Sequential Monte Carlo algorithms for agent-based models of disease transmission

Nianqiao (Phyllis) Ju
Ph.D. student at Harvard University

Feb. 11th, 2021 @ Purdue University

‘不积跬步 无以至千里 不积小流 无以成江海’ – 荀子《劝学》

Unless you collect little steps,
you can never journey a thousand miles;
Unless you gather tiny streams,
you can never make a river or a sea.
– Xun Zi, *Encouraging Learning*

Many **large-scale** phenomena emerge from **individual-level** interactions.

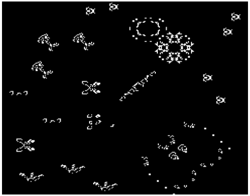

Conway

city

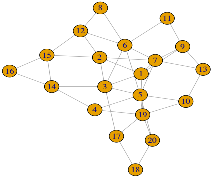

social network

photo credit: `withgoogle.com` & `unsplash.com`

Agent-based models

Agent-based models are flexible, interpretable and used in many fields:

- social sciences,
- demographics,
- ecology,
- macroeconomics.

Books on agent-based models

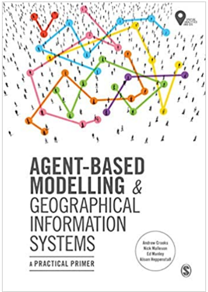

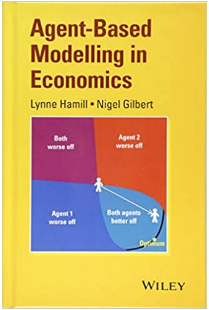

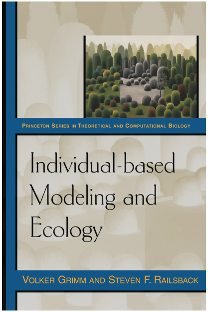

Softwares for agent-based models

Anylogic

Netlogo

Repast

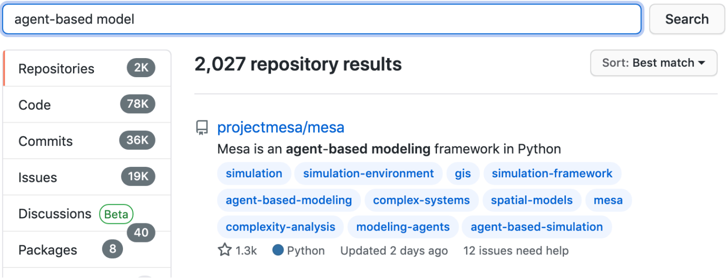

Agent-based models and COVID

Google Scholar agent-based model AND covid

Articles About 3,980 results (0.05 sec)

Any time
Since 2021
Since 2020
Since 2017
Custom range...

Sort by relevance
Sort by date

include patents
include citations

Create alert

[HTML] COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions — [HTML] sciencedirect.com
PCL Silva, PVC Batista, HS Lima, MA Alves… - Chaos, Solitons & …, 2020 - Elsevier
The COVID-19 pandemic due to the SARS-CoV-2 coronavirus has directly impacted the public health and economy worldwide. To overcome this problem, countries have adopted different policies and non-pharmaceutical interventions for controlling the spread of the …
Cited by 45 Related articles All 20 versions

Covasim: an agent-based model of COVID-19 dynamics and interventions — [PDF] medrxiv.org
CC Kerr, RM Stuart, D Mistry, RG Abeysuriya, G Hart… - medRxiv, 2020 - medrxiv.org
The COVID-19 pandemic has created an urgent need for models that can project epidemic trends, explore intervention scenarios, and estimate resource needs. Here we describe the methodology of Covasim (COVID-19 Agent-based Simulator), an open-source model …
Cited by 30 Related articles All 5 versions

Letter | Published: 14 July 2020

A stochastic agent-based model of the SARS-CoV-2 epidemic in France

Nicolas Hoertel, Martin Blachier, Carlos Blanco, Mark Olfson, Marc Massetti, Marina Sánchez Rico, Frédéric Limosin & Henri Leleu

Nature Medicine **26**, 1417–1421(2020) | Cite this article

17k Accesses | **14** Citations | **316** Altmetric | Metrics

Agent-based models and COVID: current practice

1. Build an agent-based model;

2. Find model parameters from prior studies or estimate them through simulation-based optimization;

3. Predict outcomes by simulating from the agent-based model.

From simulations to inference

- Agent-based models are ubiquitous and have been used as a **simulation** paradigm or for **model–based predictions**;

- **statistical inference** for these models has not received as much attention.

Outline

1. Scientific problem: inference in **agent-based SIS model**.
   - motivation,
   - hidden Markov model.

2. Computational challenge: design of sequential Monte Carlo algorithm.
   - intractable likelihood,
   - bootstrap particle filters are inefficient.

Compartmental models in epidemiology

Compartments: susceptible, infected, recovered, etc.

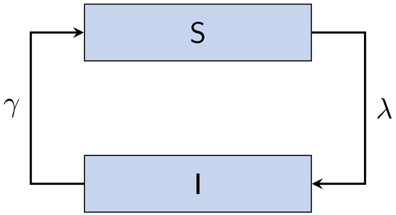

SIS model: common cold and influenza.

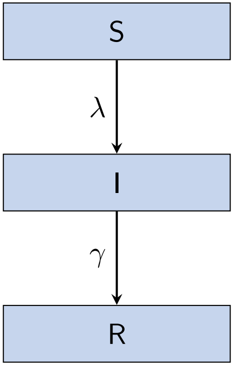

SIR model: smallpox, HIV.

