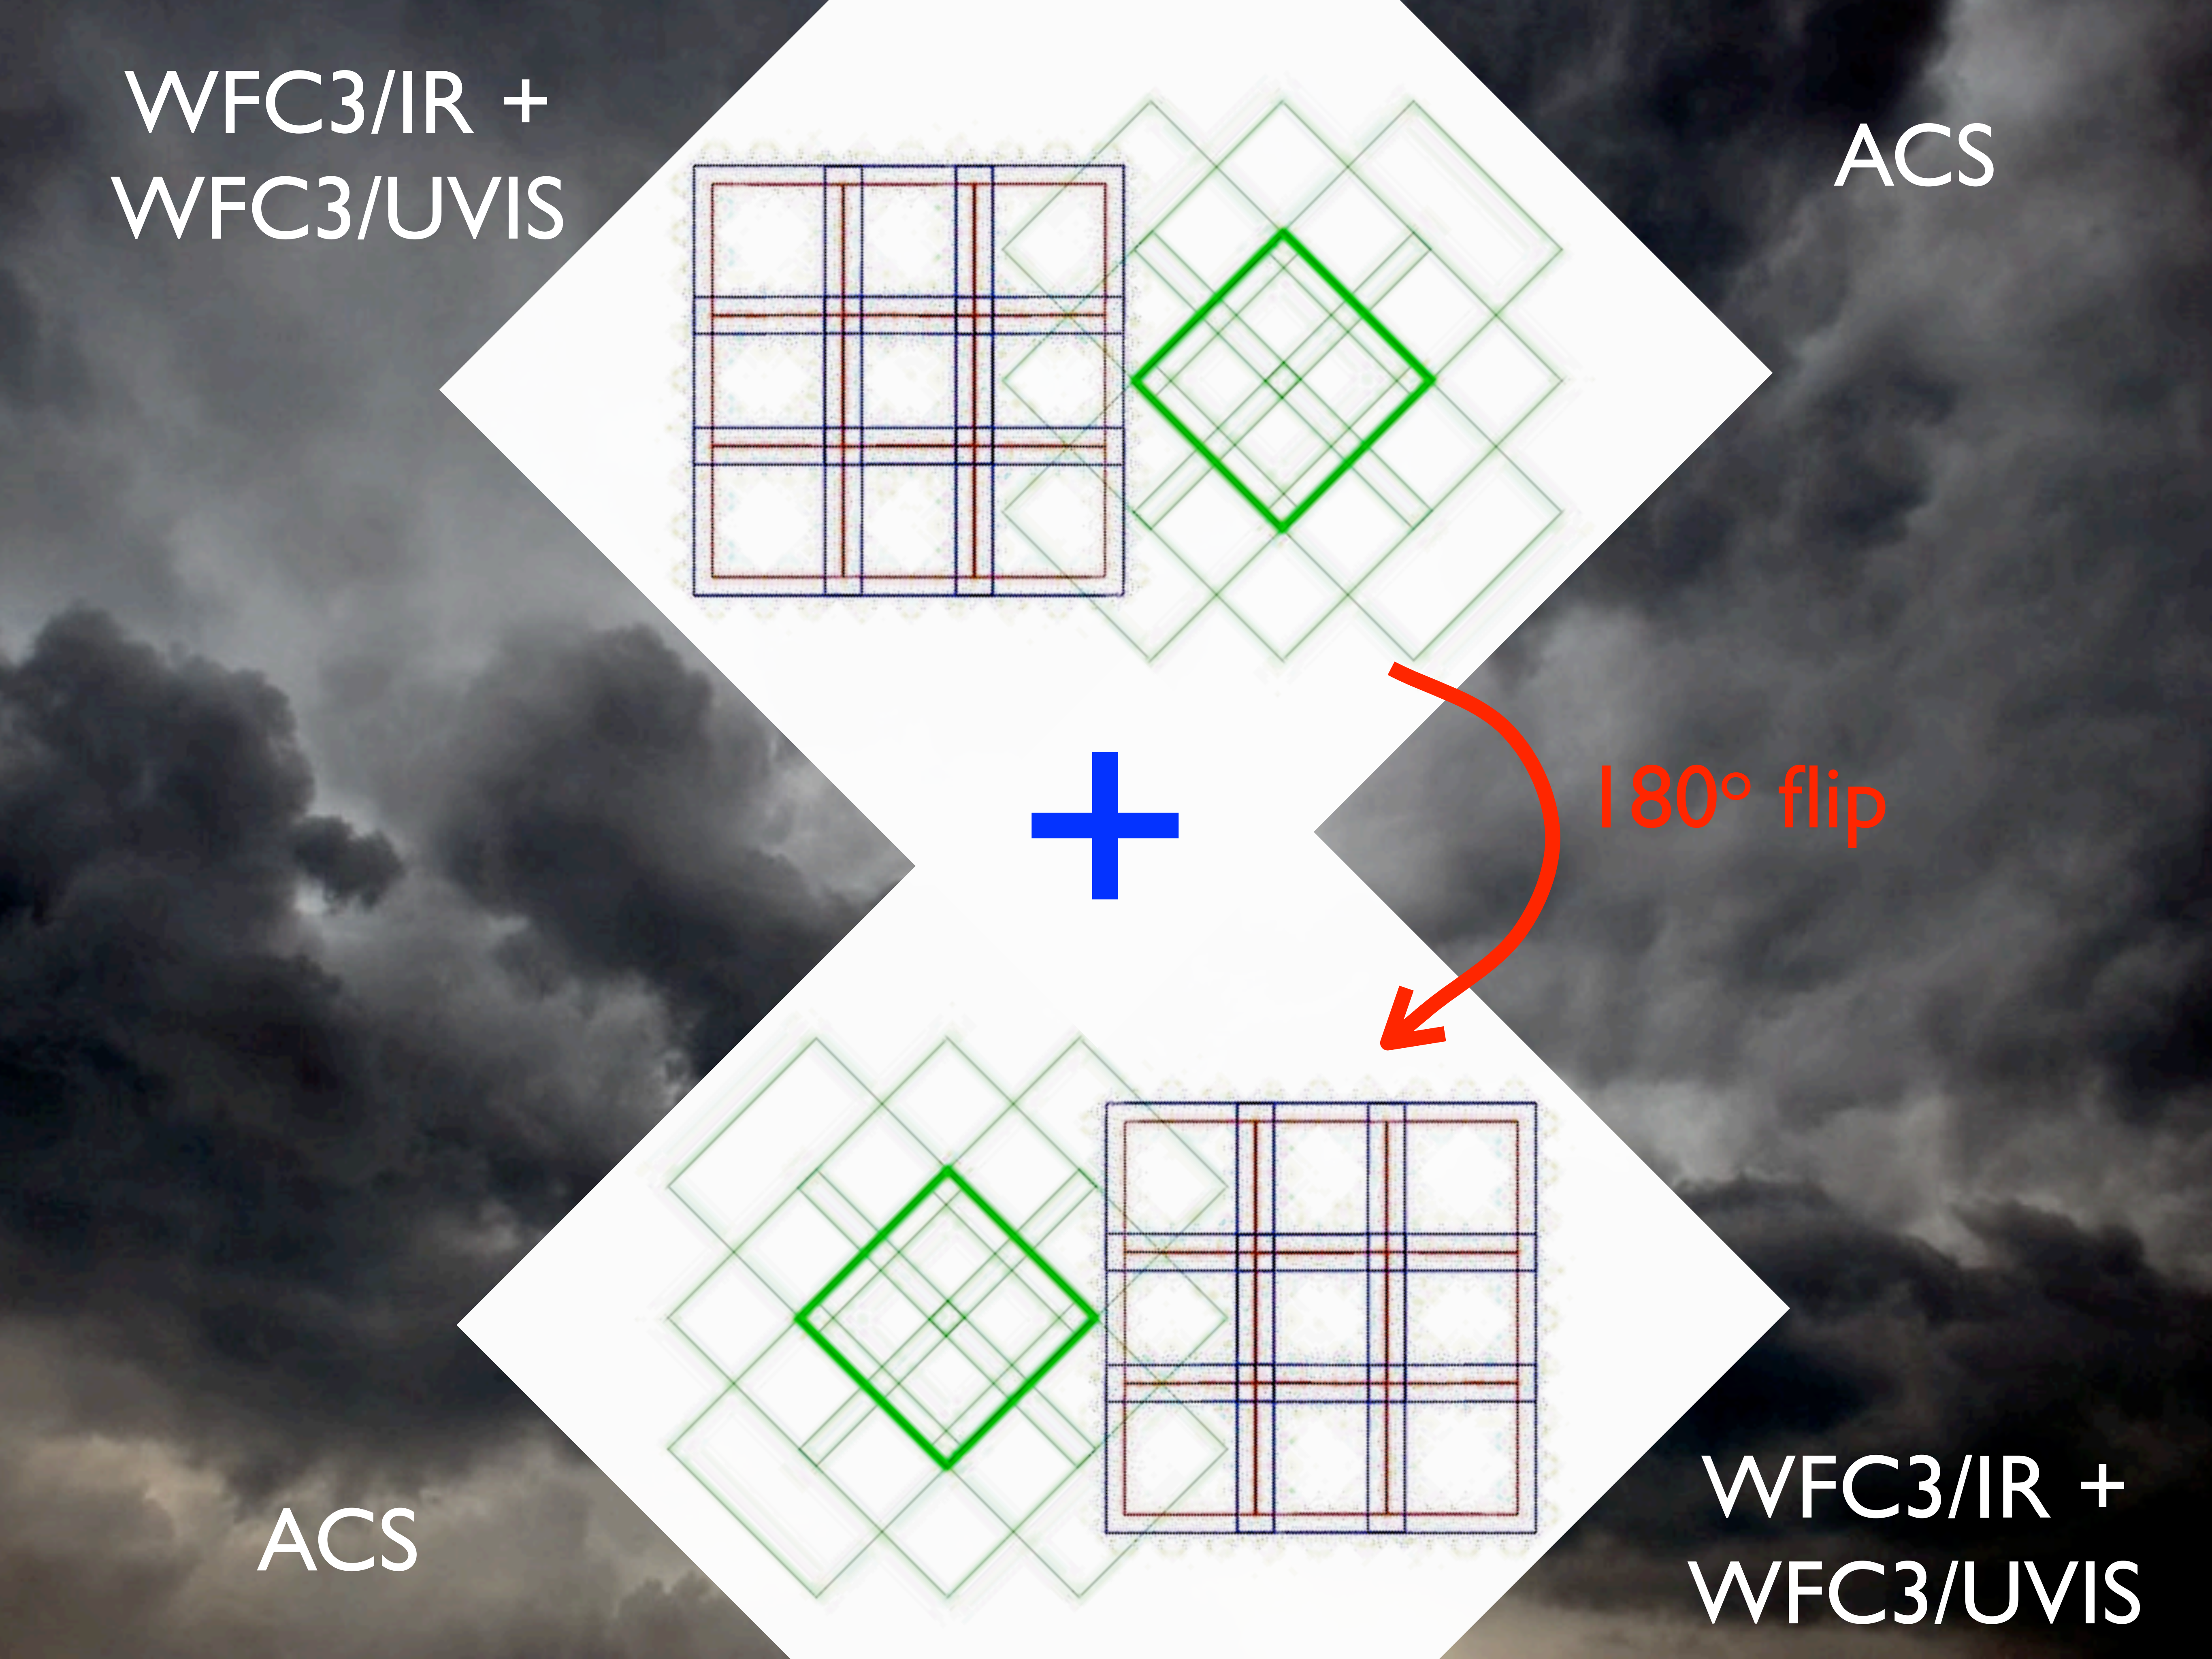
WFC3/IR +
WFC3/UVIS
ACS
180° flip
ACS
WFC3/IR +
WFC3/UVIS

6-Filter HST Tiling of M31

Outline

- Pipeline Architecture
- Challenges and solutions for running on the Amazon Cloud (EC2)
- Necessary improvements for larger data sets.
- Example application: M31 HST Survey

Pipeline Goals

- Break down data reduction into many individual tasks.
- Run tasks on lots of data in parallel.
- Group data to optimize measurements
- Efficient to reduce same data multiple ways.
- Chain of tasks can be restarted from any point of failure

Pipeline Architecture

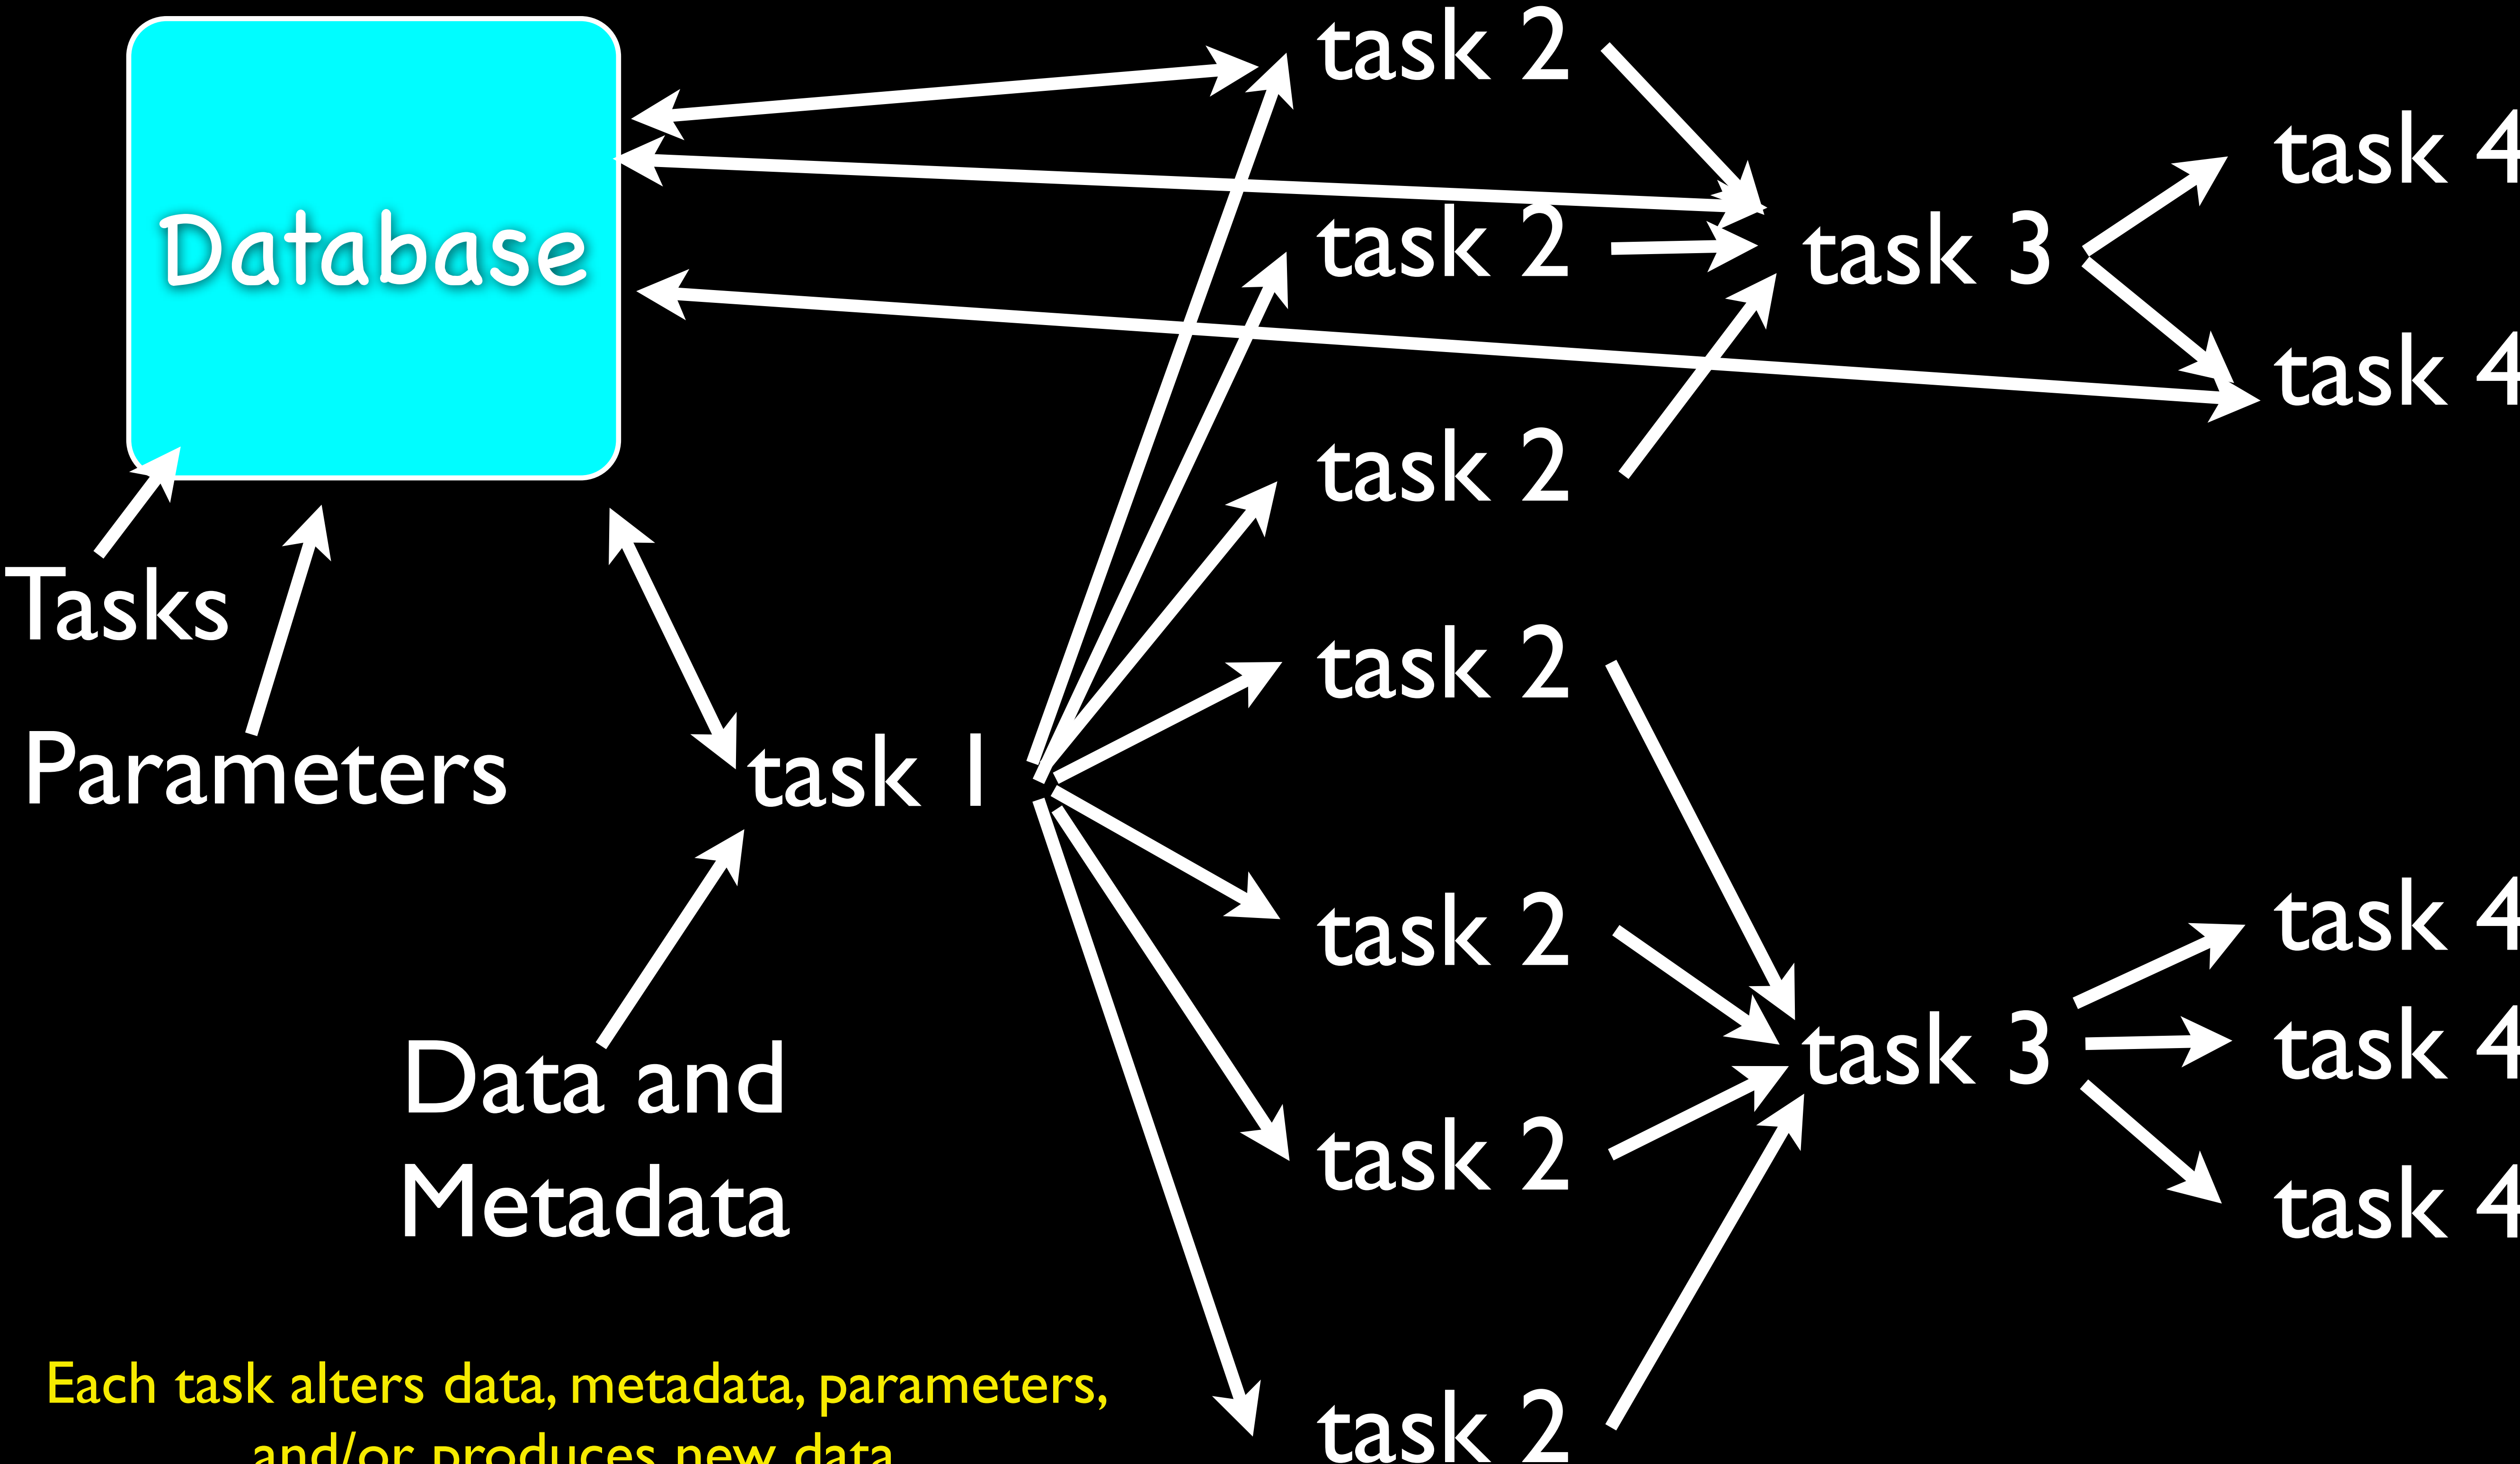

Each task alters data, metadata, parameters, and/or produces new data

Can branch off separate reduction with different parameters at any point

Cluster Computing

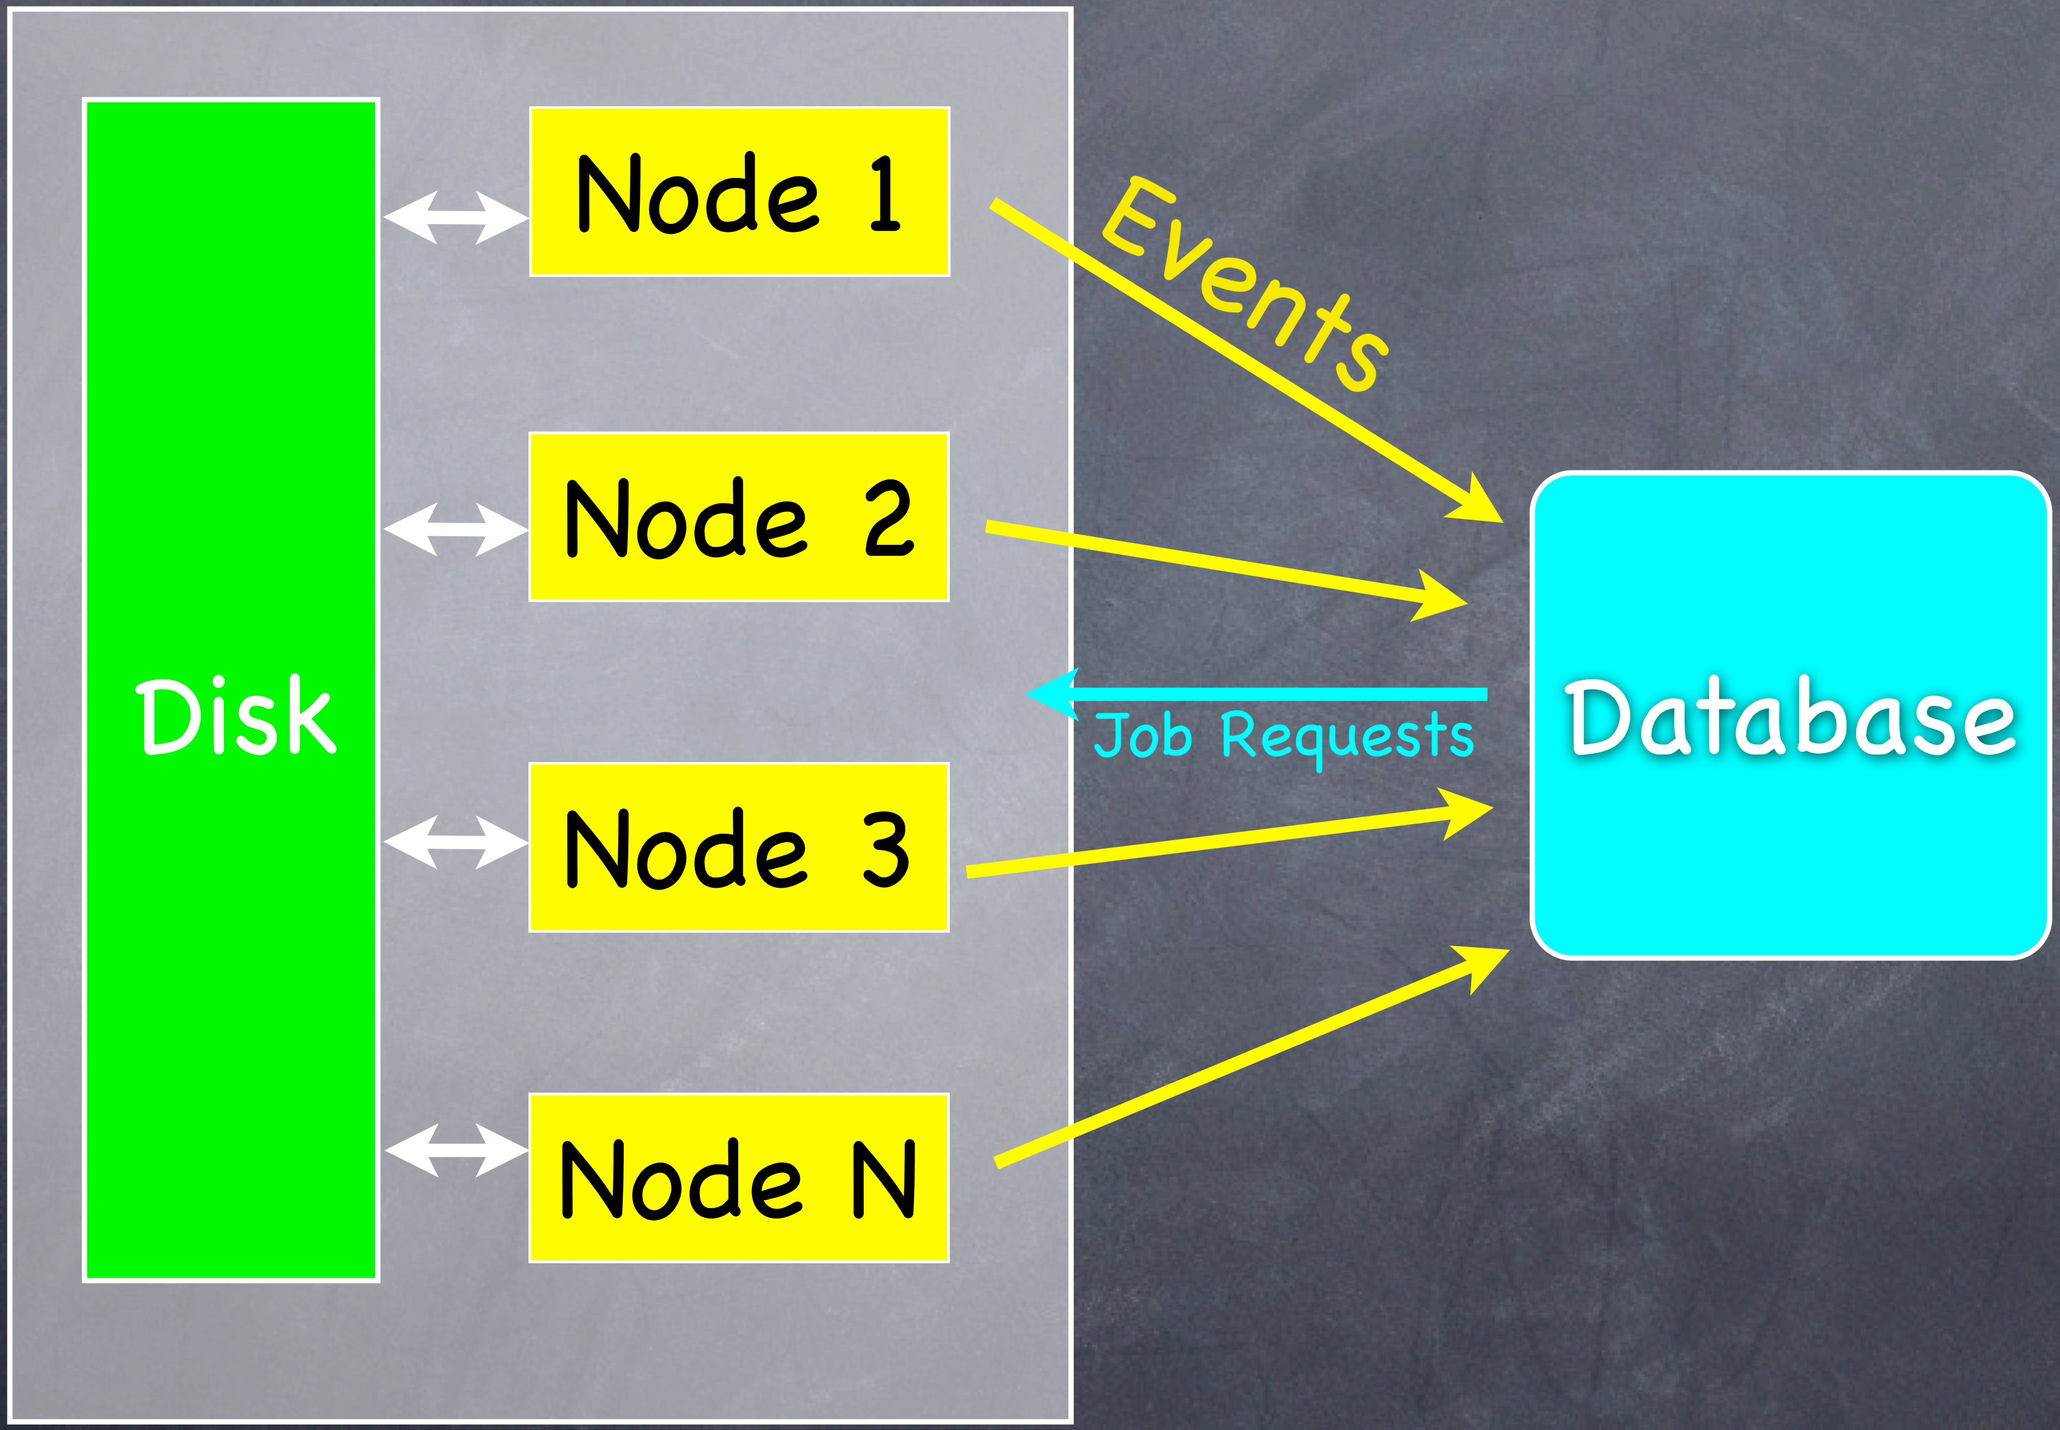

Why Move to EC2?

- Root privileges for installing and maintaining software
- No end-of-life/contract problems
- Scaling up or down based on data amounts
- Easy to change machine specifications
- Minimal unscheduled down time (hackers, upgrades, failures)
- Amazon Research Grants

