Sequential Monte Carlo for detecting and deblending objects in astronomical images

Tim White   Jeffrey Regier

Department of Statistics, University of Michigan

1 Background

Problem:

Astronomical cataloging is the task of inferring the properties of stars and galaxies in astronomical images. This task is difficult for images that contain many overlapping light sources.

The challenge of deblending sources in such an image is that the true number of sources is unknown, so the properties of the overlapping sources are ambiguous.

The Messier 15 globular cluster, imaged by Hubble

Existing methods:

Probabilistic cataloging methods infer a posterior distribution over all possible latent variable catalogs, assigning a higher probability to those that are more plausible given the pixel intensities of the image.

- Markov chain Monte Carlo samples catalogs from the posterior distribution. Requires transdimensional sampling.
- Variational inference optimizes the parameters of an approximation to the posterior. Requires nonconvex optimization.

Our contribution:

We propose a probabilistic cataloging method based on sequential Monte Carlo (SMC). Our algorithm, which we call SMC-Deblender, avoids transdimensional sampling by running multiple Markov chains in parallel over blocks of catalogs that are stratified by source count.

2 Statistical model

Image x with height of H pixels and width of W pixels.

Source count $s \sim \text{Uniform}\{0, 1, 2, ..., D\}$.

Given s, sample locations $u_1, ..., u_s \overset{\text{iid}}{\sim} \text{Uniform}([0, H] \times [0, W])$.

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

Catalog $z = \{s, \{u_j\}_{j=1}^s, \{f_j\}_{j=1}^s\}$.

Intensity at pixel (h, w) is $x_{hw} \mid z \sim \text{Poisson}(\gamma + \sum_{j=1}^s f_j \, \psi((h, w) - u_j))$,

where γ is background intensity, $\psi(\cdot)$ is point spread function.

Posterior distribution over catalogs is $p(z \mid x) \propto p(z)p(x \mid z)$.

3 Deblending with sequential Monte Carlo

Goal:

Generate a collection of weighted catalogs from $p(z \mid x)$, the posterior distribution of possible catalogs for a given image. This posterior is intractable — we cannot sample from it directly.

Approach:

- Sample catalogs from a tractable initial proposal distribution (e.g., the prior), allocating an equal number of catalogs to each candidate source count s. Preserve this block structure in the propagation and resampling steps below.
- Propagate the catalogs through a sequence of intermediate distributions $p(z)p(x \mid z)^\tau$ that connect the prior ($\tau = 0$) to the posterior ($\tau = 1$), and update the catalog weights accordingly.
- In each iteration, resample the catalogs if necessary to prevent the weight of any one catalog from dominating the others.

Likelihood-tempered SMC sampler, stratified by source count (SMC-Deblender)

Input: Image x; likelihood $p(x \mid z)$; prior $p(z)$;
method to construct invariant kernel $M_\tau(\cdot \mid \cdot)$ for $\tau \in [0,1]$;
number of blocks B (indexed by b);
number of catalogs per block N (indexed by k);
effective sample size threshold ESS_{min}.

Step $t \leftarrow -1$. Temperature $\tau_t \leftarrow 0$. Unnormalized weights $w_{bk}^{(t)} \leftarrow 1$.
Intra-block normalized weights $\widetilde{W}_{bk}^{(t)}$. Inter-block normalized weights $W_{bk}^{(t)}$.
Intra-block effective sample size $\text{ESS}_b^{(t)}$.

while $\tau_t < 1$ **do**
 $t \leftarrow t + 1$.

 if $t = 0$ **then**
  Initialize:
   Sample catalogs $z_{bk}^{(t)} \overset{\text{iid}}{\sim} p(z)$.
    (catalogs in the same block have the same source count s)

 if $t \geq 1$ **then**
  Resample:
   for block $b \in \{1, ..., B\}$ **do**
    if $\text{ESS}_b^{(t-1)} < \text{ESS}_{\text{min}}$ **then**
     Resample catalogs $\{z_{bk}^{(t-1)}\}_{k=1}^N$ using intra-block weights $\{\widetilde{W}_{bk}^{(t-1)}\}_{k=1}^N$.
     Reset inter-block weights and intra-block weights in block b.

  Propagate:
   Generate new catalogs $z_{bk}^{(t)}$ using kernel $M_{\tau_{t-1}}(\cdot \mid z_{bk}^{(t-1)})$.
    (source count s of each catalog remains unchanged)

 Temper:
  Update temperature $\tau_t \leftarrow \tau_{t-1} + \delta$, where $\delta \in [0, 1 - \tau_{t-1}]$.

