Sequential Monte Carlo for detecting and deblending objects in astronomical images

Tim White and Jeffrey Regier

Department of Statistics, University of Michigan

Michigan Student Symposium for Interdisciplinary Statistical Sciences
March 28th, 2024

Background

Astronomical cataloging for large sky surveys

The Messier 15 globular cluster, imaged by Hubble/SDSS

Astronomical cataloging for large sky surveys

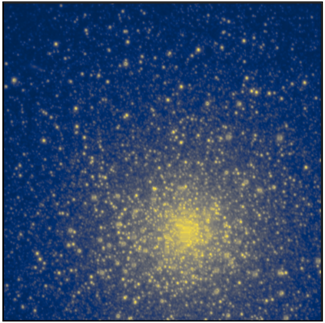

The Messier 15 globular cluster, imaged by Hubble/SDSS

Astronomical cataloging for large sky surveys

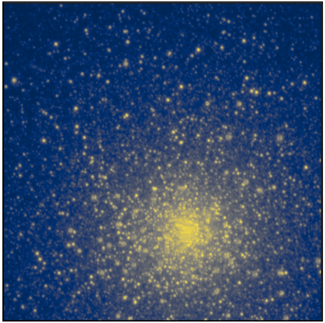

The Messier 15 globular cluster, imaged by Hubble/SDSS

- **Astronomical cataloging** is the task of inferring the properties of stars, galaxies, and other objects in astronomical images

Astronomical cataloging for large sky surveys

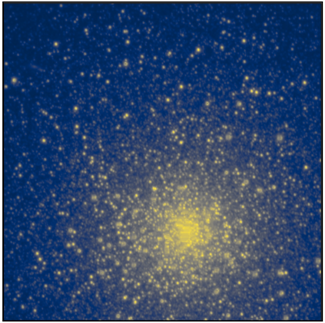

The Messier 15 globular cluster, imaged by Hubble/SDSS

- Learn about the **age** and **large-scale structure** of the universe

Astronomical cataloging for large sky surveys

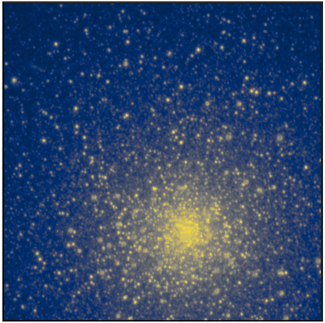

The Messier 15 globular cluster, imaged by Hubble/SDSS

- Learn about the **age** and **large-scale structure** of the universe
- Constrain models of **dark matter** and **dark energy**

Astronomical cataloging for large sky surveys

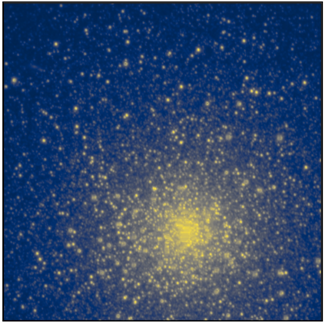

The Messier 15 globular cluster, imaged by Hubble/SDSS

- Learn about the **age** and **large-scale structure** of the universe
- Constrain models of **dark matter** and **dark energy**
- Identify locations and characteristics of special objects (e.g., **quasars**)

Blending in astronomical images

- Objects that overlap visually from the perspective of a telescope are called **blends**

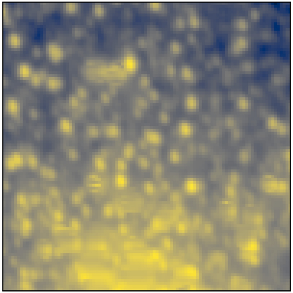

100 x 100 pixel subregion of Messier 15

Blending in astronomical images

- Objects that overlap visually from the perspective of a telescope are called **blends**
- **Deblending** is challenging
  - True source count is unknown
  - Properties are ambiguous