Equation-based models

$$S_{t+1} = S_t - \lambda I_t S_t / N + \gamma I_t,$$
$$I_{t+1} = I_t + \lambda I_t S_t / N - \gamma I_t.$$

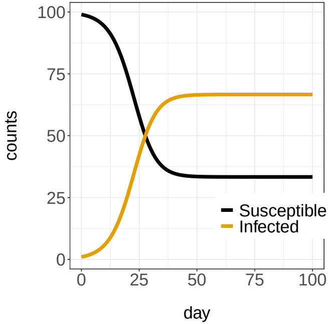

Assumptions are unrealistic:

- the network is fully connected,
- agents are homogeneous.

Network

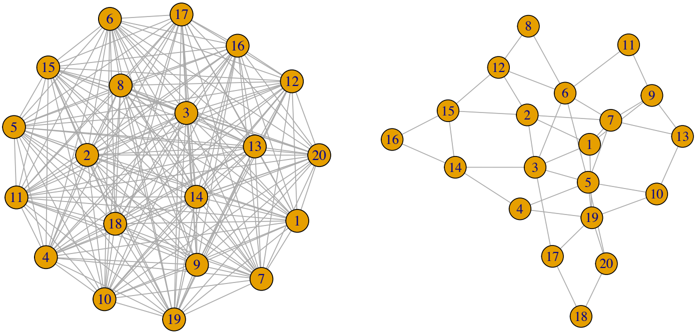

fully connected network v.s. small world network

Heterogeneous agents

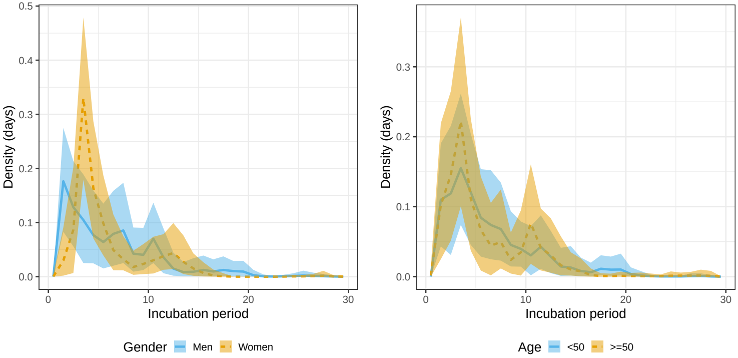

Gender-specific (left) and age-specific (right) distributions of the Covid-19 incubation period (Zhao, Ju, Bacallado & Shah, 2020, to appear on AoAS).

We need agent-based models because:

- they allow realistic assumptions on network structure and agent heterogeneity;
- we can incorporate prior individual-level information into the models;
- they allow individual-level policy making.

Outline

1. Scientific problem: inference in agent-based SIS model.
   - motivation,
   - **hidden Markov model**.
2. Computational challenge: design of sequential Monte Carlo algorithm.
   - intractable likelihood,
   - bootstrap particle filters are inefficient.

Agent-based SIS model: a hidden Markov model

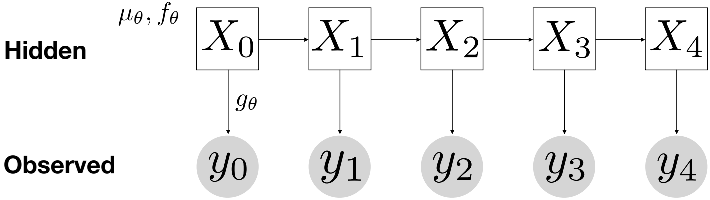

1. $X_t = (X_t^1, \ldots, X_t^N)$ is the state of the population at time $t \in [0 : T]$;

2. The time evolution of the population is
$X_0 \sim \mu_\theta, \quad X_t | X_{t-1} = x_{t-1} \sim f_\theta(\cdot | x_{t-1}), \quad t \in [1 : T].$

3. $Y_t \sim g_\theta(\cdot \mid X_t)$ is the number of reported infections at time t.

Agent-based SIS model: a hidden Markov model

- $X_t^n = \begin{cases} 0 & \text{if agent } n \text{ is susceptible} \\ 1 & \text{if agent } n \text{ is infected} \end{cases}$ at time t.

- Each agent has covariates are $w^n \in \mathbb{R}^d$ and a set of neighbors $\mathcal{N}(n)$.

- Parameters are $\theta = \{\beta_0, \beta_\lambda, \beta_\gamma, \rho\} \in \mathbb{R}^{3d} \times (0,1)$.

- Initial infection probabilities $\alpha_0^n = (1+\exp(-\beta_0^\top w^n))^{-1}$,
  infection rates $\lambda^n = (1+\exp(-\beta_\lambda^\top w^n))^{-1}$,
  recovery rates $\gamma^n = (1+\exp(-\beta_\gamma^\top w^n))^{-1}$.

