Waste-free sequential Monte Carlo for light source detection in astronomical images

Jaylin Lowe and Tim White

STATS 608 final project

April 22nd, 2024

Background

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, \ldots \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, \ldots \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.
- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, \ldots \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.
- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.

  1. $t = 0$: Sample $z_0^{1:N} \sim \pi_0(z)$ and assign weights $w_0^{1:N}$.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.

- **Goal:** Sample $z_1, z_2, ... \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.

- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.

  1. $t = 0$: Sample $z_0^{1:N} \sim \pi_0(z)$ and assign weights $w_0^{1:N}$.
  2. For $t \in \{1, ..., T\}$:

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, \ldots \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.
- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.
  1. $t = 0$: Sample $z_0^{1:N} \sim \pi_0(z)$ and assign weights $w_0^{1:N}$.
  2. For $t \in \{1, \ldots, T\}$:
     1. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, \ldots \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.
- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.

  1. $t = 0$: Sample $z_0^{1:N} \sim \pi_0(z)$ and assign weights $w_0^{1:N}$.
  2. For $t \in \{1, \ldots, T\}$:
     1. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.
     2. **Mutate** particles $z_t^{1:N} \sim \texttt{MetropolisHastings}(z_{t-1}^{A_t^{1:N}}, dz_t)$.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
    - We can evaluate $\gamma(z)$ at a particular z.
    - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, \ldots \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.
- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.
    1. $t = 0$: Sample $z_0^{1:N} \sim \pi_0(z)$ and assign weights $w_0^{1:N}$.
    2. For $t \in \{1, \ldots, T\}$:
        1. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.
        2. **Mutate** particles $z_t^{1:N} \sim \texttt{MetropolisHastings}(z_{t-1}^{A_t^{1:N}}, dz_t)$.
        3. **Update weights** $w_t^{1:N}$ of mutated particles.

Sequential Monte Carlo (SMC) samplers

- $\pi(z) = \frac{\gamma(z)}{L}$: Target density over latent random variables z.
  - We can evaluate $\gamma(z)$ at a particular z.
  - We cannot draw samples from π.
- **Goal:** Sample $z_1, z_2, ... \sim \pi(z)$ so we can approximate $E_\pi[g(z)]$.
- **Approach:** Define sequence of auxiliary targets $\pi_0 \to \cdots \to \pi_T = \pi$.
  1. $t = 0$: Sample $z_0^{1:N} \sim \pi_0(z)$ and assign weights $w_0^{1:N}$.
  2. For $t \in \{1, ..., T\}$:
     1. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.
     2. **Mutate** particles $z_t^{1:N} \sim \texttt{MetropolisHastings}(z_{t-1}^{A_t^{1:N}}, dz_t)$.
     3. **Update weights** $w_t^{1:N}$ of mutated particles.
- **Output:** $\{z_T^{1:N}, w_T^{1:N}\} \sim \pi(z)$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.
  2. While $\tau_t < 1$:

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.
  2. While $\tau_t < 1$:
     1. **Increase temperature** τ_t.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.
  2. While $\tau_t < 1$:
     1. **Increase temperature** τ_t.
     2. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.
  2. While $\tau_t < 1$:
     1. **Increase temperature** τ_t.
     2. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.
     3. **Mutate** particles $z_t^{1:N} \sim \texttt{MetropolisHastings}(z_{t-1}^{A_t^{1:N}}, dz_t)$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.
  2. While $\tau_t < 1$:
     1. **Increase temperature** τ_t.
     2. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.
     3. **Mutate** particles $z_t^{1:N} \sim \texttt{MetropolisHastings}(z_{t-1}^{A_t^{1:N}}, dz_t)$.
     4. **Update weights:** $w_t^{1:N} \leftarrow p(x \mid z_t^{1:N})^{\tau_t - \tau_{t-1}}$.

SMC samplers for Bayesian latent variable models

- Latent z with **prior** $p(z)$; observed x with **likelihood** $p(x \mid z)$.
- Target is the **posterior**: $\pi(z) := p(z \mid x) = \frac{p(z)p(x|z)}{p(x)} =: \frac{\gamma(z)}{L}$.
- Define $\pi_t(z) \propto p(z)p(x \mid z)^{\tau_t}$, where $0 = \tau_0 < \cdots < \tau_T = 1$.
  1. $t = 0$: Sample $z_0^{1:N} \sim p(z)$ and assign weights $w_0^{1:N}$.
  2. While $\tau_t < 1$:
     1. **Increase temperature** τ_t.
     2. **Resample** indices $A_t^{1:N} \leftarrow \texttt{resample}(N, w_{t-1}^{1:N})$.
     3. **Mutate** particles $z_t^{1:N} \sim \texttt{MetropolisHastings}(z_{t-1}^{A_t^{1:N}}, dz_t)$.
     4. **Update weights:** $w_t^{1:N} \leftarrow p(x \mid z_t^{1:N})^{\tau_t - \tau_{t-1}}$.
- **Output:** $\{z_T^{1:N}, w_T^{1:N}\} \sim p(z \mid x)$.