EC2 Computing

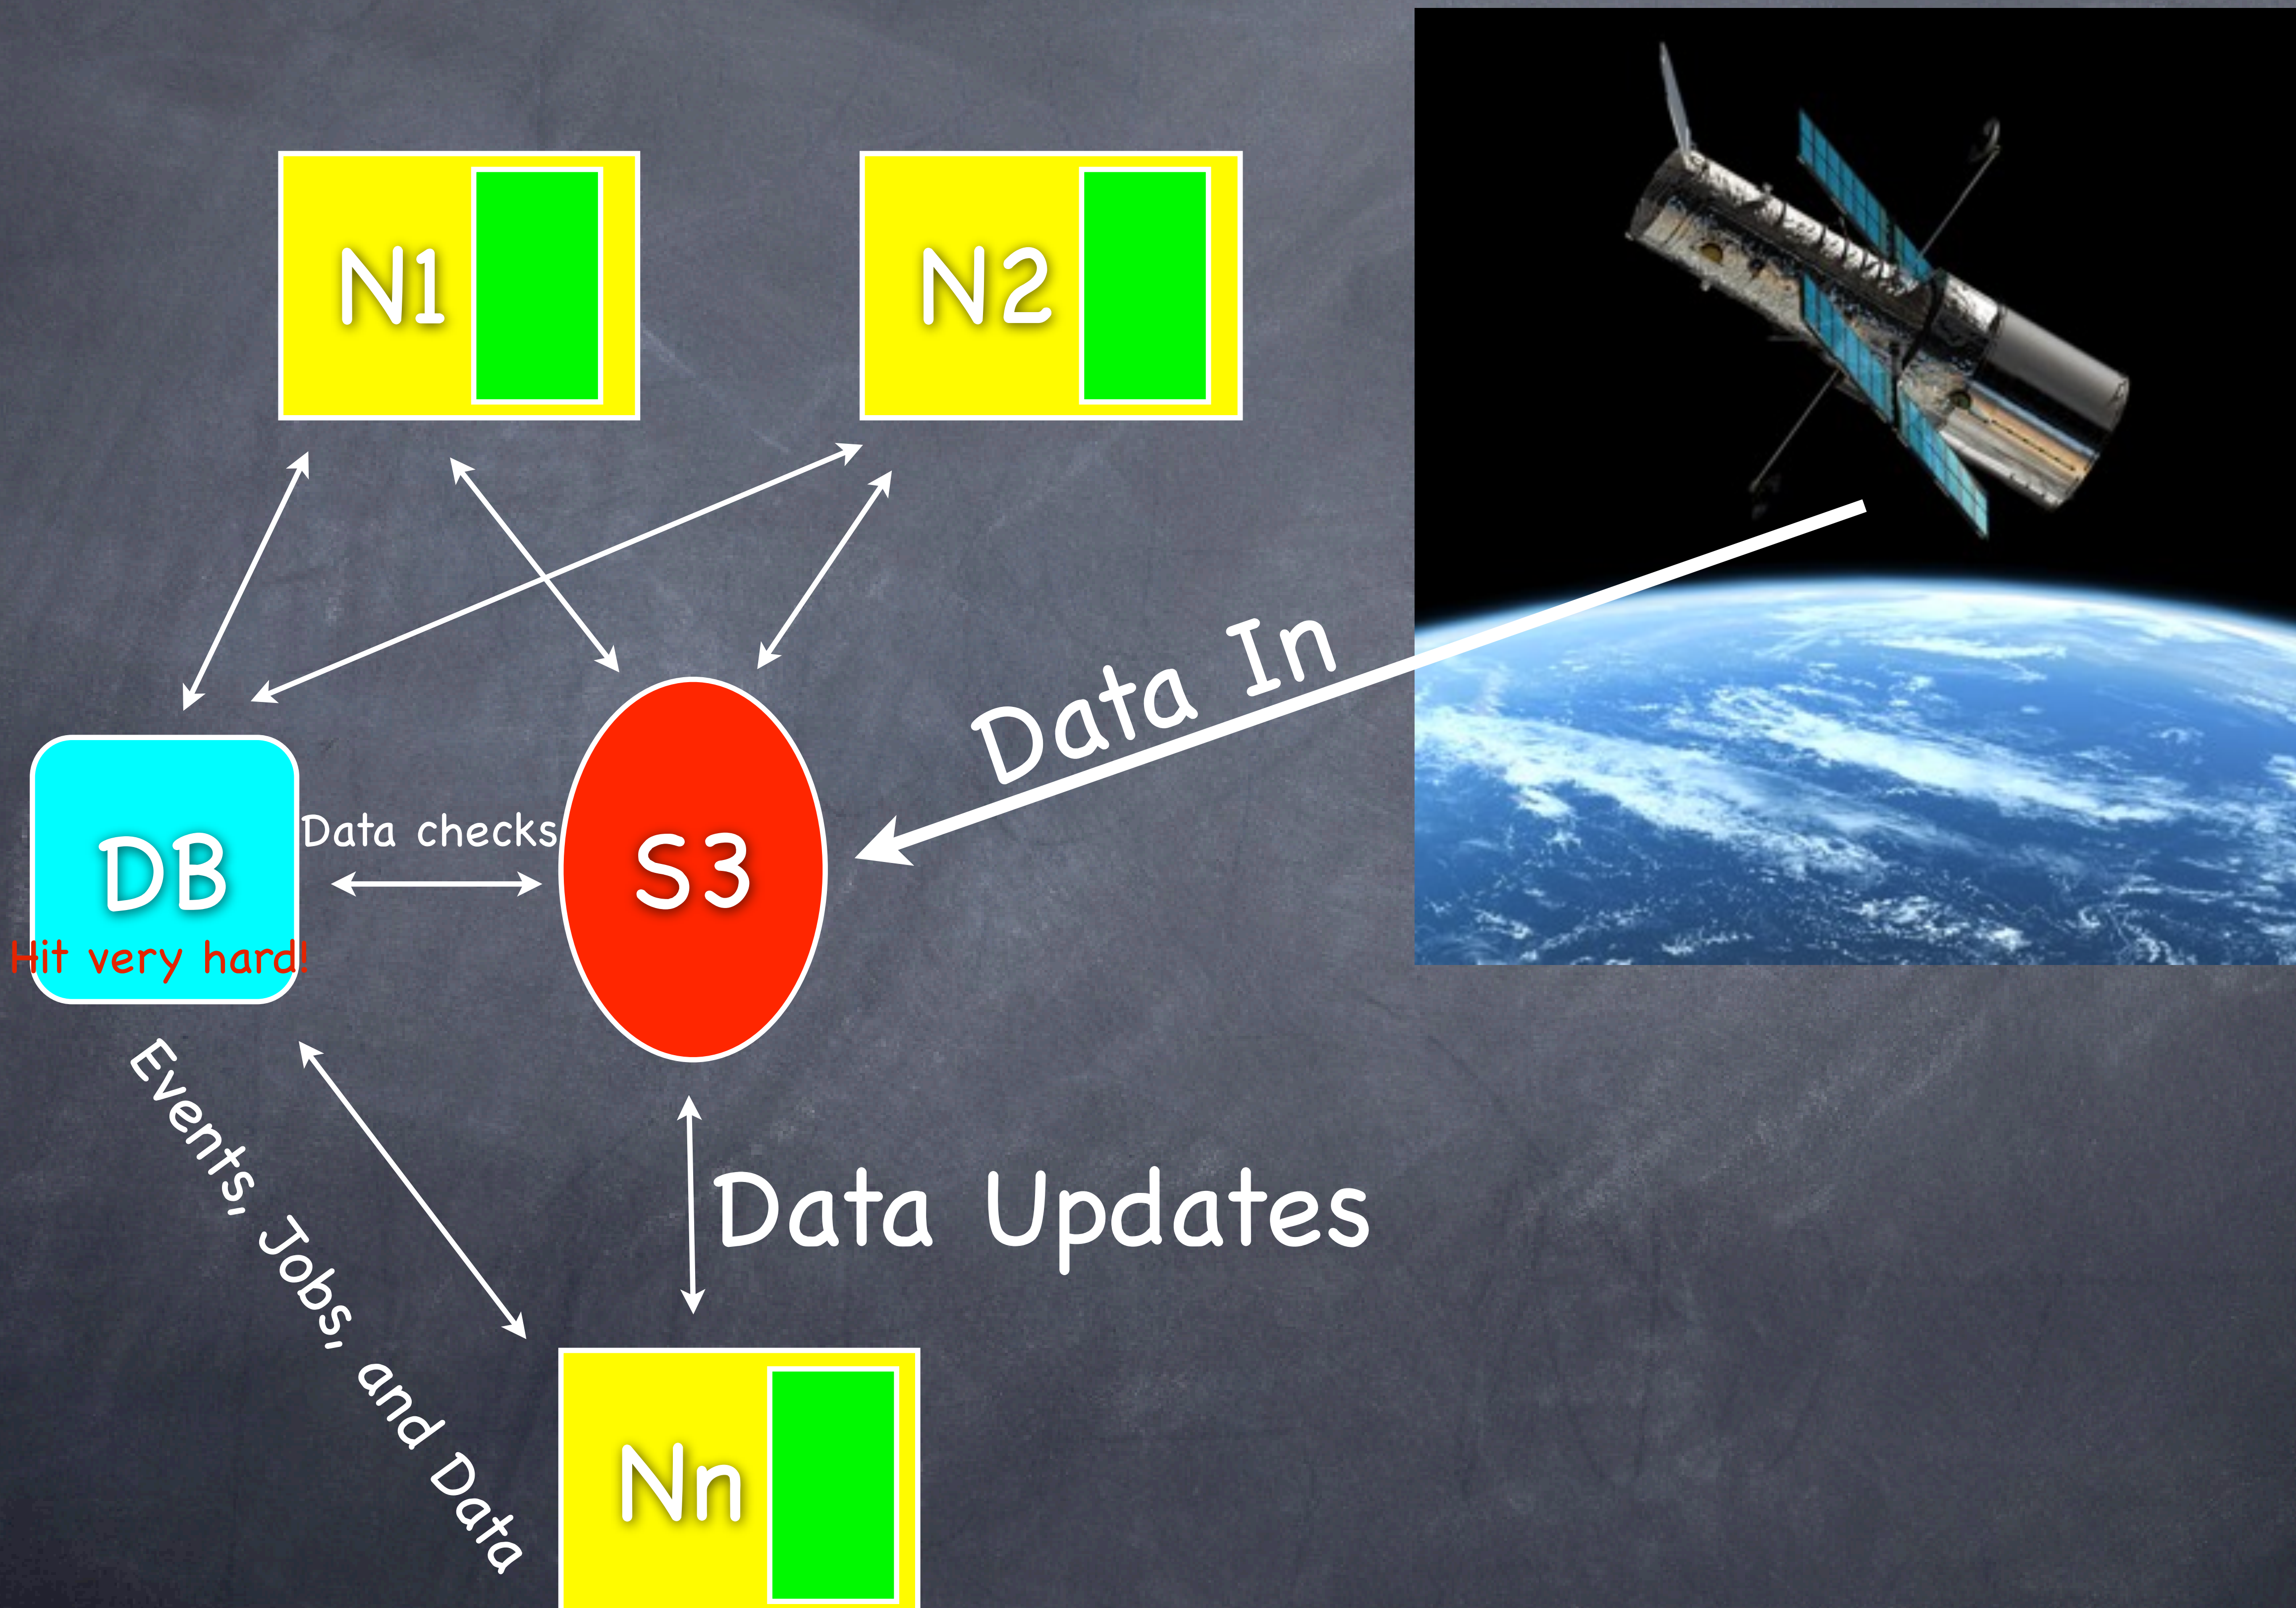

Initial Challenges

- Security: launch instances through DB server
- Disk storage: separate for each instance
- All data on Instances vanishes when they are terminated
- Job queuing and node selection done by server

Security

- Only allow instances to be passed a single number (the event code)
- Nodes only let single number through the firewall.
- Node then initiates a connection to the server -- more complex information can be exchanged.

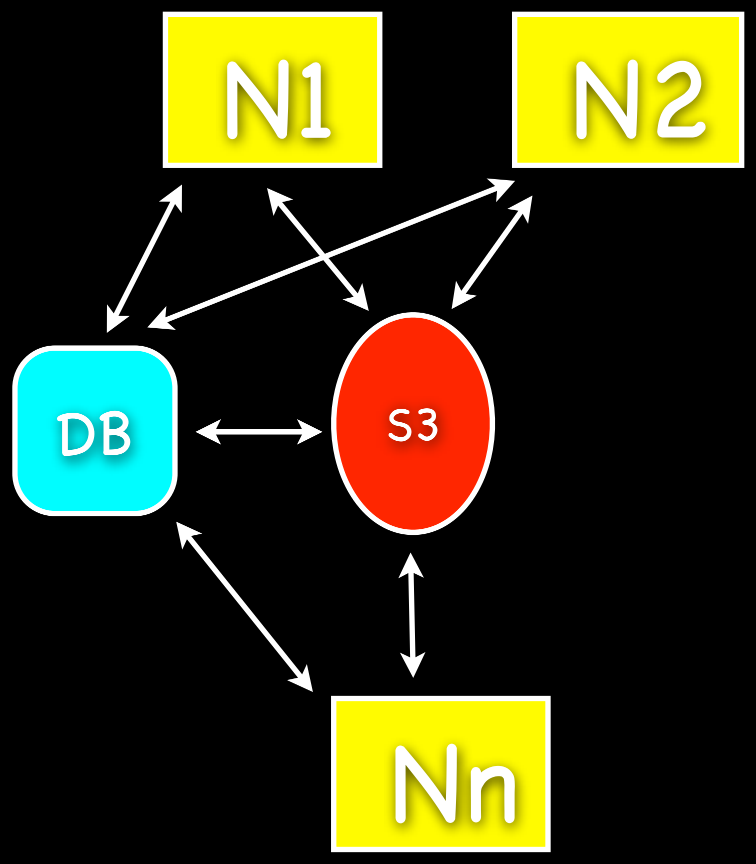

File Handling

- Each file's metadata is entered into the DB
- A copy of each file is put on S3
- When a file is requested, server syncs S3 and node to the newest version
- Requires careful scripting

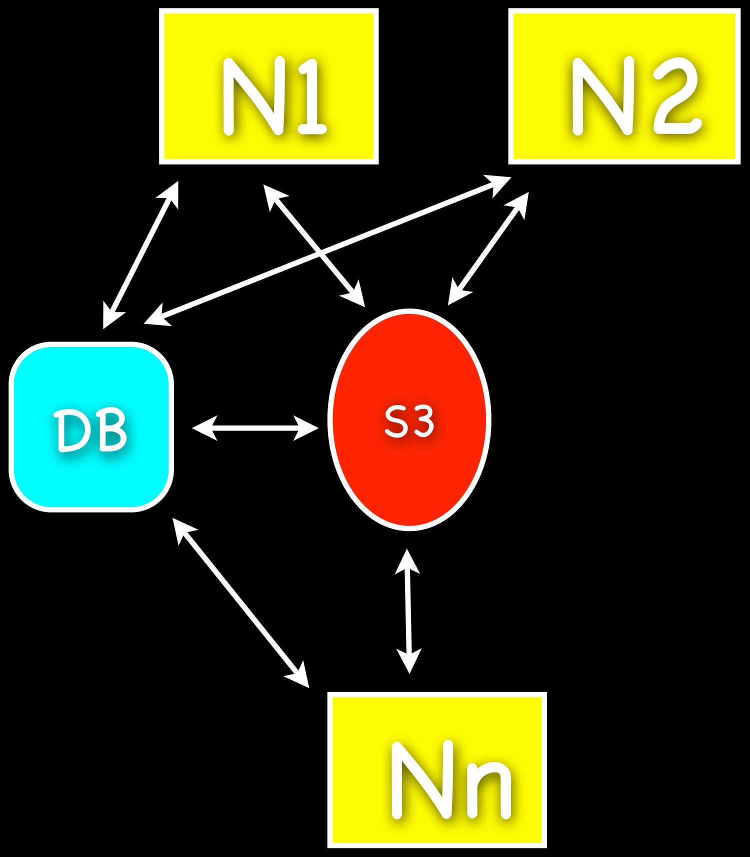

Tools Produced

- Automating start up multiple nodes with a single command to the database
- Web interface to keep track of which nodes are doing which jobs, as well as finding, and fixing failures
- Task launching regulated and coordinated across nodes
- File synchronization across many nodes through S3.

peline1.astro.washington.edu/pipelines/794?showTree=1

Event(144134): completed
Event(144127): "unsorted_image"
Job(85422): completed (node 3182 deleted) : for Task(15386): build/pipe_2.2/tag_image.pl
Event(144135): "done_tagging"
Job(85423): completed (node 3182 deleted) : for Task(15377): build/pipe_2.2/astrodrizzle.pl
Event(144139): "astrodrizzle_done"
Job(85425): completed (node 3182 deleted) : for Task(15393): build/pipe_2.2/find_reference.pl
Event(144141): "tagged_reference"
Job(85426): completed (node 3182 deleted) : for Task(15376): build/pipe_2.2/get_galaxy.pl
Event(144144): "galaxy_retrieved"
Job(85429): completed (node 3182 deleted) : for Task(15366): build/pipe_2.2/make_overlays.pl
Event(144153): completed
Event(144146): completed
Job(85427): completed (node 3182 deleted) : for Task(15385): build/pipe_2.2/count_images.pl
Event(144143): "prep"
Job(85428): completed (node 3182 deleted) : for Task(15374): build/pipe_2.2/prep_image.pl
Event(144156): completed
Event(144145): "prep"
Job(85430): completed (node 3182 deleted) : for Task(15374): build/pipe_2.2/prep_image.pl
Event(144157): completed
Event(144147): "prep"
Job(85431): completed (node 3182 deleted) : for Task(15374): build/pipe_2.2/prep_image.pl
Event(144154): completed
Event(144148): "prep"
Job(85432): completed (node 3182 deleted) : for Task(15374): build/pipe_2.2/prep_image.pl
Event(144155): completed
Event(144149): "prep"
Job(85433): completed (node 3182 deleted) : for Task(15374): build/pipe_2.2/prep_image.pl
Event(144158): completed
Event(144150): "prep"
Job(85434): completed (node 3182 deleted) : for Task(15374): build/pipe_2.2/prep_image.pl
Event(144159): "images_prepped"
Job(85436): completed (node 3182 deleted) : for Task(15369): build/pipe_2.2/shiftfind.pl
Event(144162): "shifts_done"
Job(85437): completed (node 3182 deleted) : for Task(15383): build/pipe_2.2/write_dolphot_pars.pl
Event(144164): "parameters_done"
Job(85438): completed (node 3182 deleted) : for Task(15392): build/pipe_2.2/run_dolphot.pl
Event(144192): "dolphot_done"
Job(85452): completed (node 3182 deleted) : for Task(15358): build/pipe_2.2/phot_fitstable.pl
Event(144195): "phot_fitstable_done"
Job(85453): completed (node 3182 deleted) : for Task(15391): build/pipe_2.2/cull_photometry.pl
Event(144197): "culling_done"
Job(85454): completed (node 3182 deleted) : for Task(15365): build/pipe_2.2/make_cmd.pl
Event(144200): "new_cmds_done"

Unexpected problems

- File synchronization issues -- colliding requests.
  - More careful task scripting
- Server overload -- too many requests
  - DB indexing
  - More powerful server node
- File transfer size limits and errors
  - Increased checks that transfers succeed.
- Too expensive
  - Move to “Spot” instances

Hourly Rates

type	memory	cpu	on-demand	reserved 1yr min.	typical spot
Average	8	2	0.14	0.09	0.02
more memory	15	2	0.18	0.10	0.02
more CPU	15	4	0.28	0.18	0.03
high memory	61	8	0.70	0.38	0.07
high CPU	30	16	0.84	0.51	0.13

On demand $0.28 per hour

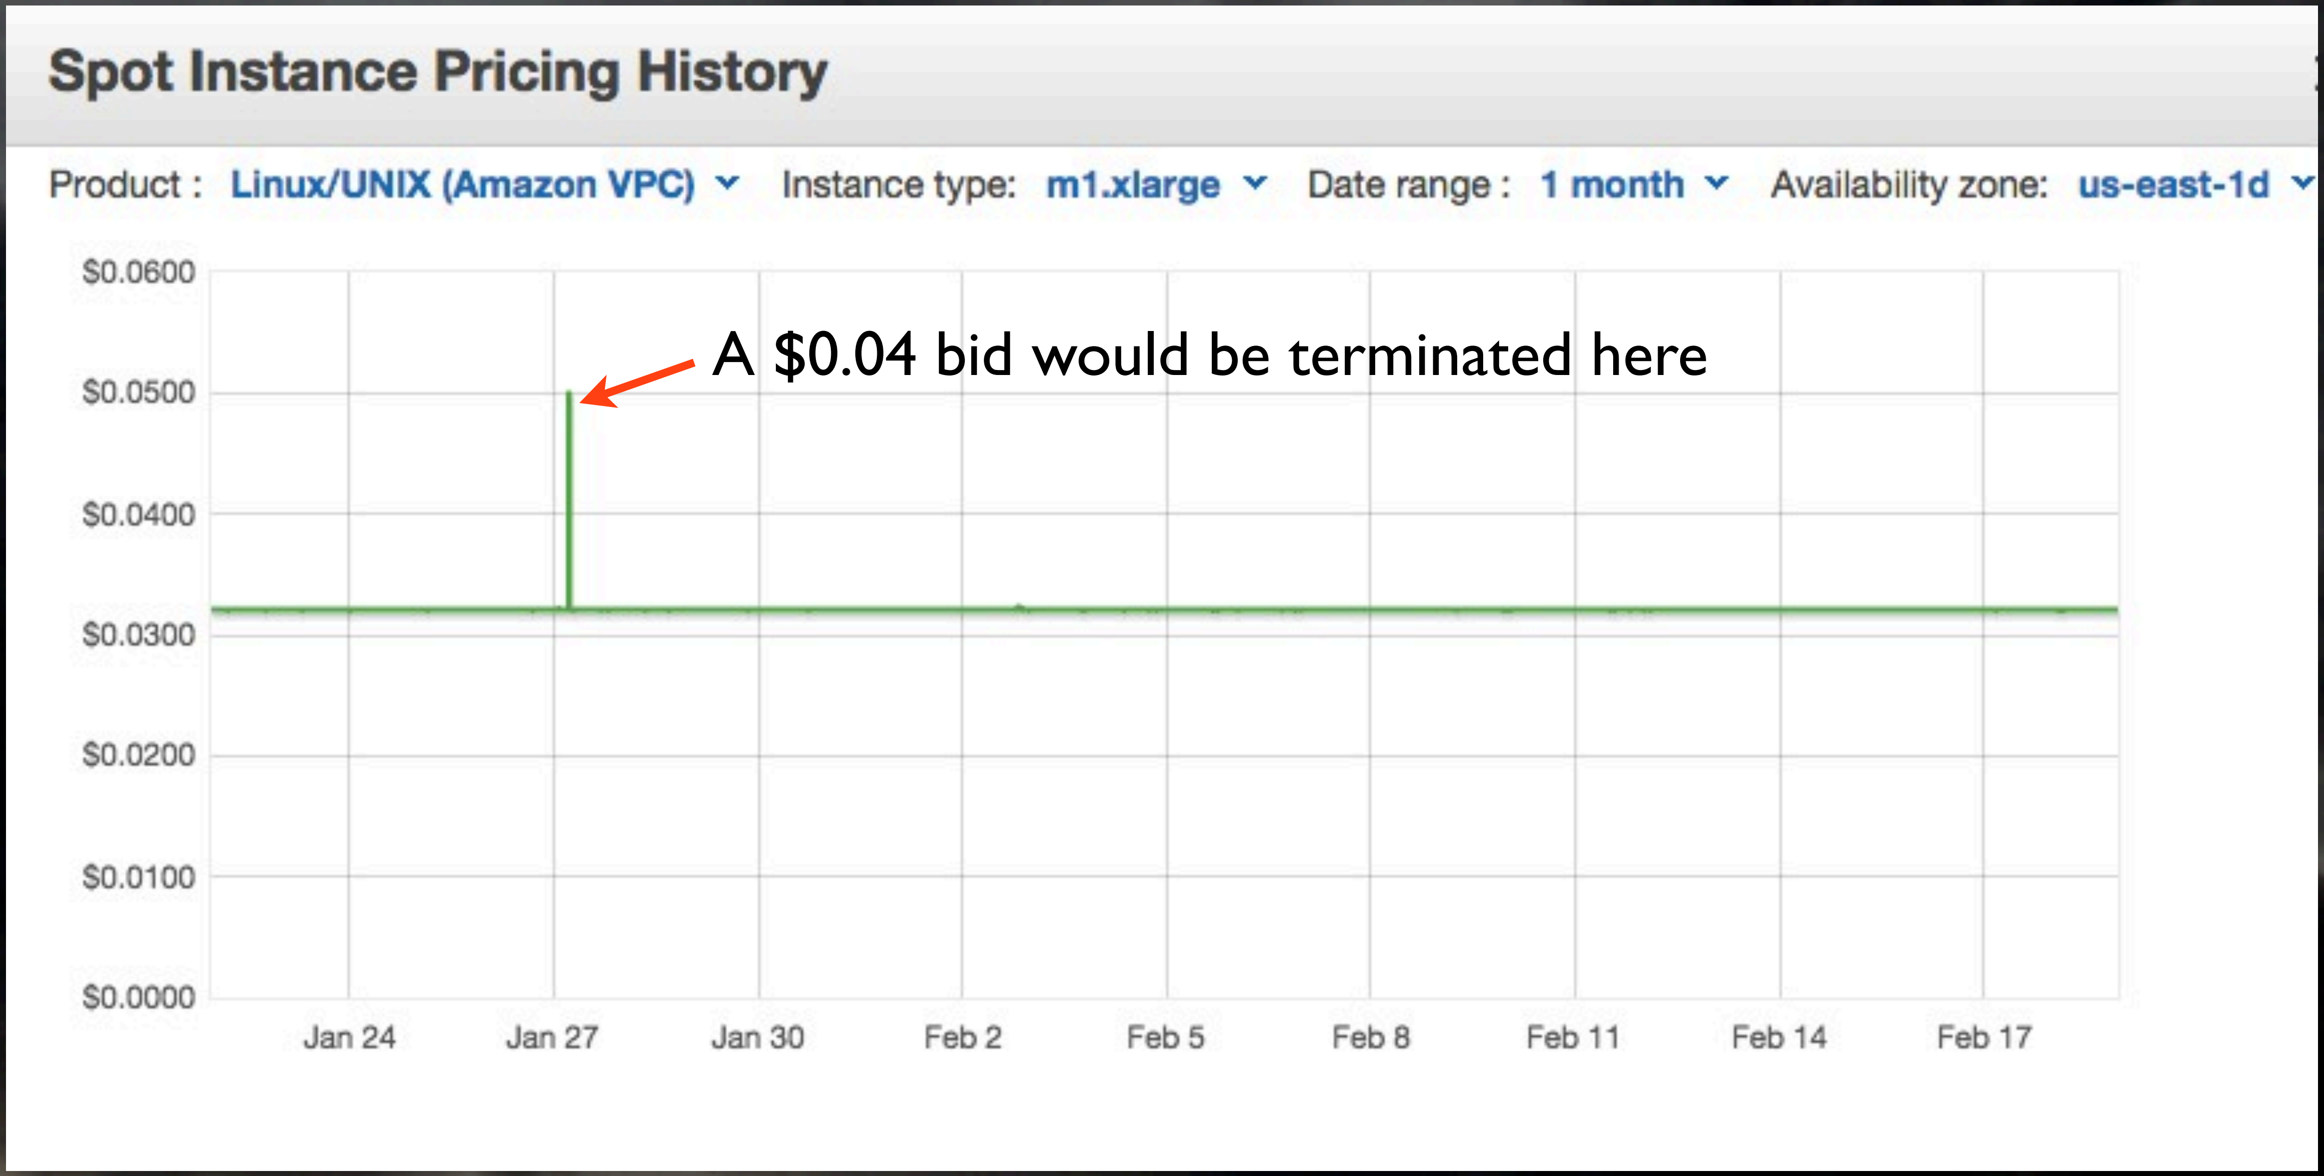

Prices can jump unexpectedly

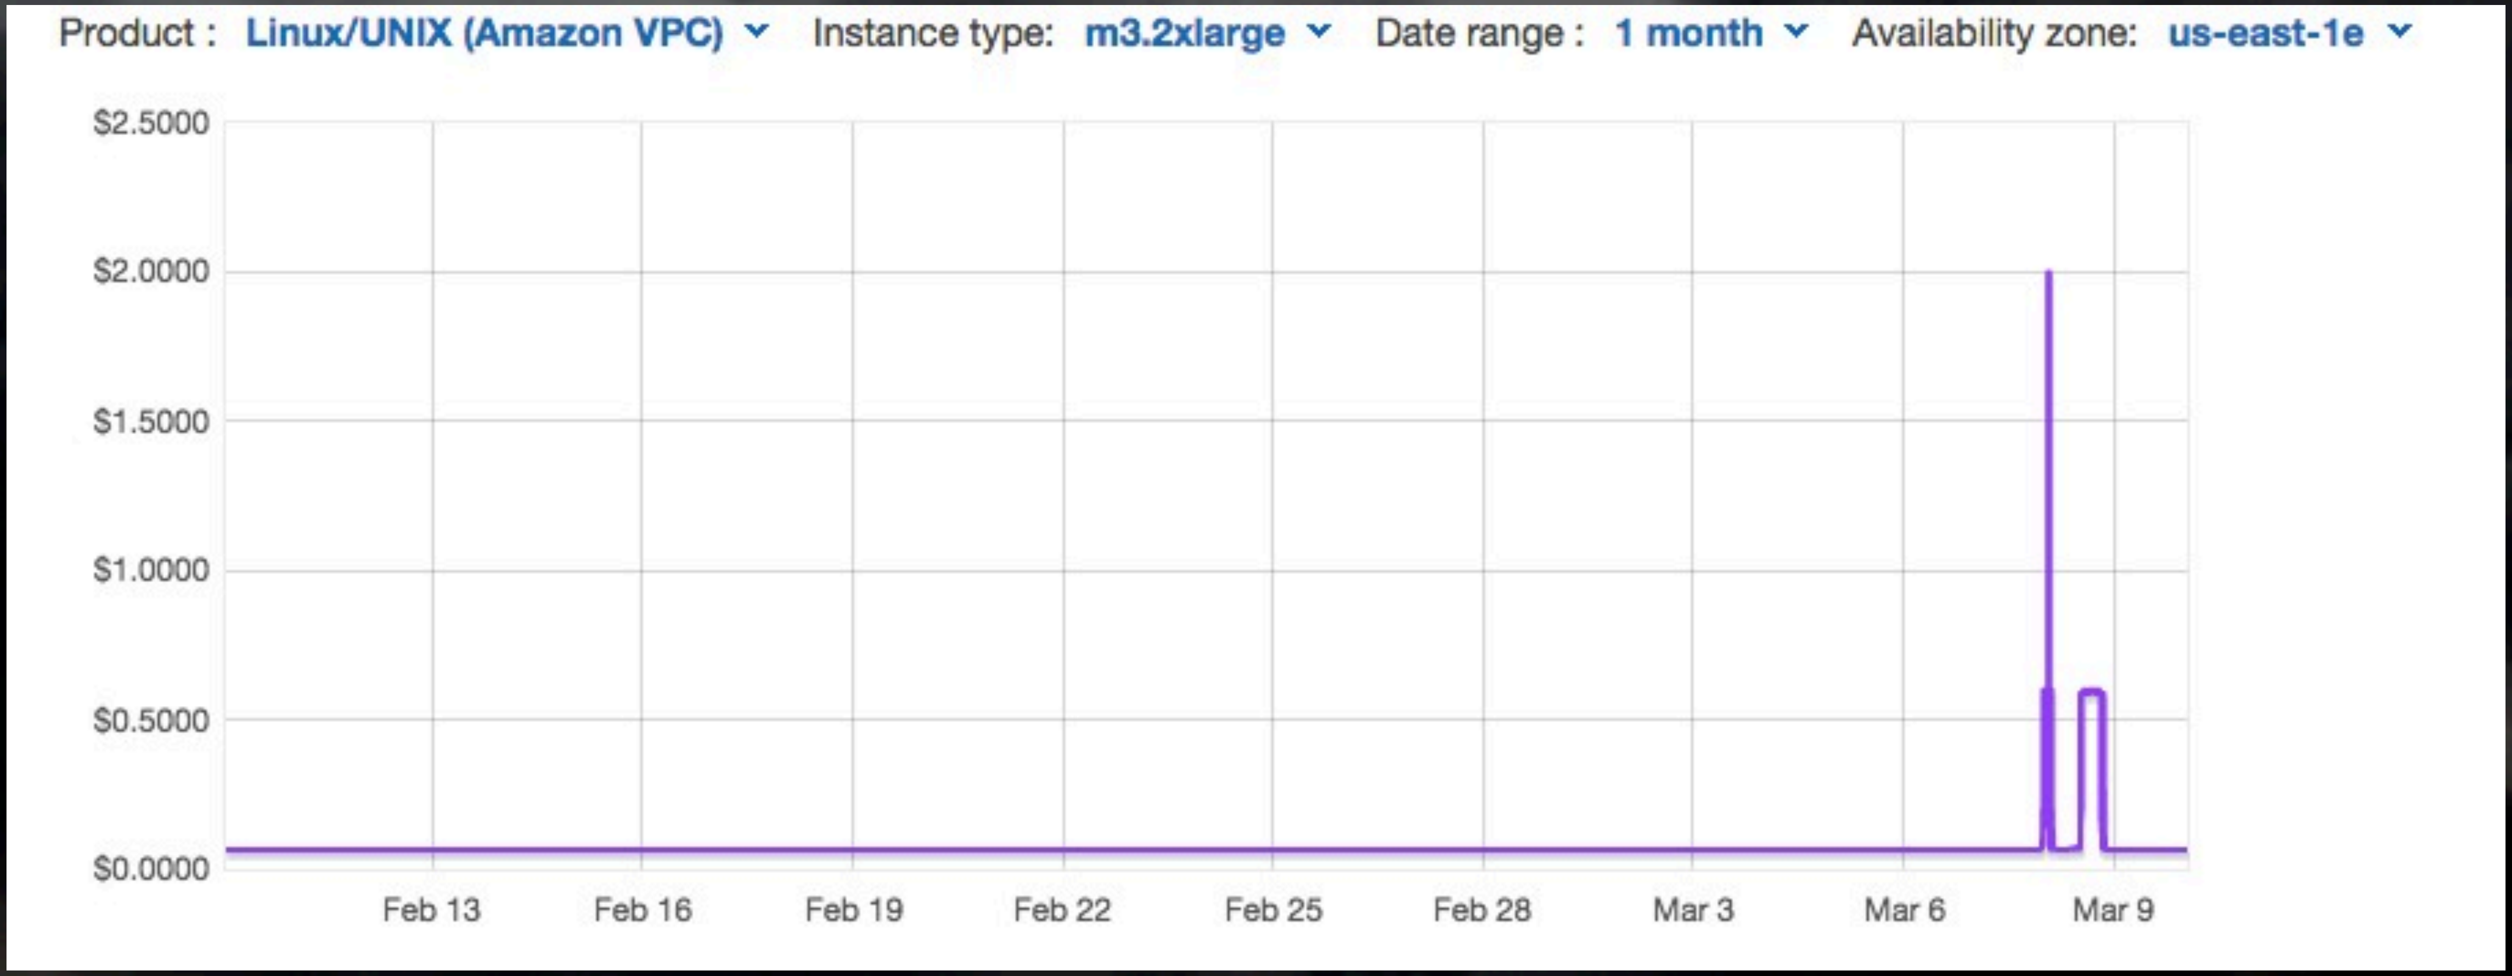

Continuing Improvements

- Optimizing node use (ratio of high to low memory nodes, and putting the right jobs on the right type of node).
- Still occasional communication problems, intermittent and very difficult to trace.
- Server always on… $$$
- Recovery for vanishing instances.

Improvements needed for Scaling

- automated termination of idle nodes
- automated adding of nodes (of appropriate type) when queue is long
- storing and using technical requirements for each task
- high-powered server instance (off cloud?)

PHAT

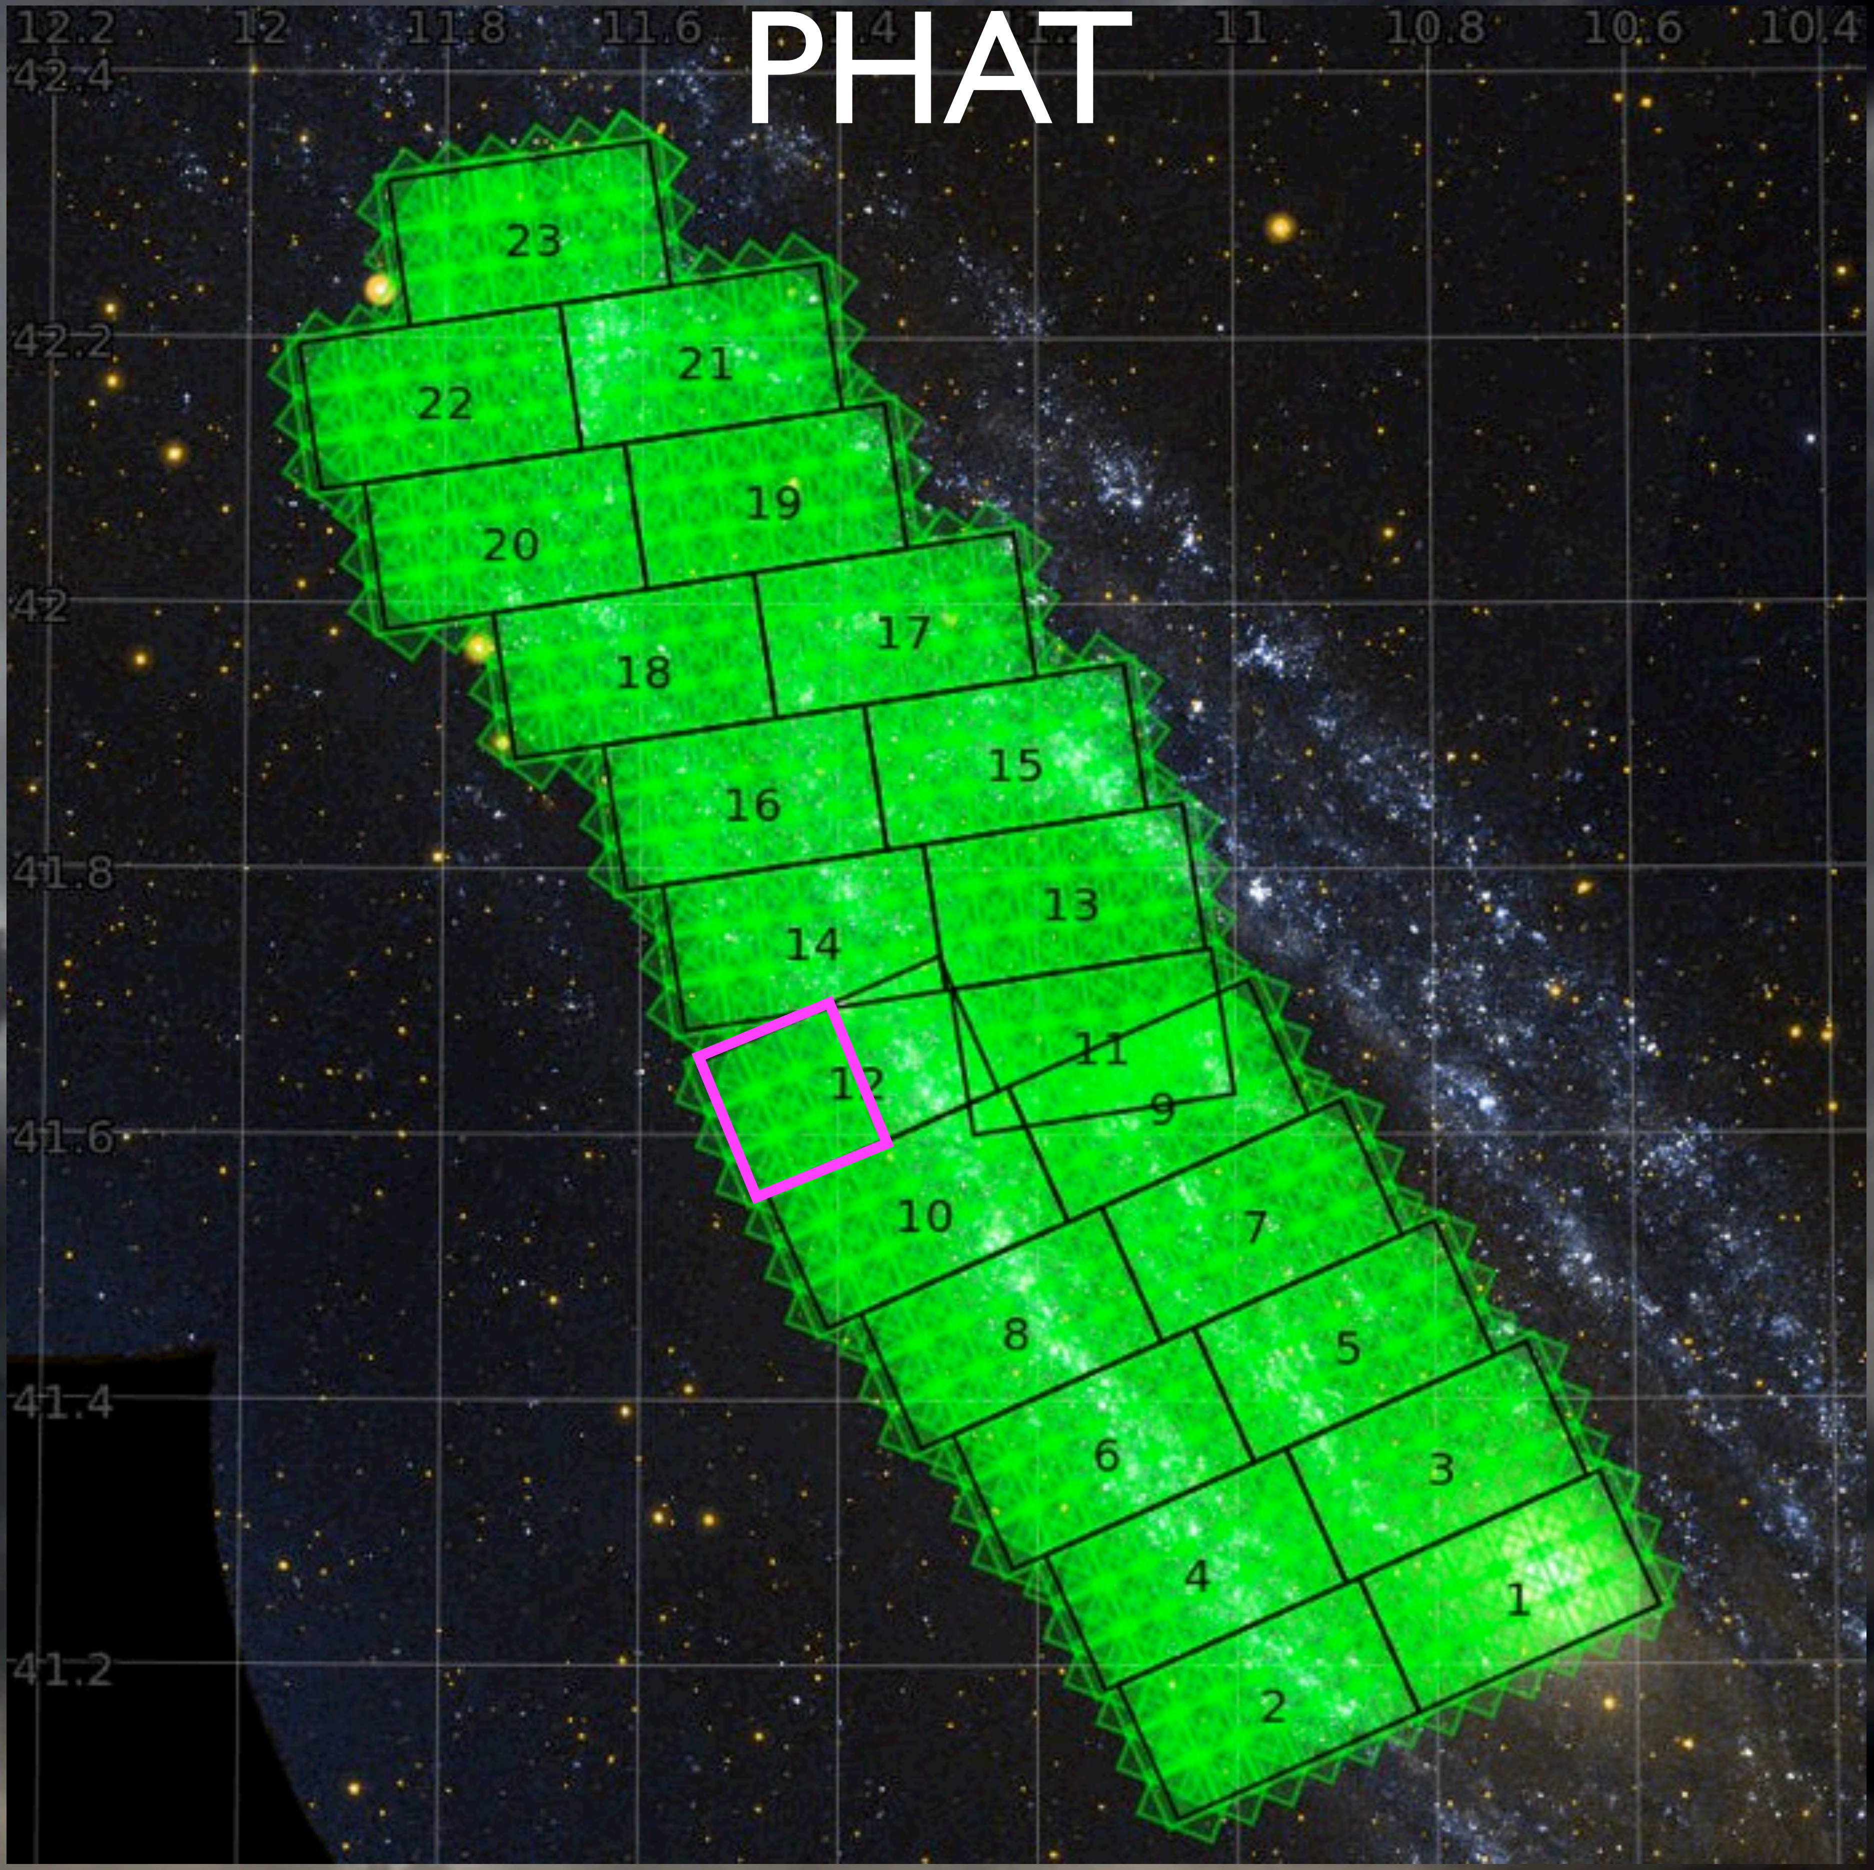
12.2
12
11.8
11.6
11
10.8
10.6
10.4
42.4
42.2
42
41.8
41.6
41.4
41.2
23
22
21
20
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1