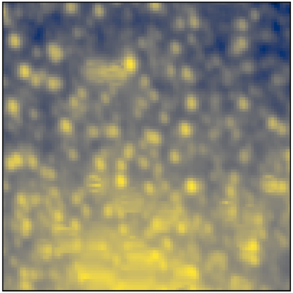

100 x 100 pixel subregion of Messier 15

Blending in astronomical images

- Objects that overlap visually from the perspective of a telescope are called **blends**
- **Deblending** is challenging
  - True source count is unknown
  - Properties are ambiguous
- This is an inverse problem
  - Given: Pixelated image of blended light sources
  - Goal: Infer source count and properties of each source

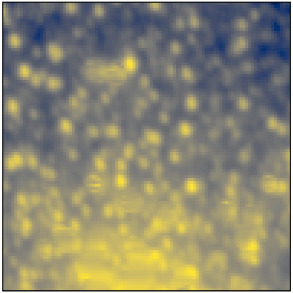

100 x 100 pixel subregion of Messier 15

Blending has become more prevalent

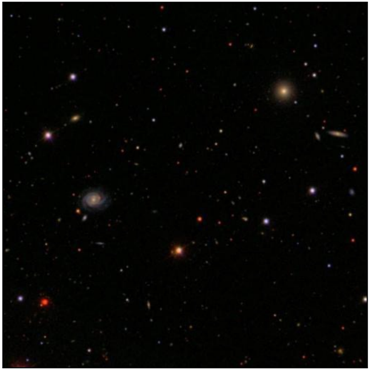

Sloan Digital Sky Survey (SDSS)

First light in 2000, terabytes of data

Blending has become more prevalent

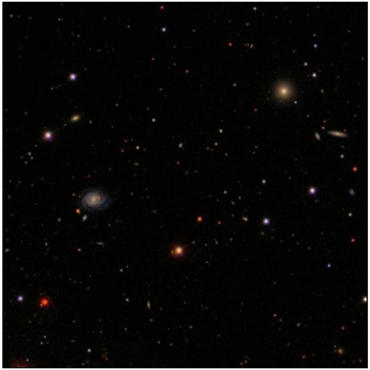

Sloan Digital Sky Survey (SDSS)

First light in 2000, terabytes of data

Legacy Survey of Space and Time (LSST)

First light in 2025, petabytes of data

Existing approaches to cataloging

Non-probabilistic cataloging

- Use peak-finding algorithms to make single-catalog estimates
- Not statistically calibrated, struggles with crowded images

Methods:

→ **Source Extractor**

Existing approaches to cataloging

Non-probabilistic cataloging

- Use peak-finding algorithms to make single-catalog estimates
- Not statistically calibrated, struggles with crowded images

Methods:

→ **Source Extractor**

Probabilistic cataloging

- Infer a posterior distribution over all possible catalogs
- Statistically calibrated, captures ambiguity of blends

Methods:

→ **Markov chain Monte Carlo**
Samples catalogs from the posterior

→ **Variational inference**
Optimizes an approximate posterior

Our contribution

- We propose **SMC-Deblender**, a probabilistic cataloging method based on **sequential Monte Carlo (SMC)** rather than Markov chain Monte Carlo (MCMC) or variational inference (VI)

Our contribution

- We propose **SMC-Deblender**, a probabilistic cataloging method based on **sequential Monte Carlo (SMC)** rather than Markov chain Monte Carlo (MCMC) or variational inference (VI)

- Tasks for the remainder of this talk:

Our contribution

- We propose **SMC-Deblender**, a probabilistic cataloging method based on **sequential Monte Carlo (SMC)** rather than Markov chain Monte Carlo (MCMC) or variational inference (VI)

- Tasks for the remainder of this talk:

  1. **Describe the SMC-Deblender algorithm**

Our contribution

- We propose **SMC-Deblender**, a probabilistic cataloging method based on **sequential Monte Carlo (SMC)** rather than Markov chain Monte Carlo (MCMC) or variational inference (VI)