Waste-free sequential Monte Carlo

- Mutation step is crucial to algorithm's performance but sensitive to design of MCMC kernel.
  - i.e., number of iterations, proposal covariance.

Waste-free sequential Monte Carlo

- Mutation step is crucial to algorithm's performance but sensitive to design of MCMC kernel.
  - i.e., number of iterations, proposal covariance.

Waste-free sequential Monte Carlo

Hai-Dang Dau | **Nicolas Chopin**

CREST-ENSAE and Institut Polytechnique de Paris, Palaiseau Cedex, France

Correspondence
Nicolas Chopin, CREST-ENSAE and Institut Polytechnique de Paris, 5 avenue Henry Le Chatelier, 91764 Palaiseau Cedex, France.
Email: nicolas.chopin@ensae.fr

Funding information
Labex Ecodec, Grant/Award Number: Ecodec/ANR-11-LABX-0047

Abstract
A standard way to move particles in a sequential Monte Carlo (SMC) sampler is to apply several steps of a Markov chain Monte Carlo (MCMC) kernel. Unfortunately, it is not clear how many steps need to be performed for optimal performance. In addition, the output of the intermediate steps are discarded and thus wasted somehow. We propose a new, waste-free SMC algorithm which uses the outputs of all these intermediate MCMC steps as particles. We establish that its output is consistent and asymptotically normal. We use the expression

Waste-free sequential Monte Carlo

- Mutation step is crucial to algorithm's performance but sensitive to design of MCMC kernel.
    - i.e., number of iterations, proposal covariance.

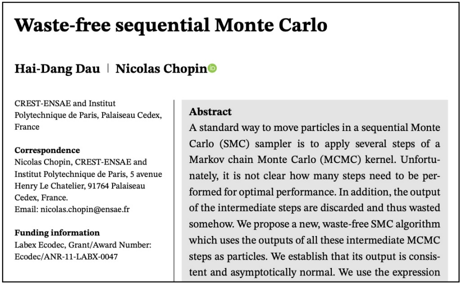

Waste-free sequential Monte Carlo

Hai-Dang Dau | Nicolas Chopin

CREST-ENSAE and Institut Polytechnique de Paris, Palaiseau Cedex, France

Correspondence
Nicolas Chopin, CREST-ENSAE and Institut Polytechnique de Paris, 5 avenue Henry Le Chatelier, 91764 Palaiseau Cedex, France.
Email: nicolas.chopin@ensae.fr

Funding information
Labex Ecodec, Grant/Award Number: Ecodec/ANR-11-LABX-0047

Abstract
A standard way to move particles in a sequential Monte Carlo (SMC) sampler is to apply several steps of a Markov chain Monte Carlo (MCMC) kernel. Unfortunately, it is not clear how many steps need to be performed for optimal performance. In addition, the output of the intermediate steps are discarded and thus wasted somehow. We propose a new, waste-free SMC algorithm which uses the outputs of all these intermediate MCMC steps as particles. We establish that its output is consistent and asymptotically normal. We use the expression

- **Benefits:** (1) ↓ effort required to tune MCMC kernel.
  (2) ↓ asymptotic variance of Monte Carlo estimates.

Guiding questions

1. **How does waste-free SMC work?**
   - Why is it supposed to have favorable properties?

Guiding questions

1. **How does waste-free SMC work?**
   - Why is it supposed to have favorable properties?

2. **Does it outperform a "standard" SMC sampler on a challenging transdimensional inference task?**
   - Are its estimates more accurate? Less variable? Is it faster?