Measure 3 cameras at once

Multiple orientations
Multiple pixel scales
Multiple distortions
Multiple point spread functions
Andy Dolphin, DOLPHOT 2.0 (http://americano.dolphinsim.com/dolphot/)

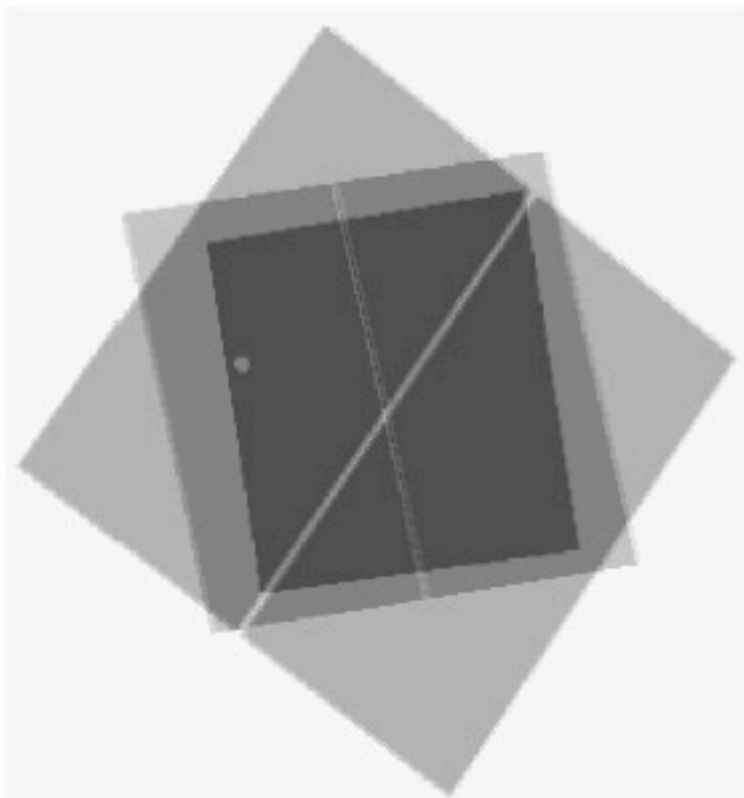

8 GB of memory: $0.2/hr

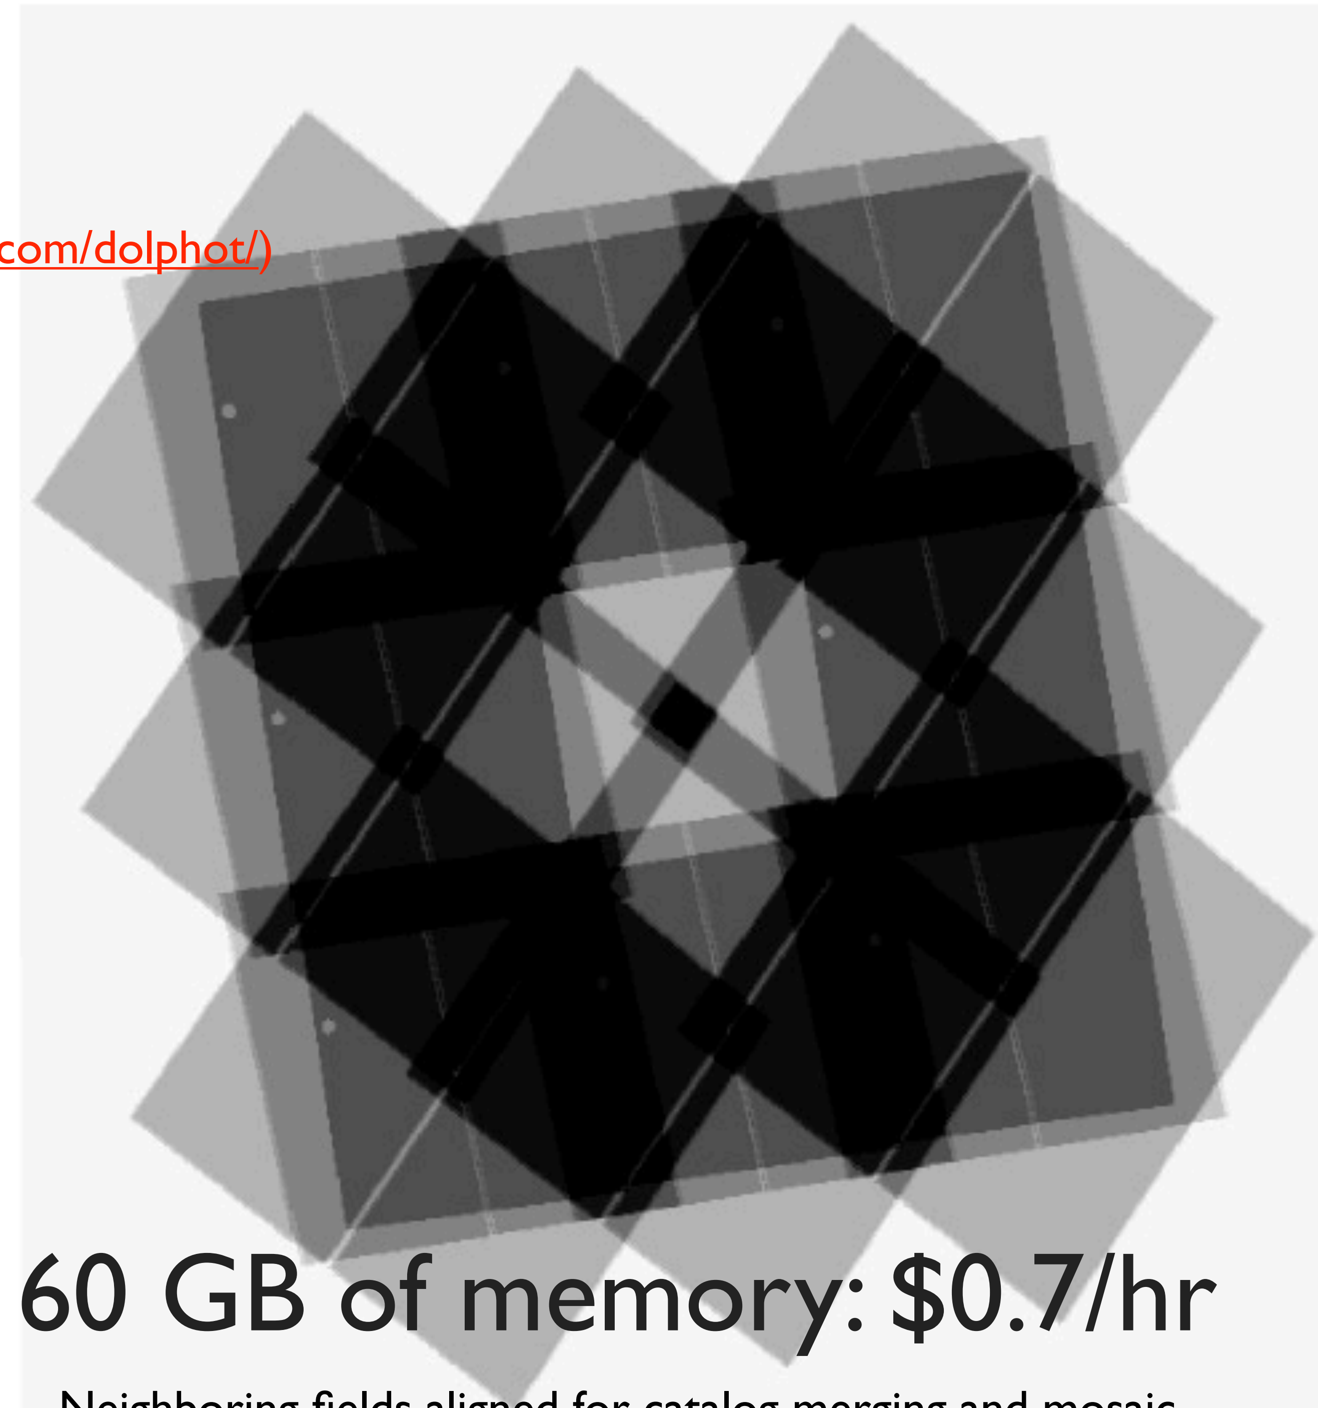

60 GB of memory: $0.7/hr

Neighboring fields aligned for catalog merging and mosaic

F475W, F814W, F160W

F475W, F814W, F160W

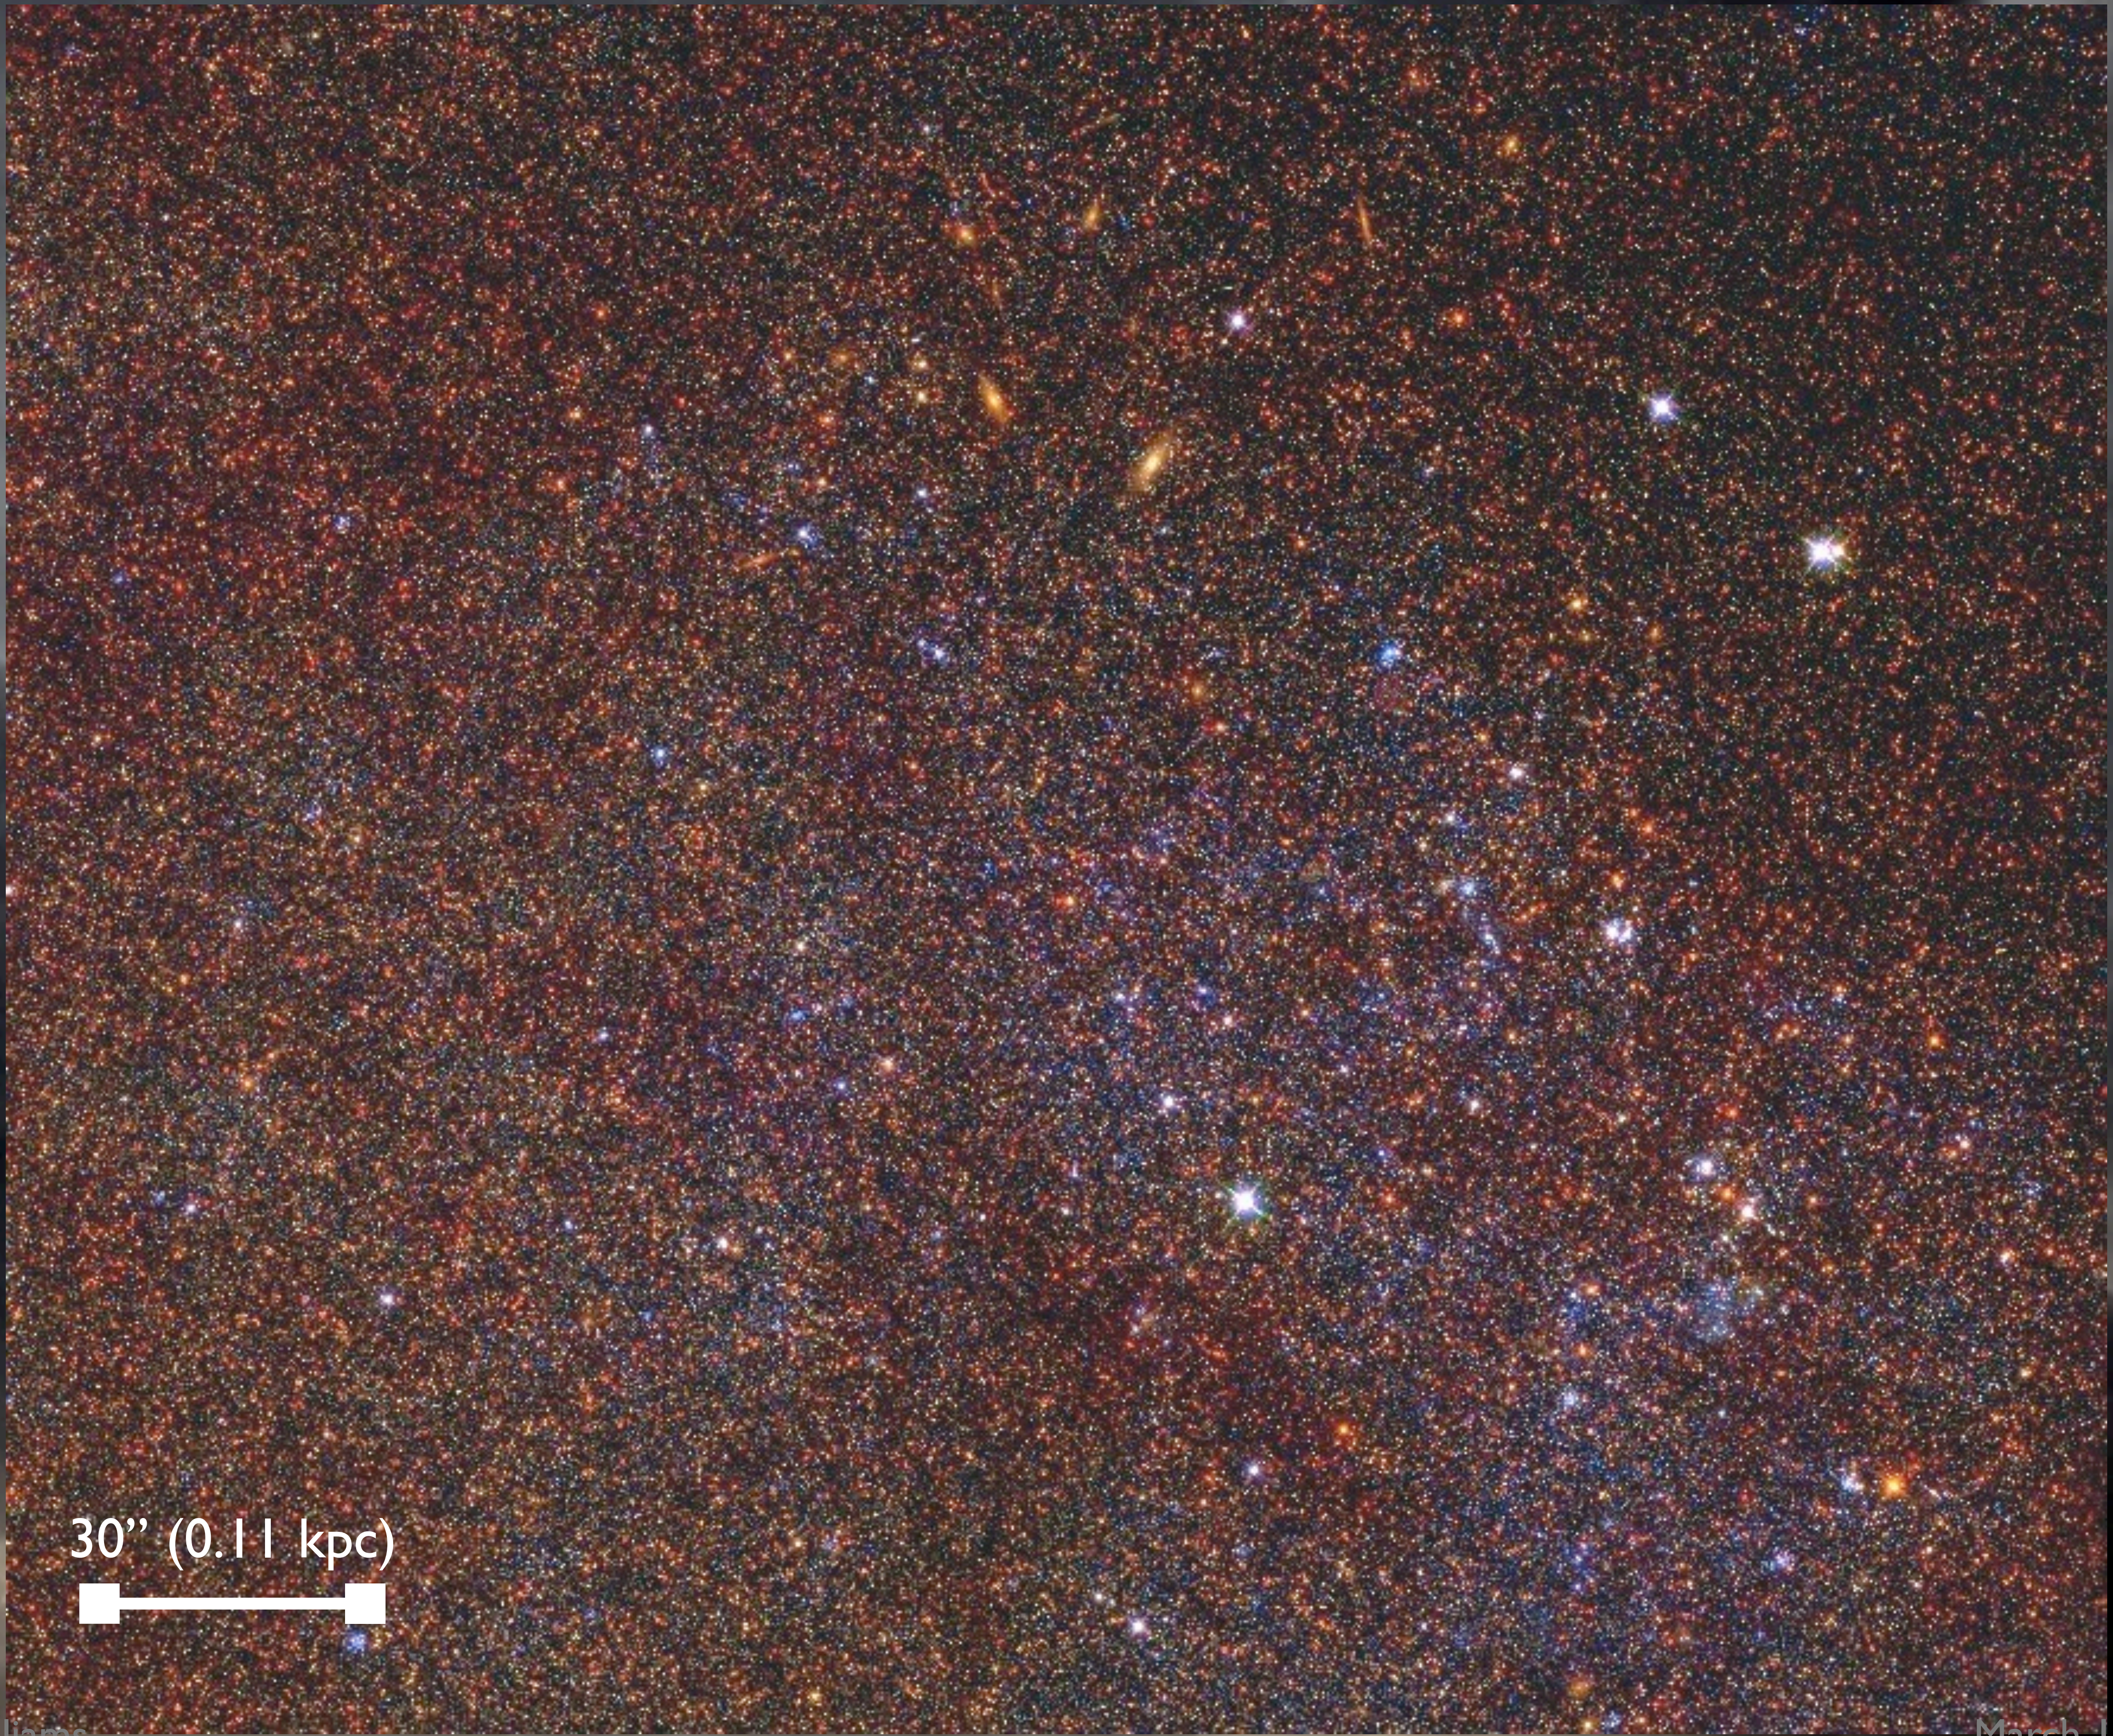
30” (0.11 kpc)

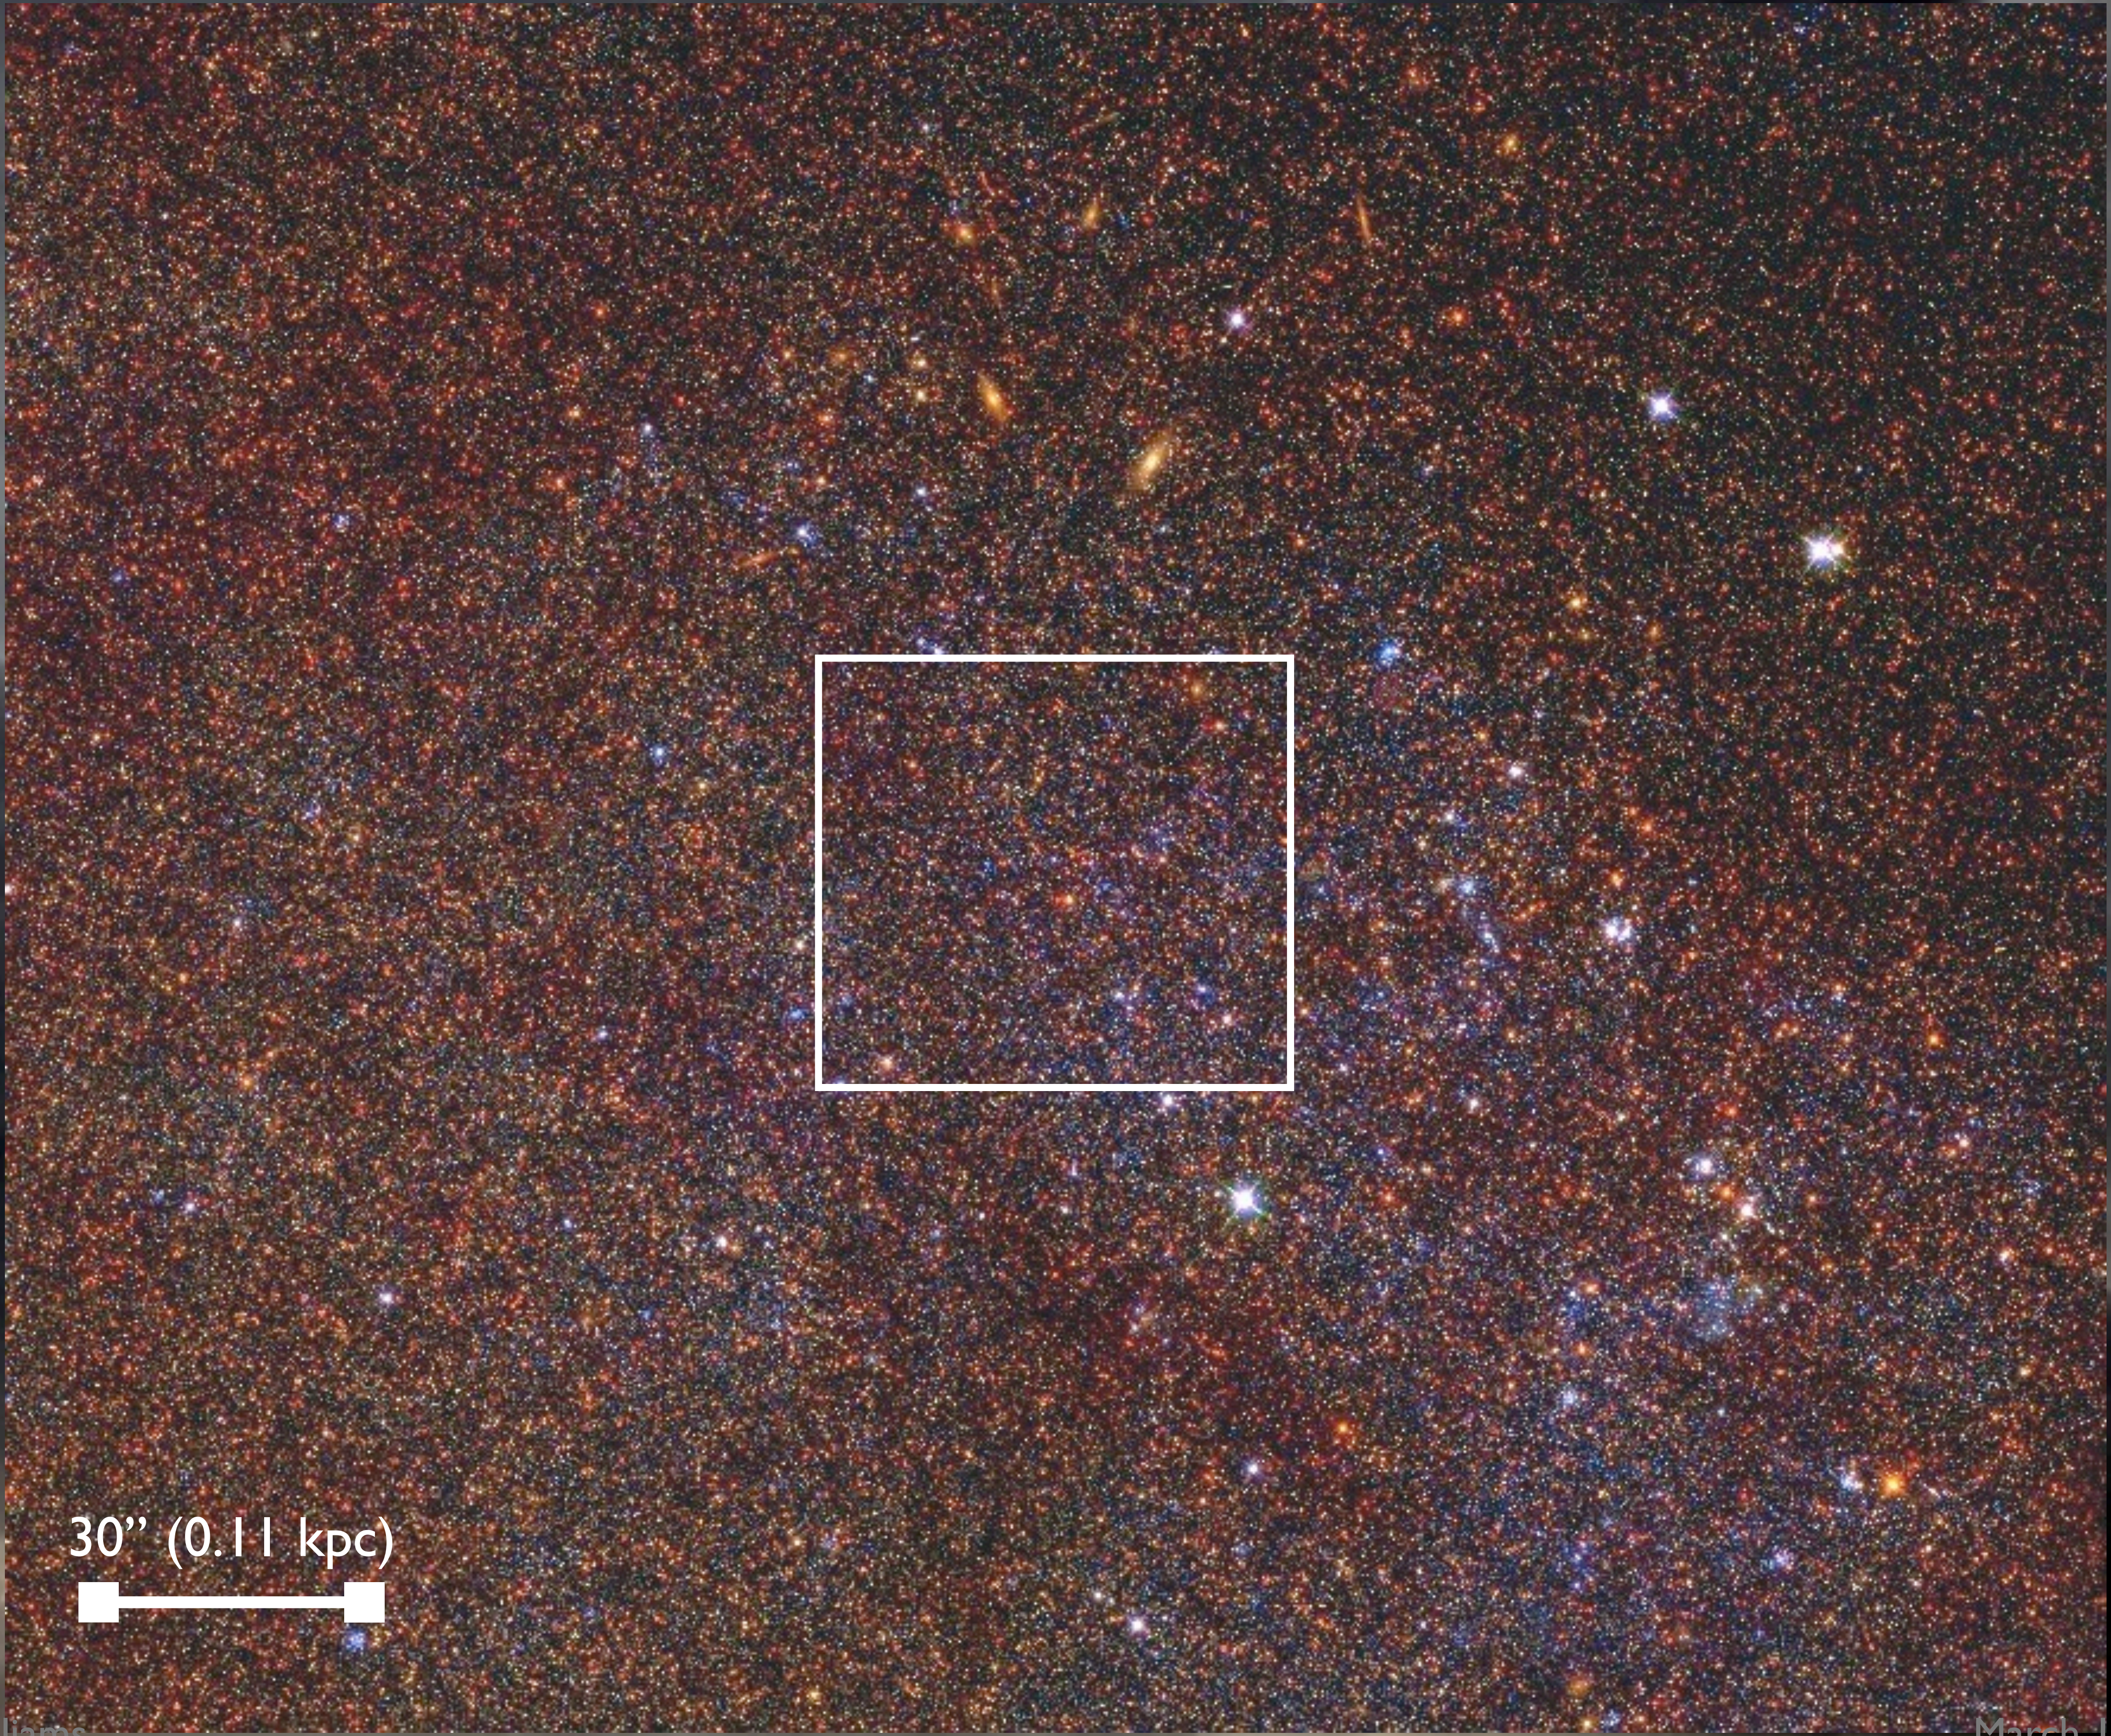
30" (0.11 kpc)