- Tasks for the remainder of this talk:
  1. **Describe the SMC-Deblender algorithm**
  2. **Assess the algorithm's performance on synthetic images**

Our contribution

- We propose **SMC-Deblender**, a probabilistic cataloging method based on **sequential Monte Carlo (SMC)** rather than Markov chain Monte Carlo (MCMC) or variational inference (VI)

- Tasks for the remainder of this talk:
  1. **Describe the SMC-Deblender algorithm**
  2. **Assess the algorithm's performance on synthetic images**
  3. **Discuss the speed and scalability of the algorithm**

Methods

Statistical model

- **Image** x with a height of H pixels and a width of W pixels

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**
  - **Source count** $s \sim \text{Uniform}\{0, 1, 2, \ldots, D\}$

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**
  - **Source count** $s \sim \text{Uniform}\{0, 1, 2, \ldots, D\}$
  - Given s, **locations** $u_1, \ldots, u_s \overset{\text{iid}}{\sim} \text{Uniform}\big([0, H] \times [0, W]\big)$

    fluxes $f_1, \ldots, f_s \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**
  - **Source count** $s \sim \text{Uniform}\{0, 1, 2, \ldots, D\}$
  - Given s, **locations** $u_1, \ldots, u_s \overset{\text{iid}}{\sim} \text{Uniform}\big([0, H] \times [0, W]\big)$
    fluxes $f_1, \ldots, f_s \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$
  - **Catalog** $z = \{s, \{u_j, f_j\}_{j=1}^{s}\}$

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**
  - **Source count** $s \sim \text{Uniform}\{0, 1, 2, \ldots, D\}$
  - Given s, **locations** $u_1, \ldots, u_s \overset{\text{iid}}{\sim} \text{Uniform}\big([0, H] \times [0, W]\big)$
    fluxes $f_1, \ldots, f_s \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$
  - **Catalog** $z = \{s, \{u_j, f_j\}_{j=1}^{s}\}$
- **Likelihood**

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**
  - **Source count** $s \sim \text{Uniform}\{0, 1, 2, \dots, D\}$
  - Given s, **locations** $u_1, \dots, u_s \overset{\text{iid}}{\sim} \text{Uniform}\big([0, H] \times [0, W]\big)$

    fluxes $f_1, \dots, f_s \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$
  - **Catalog** $z = \{s, \{u_j, f_j\}_{j=1}^{s}\}$
- **Likelihood**
  - **Intensity** at pixel (h, w) is $x_{hw} \mid z \sim \text{Poisson}(\lambda_{hw})$
    - $\lambda_{hw} = \text{background intensity} + \text{sum of fluxes at pixel } (h, w)$

Statistical model

- **Image** x with a height of H pixels and a width of W pixels
- **Prior**
  - **Source count** $s \sim \text{Uniform}\{0, 1, 2, \ldots, D\}$
  - Given s, **locations** $u_1, \ldots, u_s \overset{\text{iid}}{\sim} \text{Uniform}\big([0, H] \times [0, W]\big)$

    fluxes $f_1, \ldots, f_s \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$
  - **Catalog** $z = \{s, \{u_j, f_j\}_{j=1}^{s}\}$
- **Likelihood**
  - **Intensity** at pixel (h, w) is $x_{hw} \mid z \sim \text{Poisson}(\lambda_{hw})$
    - $\lambda_{hw} =$ background intensity $+$ sum of fluxes at pixel (h, w)
- **Posterior** $p(z \mid x) \propto p(z)p(x \mid z)$

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**
  1. Sample catalogs from the prior $p(z)$ and assign them weights

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**
  1. Sample catalogs from the prior $p(z)$ and assign them weights

  While $\tau_t < 1$:

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - → Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**

1. Sample catalogs from the prior $p(z)$ and assign them weights

While $\tau_t < 1$:

2. Increase the temperature τ_t

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - → Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**

1. Sample catalogs from the prior $p(z)$ and assign them weights

While $\tau_t < 1$:

2. Increase the temperature τ_t
3. Modify the catalogs using an MCMC kernel and update their weights

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**

  1. Sample catalogs from the prior $p(z)$ and assign them weights

  While $\tau_t < 1$:

  2. Increase the temperature τ_t
  3. Modify the catalogs using an MCMC kernel and update their weights
  4. Resample the catalogs if necessary to avoid weight degeneracy

Sequential Monte Carlo

- **Goal:** Given an image x, sample catalogs from $p(z \mid x)$
- **Problem:** Cannot sample directly from $p(z \mid x)$
- **Idea:** Define a sequence of distributions $p(z)p(x \mid z)^{\tau_t}$
  - Gradually increase **temperature** τ_t from 0 (prior) to 1 (posterior)
- **Approach:**
  1. Sample catalogs from the prior $p(z)$ and assign them weights

  While $\tau_t < 1$:

  2. Increase the temperature τ_t
  3. Modify the catalogs using an MCMC kernel and update their weights
  4. Resample the catalogs if necessary to avoid weight degeneracy

  Output: Collection of **weighted catalogs** from $p(z \mid x)$

Deblending with sequential Monte Carlo

✷ SMC-Deblender

1. Sample catalogs from the prior $p(z)$ and assign them weights

While $\tau_t < 1$:

2. Increase the temperature τ_t
3. Modify the catalogs using an MCMC kernel and update their weights
4. Resample the catalogs if necessary to avoid weight degeneracy

Output: Collection of **weighted catalogs** from $p(z \mid x)$

Deblending with sequential Monte Carlo

✷ SMC-Deblender

1. Sample catalogs from the prior $p(z)$ and assign them weights
   → **Fixed number of catalogs for each source count**

While $\tau_t < 1$:

2. Increase the temperature τ_t
3. Modify the catalogs using an MCMC kernel and update their weights
4. Resample the catalogs if necessary to avoid weight degeneracy

Output: Collection of **weighted catalogs** from $p(z \mid x)$

Deblending with sequential Monte Carlo

SMC-Deblender

1. Sample catalogs from the prior $p(z)$ and assign them weights
   - → **Fixed number of catalogs for each source count**

While $\tau_t < 1$:

2. Increase the temperature τ_t
3. Modify the catalogs using an MCMC kernel and update their weights
   - → **Modify locations and fluxes but not source count**
4. Resample the catalogs if necessary to avoid weight degeneracy

Output: Collection of **weighted catalogs** from $p(z \mid x)$

Deblending with sequential Monte Carlo

✷ SMC-Deblender

1. Sample catalogs from the prior $p(z)$ and assign them weights
   - → **Fixed number of catalogs for each source count**

While $\tau_t < 1$:

2. Increase the temperature τ_t
3. Modify the catalogs using an MCMC kernel and update their weights
   - → **Modify locations and fluxes but not source count**
4. Resample the catalogs if necessary to avoid weight degeneracy
   - → **Resample catalogs within groups stratified by source count**

Output: Collection of **weighted catalogs** from $p(z \mid x)$

Advantages of SMC-Deblender

1. Does not require transdimensional sampling (unlike MCMC)

Advantages of SMC-Deblender

1. Does not require transdimensional sampling (unlike MCMC)
2. Runs in parallel over catalogs (unlike MCMC)

Advantages of SMC-Deblender

1. Does not require transdimensional sampling (unlike MCMC)
2. Runs in parallel over catalogs (unlike MCMC)
3. Does not require non-convex optimization (unlike VI)

Advantages of SMC-Deblender

1. Does not require transdimensional sampling (unlike MCMC)
2. Runs in parallel over catalogs (unlike MCMC)
3. Does not require non-convex optimization (unlike VI)
4. Directly targets $p(z \mid x)$ instead of an approximation (unlike VI)

Experiments

Experiment settings and overall performance

- 1,000 synthetic images (15 pixels $\times$ 15 pixels)
  - Up to 10 light sources in each image
- Compare **SMC-Deblender** and **Source Extractor** in terms of accuracy and calibration of estimated source counts