 Update weights:
  Update unnormalized weights $w_{bk}^{(t)} \leftarrow W_{bk}^{(t-1)} p(x \mid z_{bk}^{(t)})^{\tau_t - \tau_{t-1}}$.
  Normalize $w_{bk}^{(t)}$ over all catalogs to obtain inter-block weights $W_{bk}^{(t)}$.
  Normalize $w_{bk}^{(t)}$ over catalogs in block b to obtain intra-block weights $\widetilde{W}_{bk}^{(t)}$.
  Compute $\text{ESS}_b^{(t)} \leftarrow \left(\sum_k (\widetilde{W}_{bk}^{(t)})^2\right)^{-1}$ for block b.

Output: Set of weighted catalogs $\{\{W_{bk}^{(t)}, z_{bk}^{(t)}\}_{k=1}^N\}_{b=1}^B$.

4 Experiments

1,000 synthetic images, each containing up to 10 light sources.

Compare accuracy of estimated source counts between SMC-Deblender and SEP (Source Extractor), a popular non-probabilistic cataloging tool.

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

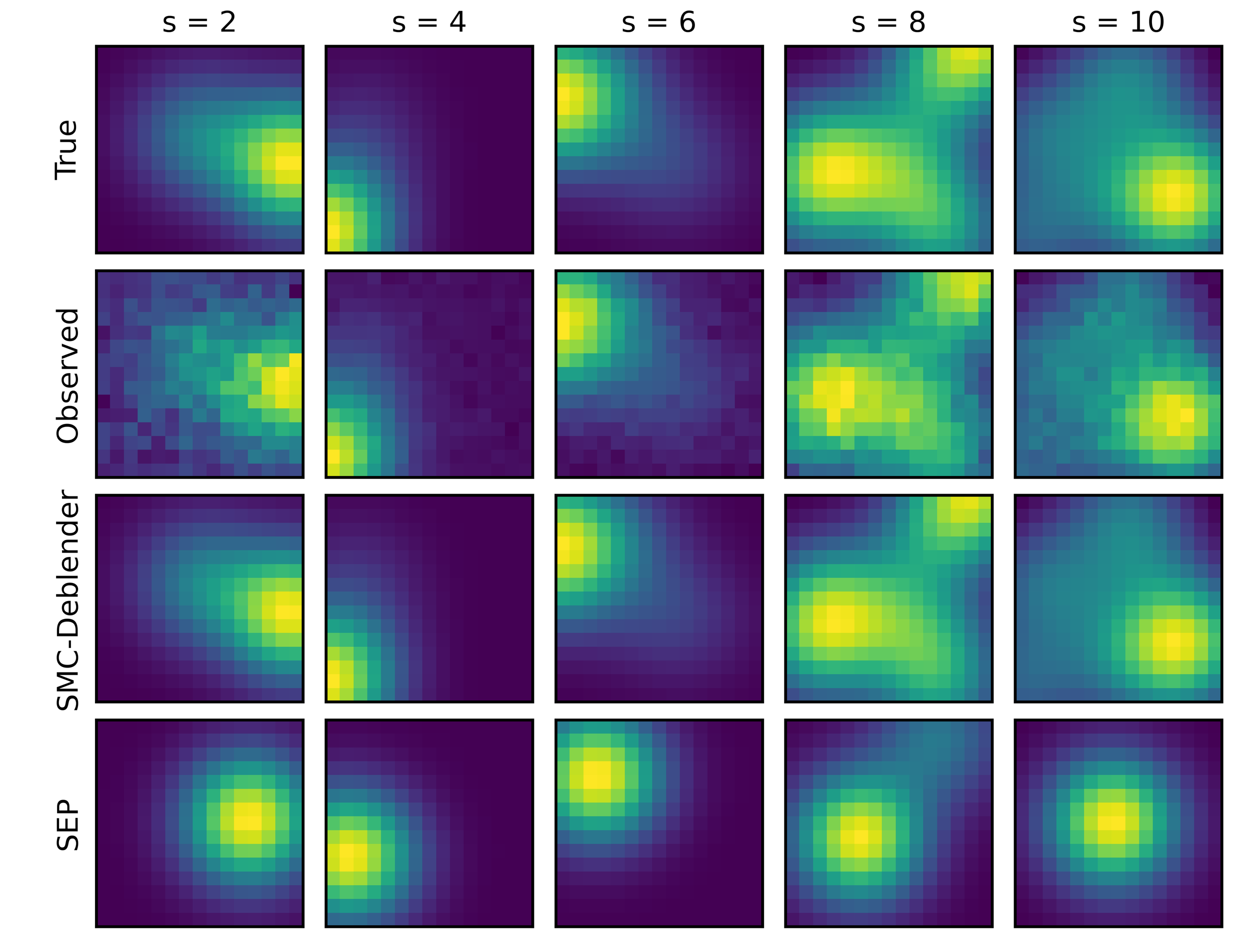

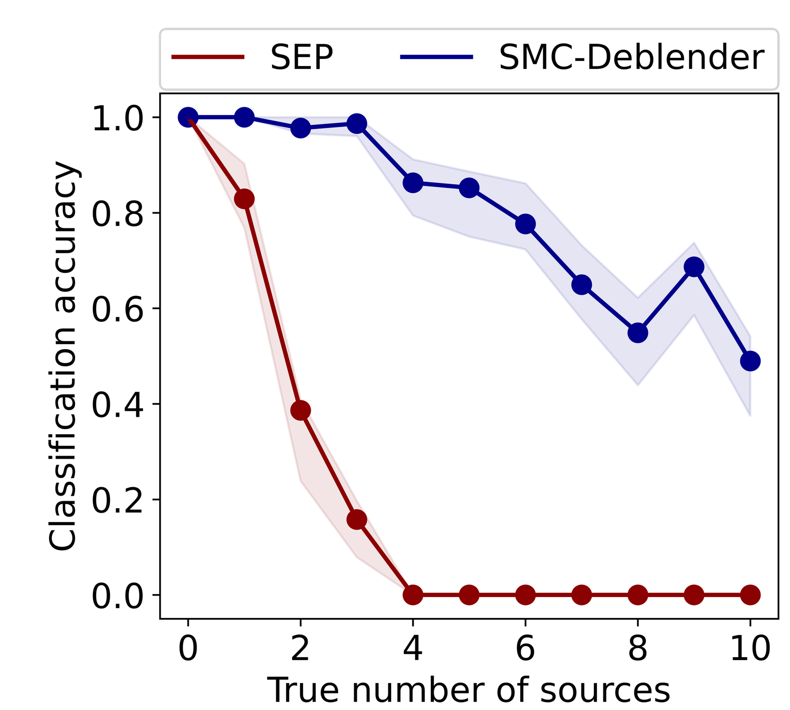

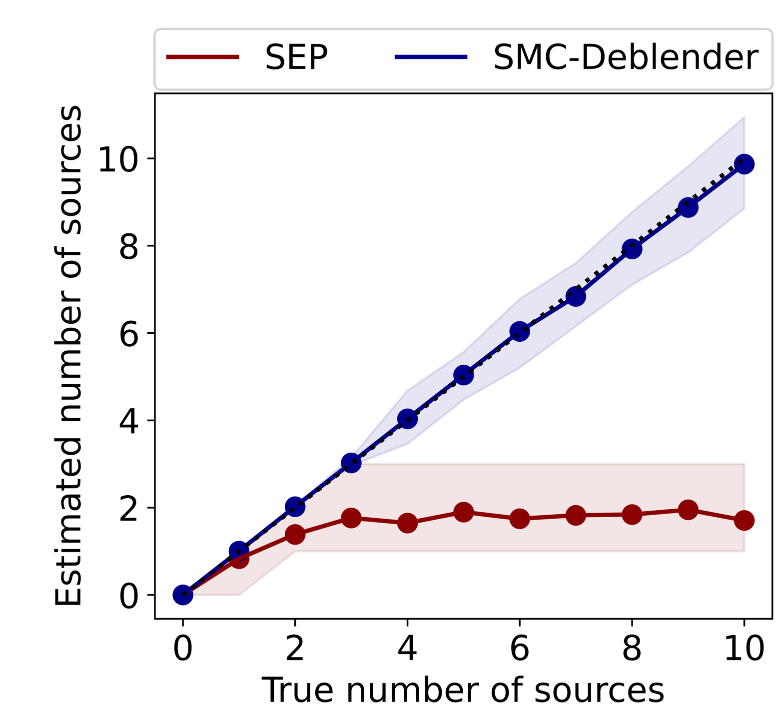

5 Next steps

- Embed the algorithm within a divide-and-conquer conditional sampling scheme to perform inference on larger images.
- Speed up the propagation step by using normalizing flows.
- Apply to other tasks in astronomy and astrophysics: (1) weak gravitational lensing, (2) galaxy clusters, (3) photometric redshift.