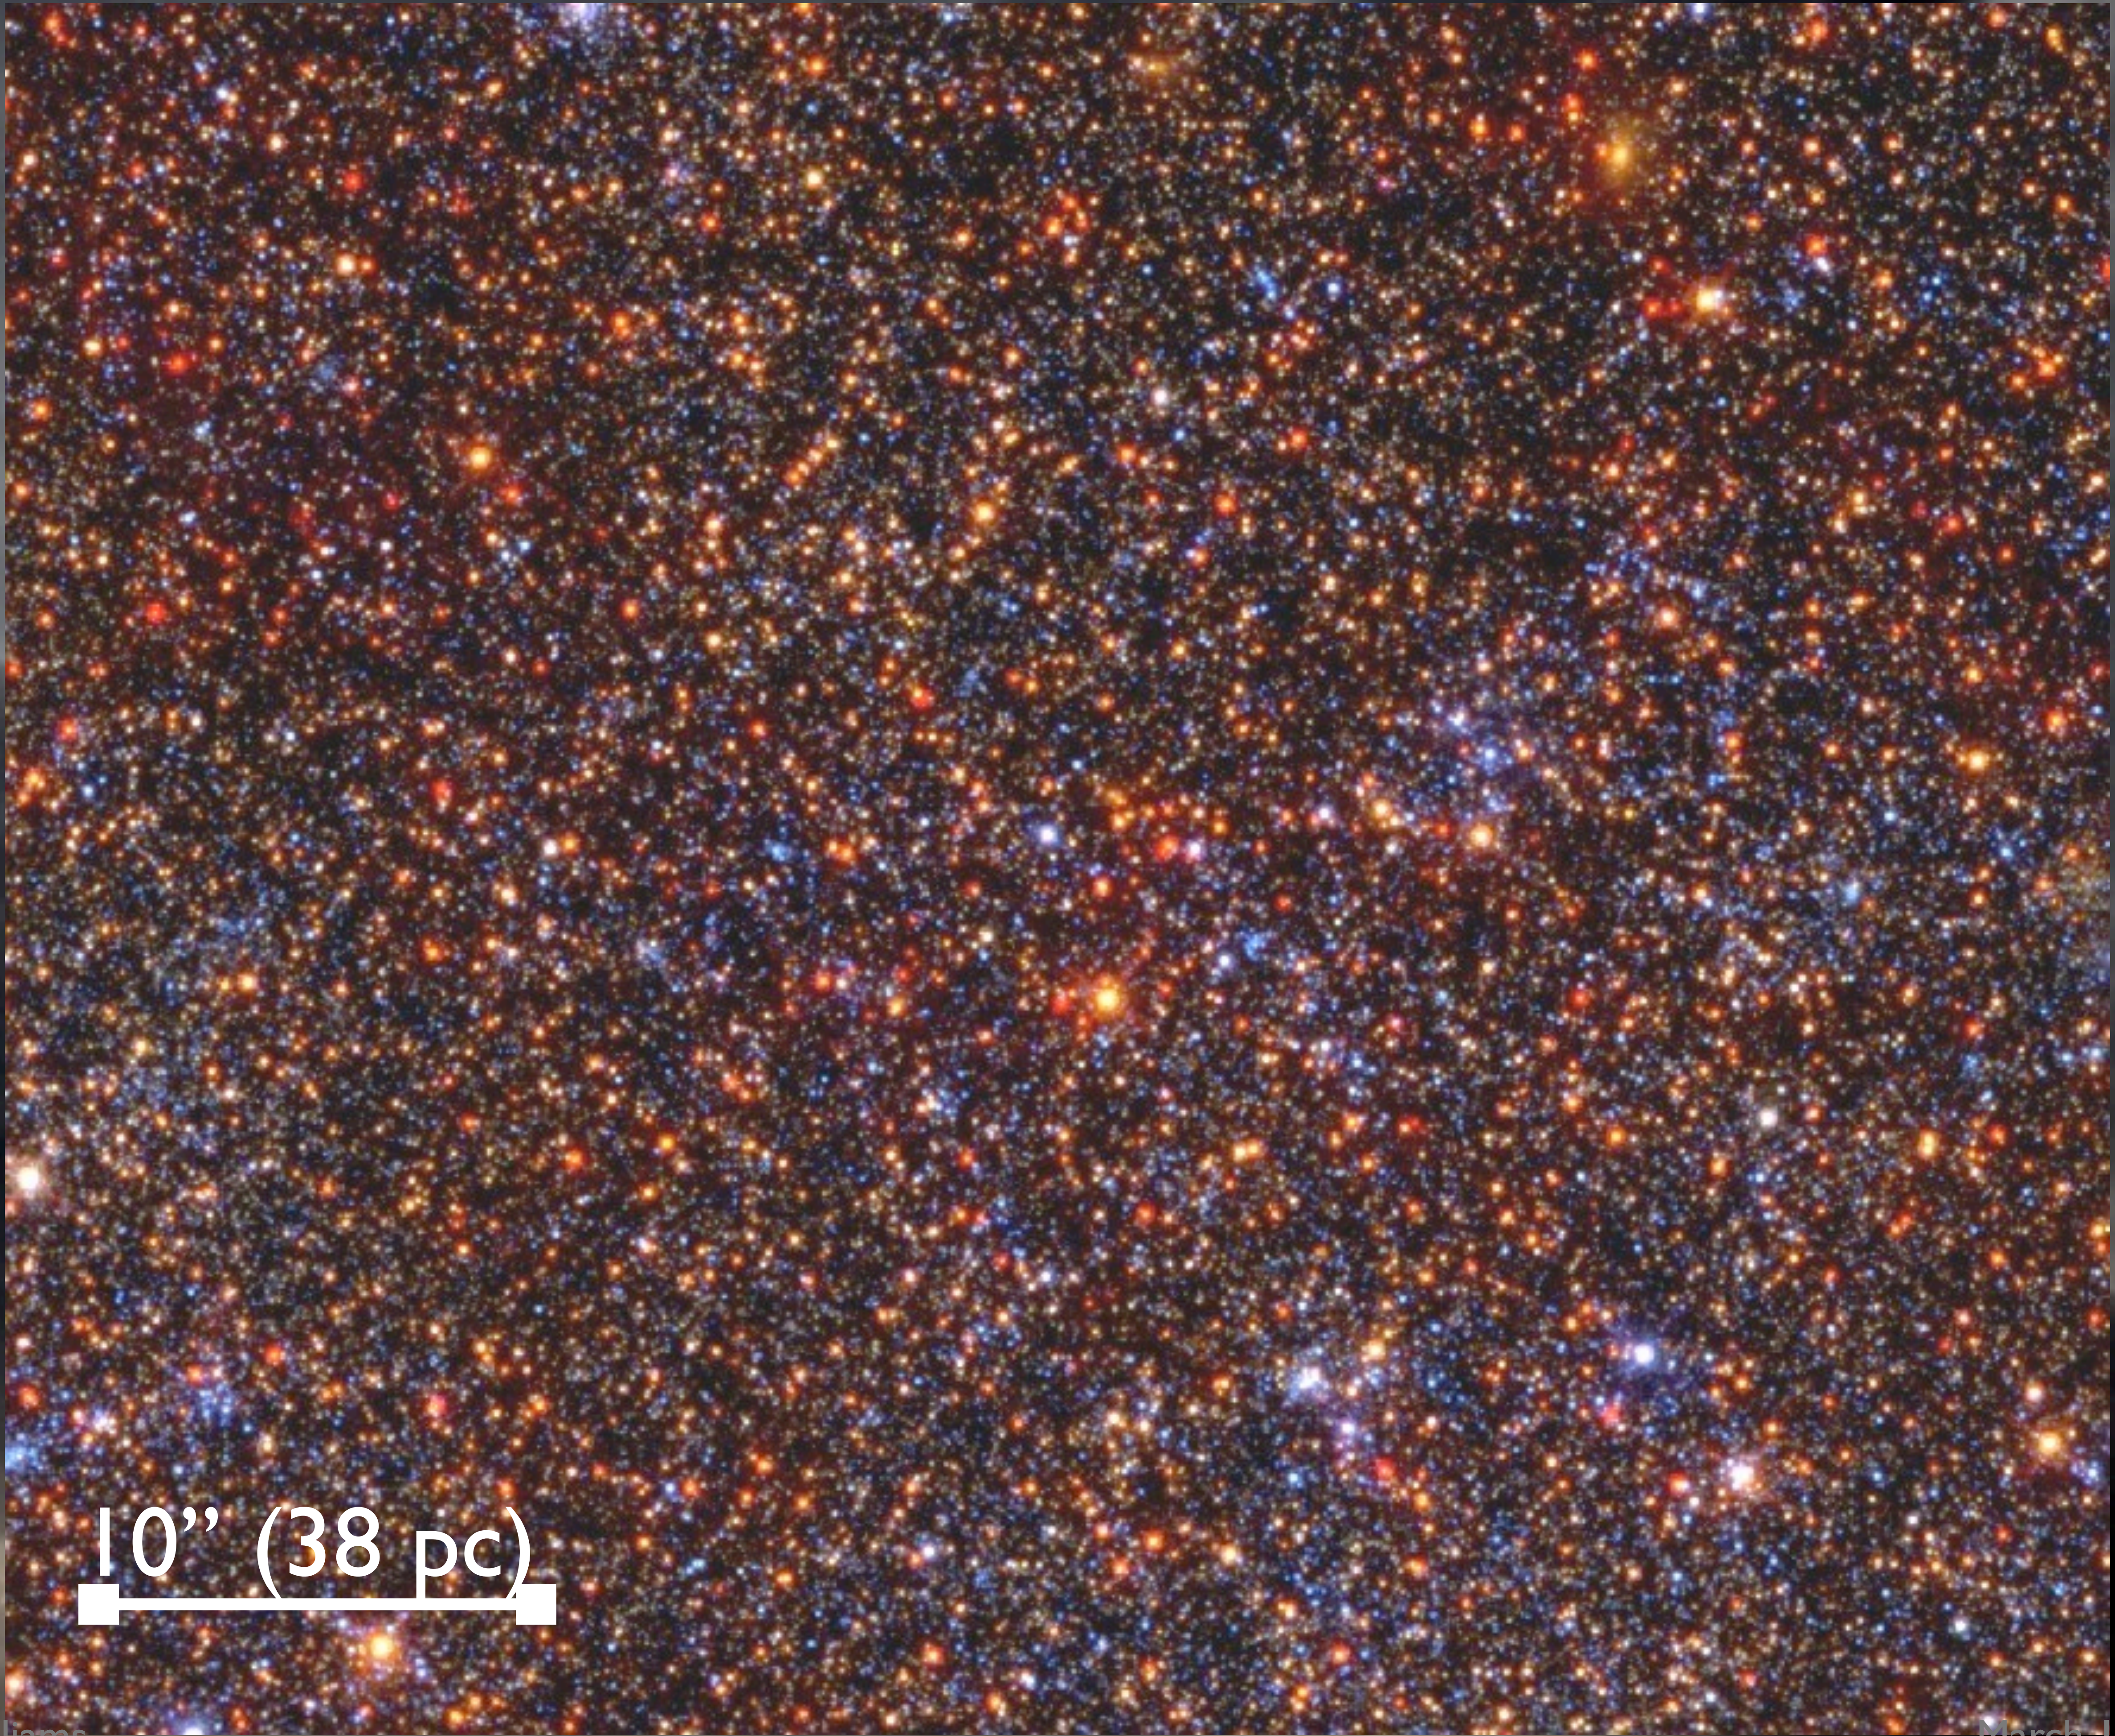
10" (38 pc)

"Crowding" varies w/ Radius

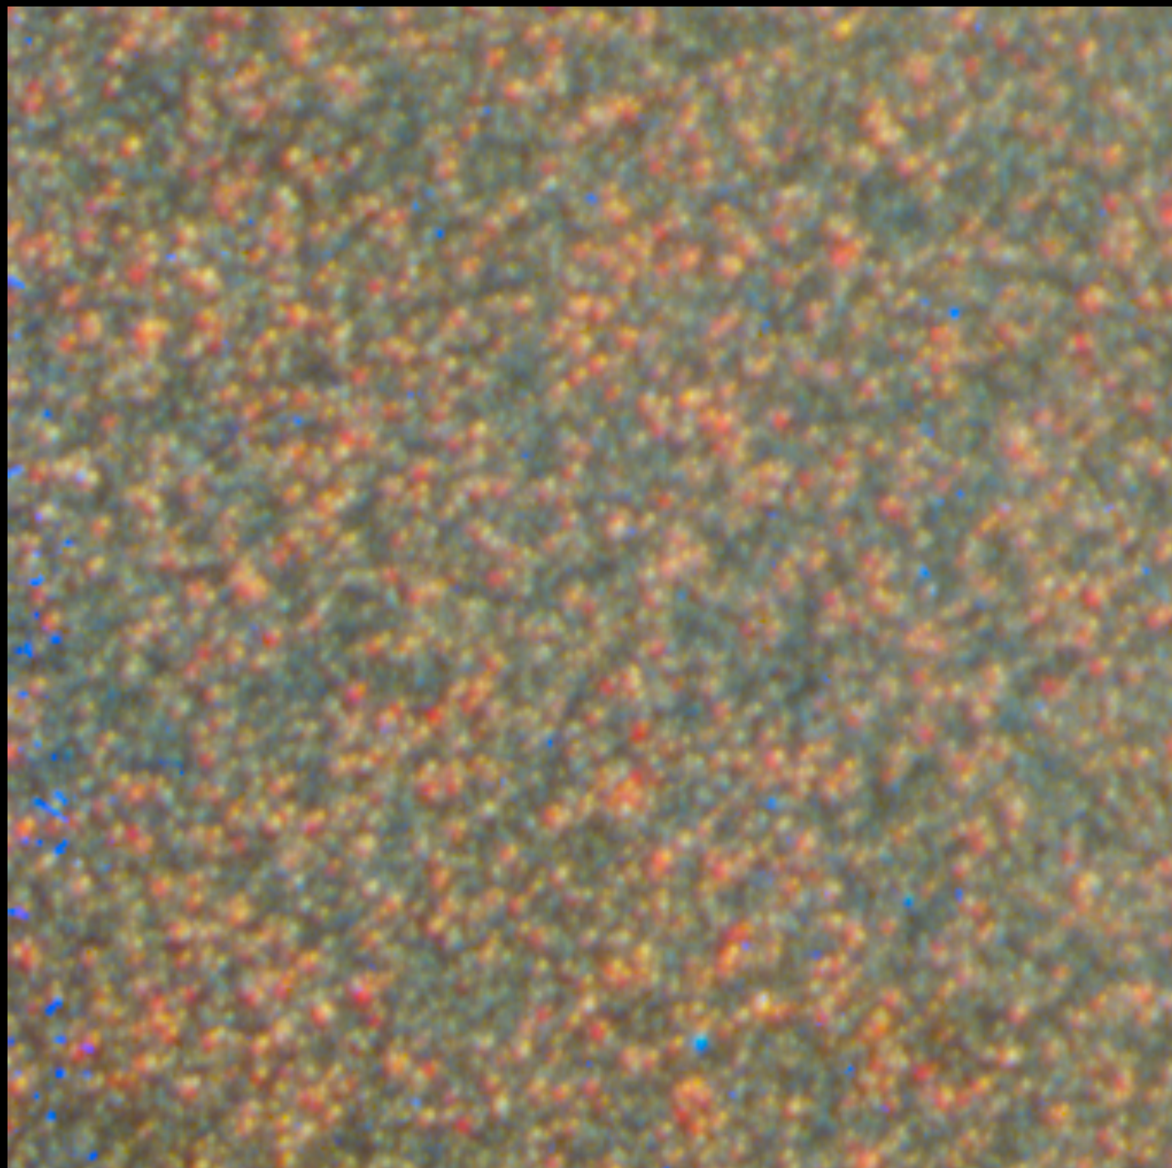

Bulge

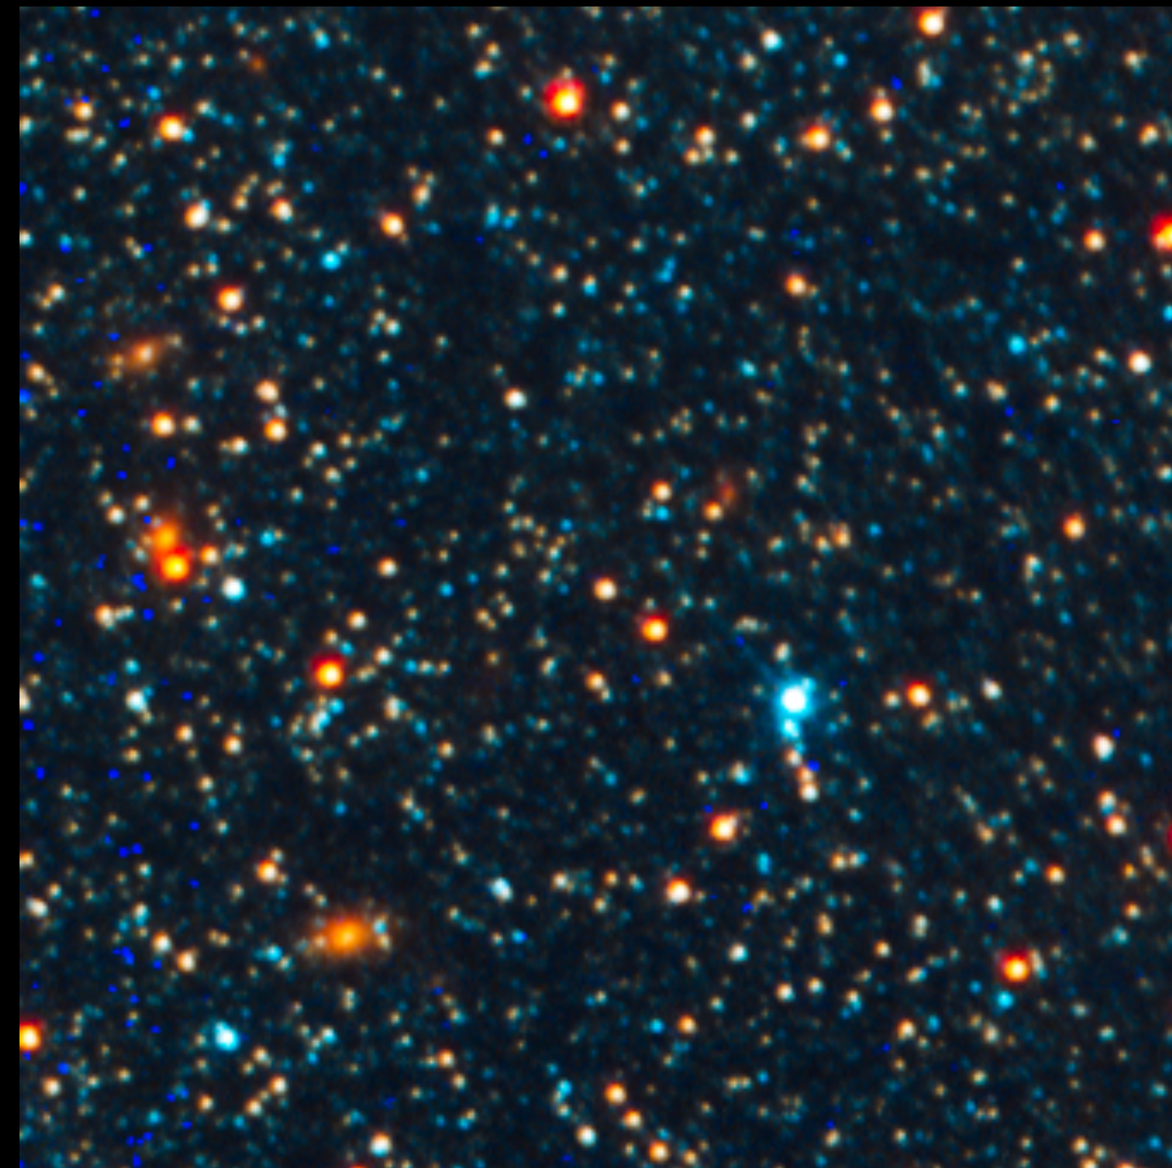

Outer Disk

Crowding Dominates Photometric Quality

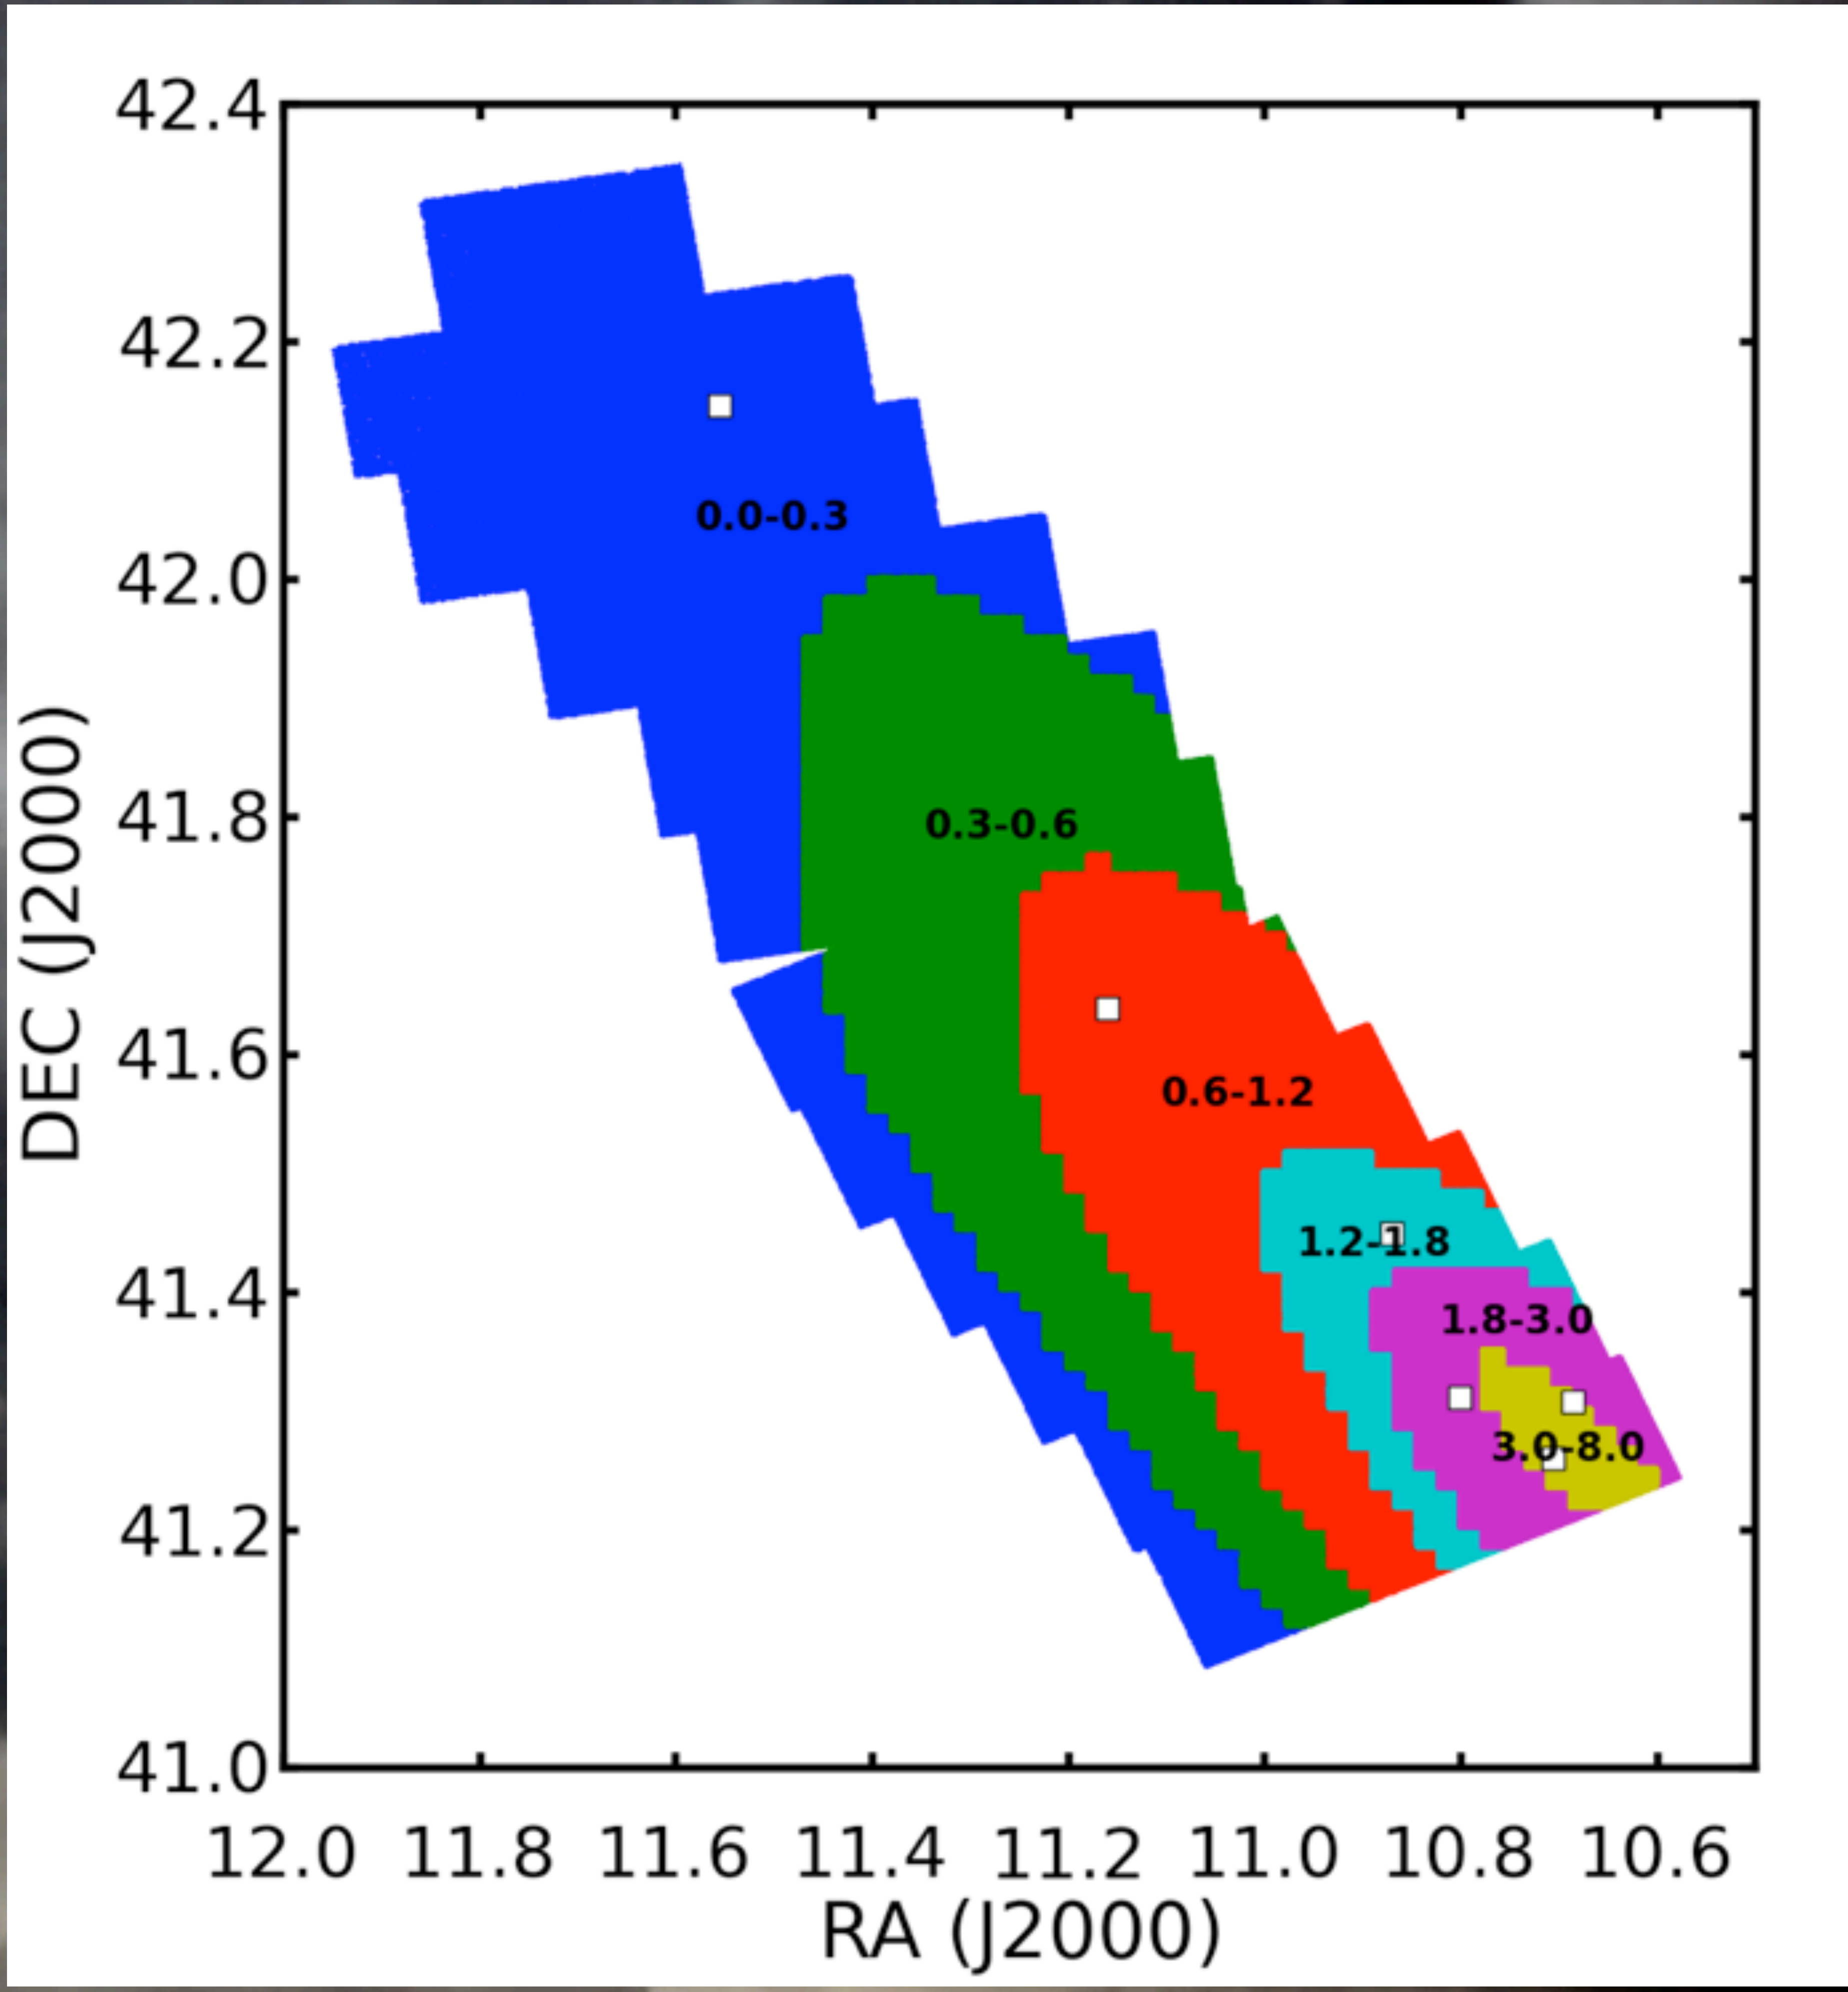

Optical-IR

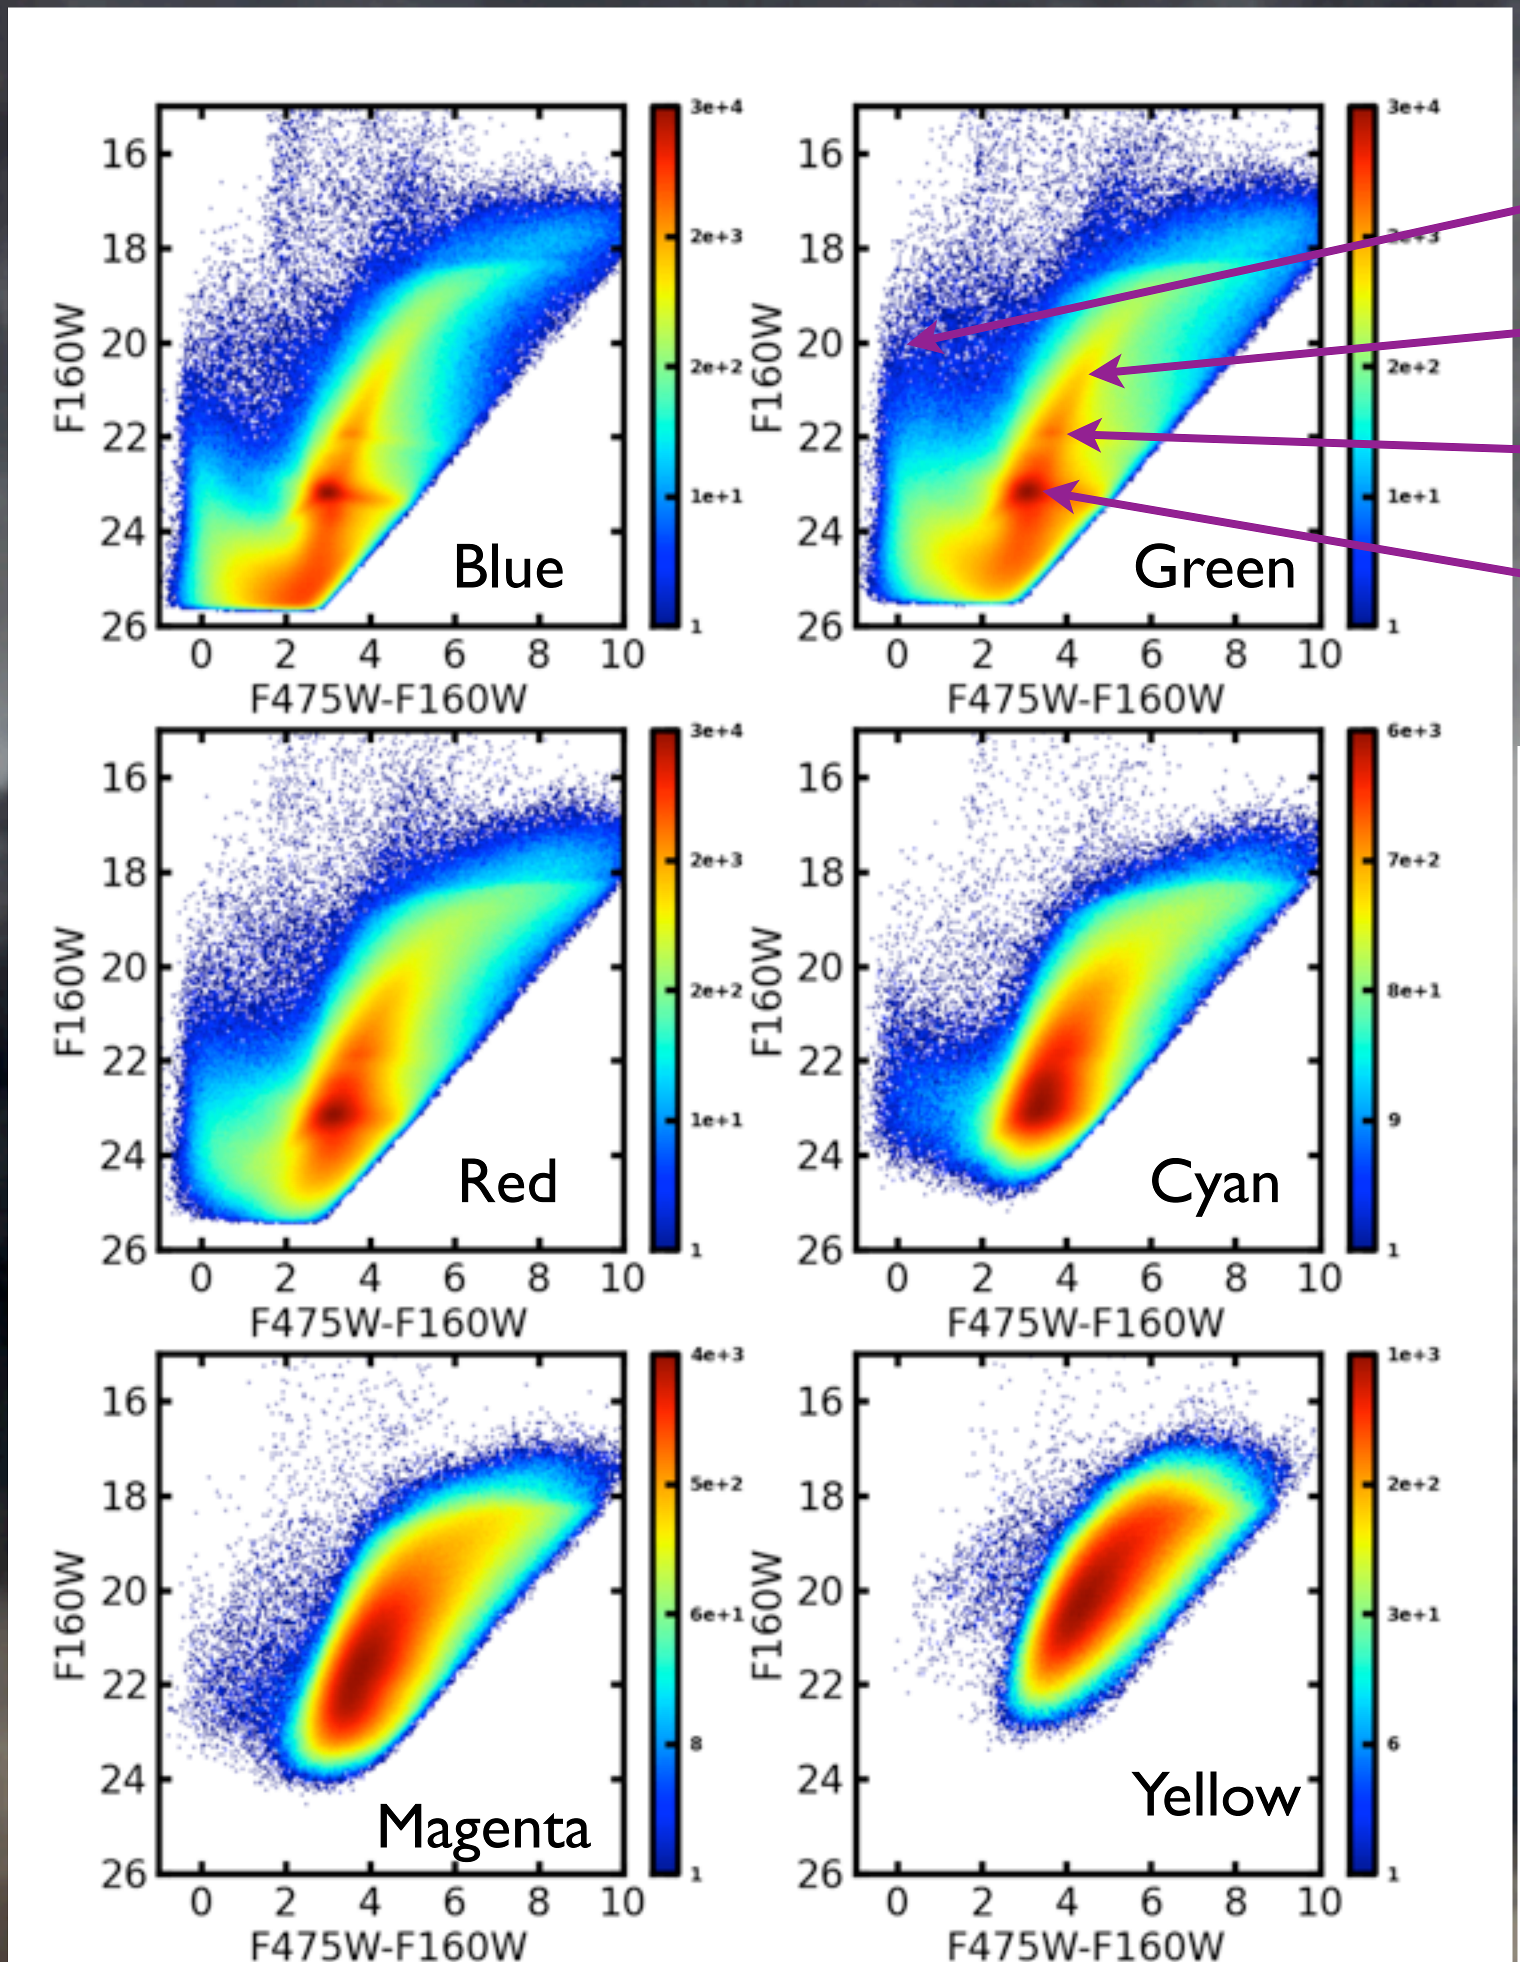

Main sequence
Red Giant Branch
Asymptotic Giant Branch Bump
Red Clump
Blue
Green
Red
Cyan
Magenta
Yellow
F160W
F475W-F160W

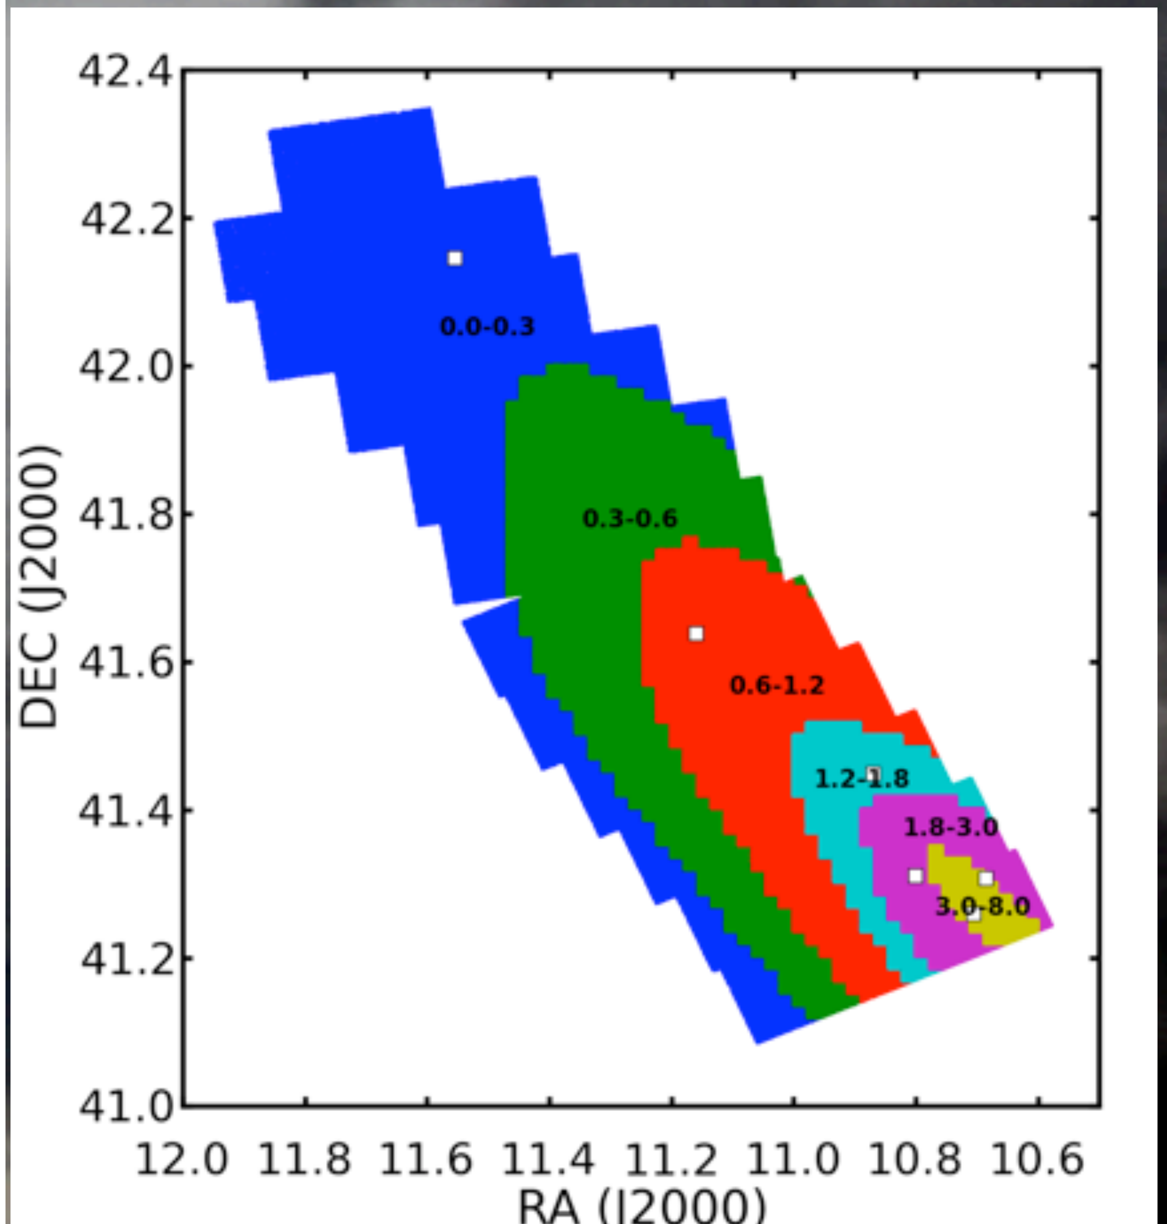

0.0-0.3
0.3-0.6
0.6-1.2
1.2-1.8
1.8-3.0
3.0-8.0
DEC (J2000)
RA (J2000)