Agent-based SIS model: a hidden Markov model

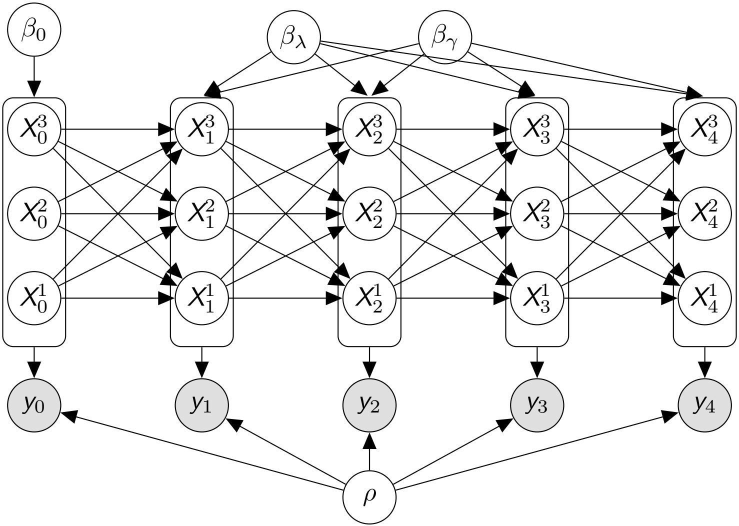

Agent-based SIS model: interactions and evolution

- Initial distribution μ: $x_0^n \sim \text{Bern}(\alpha_0^n)$ independently.
- Agents interact through a network and $\{\mathcal{N}(n)\}$ denotes the neighbors.
- The evolution is $x_t^n \mid x_{t-1} \sim \text{Bern}(\alpha^n(x_{t-1}))$ independently with

$$\alpha^n(x_{t-1}) = \begin{cases} \lambda^n \, |\mathcal{N}(n)|^{-1} \sum_{m \in \mathcal{N}(n)} x_{t-1}^m, & \text{if } x_{t-1}^n = 0 \ (\mathrm{S} \to \mathrm{I}), \\ 1 - \gamma^n, & \text{if } x_{t-1}^n = 1 \ (\mathrm{I} \to \mathrm{I}). \end{cases}$$

Agent-based SIS model: observation process

- We can only observe a proportion of the infections:

$$g_\theta(y_t \mid X_t = x_t) = \mathrm{Bin}(y_t; I(x_t), \rho),$$

where

$$I(x_t) = \sum_{n=1}^{N} x_t^n$$

is the unobserved total infection count at time t.

- Here $I(x_t)$ is a summary of x_t and it is sufficient for the observation model g_θ.

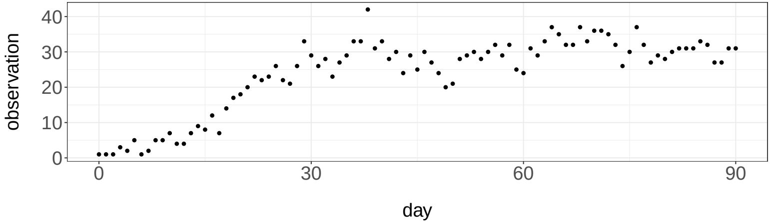

The **marginal likelihood** is $\mathcal{L}(\theta) = p_\theta(y_{0:T})$.

Frequentist: $\hat{\theta}_{\mathrm{MLE}} = \arg\max_\theta \mathcal{L}(\theta)$.

Bayesian: $p(\theta \mid y_{0:T}) \propto p(\theta) p_\theta(y_{0:T})$.

Marginal likelihood

The **marginal likelihood** of the hidden Markov model is

$$p_\theta(y_{0:T}) = \sum_{x \in \{0,1\}^{N(T+1)}} \mu(x_0) \prod_{t=1}^{T} f(x_t \mid x_{t-1}) \prod_{t=0}^{T} g(y_t \mid x_t)$$

$$= \mathbb{E}_{x \sim \mu, f}\left[\prod_{t=0}^{T} g(y_t \mid x_t)\right].$$

We will focus on computing $p(y_{0:T})$ and surpress the θ in notation.

Outline

1. Scientific problem: inference in agent-based SIS model.
   - motivation,
   - hidden Markov model.
2. Computational challenge: design of sequential Monte Carlo algorithm.
   - **intractable likelihood**
   - bootstrap particle filters are inefficient.

Forward-backward algorithm

We have a hidden Markov model with a **discrete** state-space.

The **forward-backward** algorithm allows us to exactly compute the marginal likelihood with cost of the order $(\#\text{of states})^2 \times (\#\text{of observations})$.

For our agent-based SIS model, this is $\mathcal{O}(2^{2N}T)$.

Importance sampling

We can approximate integrals using **weighted samples**.

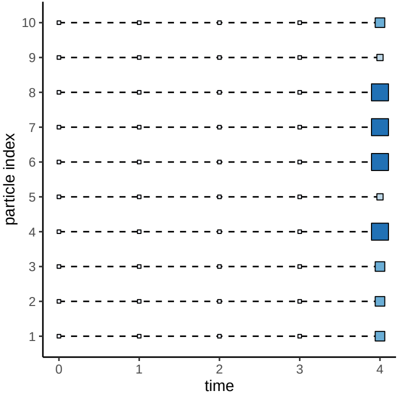

$$p(y_{0:T}) = \mathbb{E}_{x\sim\mu,f}\left[\prod_{t=0}^{T} g(y_t, X_t)\right].$$

proposal

$$x_0 \sim \mu, \quad x_t \mid x_{t-1} \sim f.$$

weight

$$w(x) = \prod_{t=0}^{T} g(y_t \mid x_t).$$

estimate

$$\hat{p}(y_{0:T}) = \frac{1}{P}\sum_{p=1}^{P} w(x^{(p)}).$$

$y_0 = (1, 2, 1, 1, 0), \alpha_0 = (0.3, 0.4, 0.5), \rho = 0.8, \lambda = (0.3, 0.4, 0.5), \gamma = (0.3, 0.2, 0.1).$

Bootstrap particle filter

To estimate $p(y_{0:T})$, we can apply importance sampling recursively and add a resampling step.