Methods

Detecting light sources in astronomical images

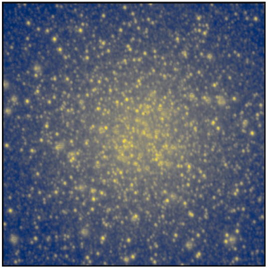

The Messier 53 globular cluster, imaged by SDSS

Detecting light sources in astronomical images

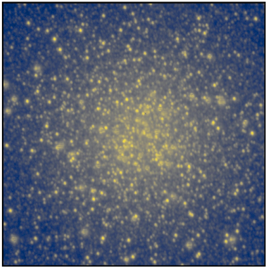

The Messier 53 globular cluster, imaged by SDSS

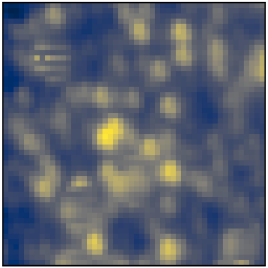

50 pixel by 50 pixel subregion of Messier 53

Detecting light sources in astronomical images

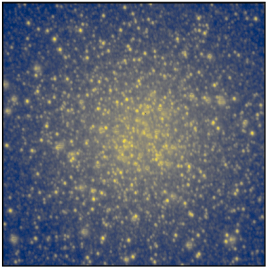

The Messier 53 globular cluster, imaged by SDSS

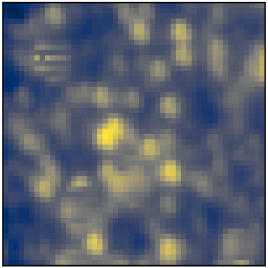

50 pixel by 50 pixel subregion of Messier 53

- **Given:** Pixelated image of blended light sources.

  Goal: Infer source count and properties of each source.

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** $s \sim \text{Uniform}\{0, 1, 2, \dots, D\}$
  - 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** $z = \{s, \{u_j, f_j\}_{j=1}^{s}\}$

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** $s \sim \text{Uniform}\{0, 1, 2, \ldots, D\}$
  - 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)

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 $x_{hw} \mid z \sim \text{Poisson}(\lambda_{hw})$
    - $\lambda_{hw} = \text{background intensity} + \text{sum of fluxes at pixel } (h, w)$
- **Posterior** $p(z \mid x) \propto p(z)p(x \mid z)$

Standard vs. waste-free SMC samplers

Standard

1. Initialize N catalogs and weights.
2. While $\tau_t < 1$:
   1. Increase temperature.
   2. Resample N catalogs.
   3. For $i \in \{1, ..., N\}$:
      Mutate ith catalog k times.
      Keep last mutation only.
   4. Update weights.

Standard vs. waste-free SMC samplers

Standard

1. Initialize N catalogs and weights.
2. While $\tau_t < 1$:
   1. Increase temperature.
   2. Resample N catalogs.
   3. For $i \in \{1, ..., N\}$:
      Mutate ith catalog k times.
      Keep last mutation only.
   4. Update weights.

Waste-free

1. Initialize N catalogs and weights.
2. While $\tau_t < 1$:
   1. Increase temperature.
   2. Resample $M << N$ catalogs, where $N = MP$.
   3. For $i \in \{1, ..., M\}$:
      Mutate ith catalog $P-1$ times.
      Keep all $P-1$ mutations.
   4. Update weights.

Several questions

1. Why is it called "waste-free"?

Several questions

1. Why is it called “waste-free”?
    - **Intermediate mutations are not discarded, and compute is not wasted on them.**

Several questions

1. Why is it called "waste-free"?
   - **Intermediate mutations are not discarded, and compute is not wasted on them.**
2. Does it have the same theoretical guarantees as standard SMC?

Several questions

1. Why is it called “waste-free”?
   - **Intermediate mutations are not discarded, and compute is not wasted on them.**
2. Does it have the same theoretical guarantees as standard SMC?
   - **Yes — e.g., unbiased estimate of $p(x)$, posterior estimates are consistent and asymptotically normal.**

Several questions

1. Why is it called "waste-free"?
   - **Intermediate mutations are not discarded, and compute is not wasted on them.**
2. Does it have the same theoretical guarantees as standard SMC?
   - **Yes — e.g., unbiased estimate of $p(x)$, posterior estimates are consistent and asymptotically normal.**
3. What value does it add?

Several questions

1. Why is it called "waste-free"?
   - **Intermediate mutations are not discarded, and compute is not wasted on them.**
2. Does it have the same theoretical guarantees as standard SMC?
   - **Yes — e.g., unbiased estimate of p(x), posterior estimates are consistent and asymptotically normal.**
3. What value does it add?
   - **Posterior estimates have smaller asymptotic variance under certain assumptions. Choice of M and P may be more robust than choice of k.**

Experiment 1

Experiment 1: Variability of posterior estimates

- Questions

Experiment 1: Variability of posterior estimates

- **Questions**
  - How variable are SMC estimates across different values of k (standard) and M and P (waste-free)?

Experiment 1: Variability of posterior estimates

- **Questions**
  - How variable are SMC estimates across different values of k (standard) and M and P (waste-free)?
  - Is waste-free SMC more robust to the choice of M and P than standard SMC is to the choice of k?