Clusters (HST vs Ground)

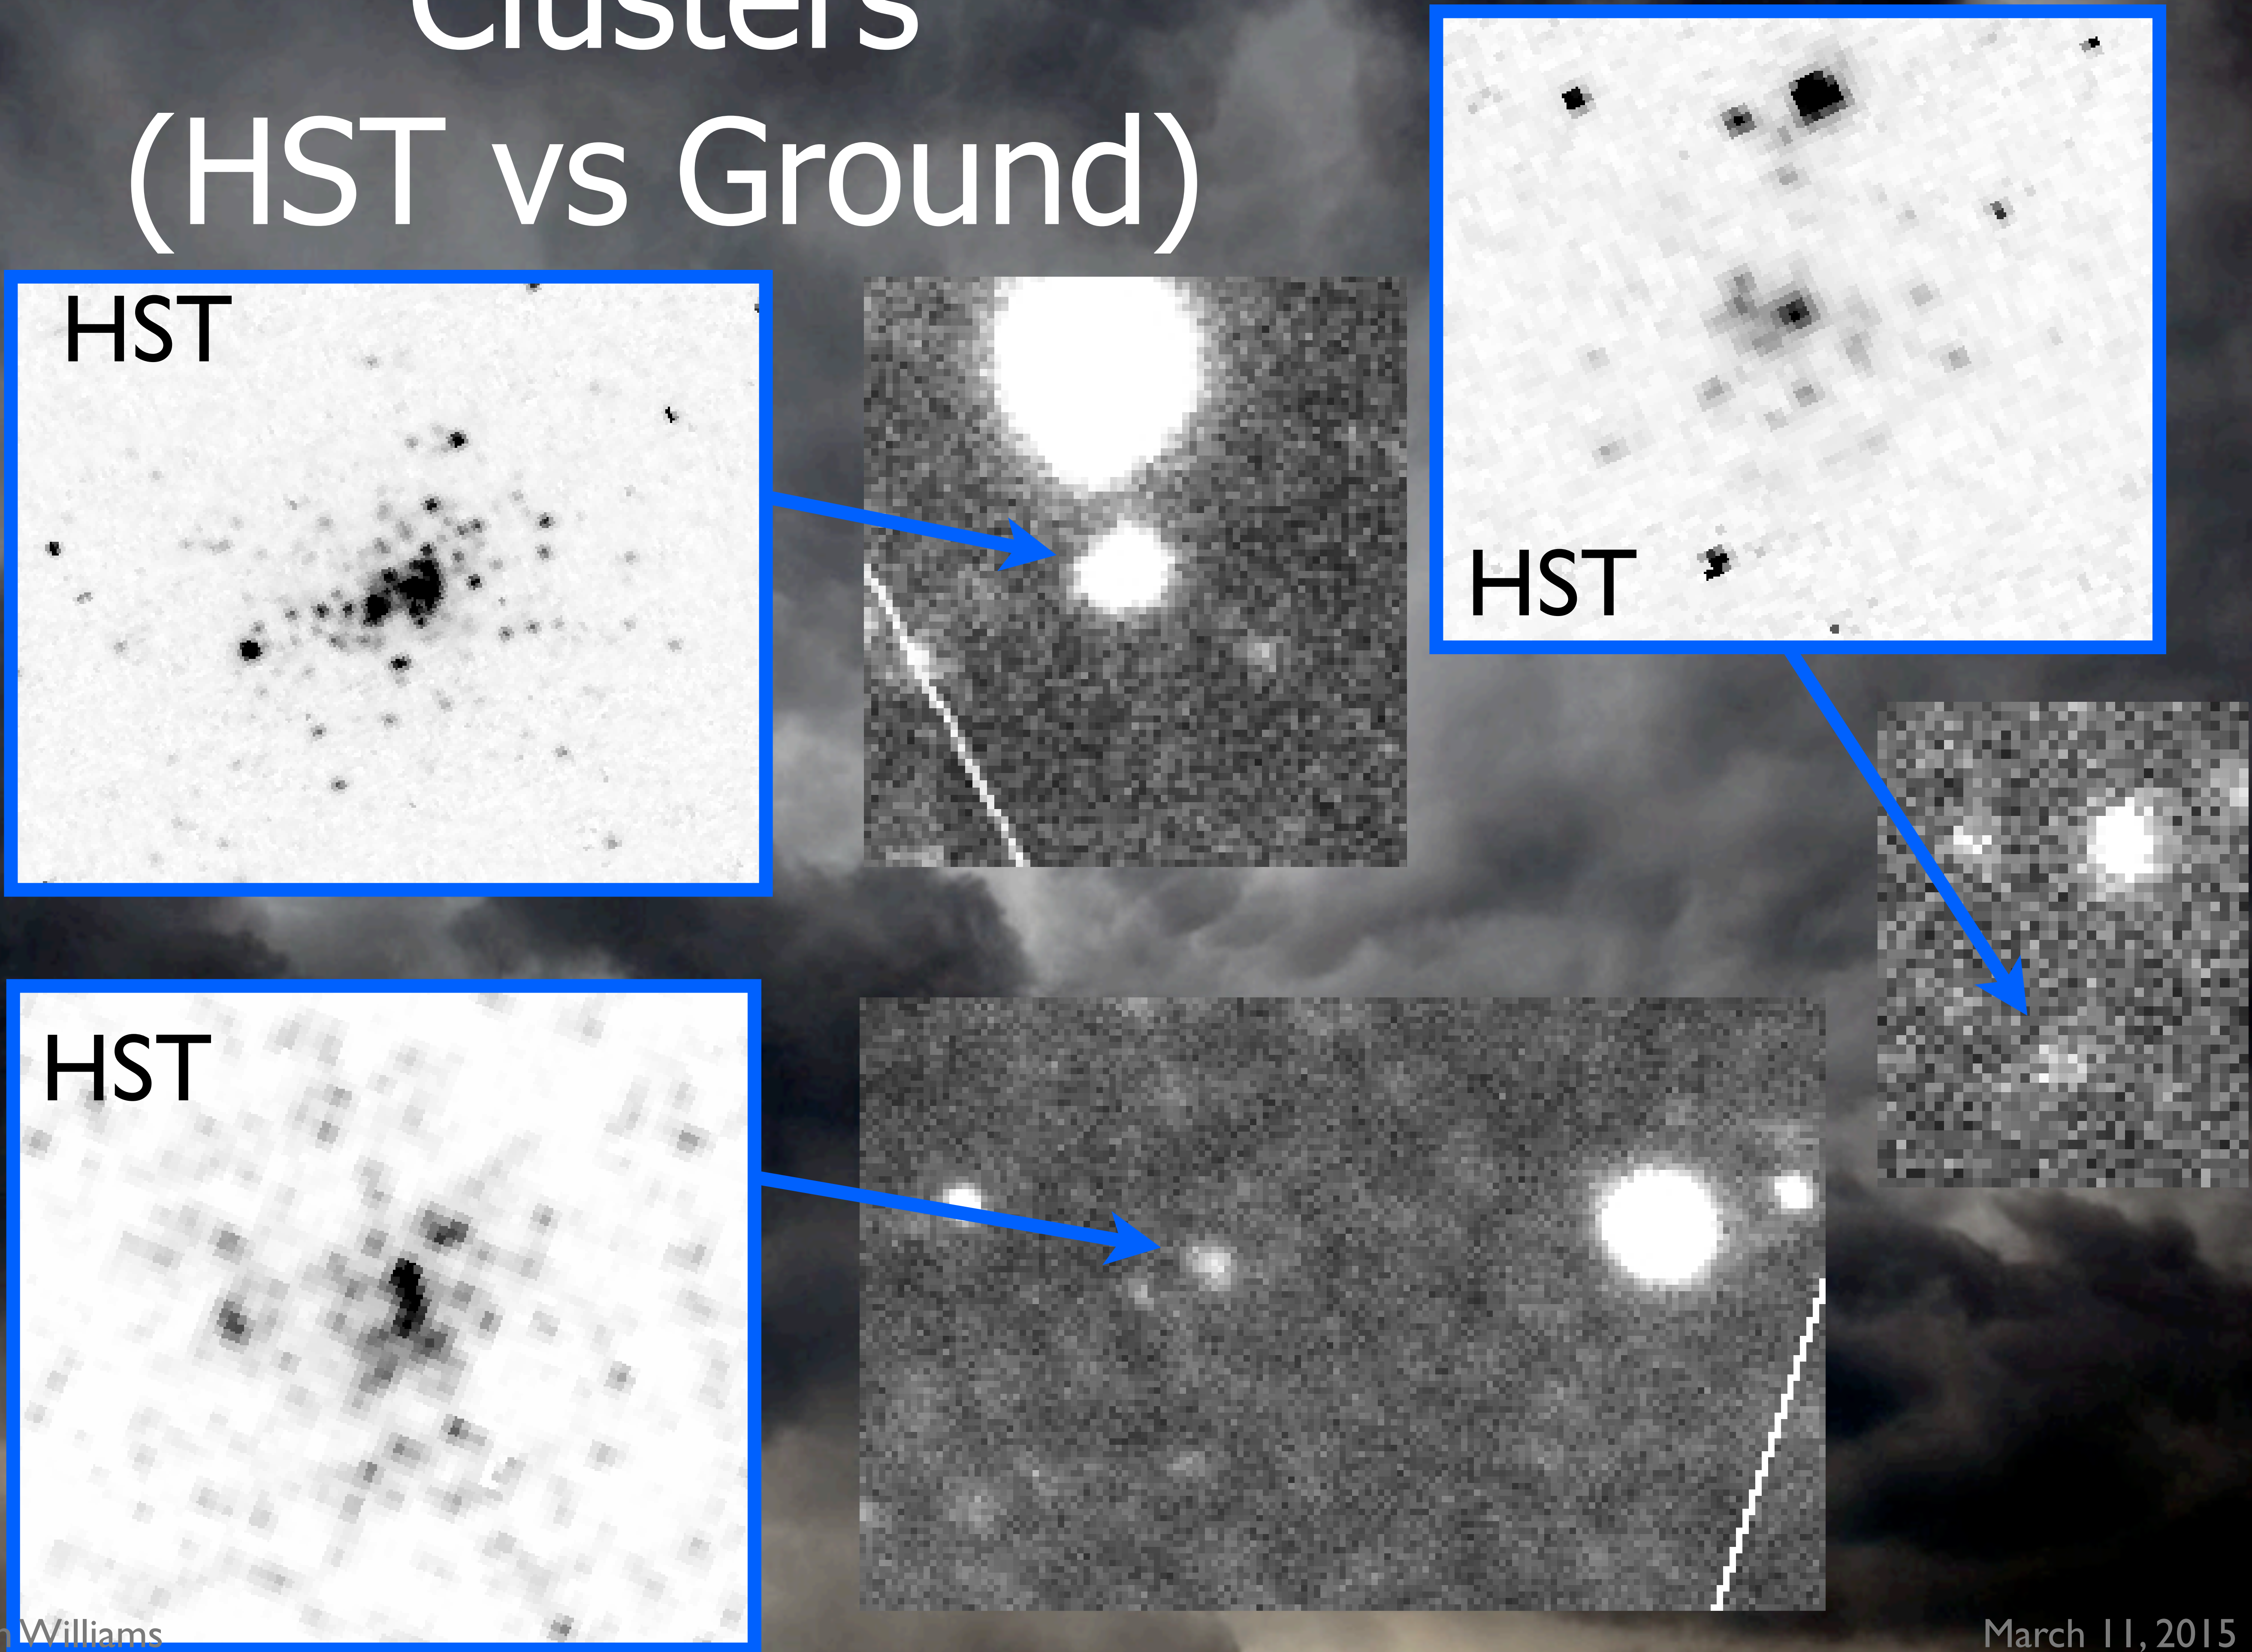

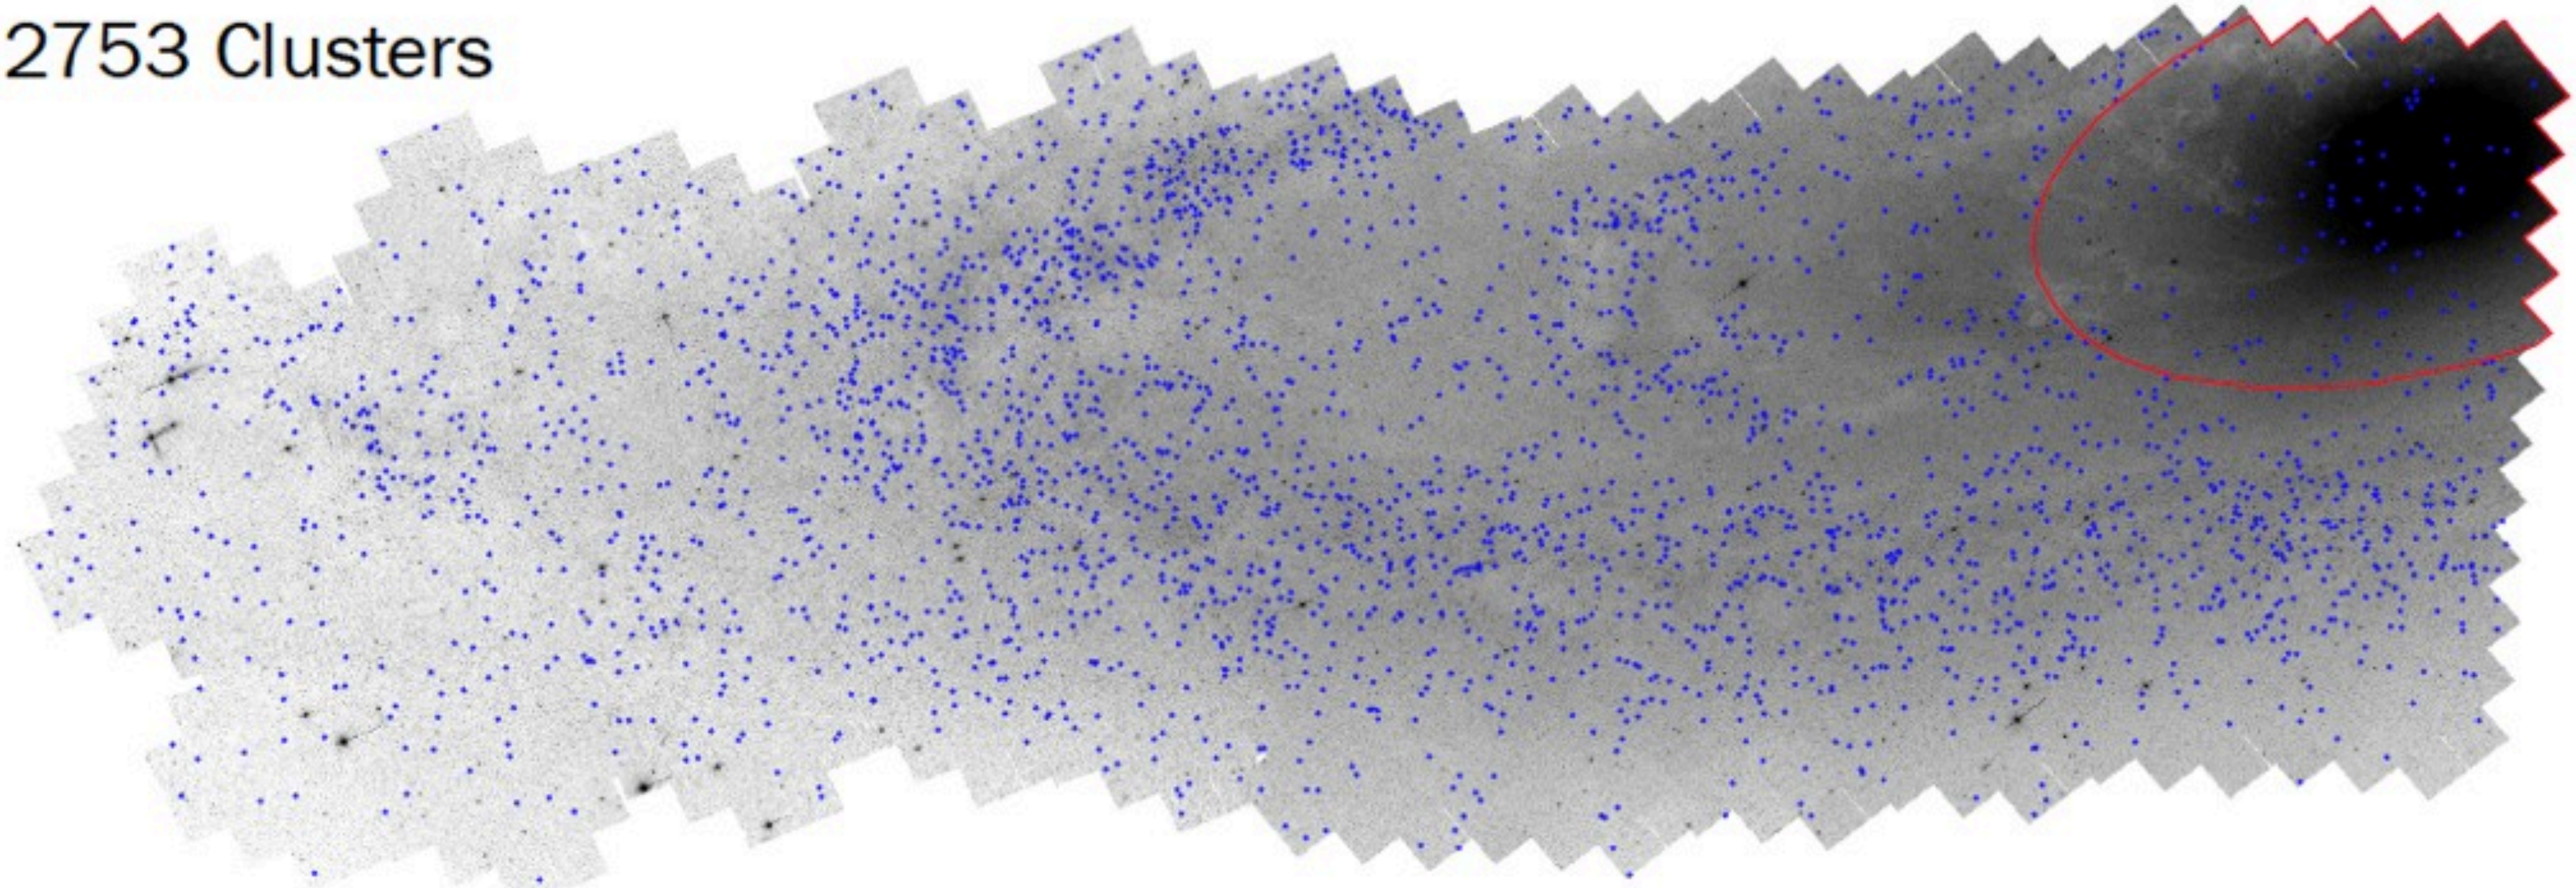

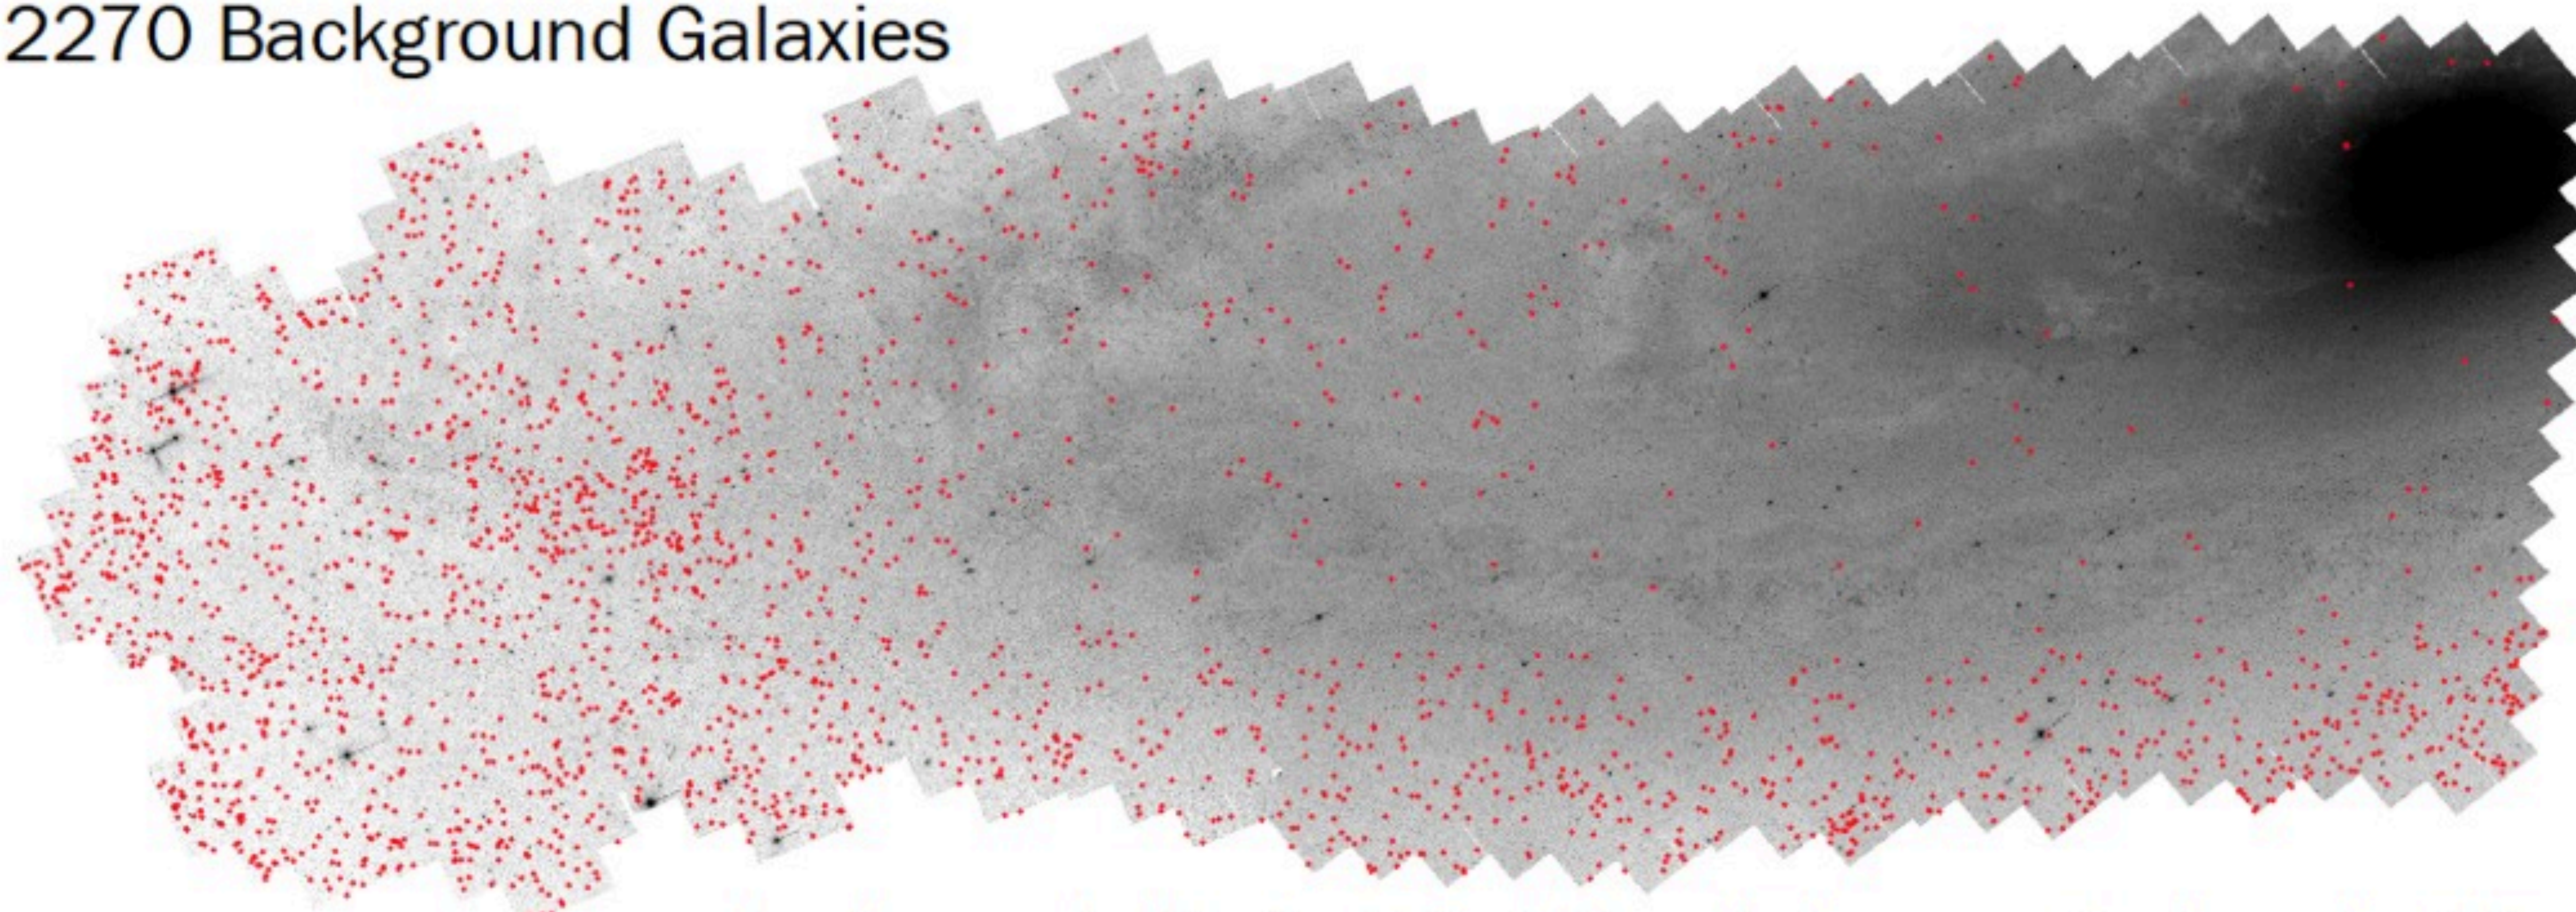

Andromeda Project Catalog: Johnson et al., submitted

6x increase in number of known clusters within PHAT footprint.

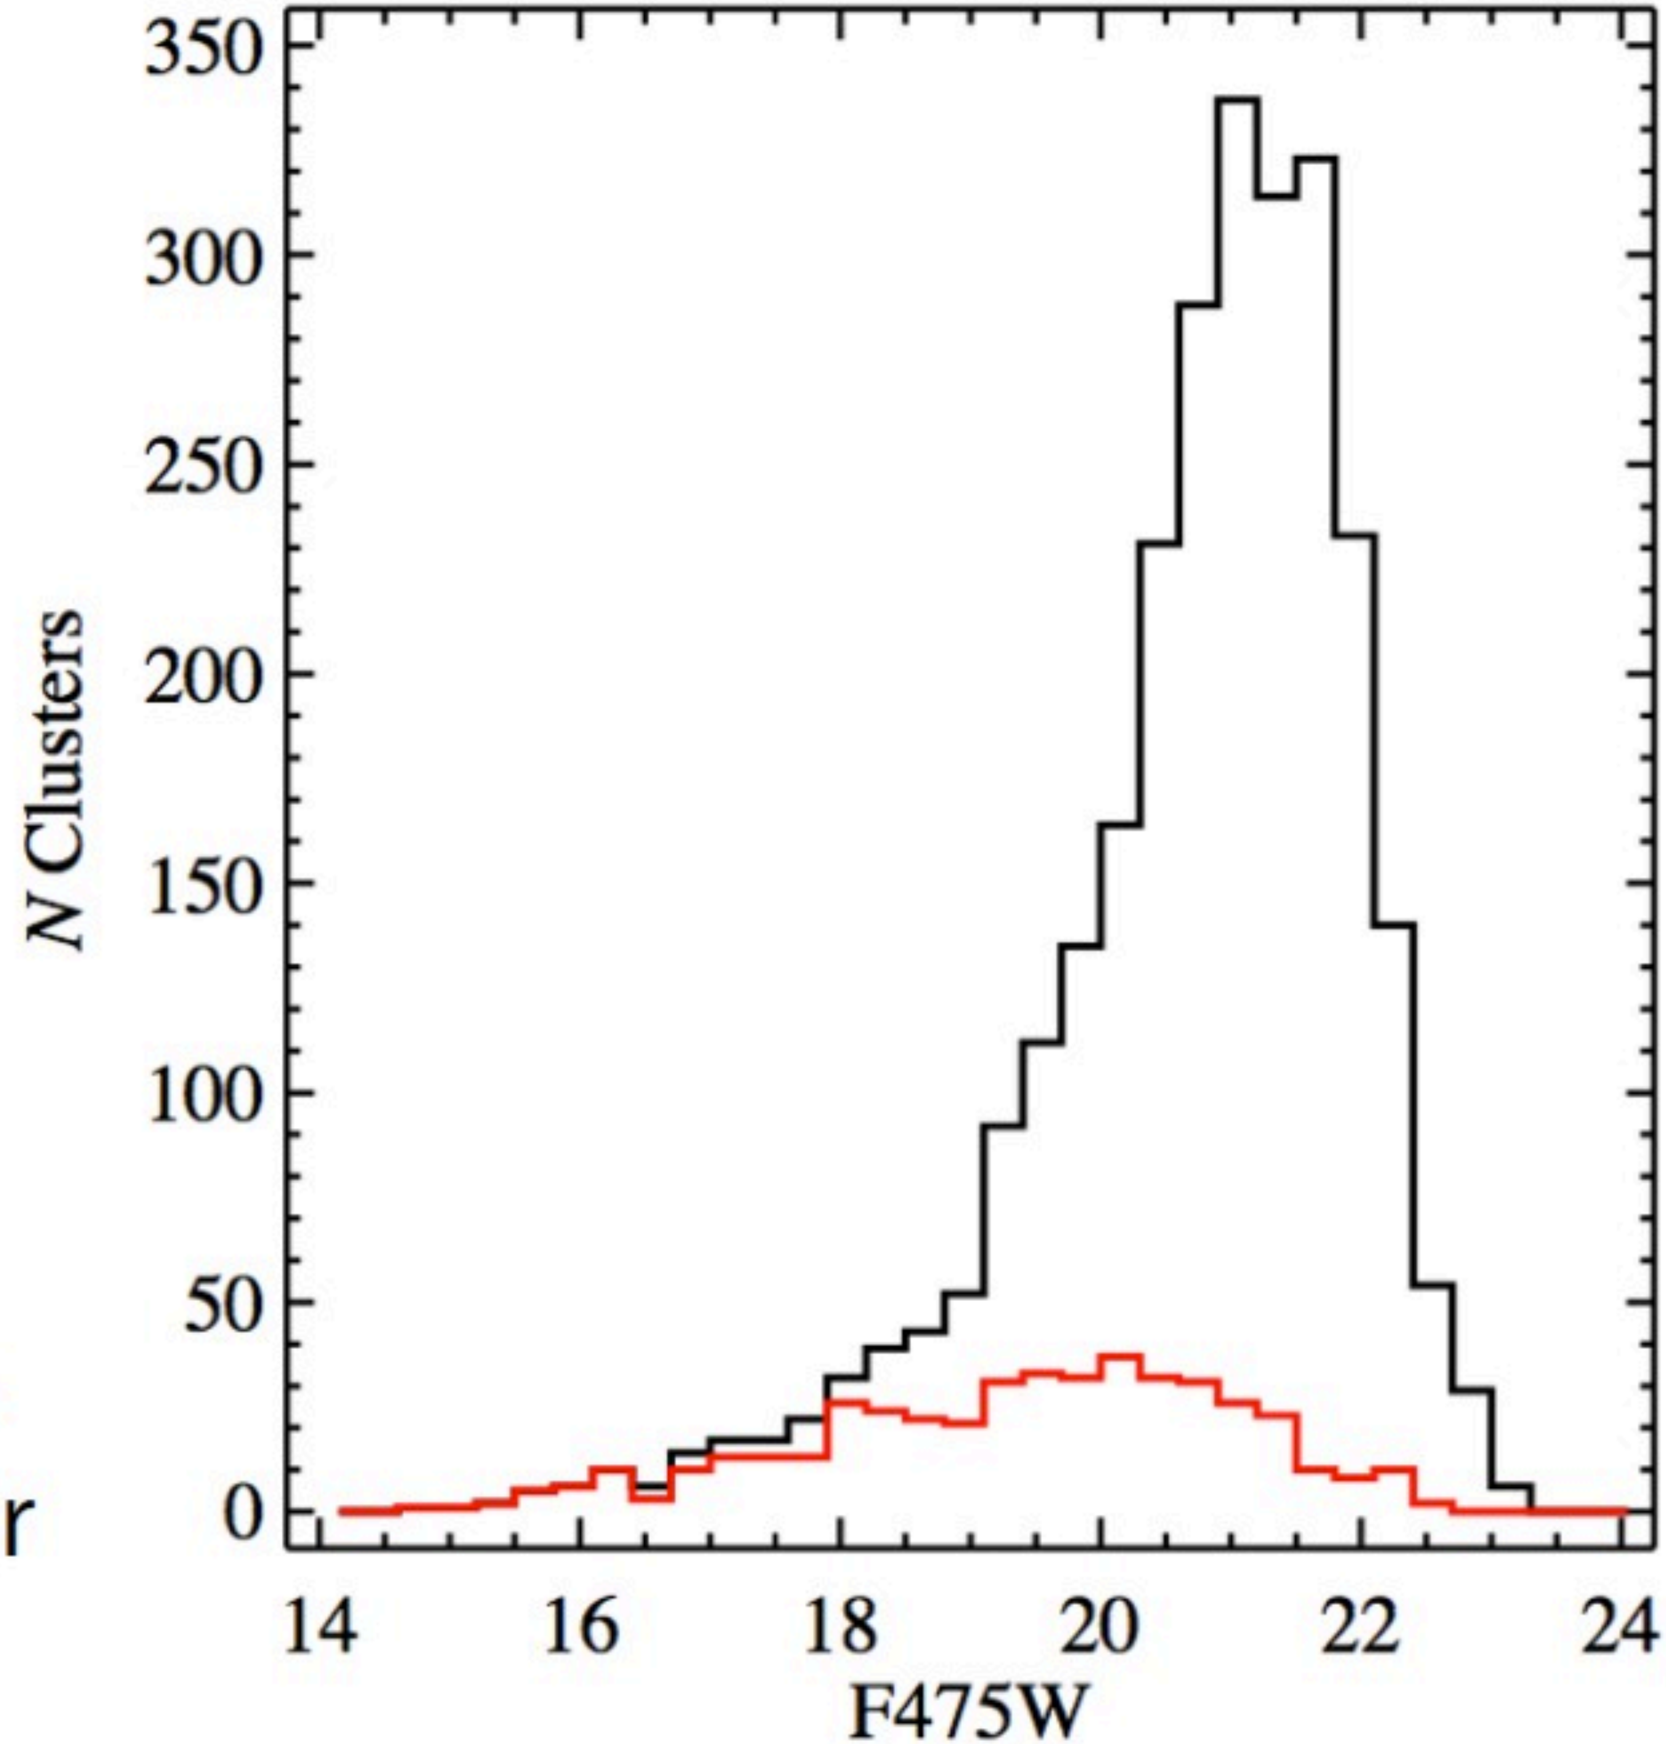

Catalog complete down to 500 $M_\odot$ for ages <100 Myr.