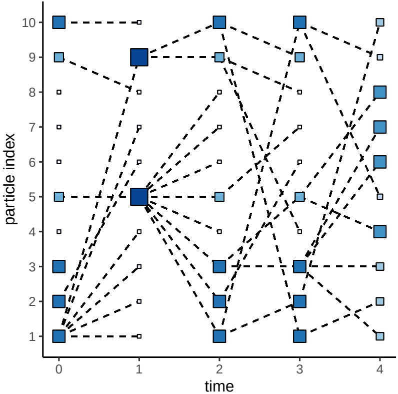

$$p(y_{0:T}) = \mathbb{E}_{x\sim\mu,f}\left[\prod_{t=0}^{T} g(y_t, x_t)\right]$$

proposal

$$x_0 \sim \mu, \quad x_t \mid x_{t-1} \sim f.$$

weight

$$w(x_t) = g(y_t \mid x_t).$$

resample with replacement according to weights.

Propose $x_0 \sim \mu$ for each particle.

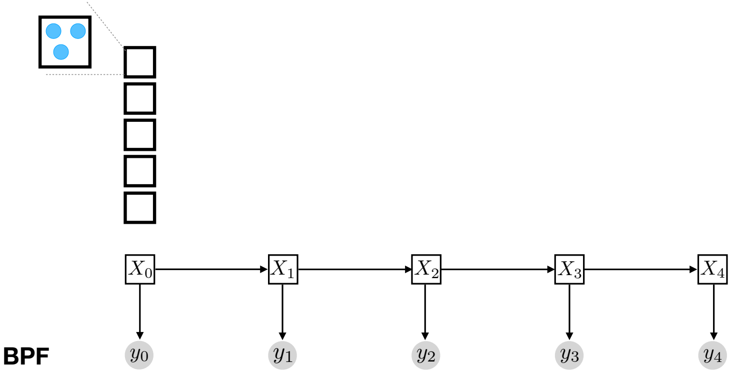

Weight $w(x_0) = g(y_0 \mid x_0)$ for each particle.

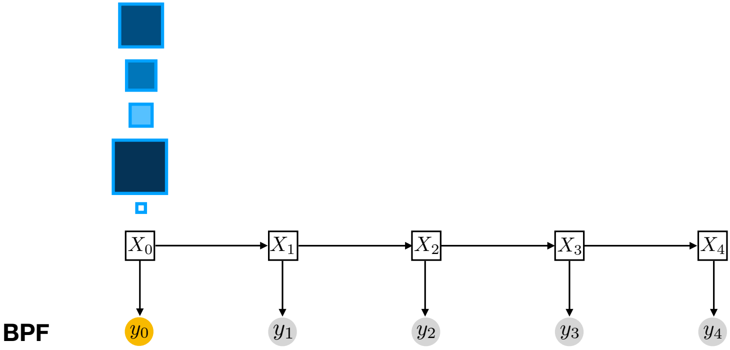

Resample according to weights, and then propose $x_1 \sim f(x_1 \mid x_0)$ for each particle.

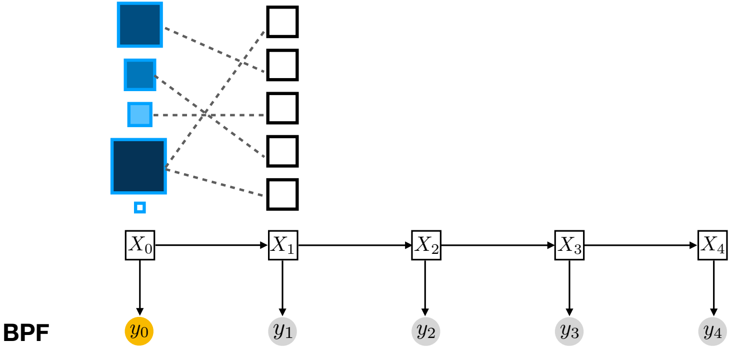

Weight $w(x_1) = g(y_1 \mid x_1)$ for each particle.

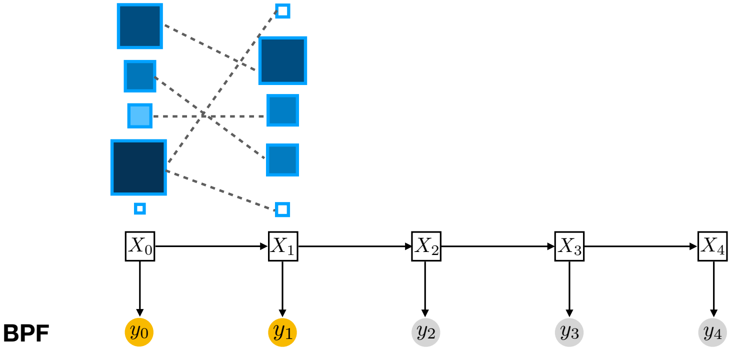

We can keep going until time T.

Bootstrap particle filter: curse of dimensionality

The bootstrap particle filter is not practical for large N because:

- need many particles to obtain small variance of $\hat{p}(y_{0:T})$,
- the marginal likelihood estimate might collapse to zero:

$$
\begin{aligned}
w(x_t) = g(y_t \mid x_t) &= \mathrm{Bin}(y_t; I(x_t), \rho) \\
&= \mathbb{1}\left(y_t \leq I(x_t)\right) \binom{I(x_t)}{y_t} \rho^{y_t} (1-\rho)^{I(x_t)-y_t}.
\end{aligned}
$$

Outline

1. Scientific problem: inference in agent-based SIS model.
   - motivation,
   - hidden Markov model.
2. Computational challenge: design of sequential Monte Carlo algorithm.
   - intractable likelihood,
   - **bootstrap particle filter is inefficient**.