Experiment 1: Variability of posterior estimates

- **Questions**
  - How variable are SMC estimates across different values of k (standard) and M and P (waste-free)?
  - Is waste-free SMC more robust to the choice of M and P than standard SMC is to the choice of k?
- **Details**

Experiment 1: Variability of posterior estimates

- **Questions**
  - How variable are SMC estimates across different values of k (standard) and M and P (waste-free)?
  - Is waste-free SMC more robust to the choice of M and P than standard SMC is to the choice of k?
- **Details**
  - Four 15×15 synthetic images with source count $\in \{2, 4, 6, 8\}$.

Standard

k	N
5	2000
25	400
50	200
100	100
200	50

Waste-free

M	P
25	400
50	200
80	125
125	80
200	50

Experiment 1: Variability of posterior estimates

- **Questions**
  - How variable are SMC estimates across different values of k (standard) and M and P (waste-free)?
  - Is waste-free SMC more robust to the choice of M and P than standard SMC is to the choice of k?
- **Details**
  - Four 15×15 synthetic images with source count $\in \{2, 4, 6, 8\}$.

Standard

k	N
5	2000
25	400
50	200
100	100
200	50

Waste-free

M	P
25	400
50	200
80	125
125	80
200	50

  - 100 runs × 5 parameter combinations × 4 images × 2 methods

Four synthetic images

Source count = 2

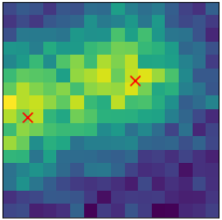

Source count = 4

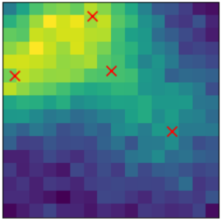

Source count = 6

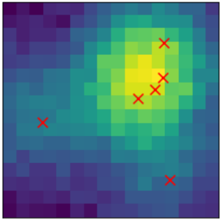

Source count = 8

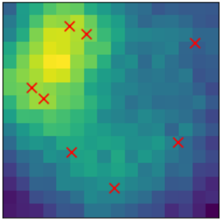

Results: Posterior mean source counts

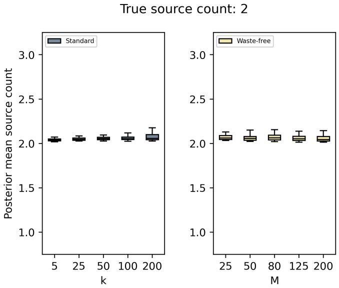

Results: Posterior mean source counts

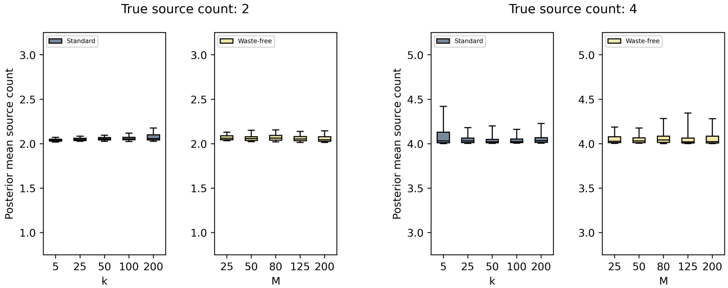

Results: Posterior mean source counts

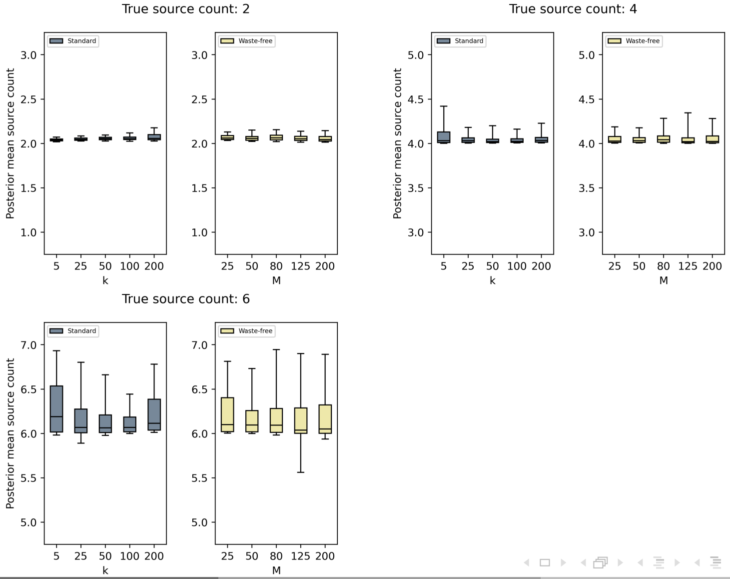