Experiment settings and overall performance

- 1,000 synthetic images (15 pixels $\times$ 15 pixels)
  - Up to 10 light sources in each image
- Compare **SMC-Deblender** and **Source Extractor** in terms of accuracy and calibration of estimated source counts

	% correct	MAE	Time per image
SMC-Deblender	79.8%	0.229	60 seconds
Source Extractor	21.0%	3.576	$<$1 second

Classification accuracy by source count

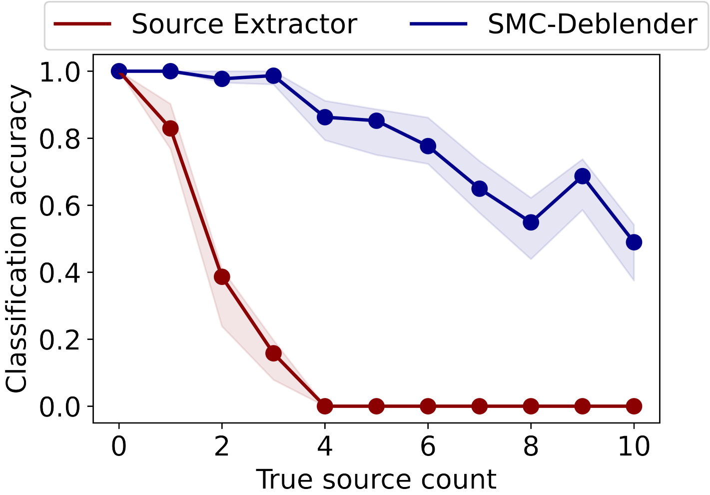

Calibration of estimated source counts

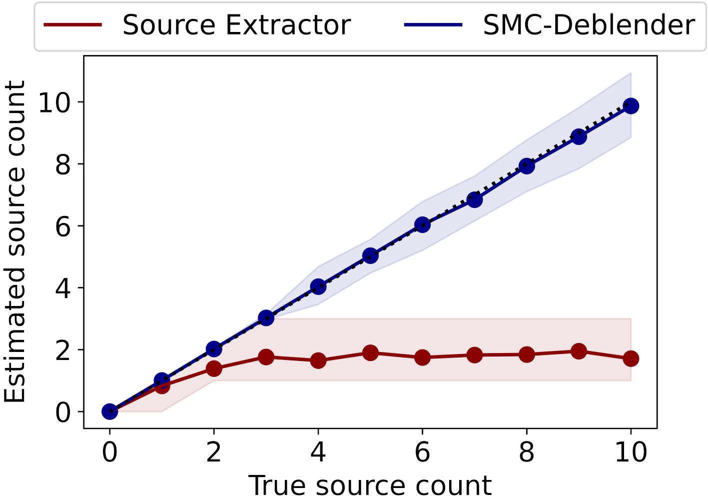

Reconstructed images

True source count = 2

Observed image

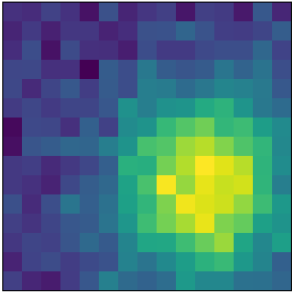

Reconstructed images

True source count = 2

Observed image

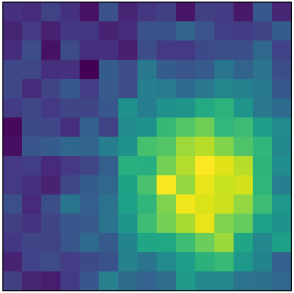

SMC-Deblender

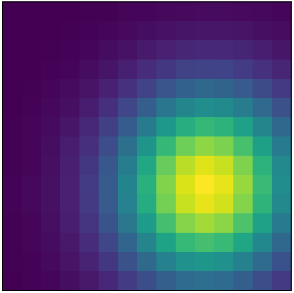

Reconstructed images

True source count = 2

Observed image

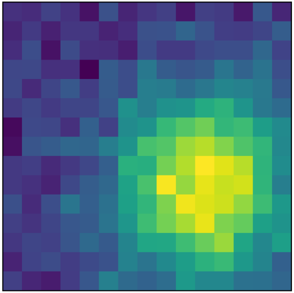

SMC-Deblender

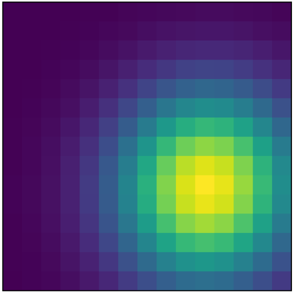

Source Extractor

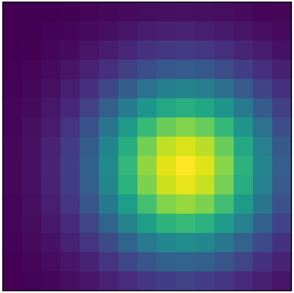

Reconstructed images

True source count $= 7$

Observed image

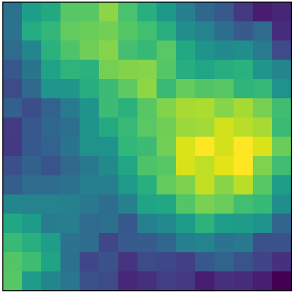

SMC-Deblender

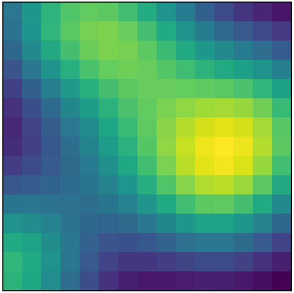

Source Extractor

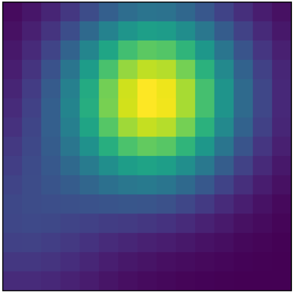

Discussion

Summary and next steps

- **Summary**
  - **SMC-Deblender** is a novel probabilistic cataloging algorithm
  - Has favorable properties compared to methods based on MCMC or VI
  - Outperformed a popular non-probabilistic method in our experiments

Summary and next steps

- **Summary**