Bootstrap particle filter: how to improve?

Three elements of sequential Monte Carlo samplers:

1. sampling from the proposal,
2. computing the weight,
3. resampling scheme.

Some proposals are better than the others, and we will improve the particle filters by **designing good proposals**.

picture credit: unsplash.com

- Bootstrap particle filter (BPF) proposes from $f(x_t \mid x_{t-1})$ and knows the past;
- Auxiliary particle filter (APF) knows the past and the present;
- Controlled sequential Monte Carlo (cSMC) knows the past, the present, and the future.

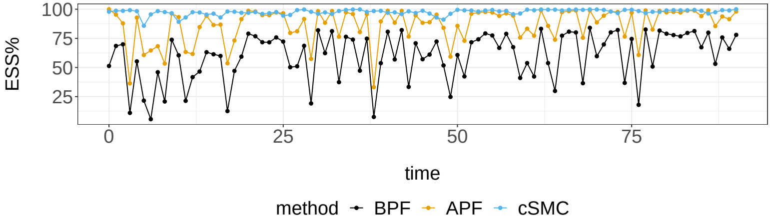

More information $\rightarrow$ better performance. But what's the price?

Auxiliary particle filter

The **one-step look-ahead proposal** of APF:

$$p(x_t \mid x_{t-1}, y_t) \propto f(x_t \mid x_{t-1}) g(y_t \mid x_t).$$

BPF proposal:

$$f(x_t \mid x_{t-1}).$$

(Pitt & Shephard, 1999; Chen et al., 2000; Johansen & Doucet, 2008)

Propose $x_0 \sim p(x_0 \mid y_0)$ for each particle.

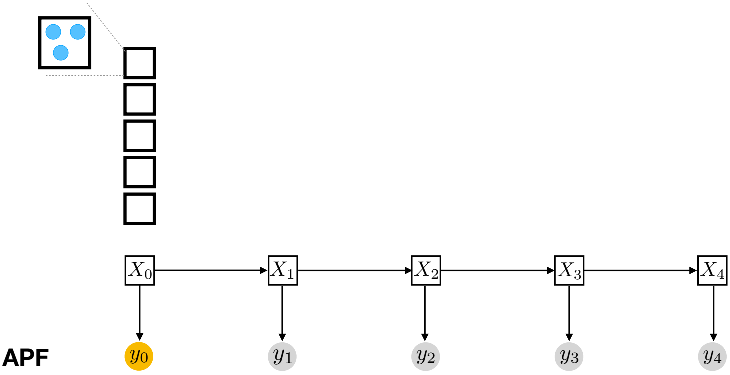

Propose $x_0 \sim p(x_0 \mid y_0)$ for each particle.

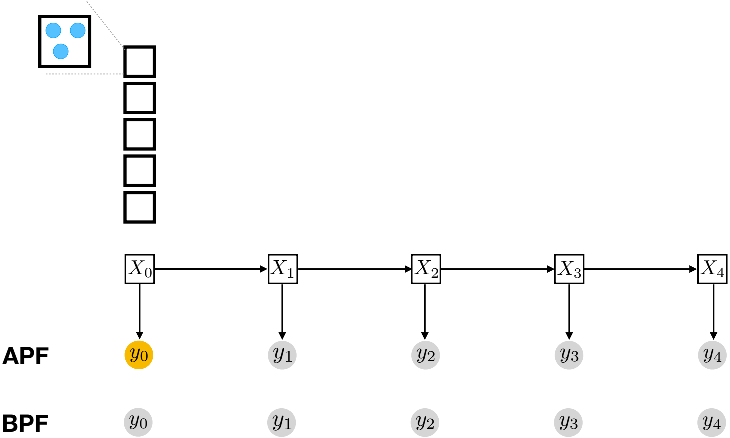

Weight $w(x_0) = g(y_1 \mid x_0)$ for each particle.

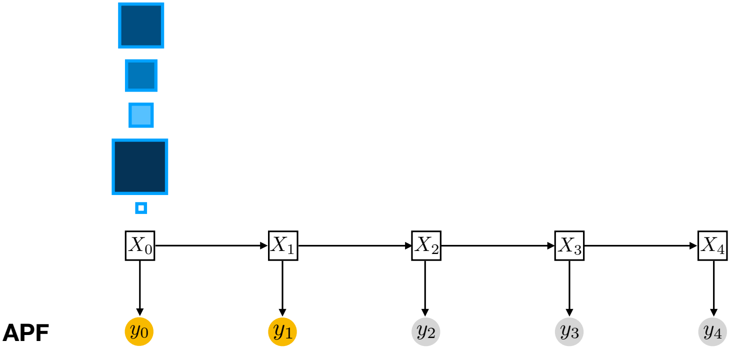

Resample according to weights, and then propose $x_1 \sim f(x_1 \mid x_0, y_1)$ for each particle.