Results: Posterior mean source counts

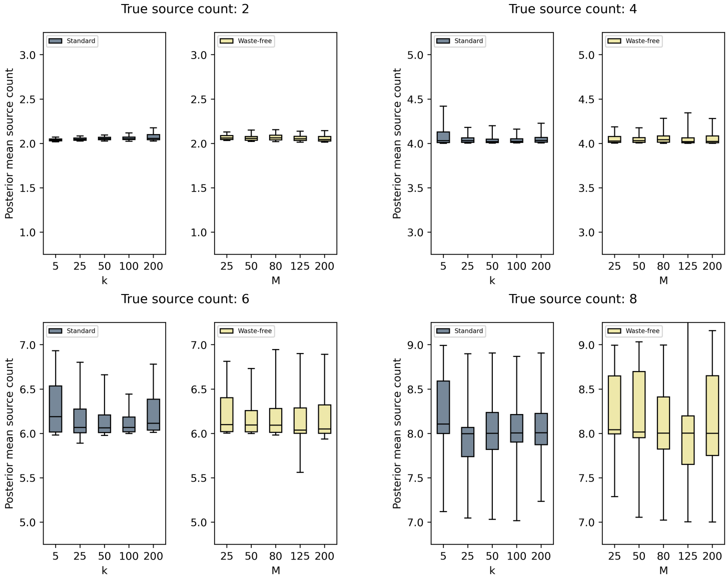

Results: Posterior mean total fluxes

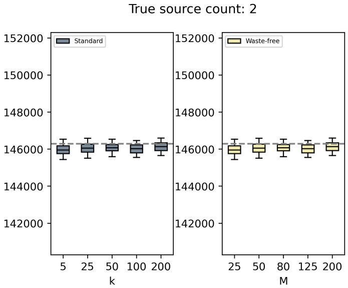

Results: Posterior mean total fluxes

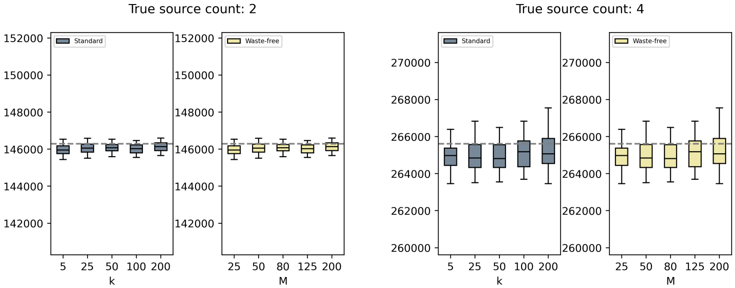

Results: Posterior mean total fluxes

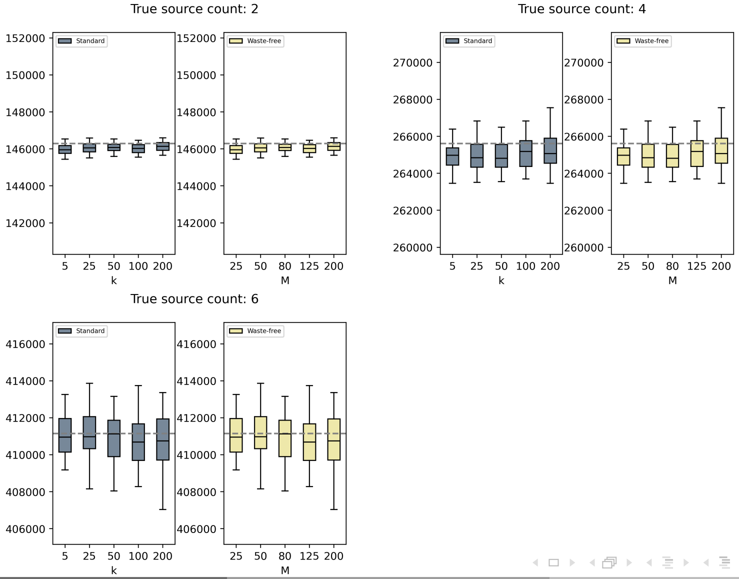

Results: Posterior mean total fluxes

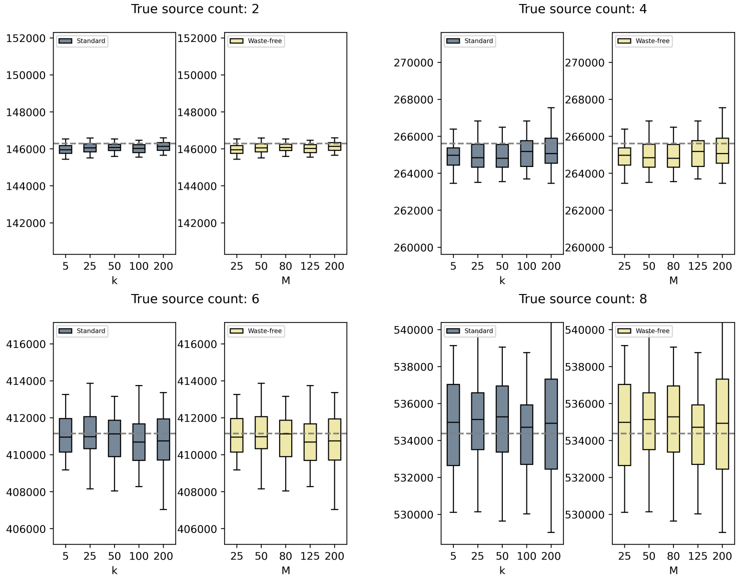

Results: $\log p(x)$

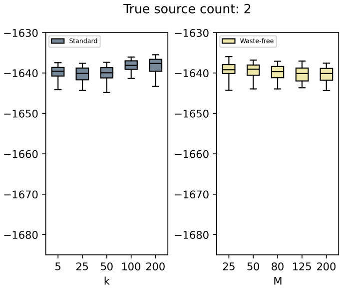

Results: $\log p(x)$

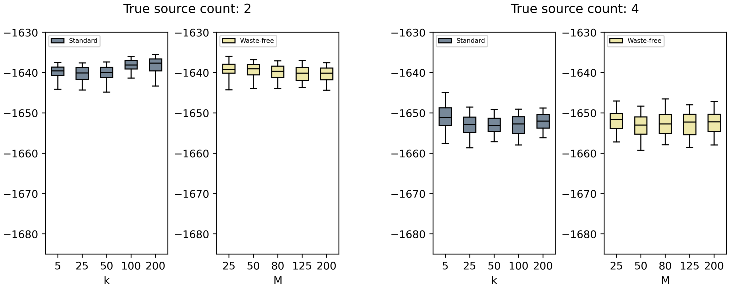

Results: $\log p(x)$

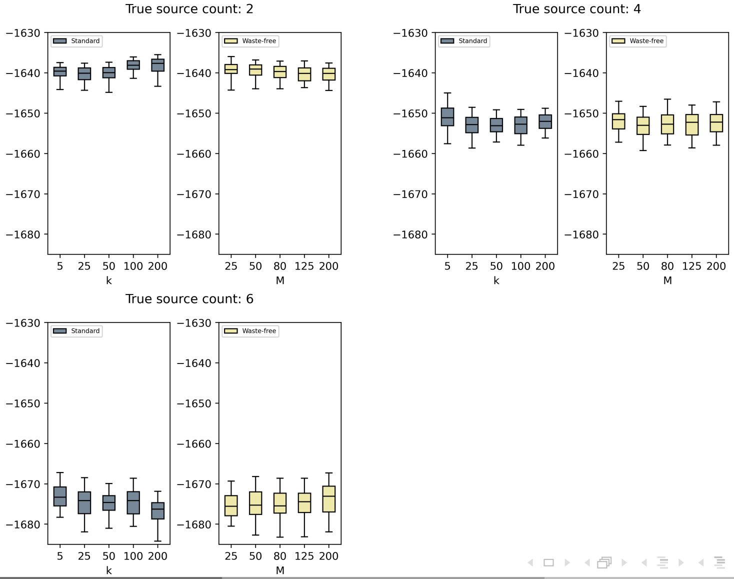

Results: $\log p(x)$

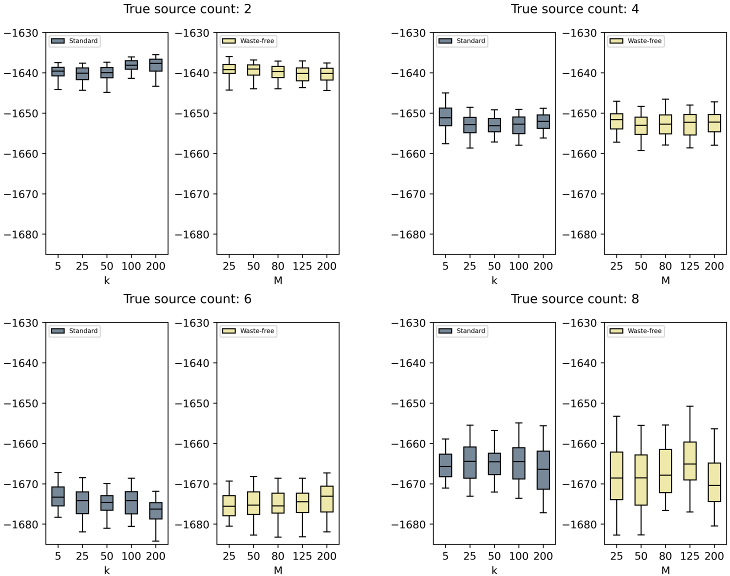

Experiment 2

Experiment 2: Calibration of posterior estimates

- Questions

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?
  - Does the waste-free method outperform the standard method?