Young clusters probe Initial Mass Function

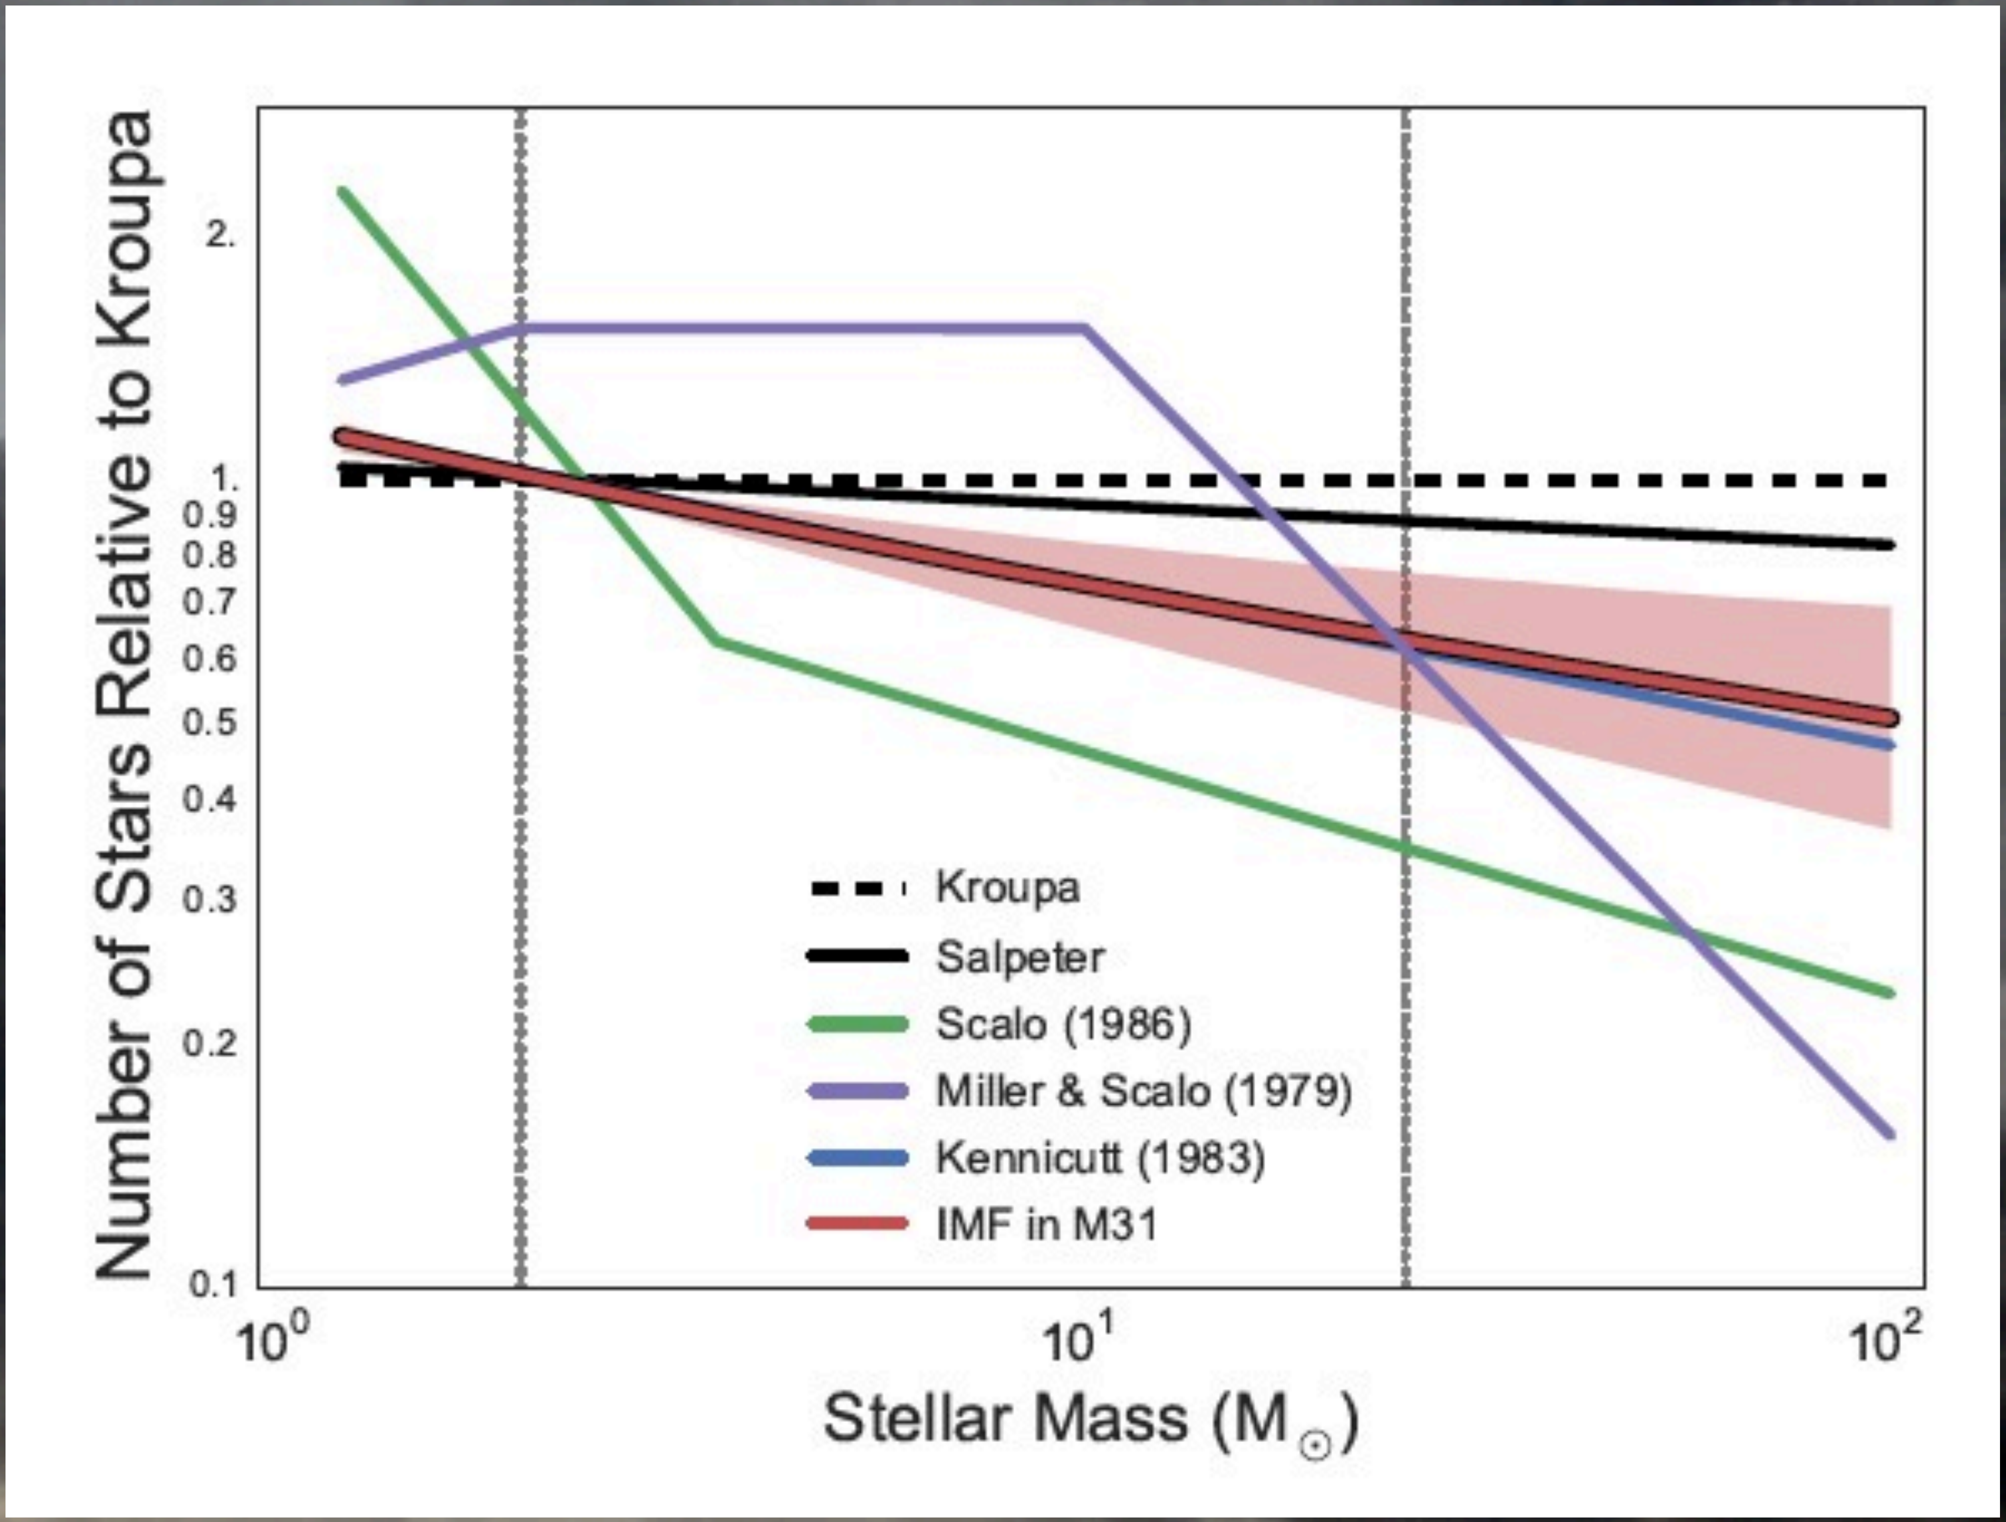

D. Weisz et al., submitted

Initial Mass Function

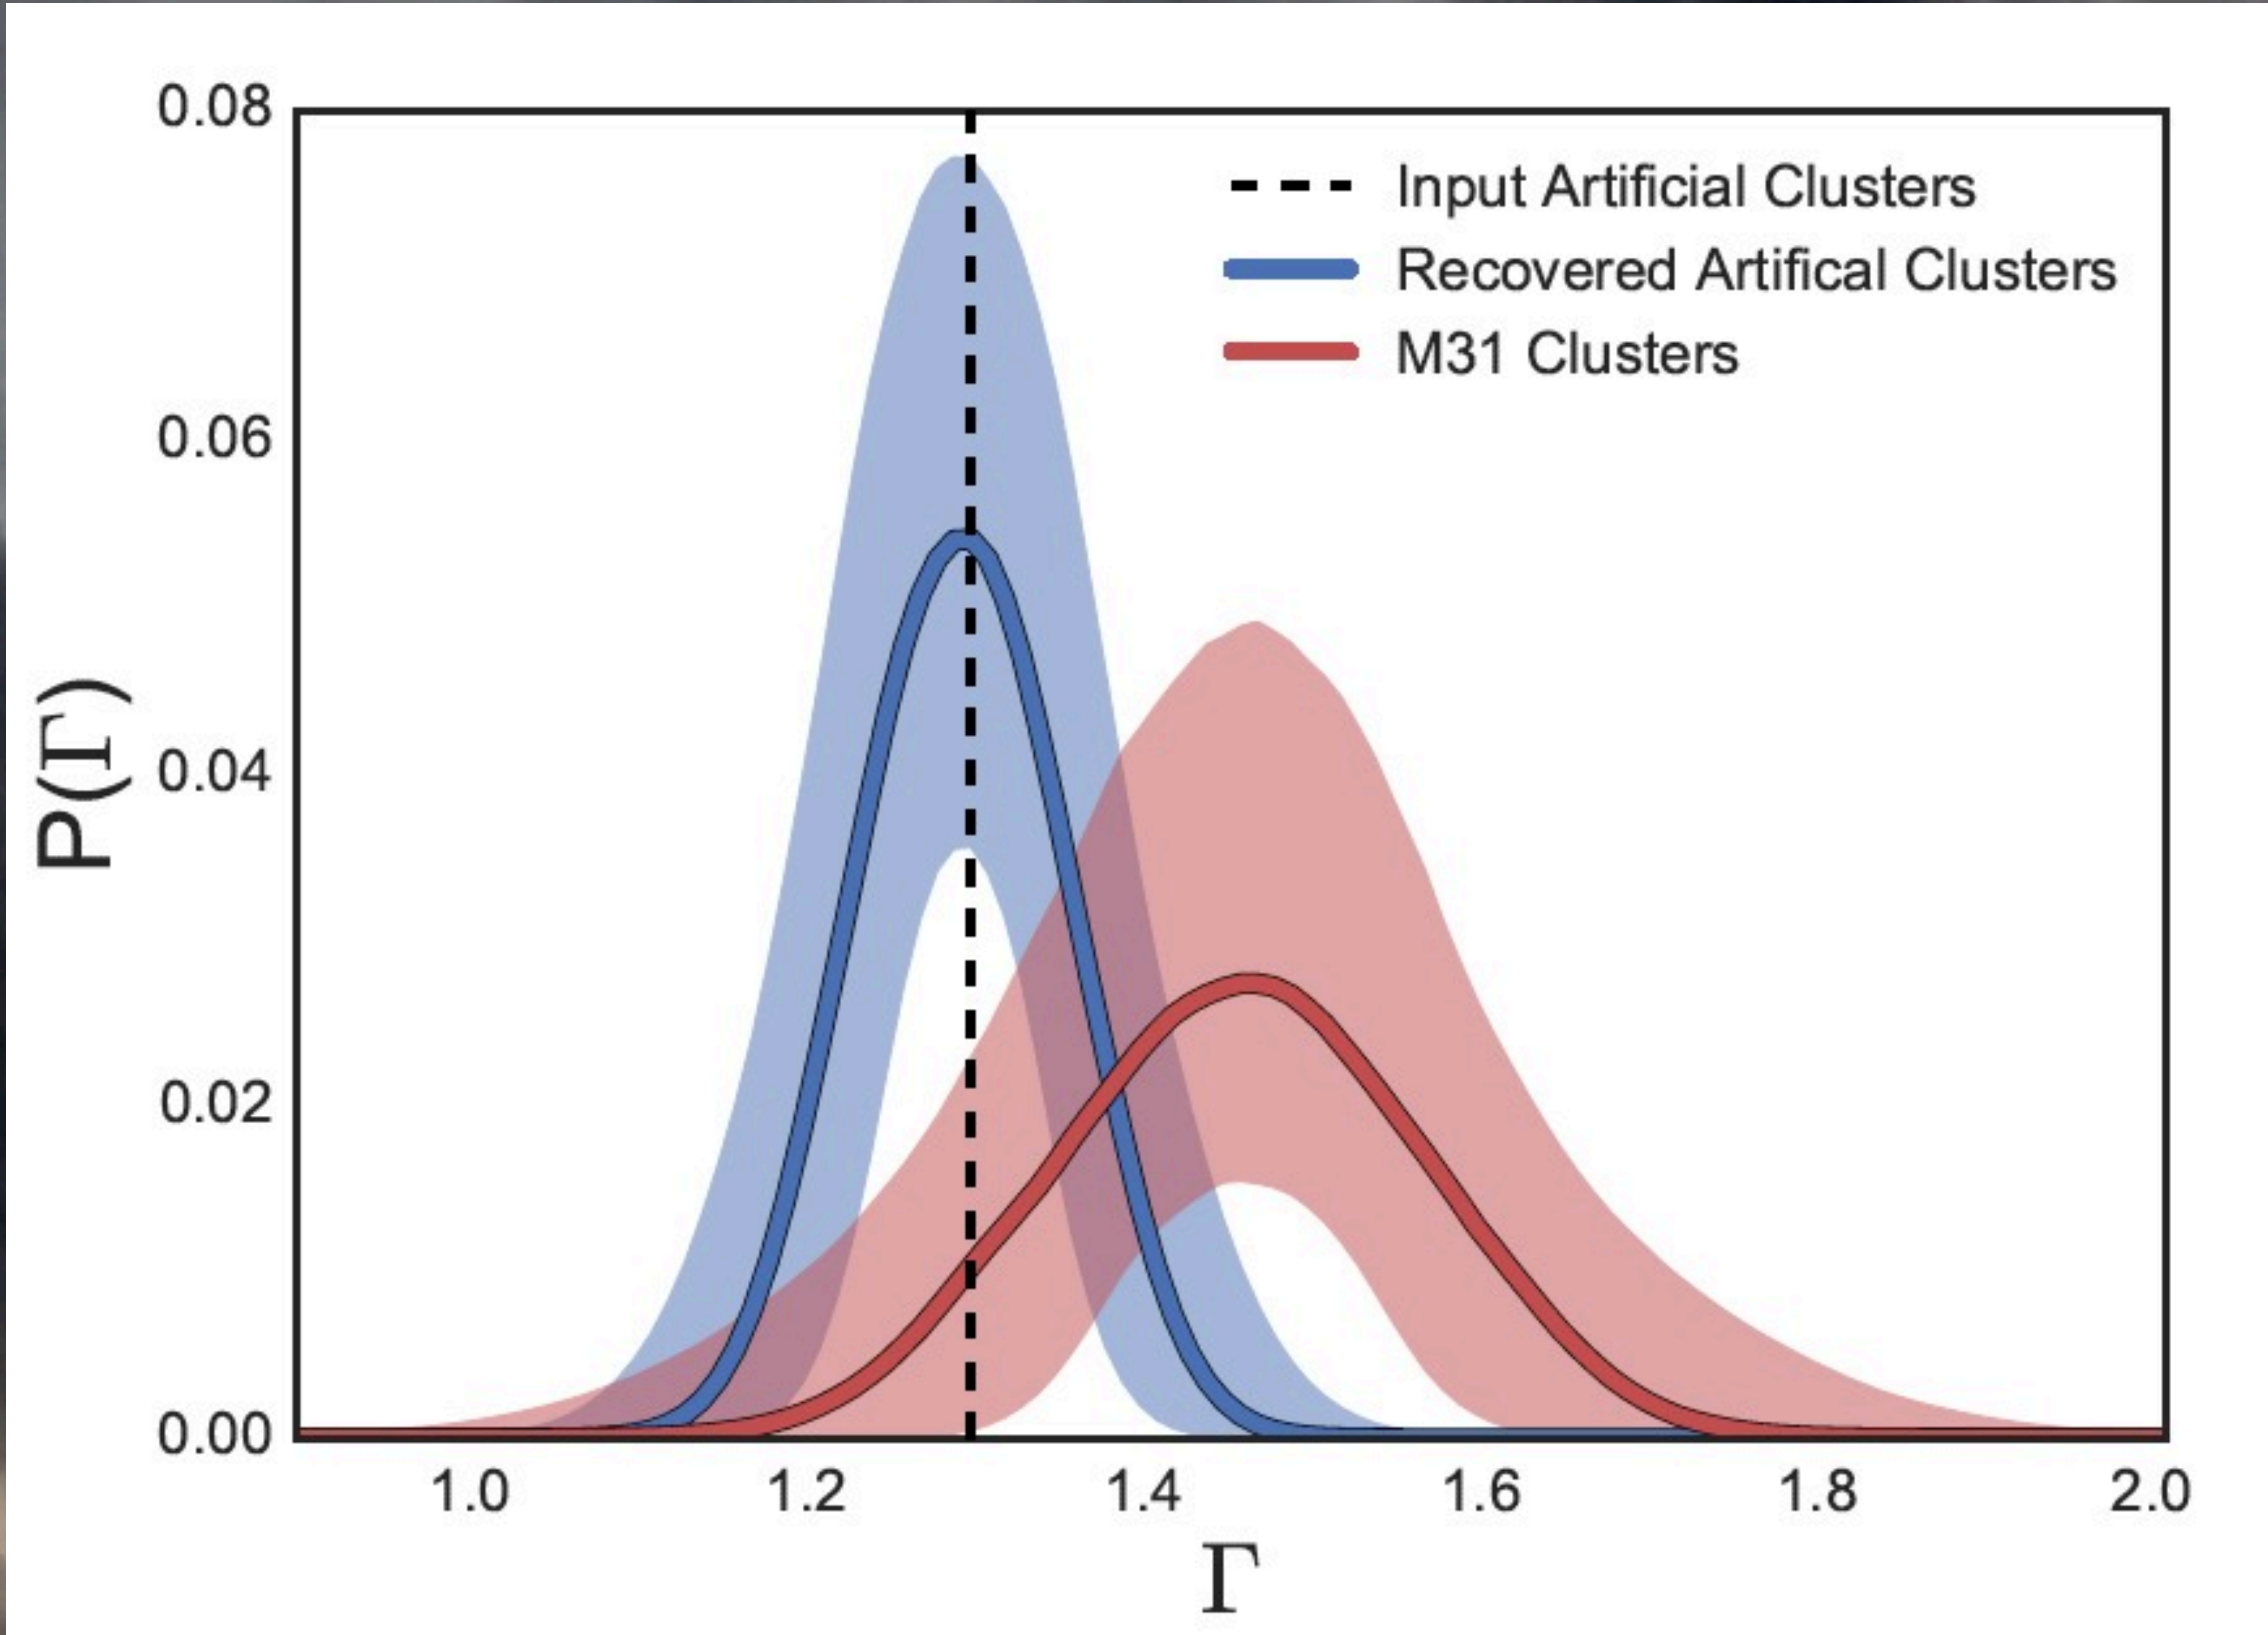

D. Weisz et al., submitted

Red Giant Branch Sensitive to Extinction

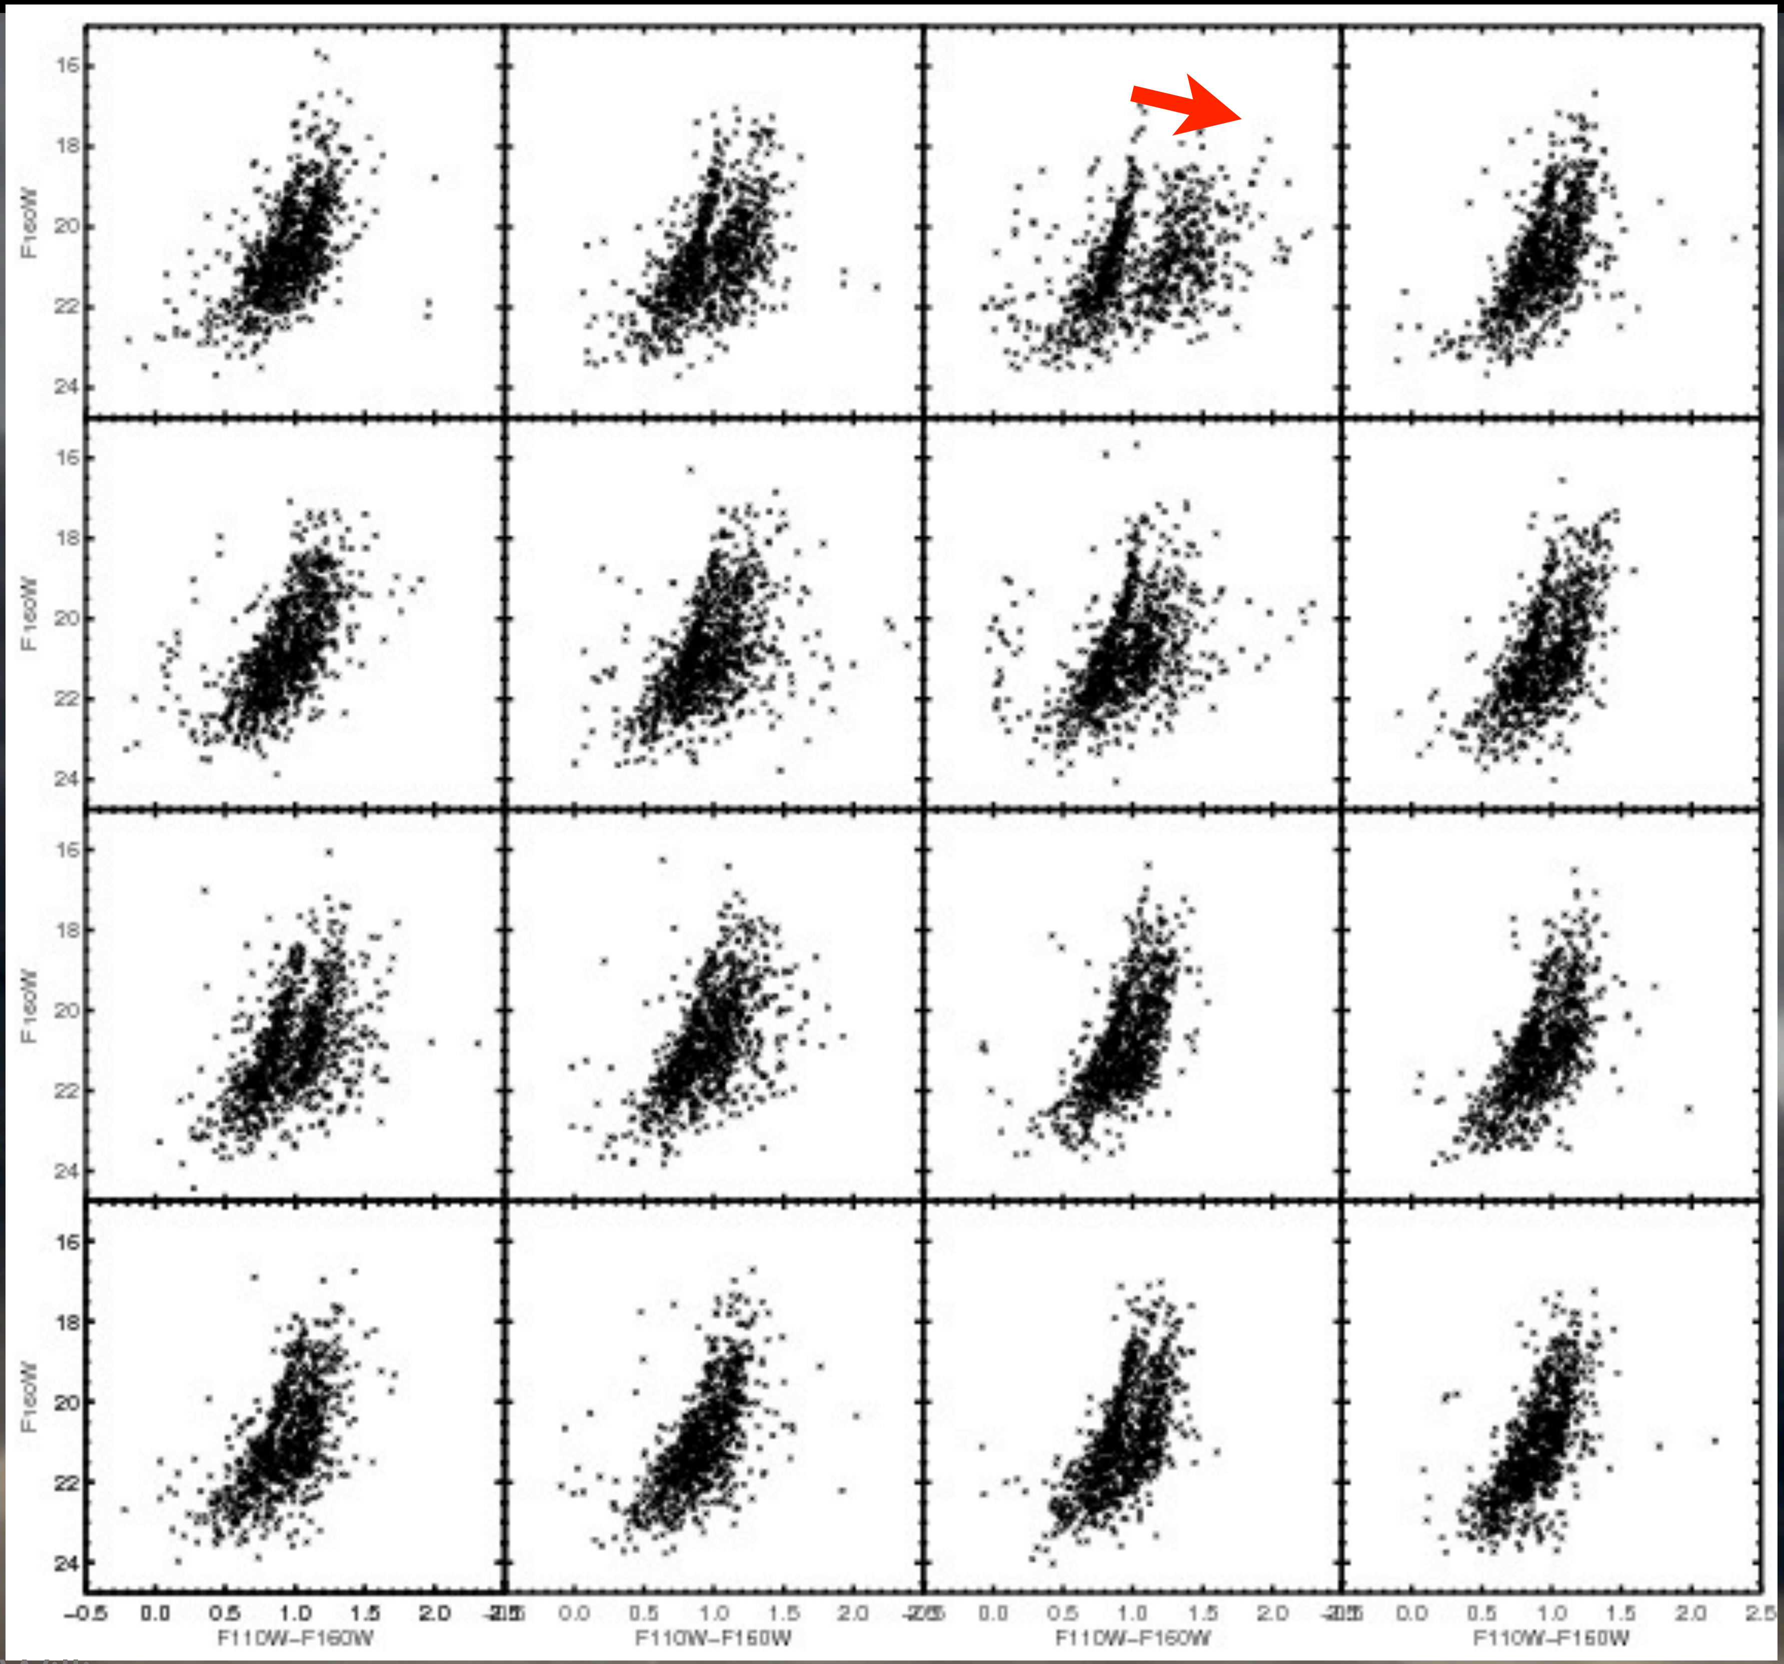

F110W-F160W CMD of Single WFC3/IR frame, subdivided in 4x4 grid

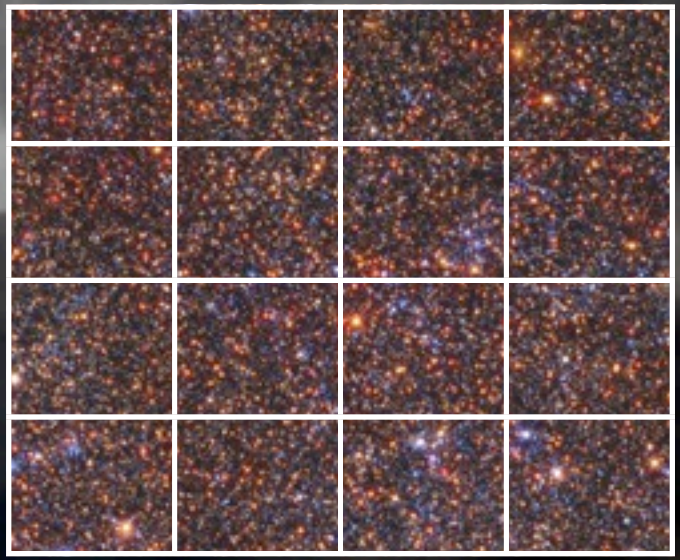

Subregions of single WFC3/IR frame

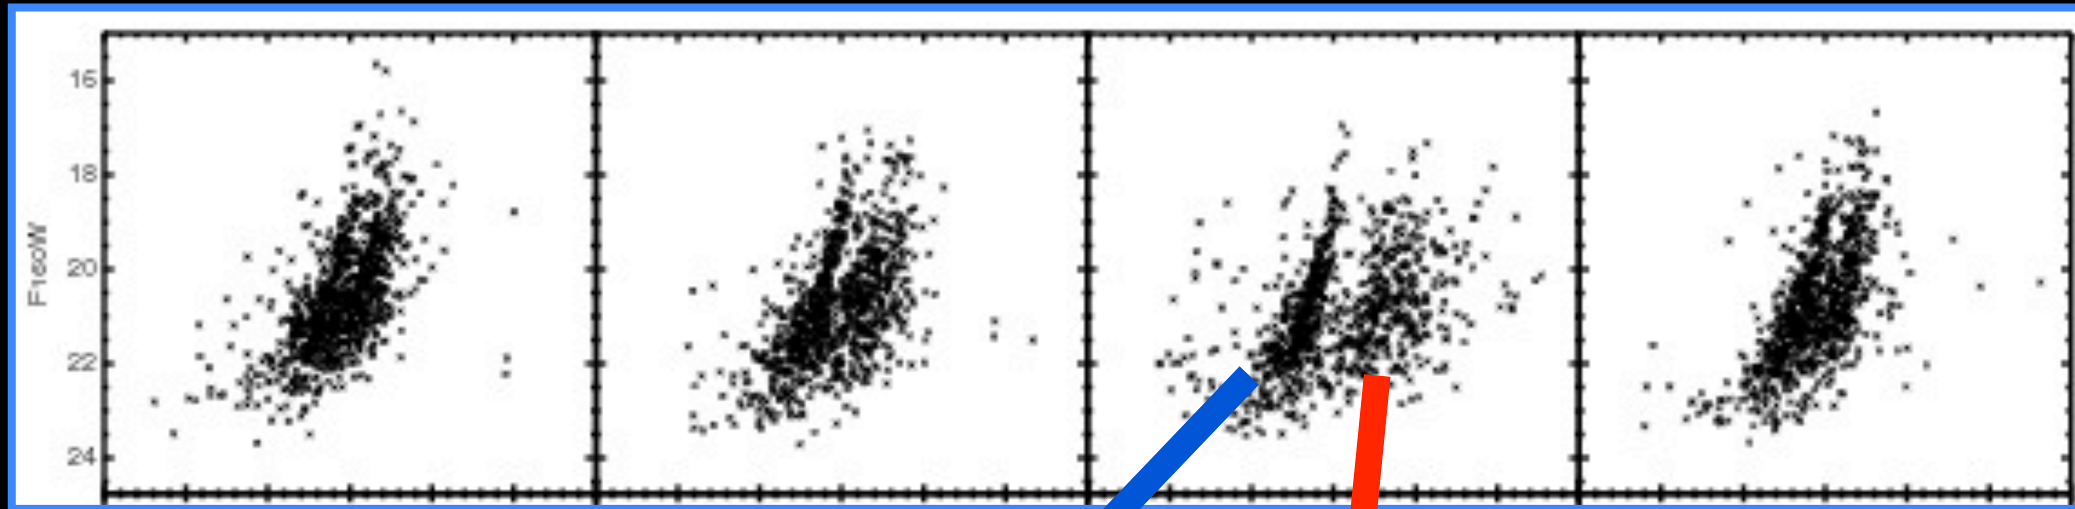

Unreddened peak

Reddened peak

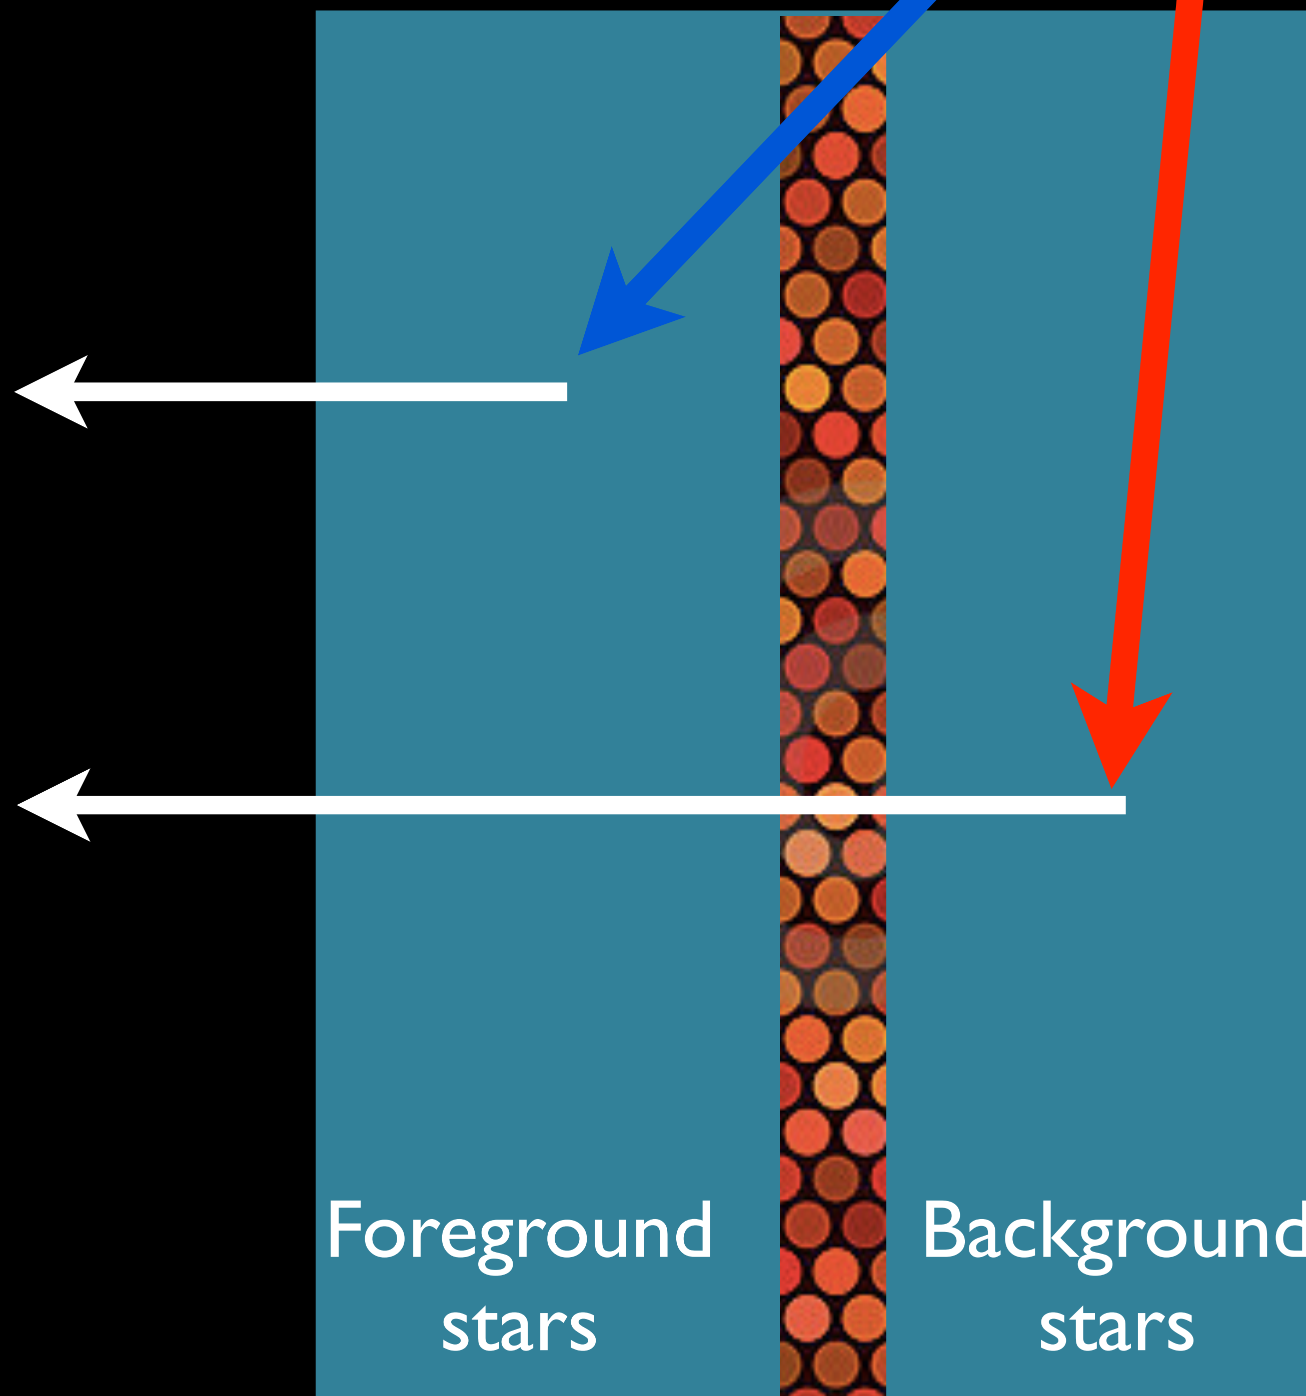

Map of the Dust

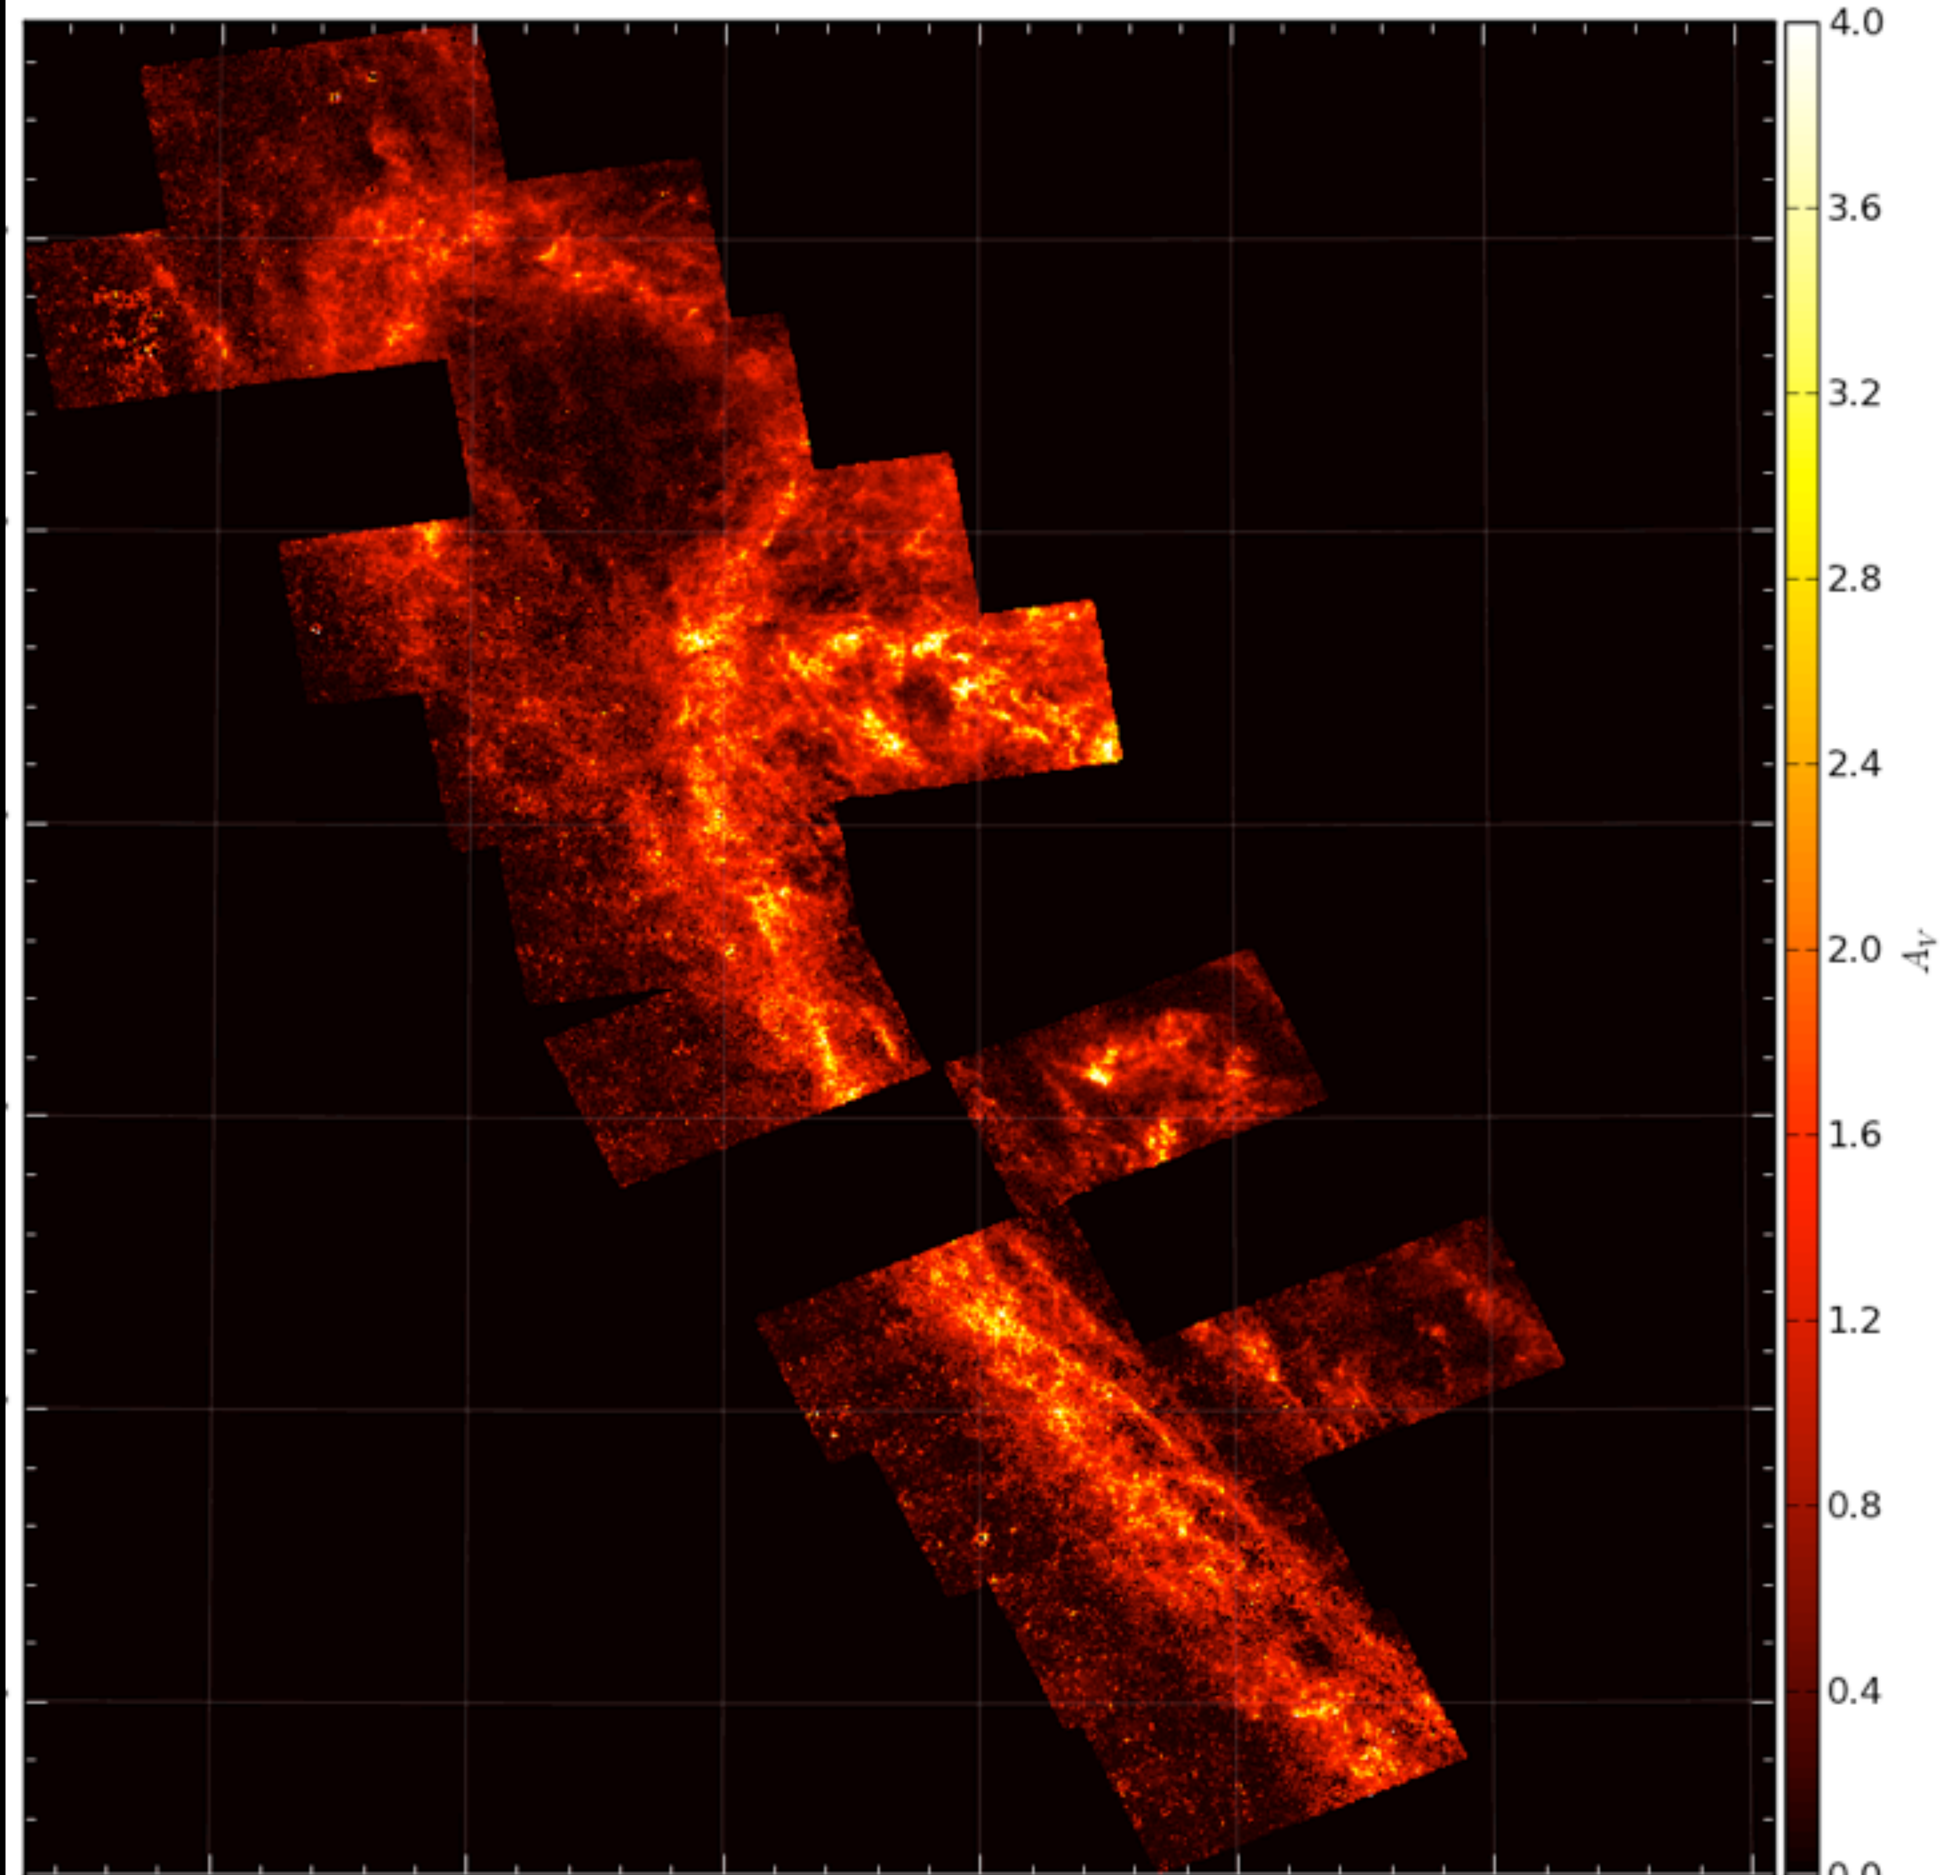

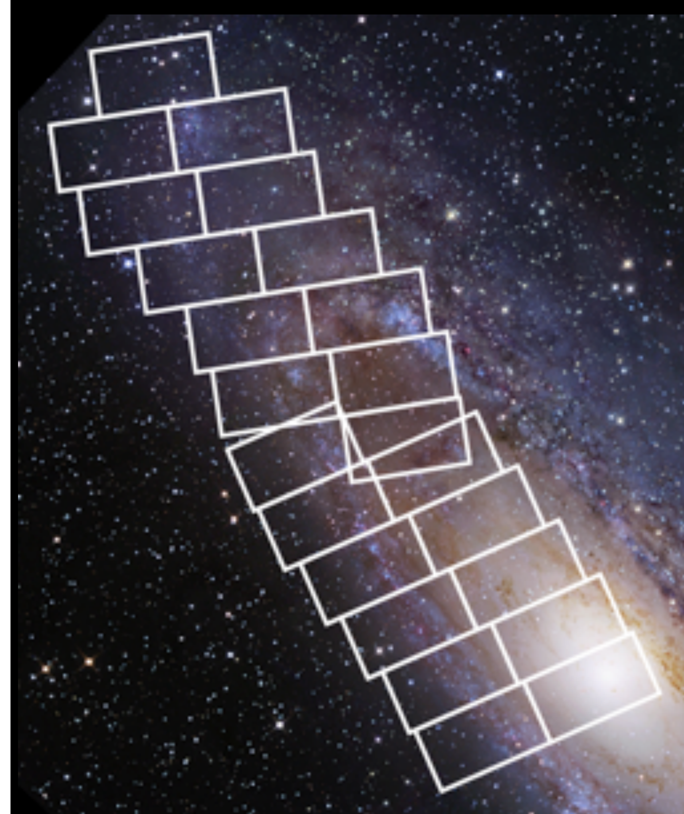

Residual structure: 1.5% flat fielding errors in WFC3/IR!