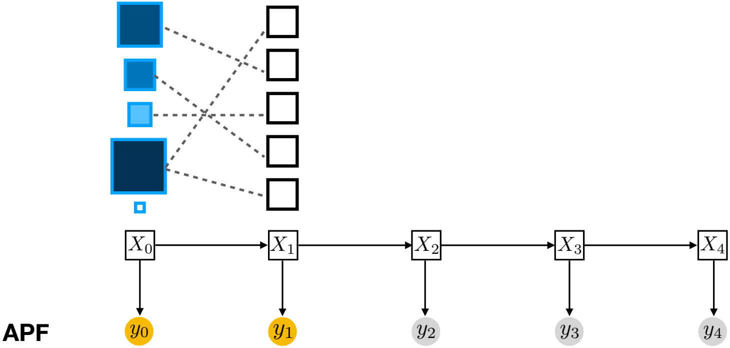

Weight $w(x_1) = g(y_2 \mid x_1)$ for each particle.

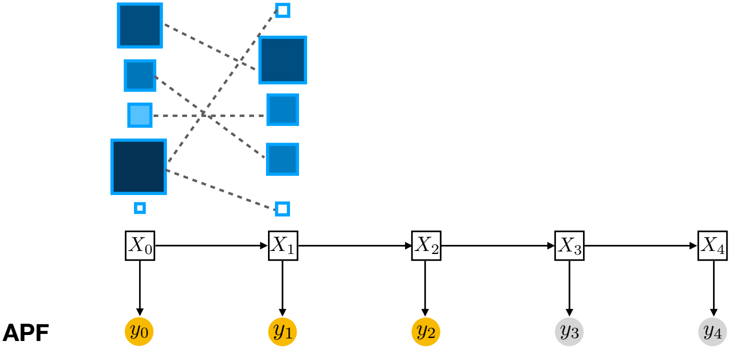

Auxiliary particle filter: the crystal ball

proposal

$$p(x_t \mid x_{t-1}, y_t)$$

weight

$$w(x_t) = p(y_{t+1} \mid x_t)$$

the **auxiliary variable**

$$I(x_t) = \sum_{n=1}^{N} x_t^n.$$

photo credit: harrypotter.fandom.com

Auxiliary variable: the wand

The weight is

$$p(y_{t+1} \mid x_t) = \sum_{x_{t+1} \in \{0,1\}^N} g(y_{t+1} \mid x_{t+1}) \; f(x_{t+1} \mid x_t)$$

$$= \sum_{I(x_{t+1})=0}^{N} p(y_{t+1} \mid I(x_{t+1})) \; p(I(x_{t+1}) \mid x_t).$$

I is the magical wand:
size of enumeration from $2^N \searrow N+1$.

Auxiliary variable: the wand

What is the distribution of $I(X_{t+1}) \mid x_t$?

1. $I(X^n_{t+1}) = \sum_{n=1}^{N} X^n_{t+1}$;
2. $X^n_{t+1} \mid X_t = x_t \sim \text{Bern}(\alpha^n(x_t))$ independently.

The sum of N independent Bernoulli variables with non-identical success probabilities follows a **Poisson binomial** distribution.

$$p(I(X_{t+1}) \mid x_t) = \text{PoisBin}(I(x_{t+1}); \alpha(x_{t-1})).$$

We can compute Poisson binomial PMF in $\mathcal{O}(N^2)$ steps with dynamic programming.

(Chen & Liu, 1997)

APF weight

The weight is $w(x_t) = p(y_{t+1} \mid x_t)$ and

$$
\begin{aligned}
p(y_{t+1} \mid x_t) &= \sum_{I(x_{t+1})=0}^{N} p(y_{t+1} \mid I(x_{t+1}))\ p(I(x_{t+1}) \mid x_t) \\
&= \sum_{I(x_{t+1})=0}^{N} \mathrm{Bin}(y_{t+1}; I(x_{t+1}), \rho)\ \mathrm{PoisBin}(I(x_{t+1}); \alpha(x_t)).
\end{aligned}
$$

We can compute $p(y_{t+1} \mid x_t)$ in $\mathcal{O}(N^2)$ steps.

APF proposal

Remember that weight and proposal go hand in hand.

compute the weight

$$p(y_{t+1} \mid x_t).$$

sample from the proposal

$$p(x_t \mid x_{t-1}, y_t).$$

APF proposal: a decomposition

The proposal is $p(x_t \mid x_{t-1}, y_t)$.

With the auxiliary variable $i_t = I(x_t)$, it becomes

$$p(x_t, i_t \mid x_{t-1}, y_t) = p(i_t \mid x_{t-1}, y_t)\ p(x_t \mid x_{t-1}, i_t).$$

First part

$$p(i_t \mid x_{t-1}, y_t) = \frac{\mathrm{PoisBin}(i_t; \alpha(x_{t-1}))\mathrm{Bin}(y_t; i_t, \rho)}{p(y_t \mid x_{t-1})}.$$

Using Poisson binomial PMF, we can sample $i_t \mid y_t, x_{t-1}$ in $\mathcal{O}(N^2)$ steps.

APF proposal

Second part has density

$$p(x_t \mid x_{t-1}, i_t) = \mathbb{I}(I(x_t) = i_t) \frac{\prod_{n=1}^{N} (\alpha^n(x_{t-1}))^{x_t^n} (1 - \alpha^n(x_{t-1}))^{1-x_t^n}}{\mathrm{PoisBin}(i_t; \alpha(x_{t-1}))}.$$

Conditional Bernoulli distribution:

$$\text{if } Z^n \sim \mathrm{Bern}(\alpha^n) \quad \text{independently,} \quad \text{then } Z \,\Big|\, \sum_{n=1}^{N} Z^n = i \sim ?$$

We can sample exactly from the **conditional Bernoulli distribution** in $\mathcal{O}(N^2)$ steps.