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?
  - Does the waste-free method outperform the standard method?
- **Details**

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?
  - Does the waste-free method outperform the standard method?
- **Details**
  - 1500 synthetic images with source count $\in \{0, 1, 2, ..., 8\}$.

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?
  - Does the waste-free method outperform the standard method?
- **Details**
  - 1500 synthetic images with source count $\in \{0, 1, 2, ..., 8\}$.
  - Standard: $k = 100$ and $N = 100$

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?
  - Does the waste-free method outperform the standard method?
- **Details**
  - 1500 synthetic images with source count $\in \{0, 1, 2, ..., 8\}$.
  - Standard: $k = 100$ and $N = 100$
  - Waste free: $M = 80$ and $P = 125$

Experiment 2: Calibration of posterior estimates

- **Questions**
  - How accurate and calibrated are SMC estimates across many images?
  - Does the waste-free method outperform the standard method?
- **Details**
  - 1500 synthetic images with source count $\in \{0, 1, 2, ..., 8\}$.
  - Standard: $k = 100$ and $N = 100$
  - Waste free: $M = 80$ and $P = 125$
  - 1 run $\times$ 1 parameter combination $\times$ 1500 images $\times$ 2 methods

Results: Accuracy of posterior mean source counts

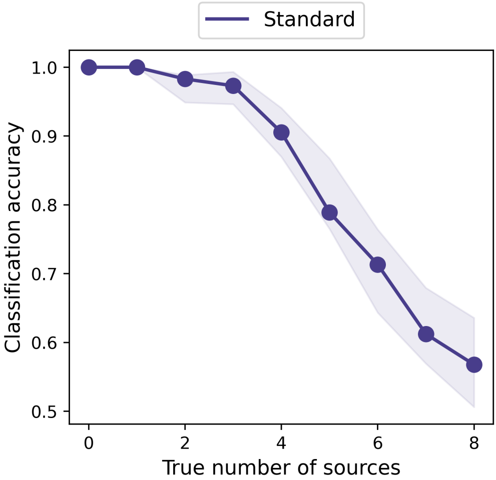

Results: Accuracy of posterior mean source counts

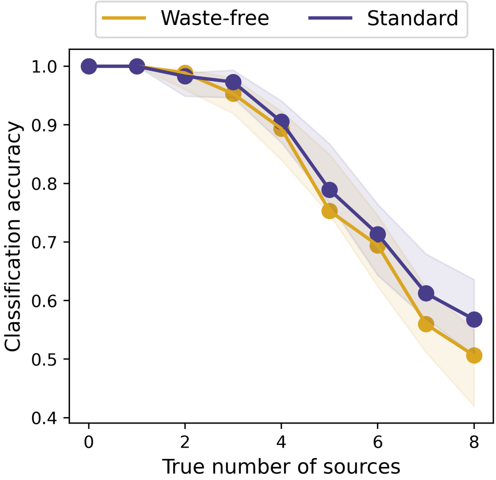

Results: Calibration of posterior mean source counts

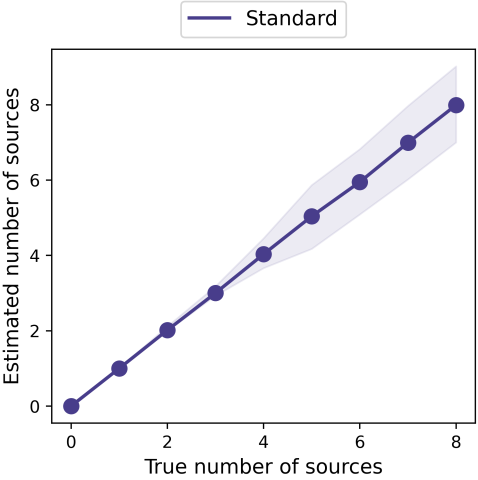

Results: Calibration of posterior mean source counts

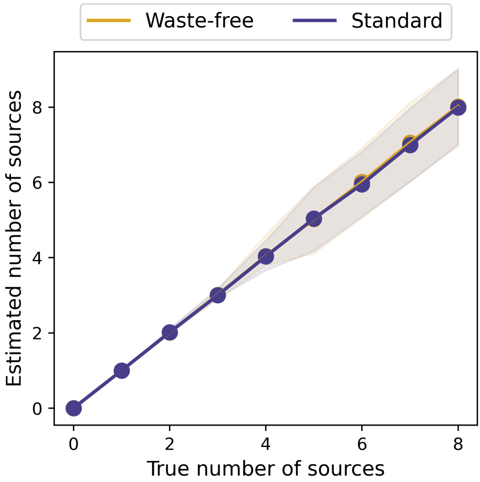

Results: MAE of posterior mean source counts

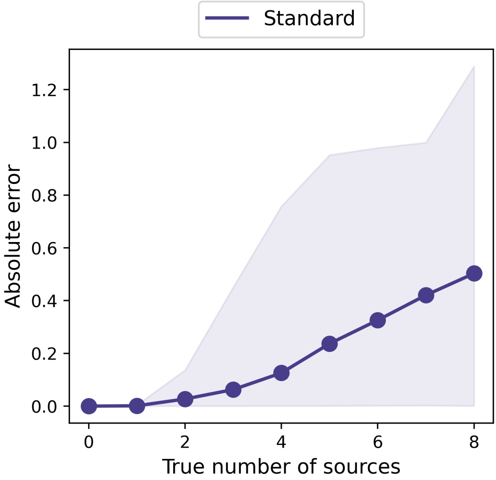