  - **SMC-Deblender** is a novel probabilistic cataloging algorithm
  - Has favorable properties compared to methods based on MCMC or VI
  - Outperformed a popular non-probabilistic method in our experiments
- **Next steps**

Summary and next steps

- **Summary**
  - **SMC-Deblender** is a novel probabilistic cataloging algorithm
  - Has favorable properties compared to methods based on MCMC or VI
  - Outperformed a popular non-probabilistic method in our experiments
- **Next steps**
  - **Divide and conquer** to perform inference on larger images

Summary and next steps

Summary

- **SMC-Deblender** is a novel probabilistic cataloging algorithm
- Has favorable properties compared to methods based on MCMC or VI
- Outperformed a popular non-probabilistic method in our experiments

Next steps

- **Divide and conquer** to perform inference on larger images
- $\downarrow$ runtime by using a **normalizing flow** instead of an MCMC kernel

Summary and next steps

- **Summary**
  - **SMC-Deblender** is a novel probabilistic cataloging algorithm
  - Has favorable properties compared to methods based on MCMC or VI
  - Outperformed a popular non-probabilistic method in our experiments

- **Next steps**
  - **Divide and conquer** to perform inference on larger images
  - $\downarrow$ runtime by using a **normalizing flow** instead of an MCMC kernel
  - Apply to other image analysis tasks in astronomy and astrophysics:
    1. **Weak gravitational lensing**
    2. **Galaxy clusters**
    3. **Photometric redshift**

Thank you!

https://linktr.ee/timwhite0