(Chen & Liu, 1997)

Reducing cost of APF

- We have seen that the proposal and weight steps can be performed exactly in $\mathcal{O}(N^2)$ operations.

- Can this be even faster?

Translated Poisson approximation

The weights $w(x_t) = p(y_{t+1} \mid x_t)$ are Poisson binomial PMFs.

Translated Poisson approximation utilizes ‘moment-matching’.

This approximation gets better as N increases:

$\|\text{PoisBin} - \text{TP}\|_{\text{TV}} \leq c_\alpha / \sqrt{N}$.

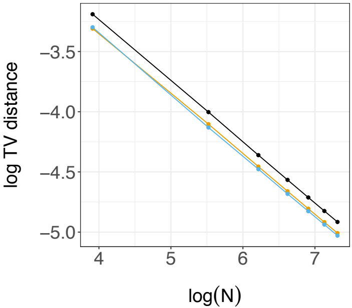

(Barbour & Ćekanavićius, 2002)

Conditional Bernoulli distribution

We can employ Markov chain Monte Carlo (MCMC) to sample from conditional Bernoulli distributions $\mathrm{CondBern}(\alpha, I)$.

Swap move has constant cost per iteration.

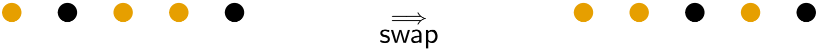

Does this chain converge fast enough when $I \propto N$?

Conditional Bernoulli distribution

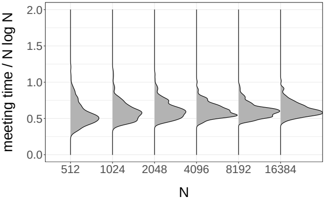

- Prove convergence via coupling;
- Combine path coupling and a partition of the state space.

Theorem from Heng, Jacob & Ju, 2020 (under review at Biometrika)

With probability at least $1 - \exp(-\nu N)$, we have

$$\|z^{(t)} - \mathrm{CondBern}(\alpha, I)\| \leq \varepsilon \text{ for all } t \geq \kappa N \log(N/\varepsilon).$$

APF

Auxiliary variable

$$i_t = I(x_t) = \sum_{n=1}^{N} x_t^n.$$

1. intermediate step in sampling

$$p(x_t, i_t \mid x_{t-1}, y_t);$$

2. instrumental variable in marginalization

$$p(y_t \mid x_{t-1}) = \sum_{i_t=0}^{N} p(y_t \mid i_t) p(i_t \mid x_{t-1}).$$

From APF to cSMC

One-step look-ahead is still a local strategy.

We want an ‘all-step’ look-ahead proposal.

photo credit: harrypotter.fandom.com

Controlled SMC: the hall of prophecy

The 'best' proposal takes **all future observations** into account: the **smoothing distribution**

$$p(x_t \mid x_{t-1}, y_{t:T}).$$

Controlled SMC: the hall of prophecy

Propose $x_0 \sim q(x_0 \mid y_{0:T}) \approx p(x_0 \mid y_{0:T})$ for each particle.

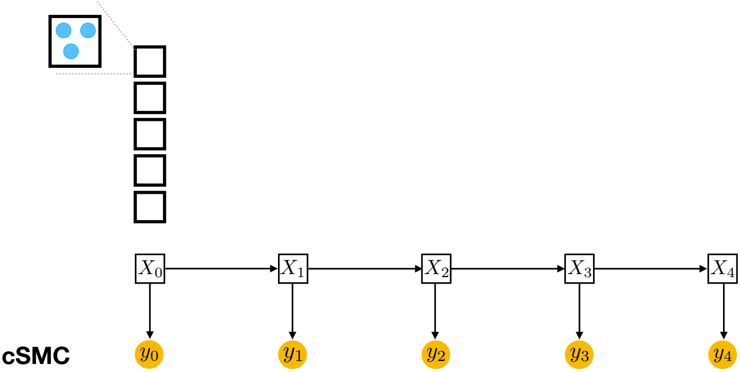

Controlled SMC: the hall of prophecy

Propose $x_0 \sim q(x_0 \mid y_{0:T}) \approx p(x_0 \mid y_{0:T})$ for each particle.

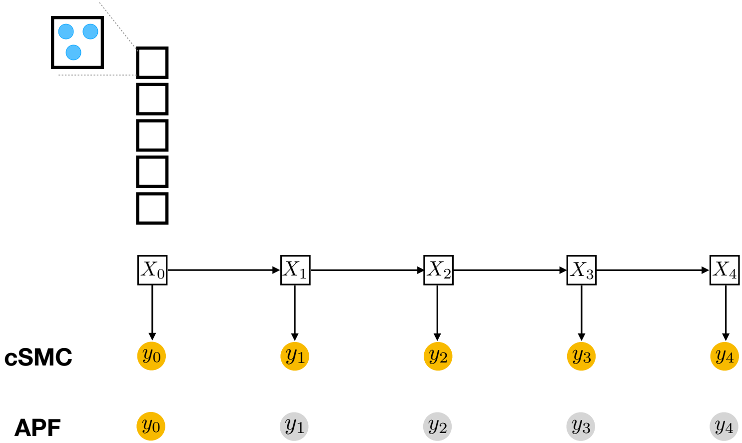

Controlled SMC: the hall of prophecy

Ideal proposal is the **smoothing distribution**

$$p(x_t \mid x_{t-1}, y_{t:T}) \propto f(x_t \mid x_{t-1}) \underbrace{g(y_t \mid x_t) p(y_{t+1:T} \mid x_t)}_{=p(y_{t:T}|x_t):=\psi_t^\star(x_t)}.$$

Construct **approximate** proposal

$$q(x_t \mid x_{t-1}, y_{t:T}) \propto f(x_t \mid x_{t-1}) \psi_t(I(x_t)),$$

through approximation of $\psi_t^\star(x_t)$.