Results: MAE of posterior mean source counts

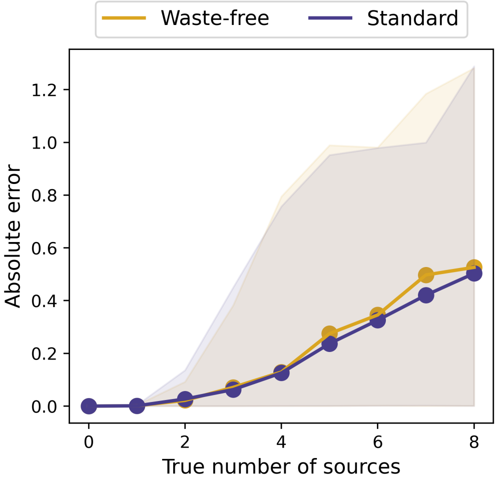

Discussion

Closing thoughts

- **Takeaway**
  - Waste-free resampling and mutation does not help (or hurt) much for this inference task.

Closing thoughts

- **Takeaway**
  - Waste-free resampling and mutation does not help (or hurt) much for this inference task.
- **Possible explanations**

Closing thoughts

- **Takeaway**
  - Waste-free resampling and mutation does not help (or hurt) much for this inference task.
- **Possible explanations**
  1. Our SMC samplers are tailored to object detection.
     - Catalogs stratified by source count to avoid transdimensional sampling.

Closing thoughts

- **Takeaway**
  - Waste-free resampling and mutation does not help (or hurt) much for this inference task.
- **Possible explanations**
  1. Our SMC samplers are tailored to object detection.
     - Catalogs stratified by source count to avoid transdimensional sampling.
  2. We used fixed proposal variances in the mutation step.
     - Could be adapted, but not obvious how to do this in this setting.

Closing thoughts

- **Takeaway**
  - Waste-free resampling and mutation does not help (or hurt) much for this inference task.
- **Possible explanations**
  1. Our SMC samplers are tailored to object detection.
     - Catalogs stratified by source count to avoid transdimensional sampling.
  2. We used fixed proposal variances in the mutation step.
     - Could be adapted, but not obvious how to do this in this setting.
  3. Original paper focused on long-chain setting (small M, large P).
     - Advantages of waste-free procedure may be exaggerated.

Thank you!

Results: Runtime

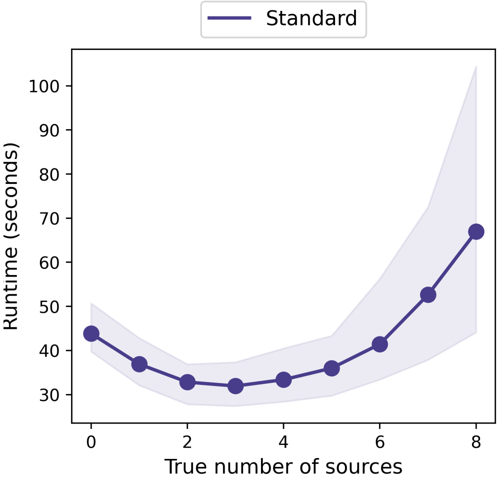

Results: Runtime