Excellent Consistency with Dust Emission

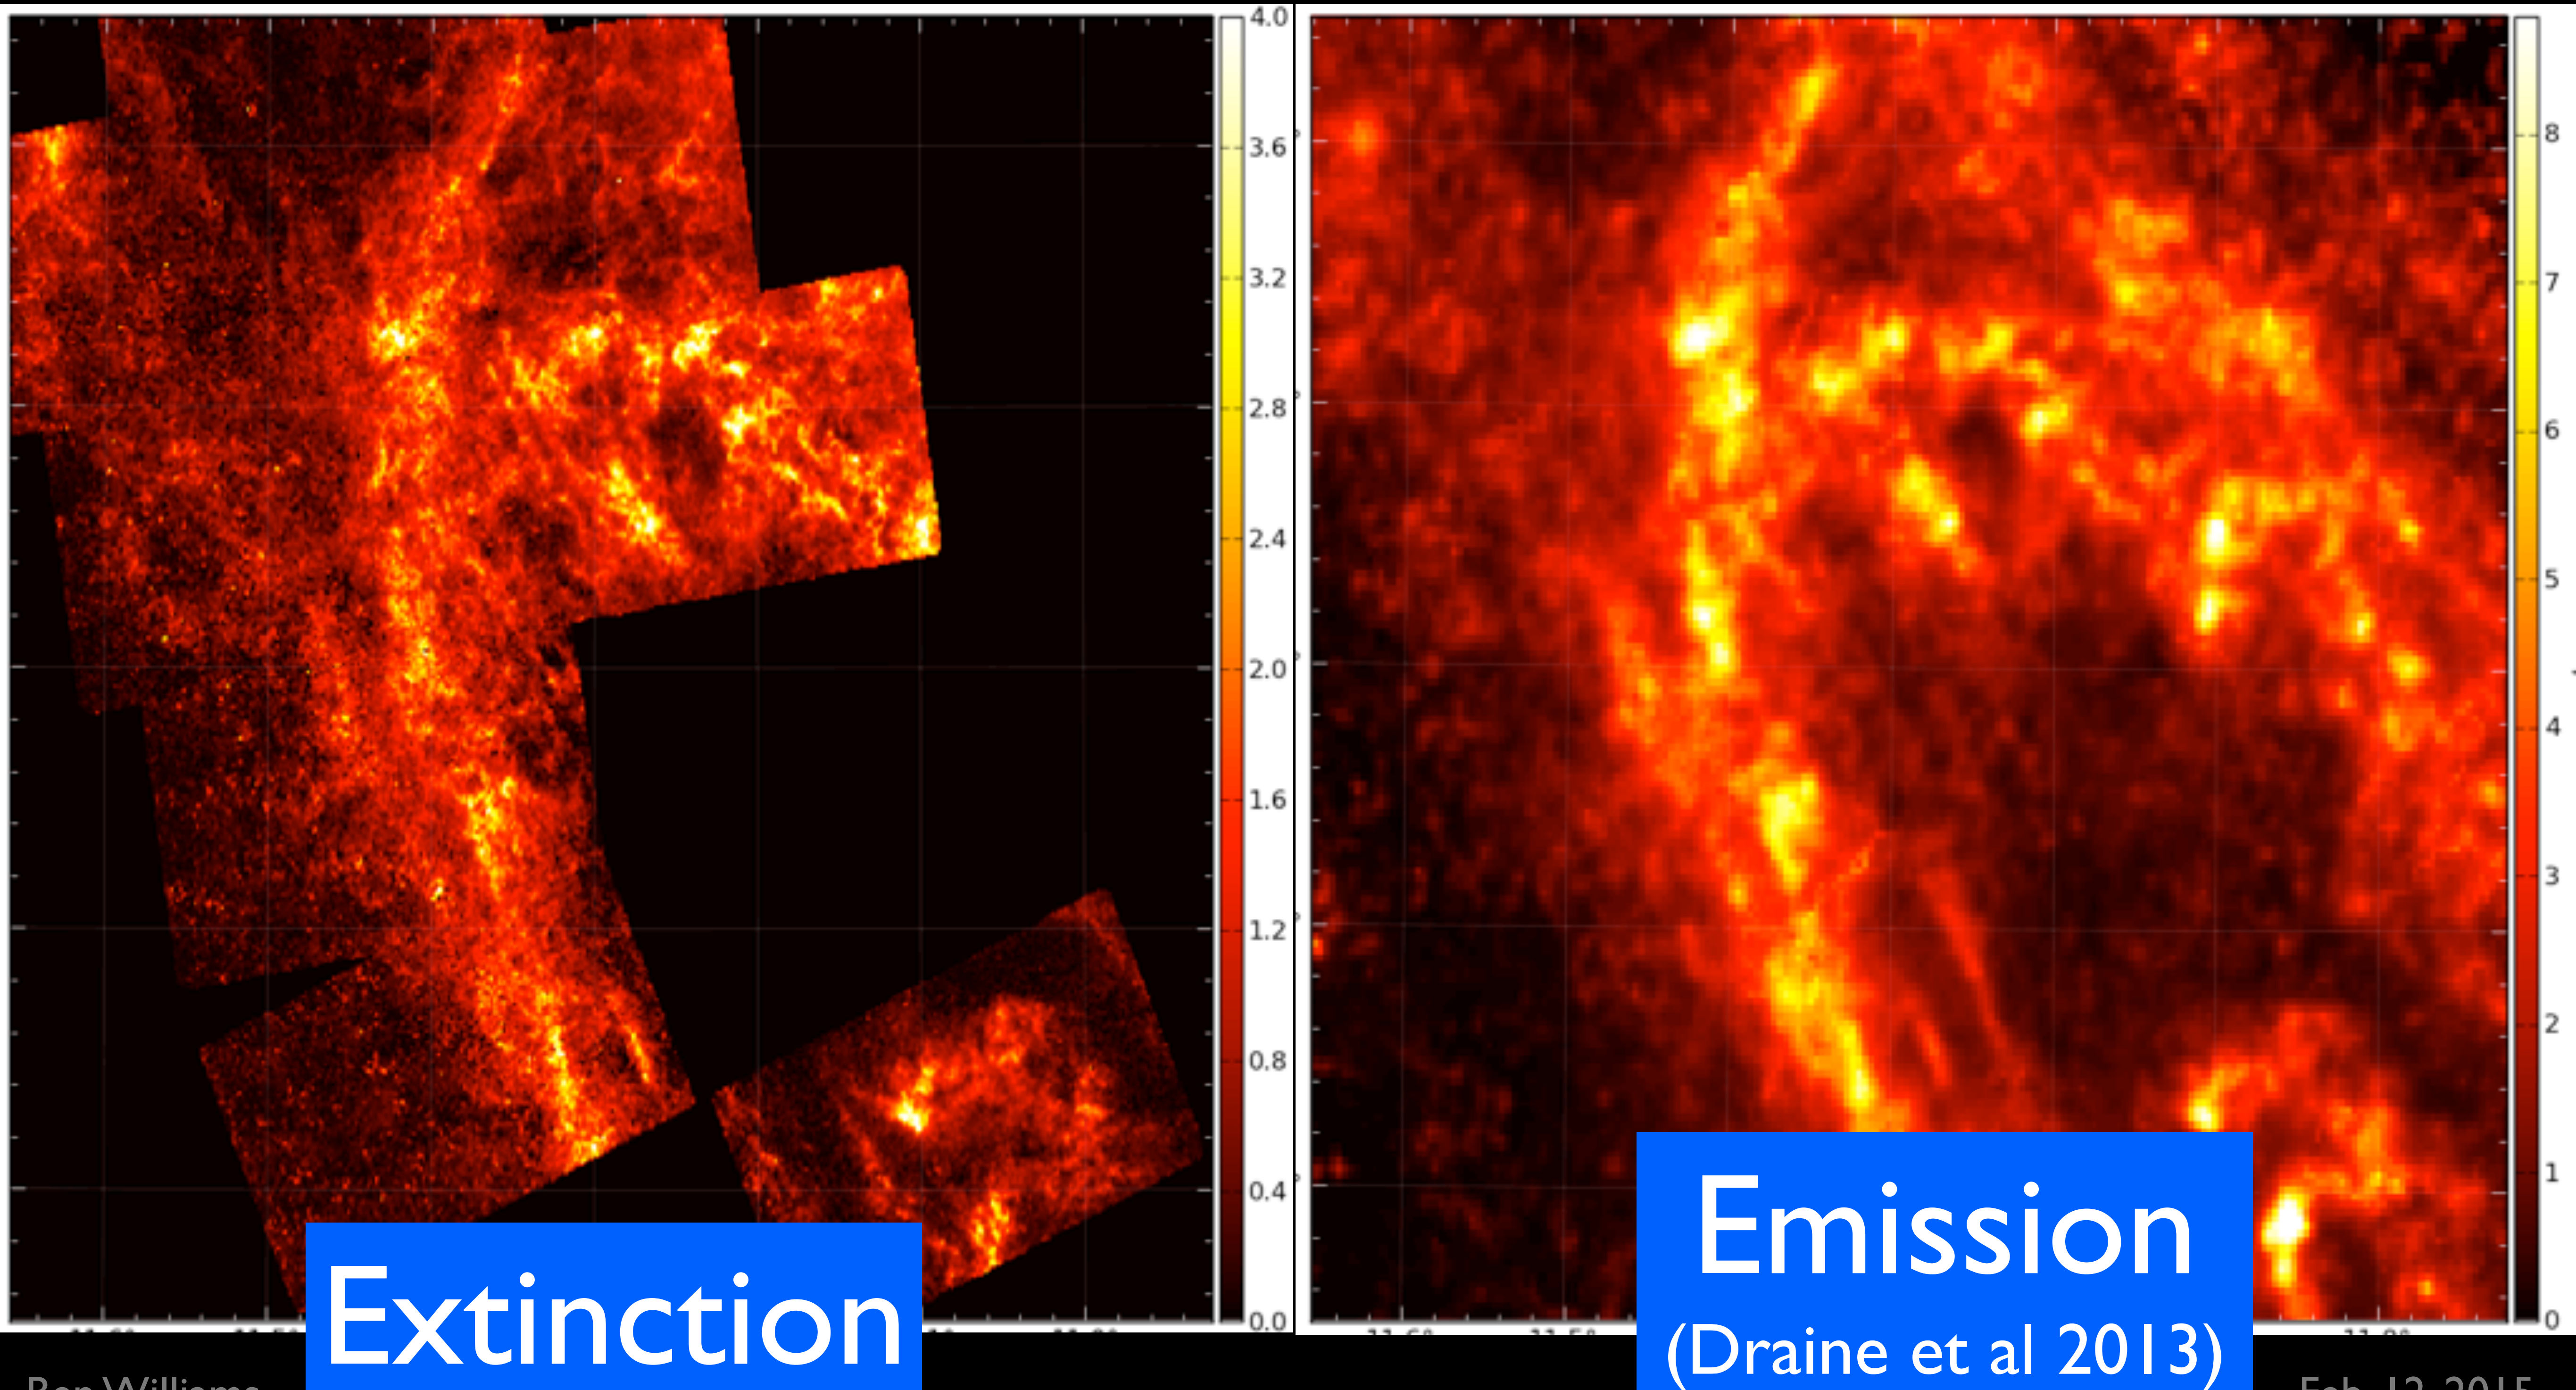

Choose dust-free locations

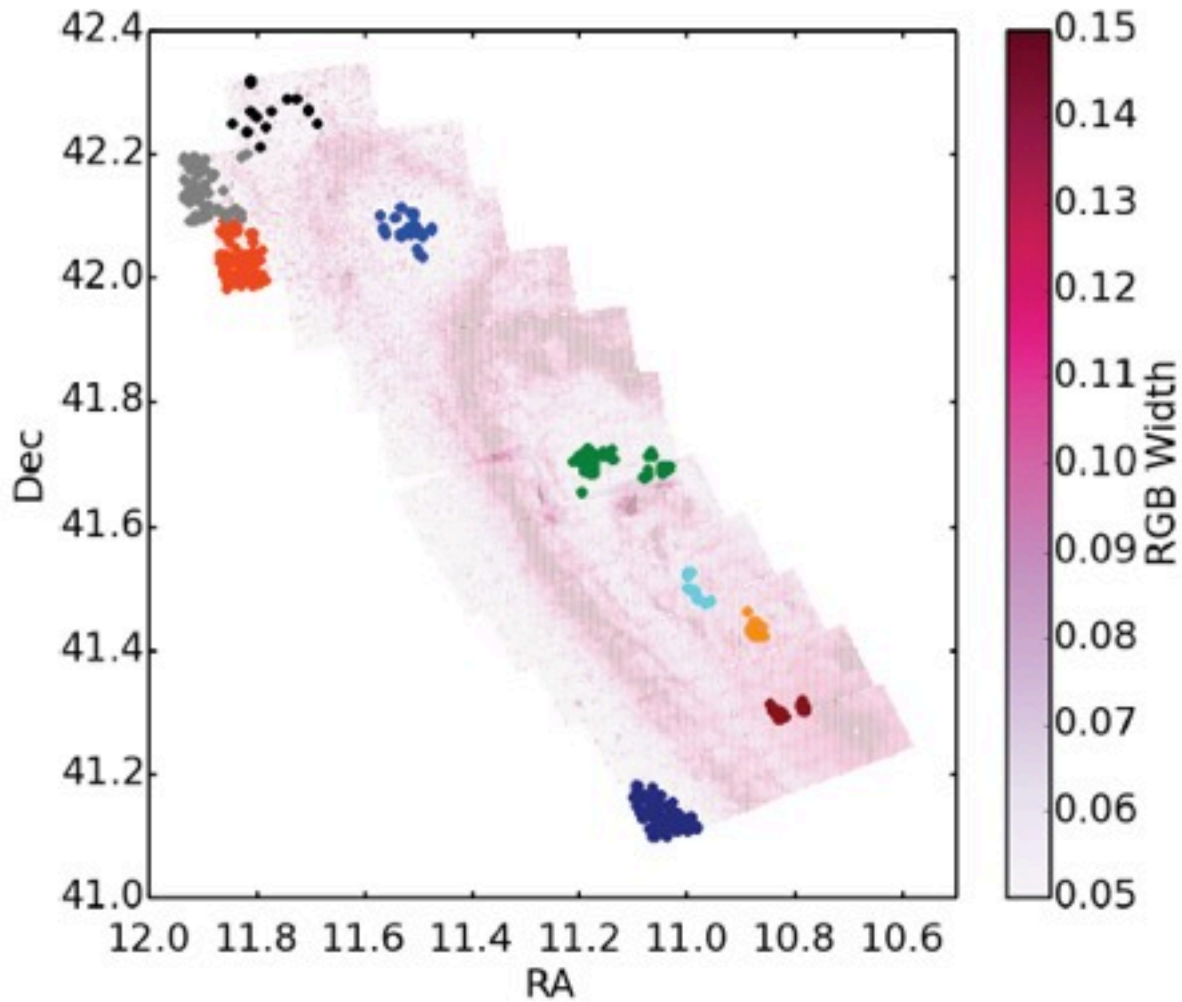

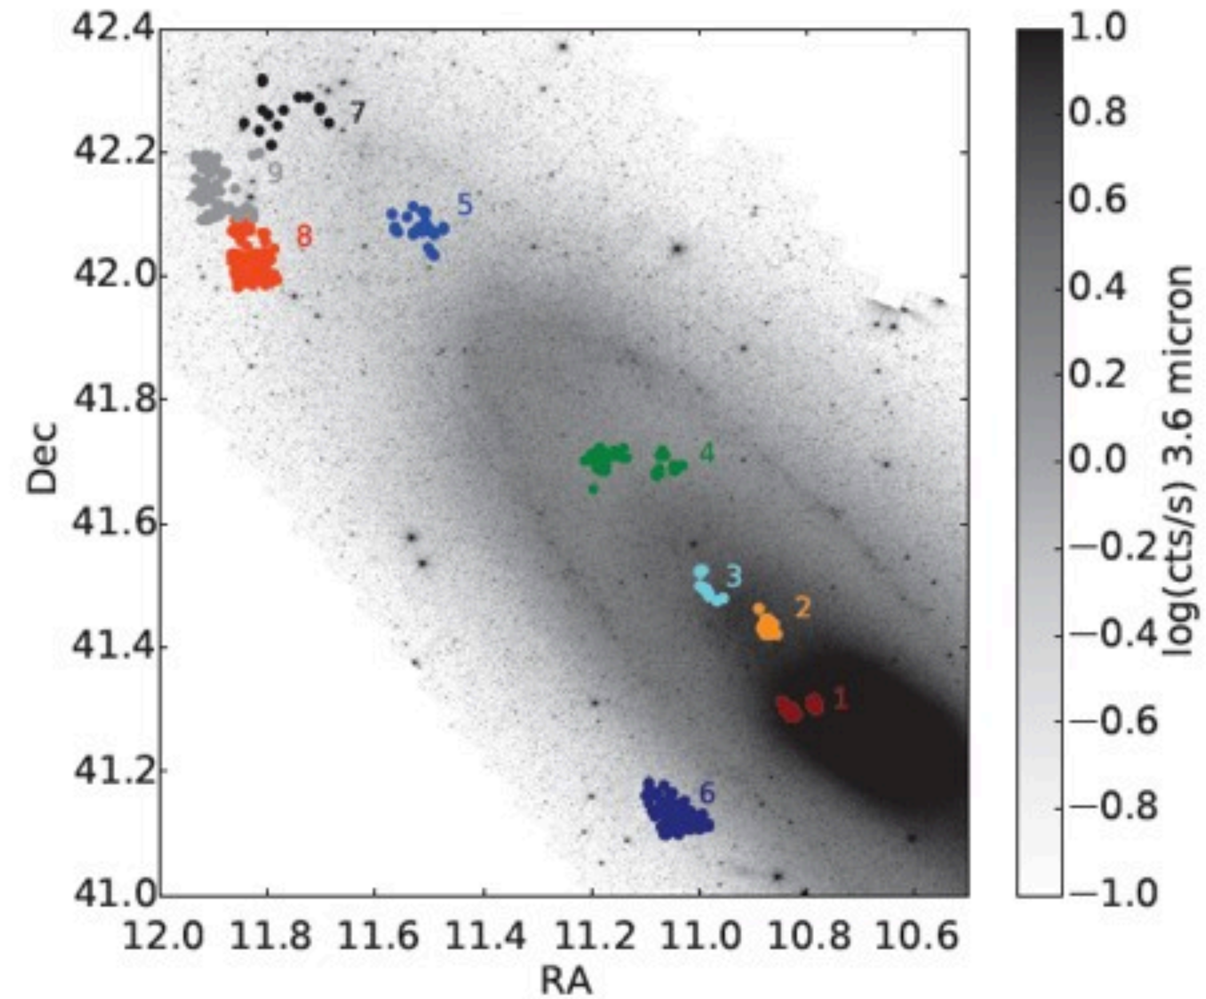

Group into radial bins

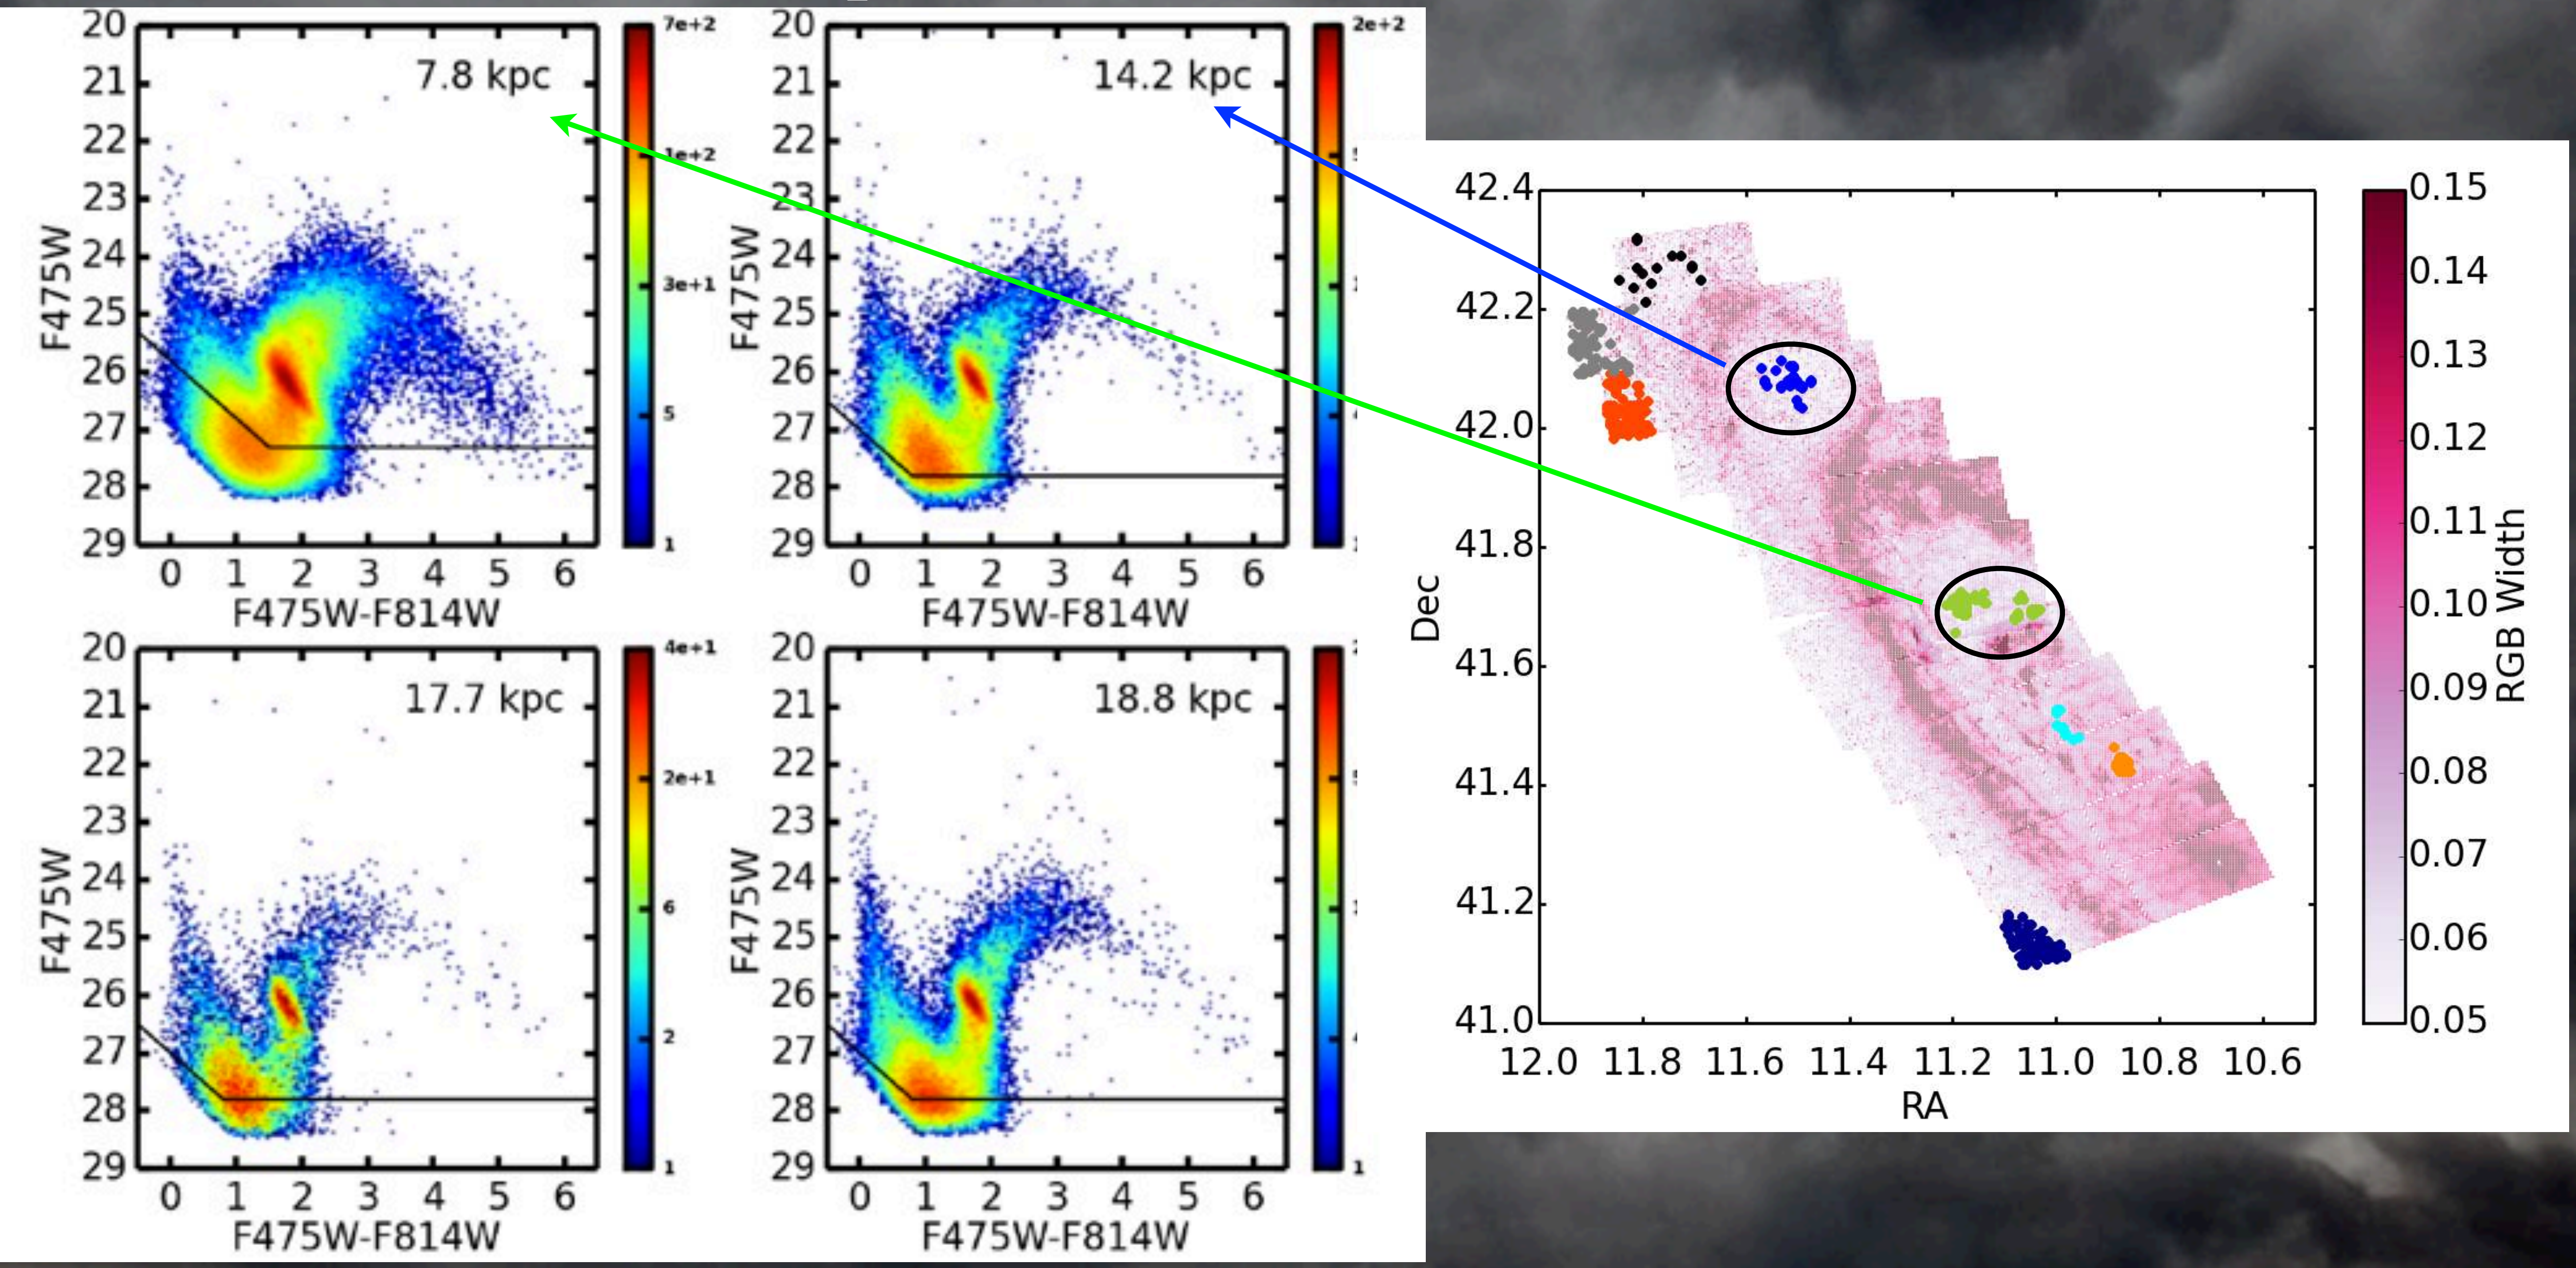

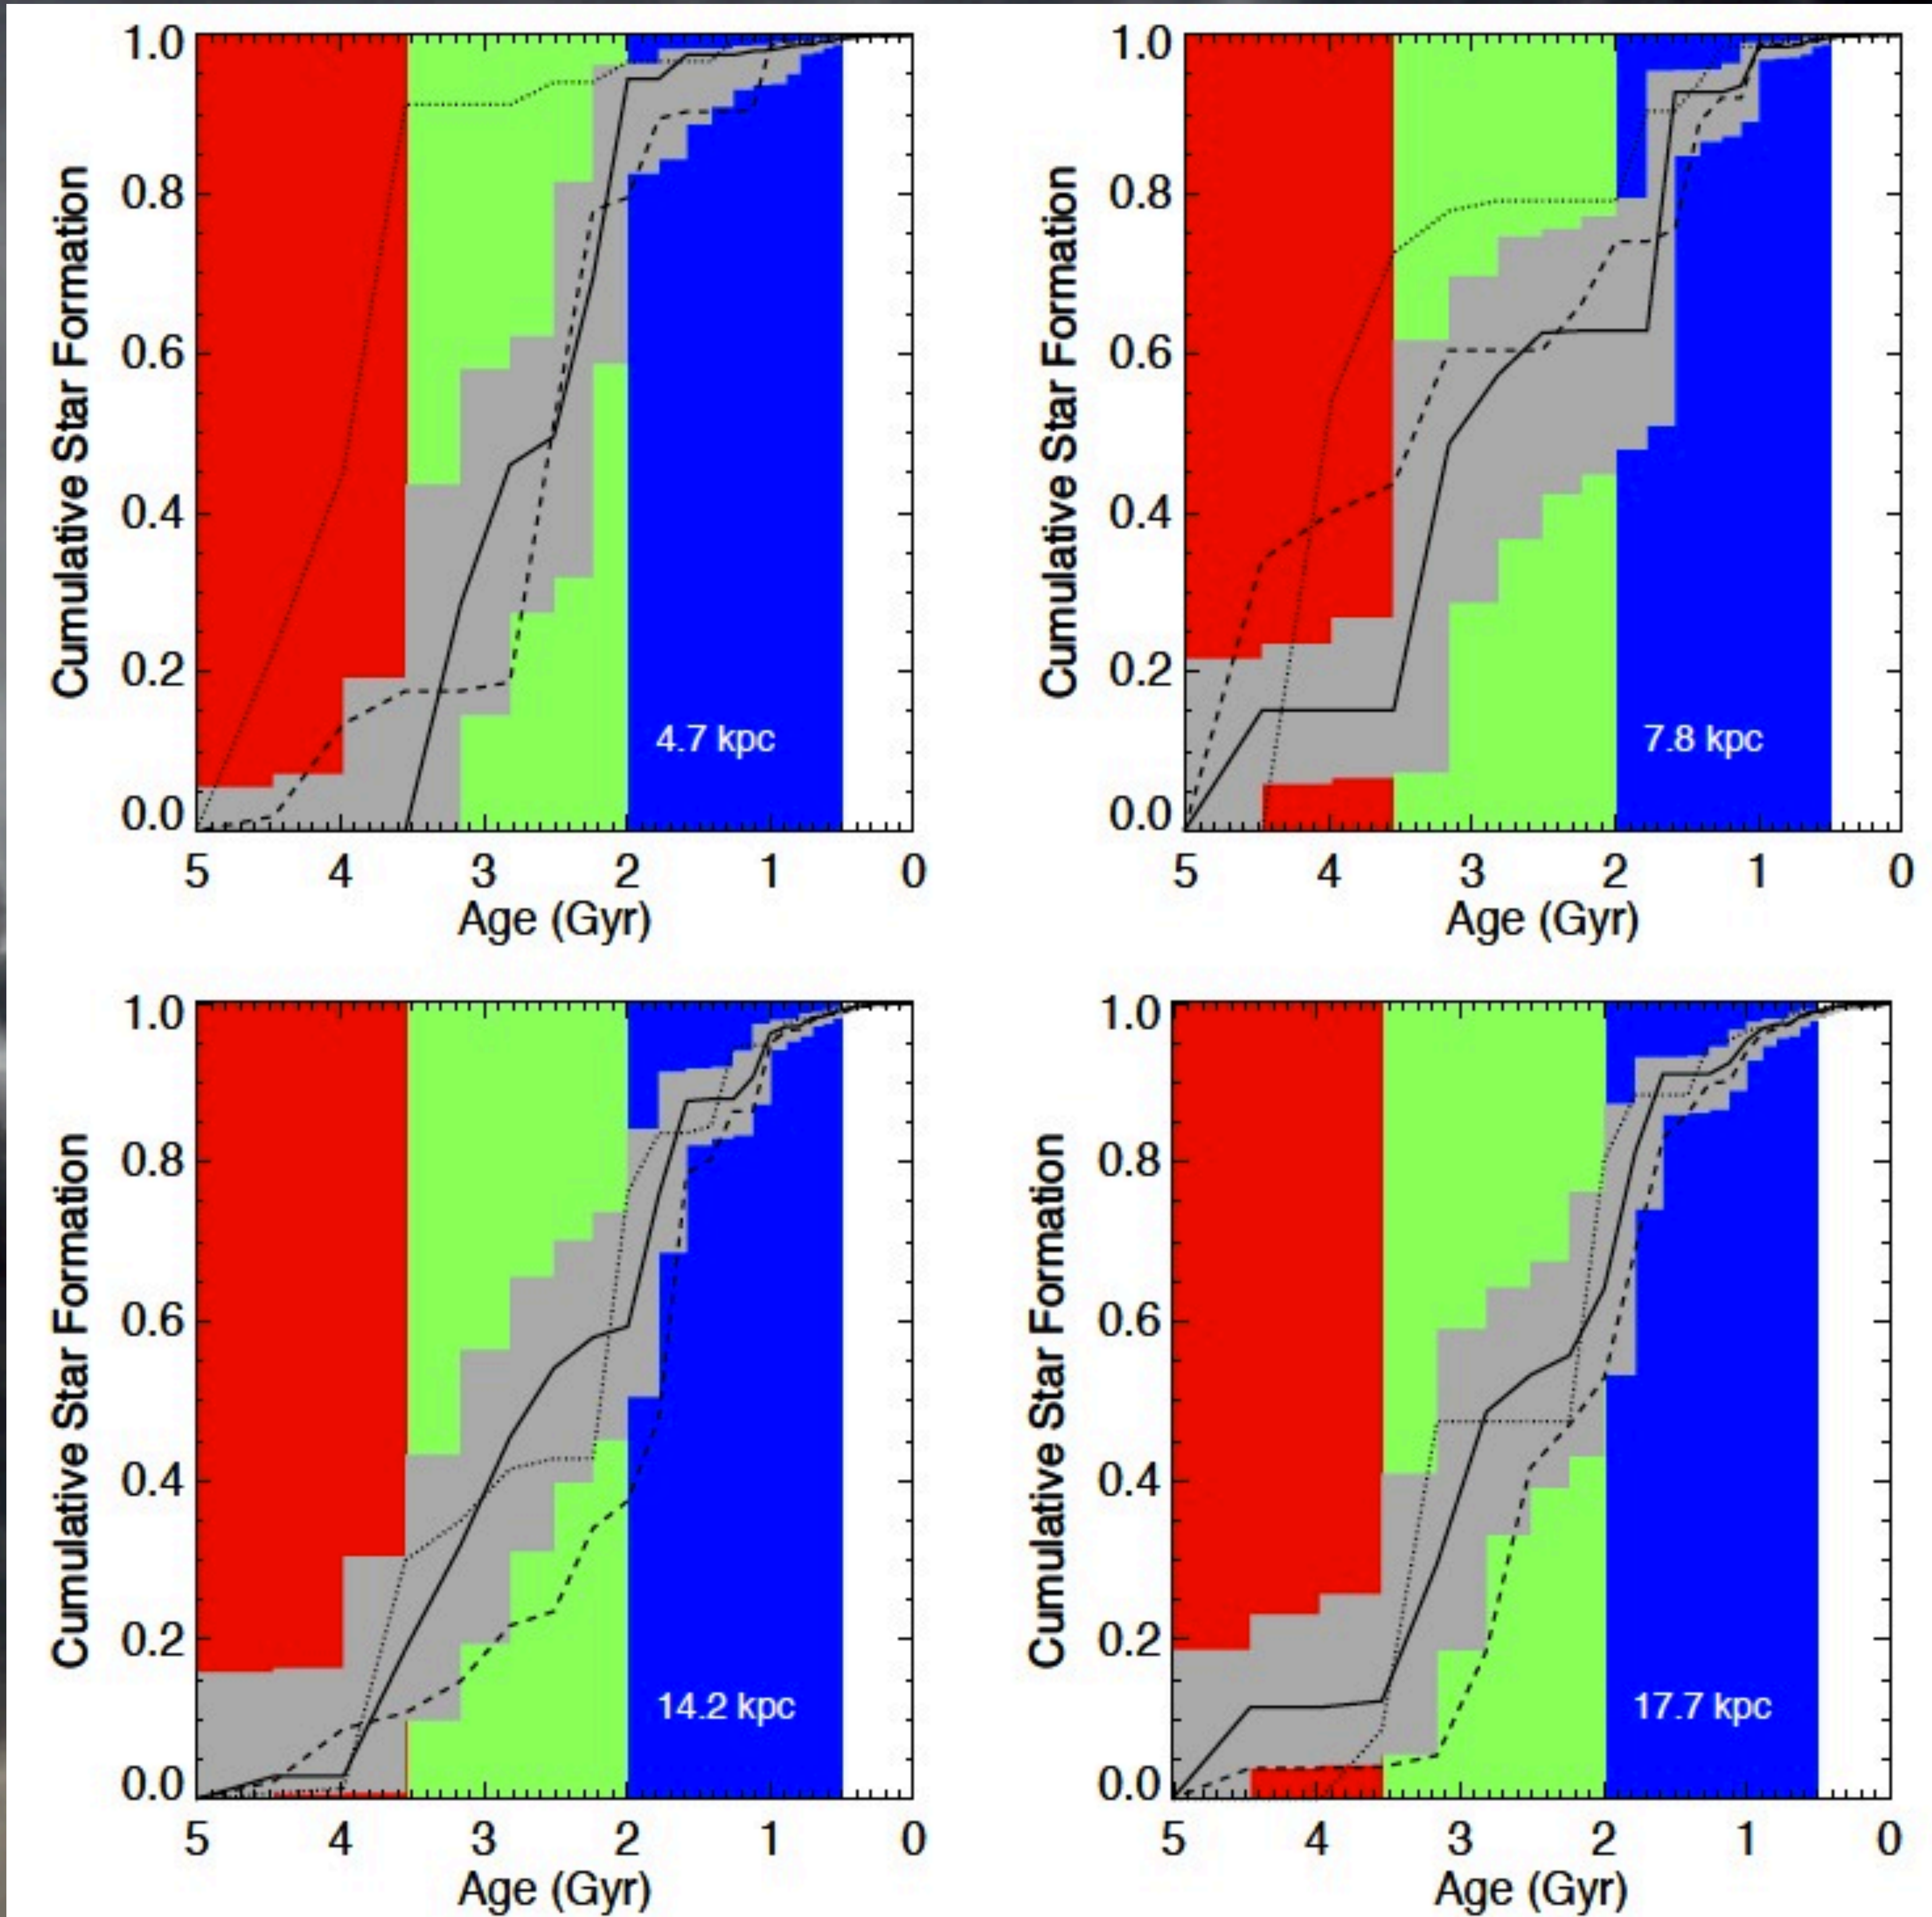

Cumulative Star Formation
Age (Gyr)
4.7 kpc
7.8 kpc
14.2 kpc
17.7 kpc

Bernard et al. 2015: Outer disk

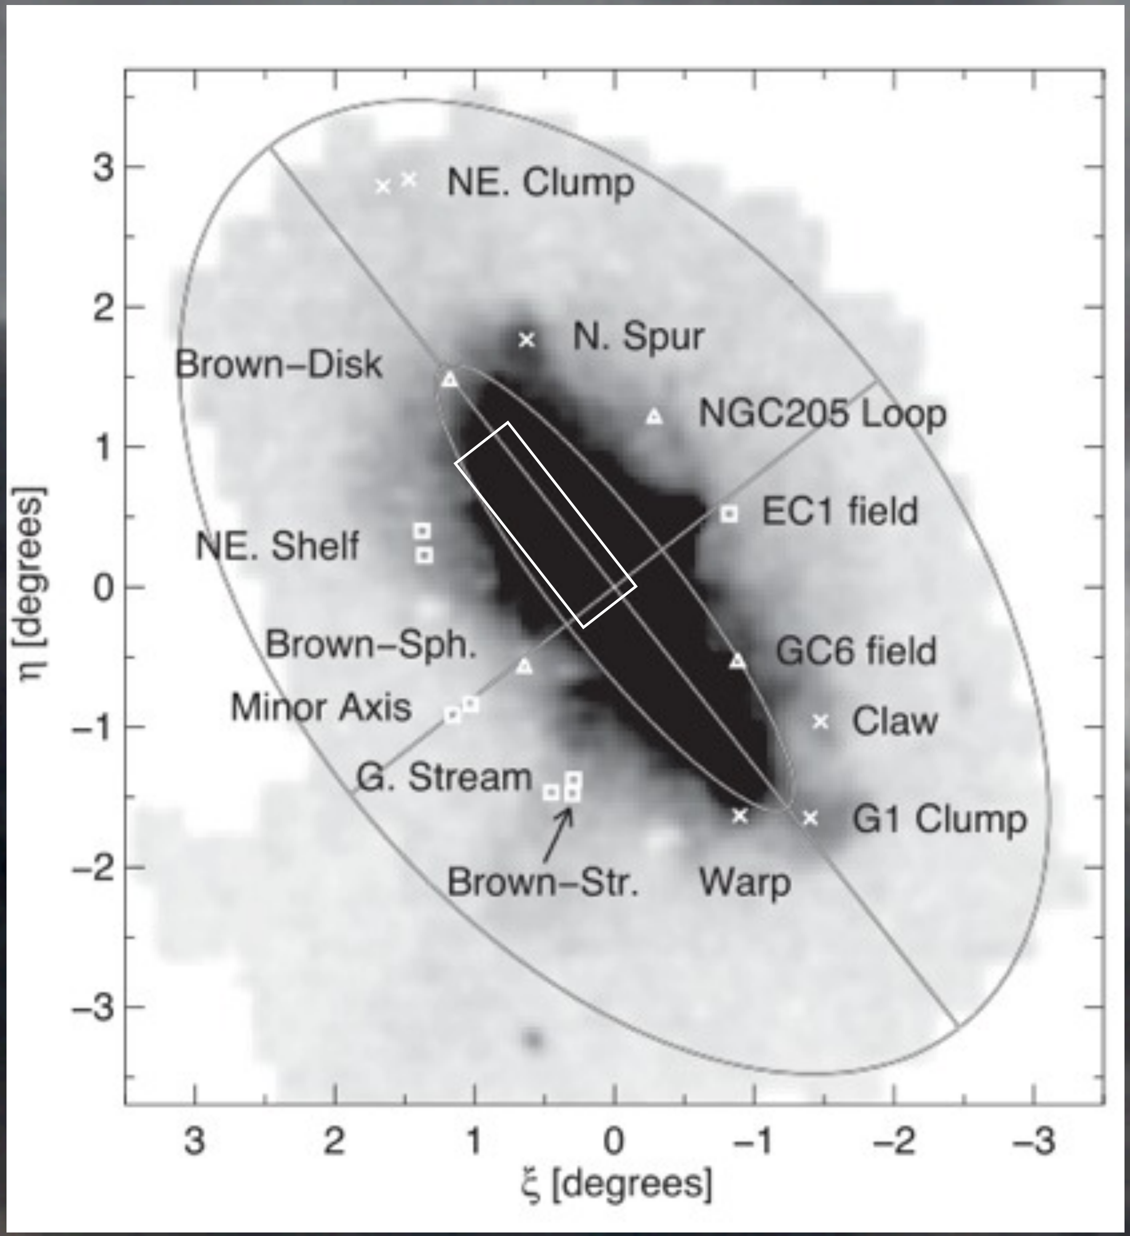

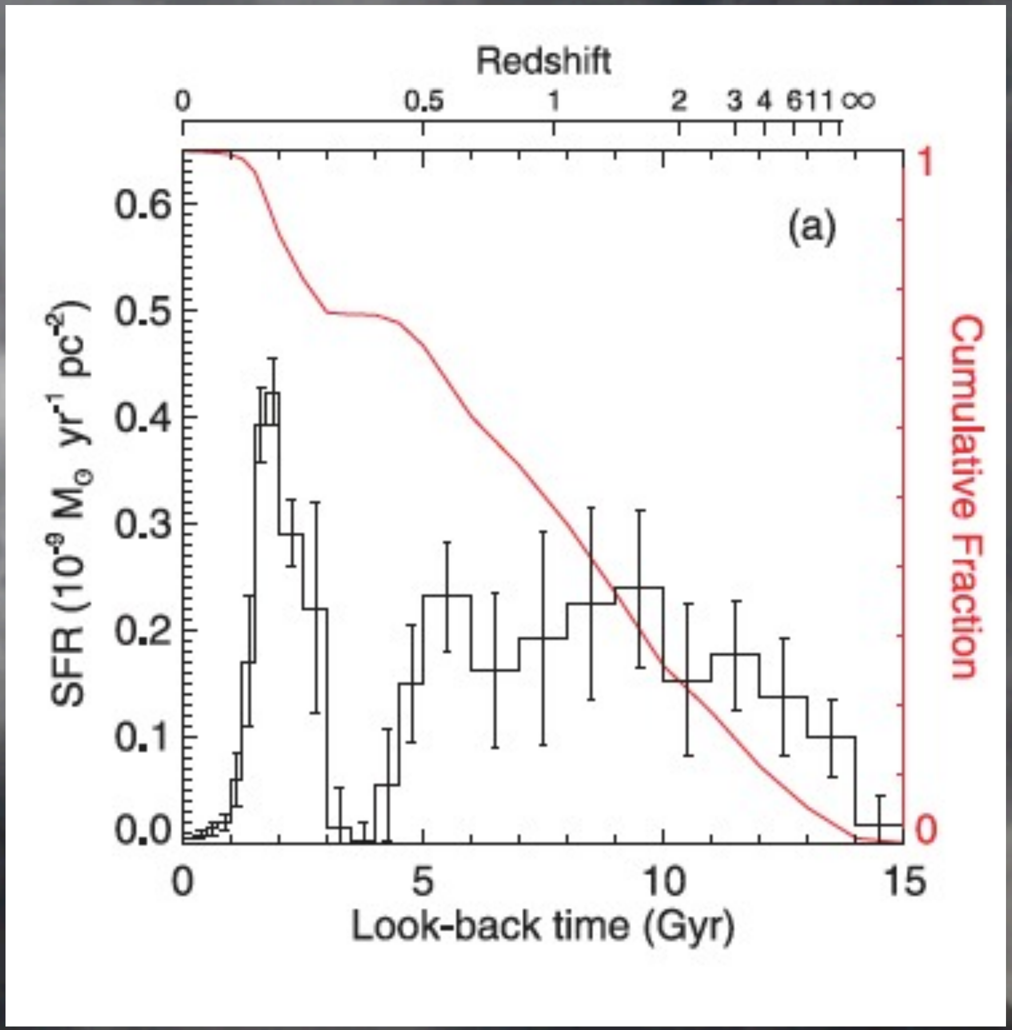

Power of 6 bands

Spectral energy distribution

↓

Temperature, gravity, metallicity, intervening dust properties

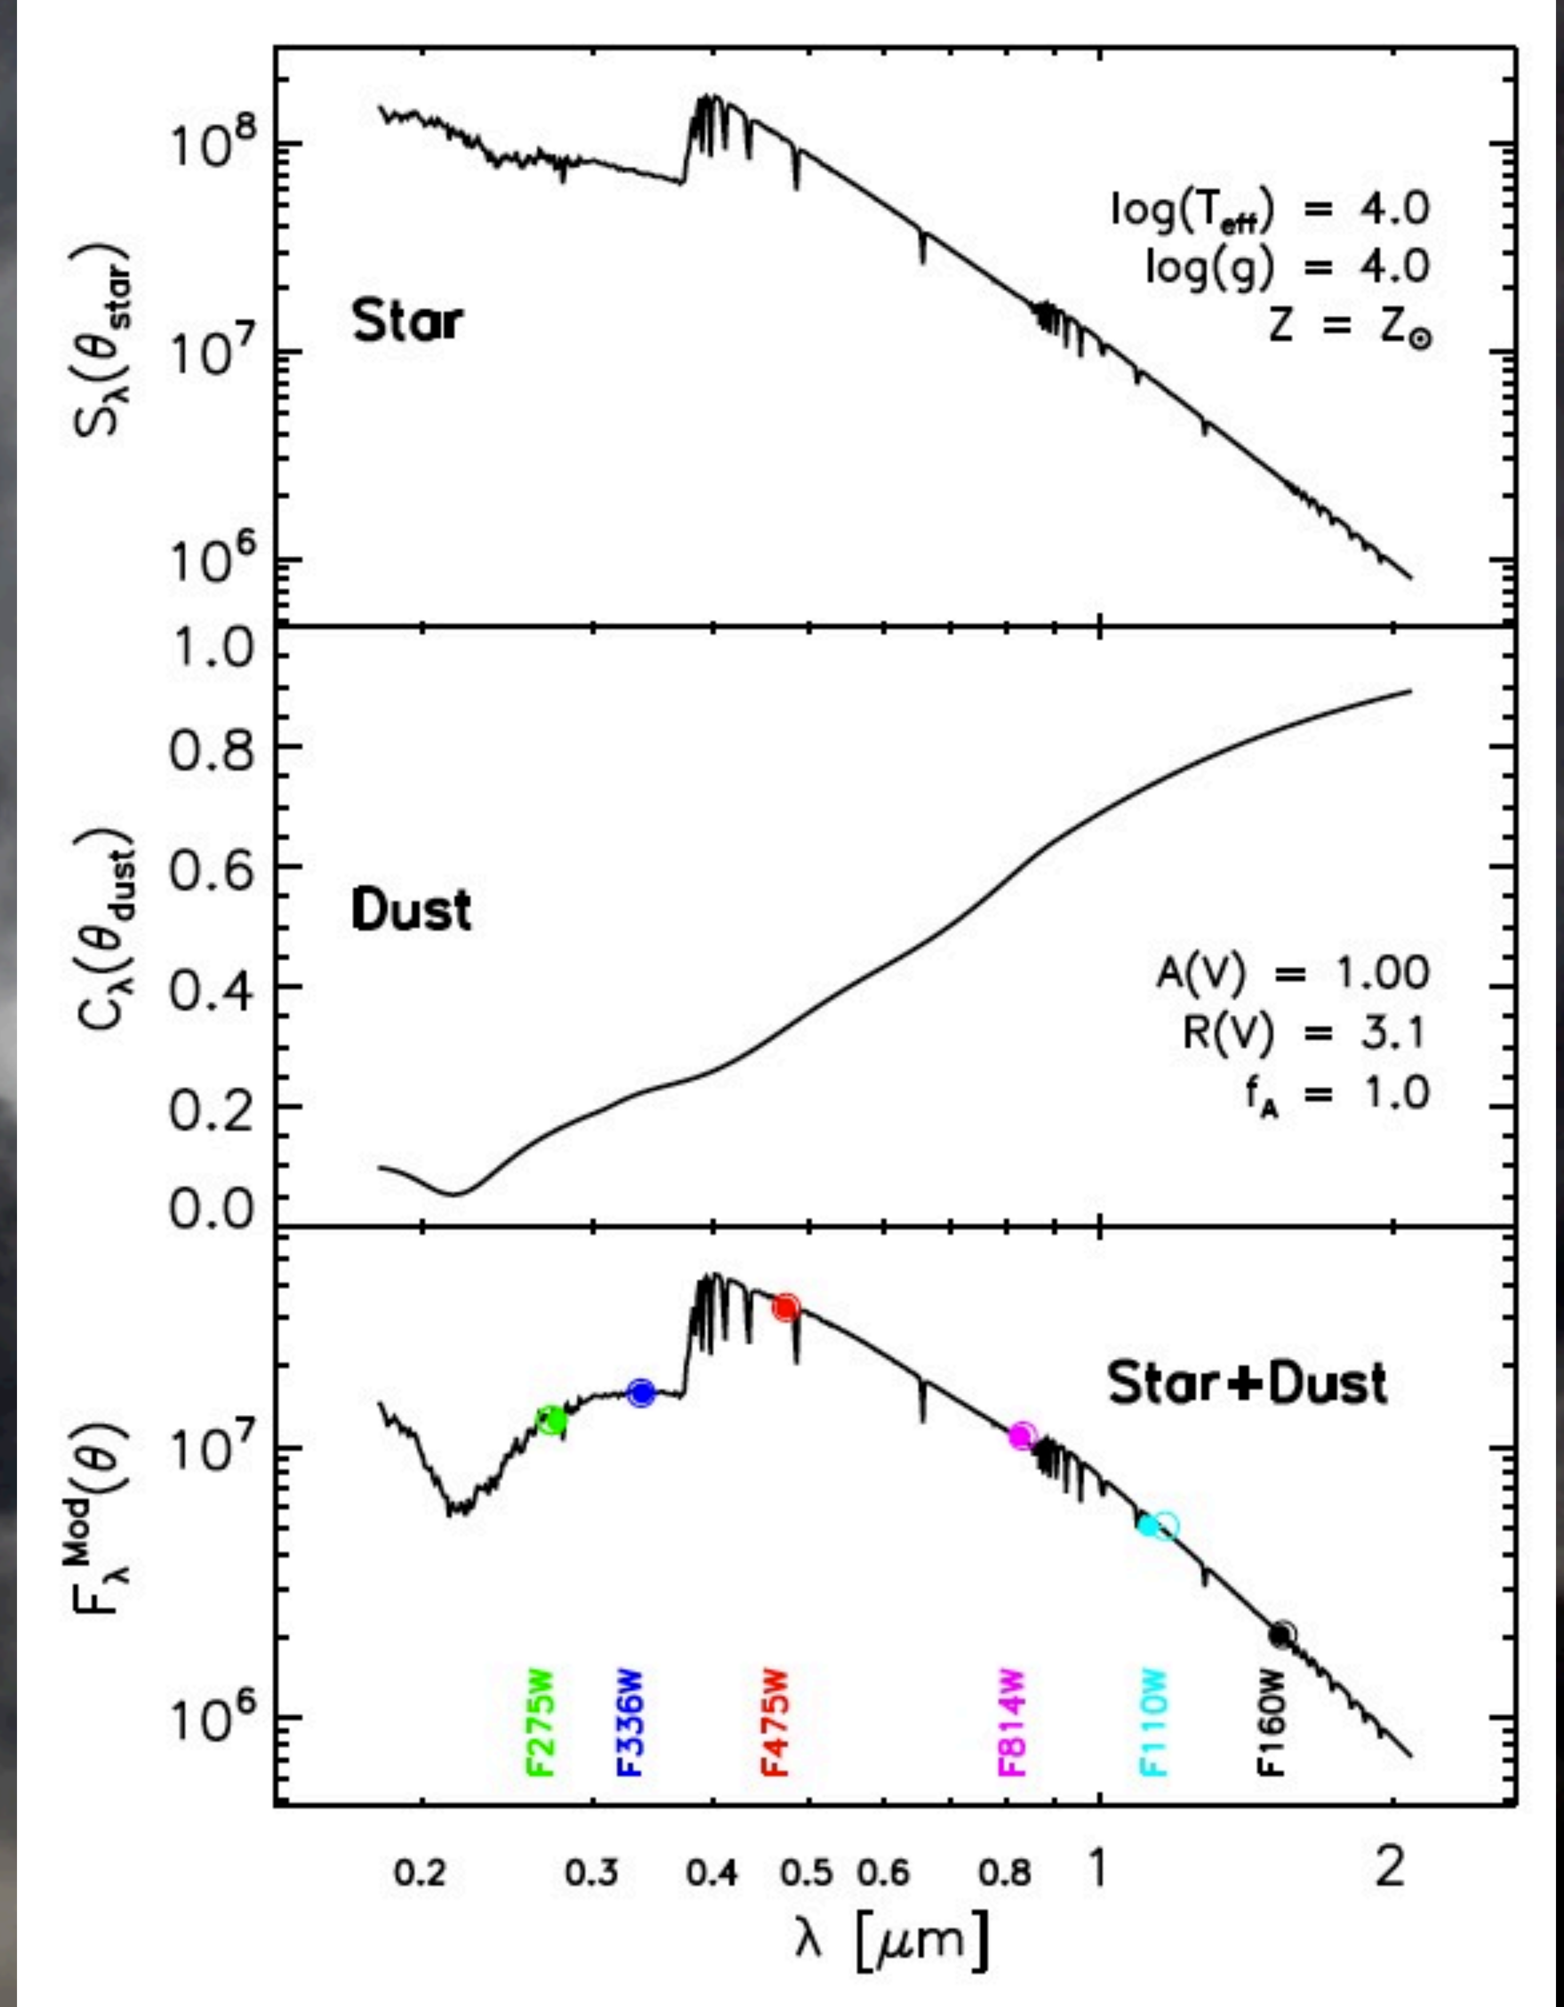

Derive True H-R Diagram

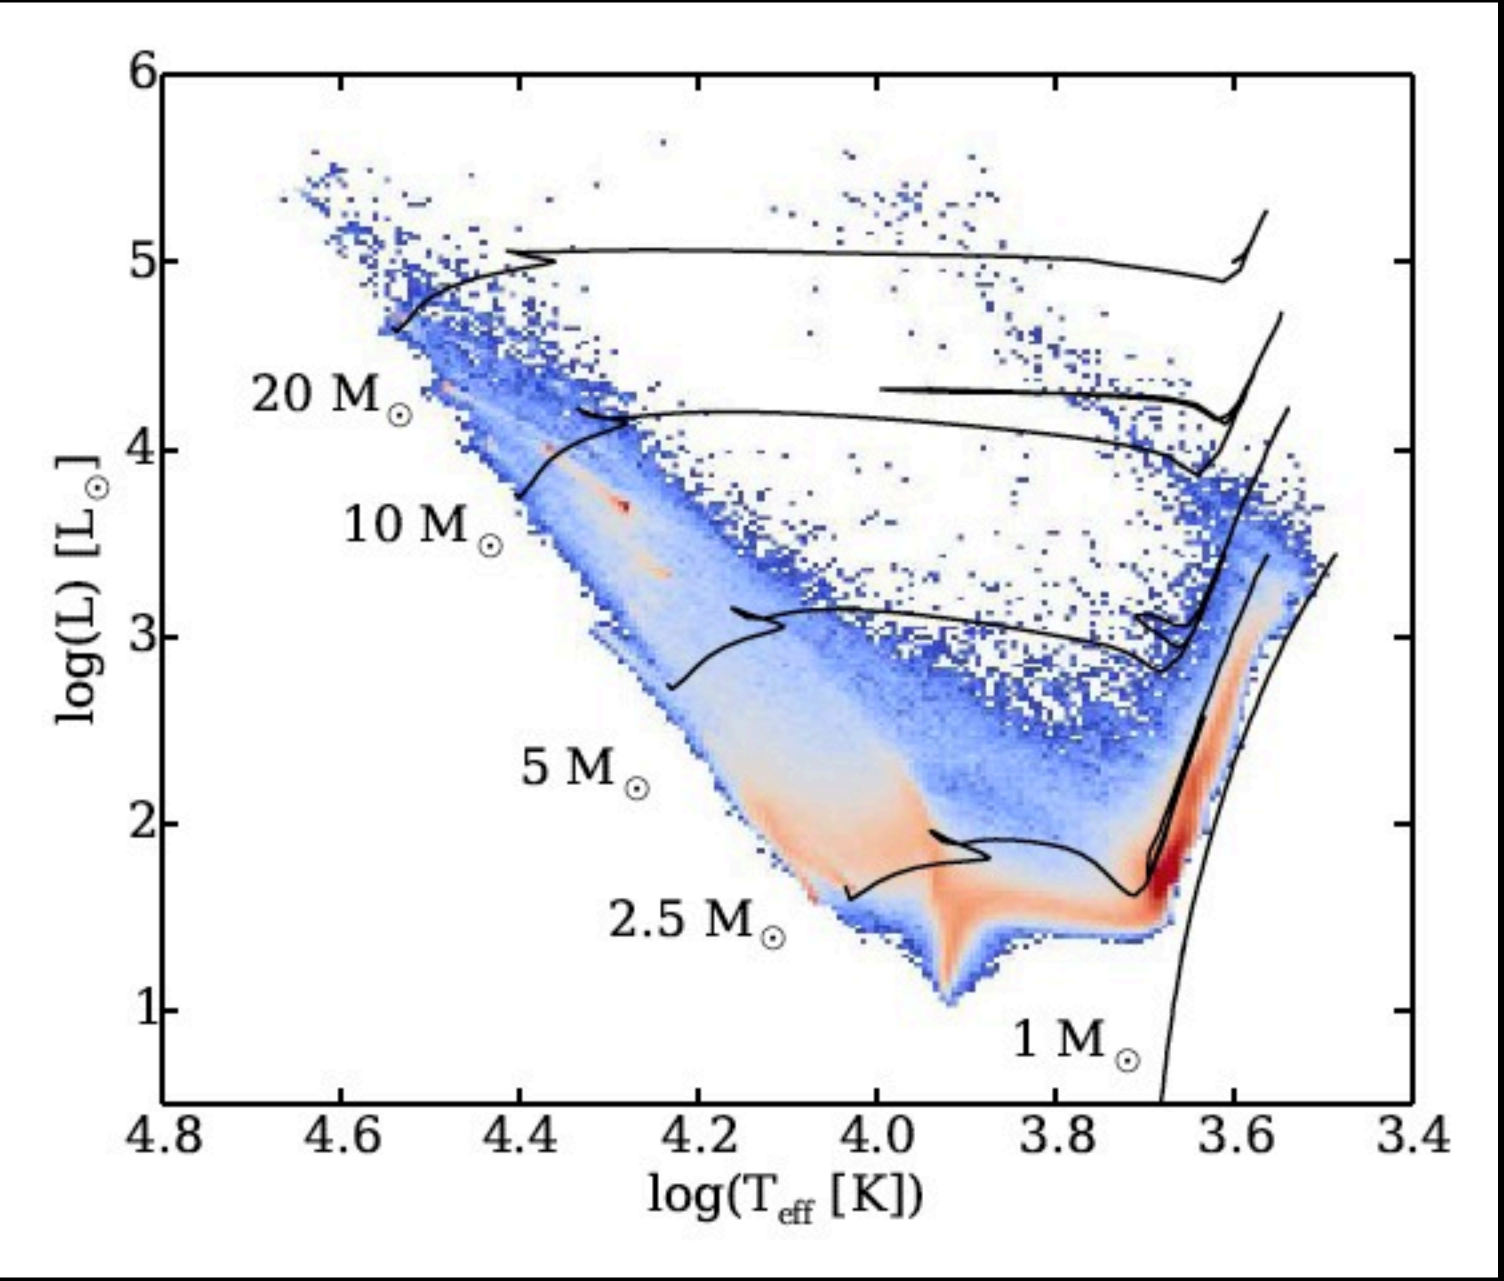

Gordon et al. in prep

Public Release

http://archive.stsci.edu/prepds/phat/

117 million stars

Paper reference:

Williams et al. 2014, ApJS, 215, 9

Collaborators

Julianne J. Dalcanton (UW)
Keith Rosema (Random Walk)
Dan Pirone (Qumulo)
Dustin Lang (CMU)
Andrew Dolphin (Raytheon)
Daniel R. Weisz (UW)
Lori Beerman (UW)
Eric F. Bell (UMich)
Luciana Bianchi (STScI)
Nell Byler (UW)
Morgan Fouesneau (MPIA)
Karoline Gilbert (STScI)
Leo Girardi (INAF)
Karl Gordon (STScI)
Dimitrios Gouliermis (MPIA
Dylan Gregersen (UUtah)
Cliff Johnson (UW)
Jason Kalirai (STScI)
Tod Lauer (NOAO)
Alexia Lewis (UW)
Antonela Monachesi (UMich)
Hans-Walter Rix (MPIA)
Philip Rosenfield (UNIPD)
Anil Seth (UUtah)
Evan Skillman (UMinn)

The PHAT Team

THANKS

NASA: Space Telescope Science Institute, GO 12055

Amazon Web Services: Research Grants