(Heng at al., 2020; Guarniero, Johansen & Lee, 2017)

Backward information filter

Backward information filter (BIF) $\quad \psi_t^\star(x_t) = p(y_{t:T} \mid x_t).$

Recursive definition

$$\psi_t^\star(x_t) = g(y_t \mid x_t) \underbrace{\sum_{x_{t+1} \in \{0,1\}^N} \psi_{t+1}^\star(x_{t+1}) f(x_{t+1} \mid x_t)}_{f(\psi_{t+1}^\star(x_{t+1})|x_t)} .$$

We seek an **recursive approximation**.

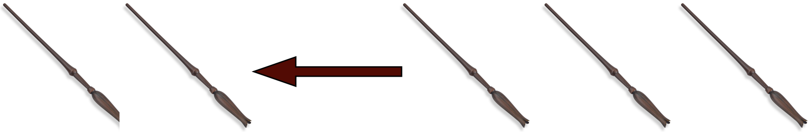

Approximating the conditional expectation

$$\begin{tikzcd}
f(\psi^{\star}_{t+1}(x_{t+1}) \mid x_t) \arrow[d, dotted, no head] & f(\cdot \mid x_t) \arrow[l] \arrow[d, dotted, no head] & \alpha(x_t) \arrow[l] \arrow[d, dotted, no head] \\
\overline{f}(\psi_{t+1}(i_{t+1}) \mid i_t) & \overline{f}(\cdot \mid i_t) \arrow[l] & \overline{\alpha}(x_t) \arrow[l]
\end{tikzcd}$$

$$\overline{\alpha}^{n}(x_t) = \begin{cases} \overline{\lambda} N^{-1} I(x_t), & \text{if } x_t^n = 0 \ (\mathrm{S} \to \mathrm{I}), \\ 1 - \overline{\gamma}, & \text{if } x_t^n = 1 \ (\mathrm{I} \to \mathrm{I}). \end{cases}$$

$$\overline{\lambda} = N^{-1} \sum_{n=1}^{N} \lambda^n, \quad \overline{\gamma} = N^{-1} \sum_{n=1}^{N} \gamma^n.$$

$\overline{\alpha}$ is in fact a function of $i_t \in [0 : N]$.

Recursive approximation

$$
\begin{tikzcd}
\psi_t^\star(x_t) \arrow[d, dotted, no head] & f(\psi_{t+1}^\star(x_{t+1}) \mid x_t) \arrow[l] \arrow[d, dotted, no head] & f(\cdot \mid x_t) \arrow[l] \arrow[d, dotted, no head] & \alpha(x_t) \arrow[l] \arrow[d, dotted, no head] \\
\psi_t(i_t) & \overline{f}(\psi_{t+1}(i_{t+1}) \mid i_t) \arrow[l] & \overline{f}(\cdot \mid i_t) \arrow[l] & \overline{\alpha}(x_t) \arrow[l]
\end{tikzcd}
$$

$$\psi_T(i_T) = \mathrm{Bin}(y_T; i_T, \rho), \quad \psi_t(i_t) = \mathrm{Bin}(y_t; i_t, \rho)\bar{f}(\psi_{t+1} \mid i_t).$$

Approximating the backward information filter costs $\mathcal{O}(N^3 T)$ to compute and $\mathcal{O}(NT)$ in storage.

cSMC: weight and sampling

After the recursive approximation steps, cSMC takes essentially the same sampling steps as in APF and has same cost of $\mathcal{O}(N^2)$ per time step.

Proposal

$$q_t(x_t, i_t \mid x_{t-1}, \theta) = \frac{f(x_t \mid x_{t-1})\psi_t(i_t)}{f(\psi_t \mid x_{t-1})}.$$

Weight

$$w_t(x_t) = \frac{g(y_t \mid x_t) f(\psi_{t+1} \mid x_t)}{\psi_t(I(x_t))}.$$

More information

Proposals to sample particles for x_0 in the three SMC methods:

BPF	$\mu(x_0)$
APF	$f(x_0 \mid y_0)$
cSMC	$q(x_0 \mid y_{0:T}) \approx p(x_0 \mid y_{0:T})$

BPF	y_0	y_1	y_2	y_3	y_4
APF	y_0	y_1	y_2	y_3	y_4
cSMC	y_0	y_1	y_2	y_3	y_4

More information, better performance & higher cost

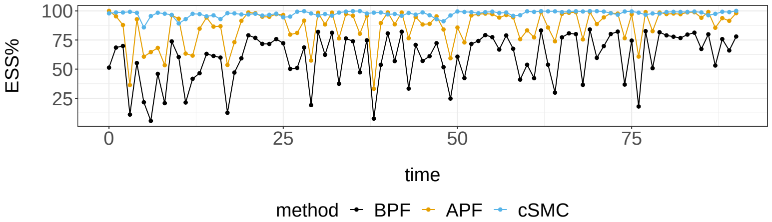

method	use of $y_{0:T}$ in proposal	run-time
BPF	past	$\mathcal{O}(NTP)$
APF	past and present	$\mathcal{O}(N^2TP)$
cSMC	past, present and future	$\mathcal{O}(N^2TP) + \mathcal{O}(N^3T)$

SMC comparison: effective sample size

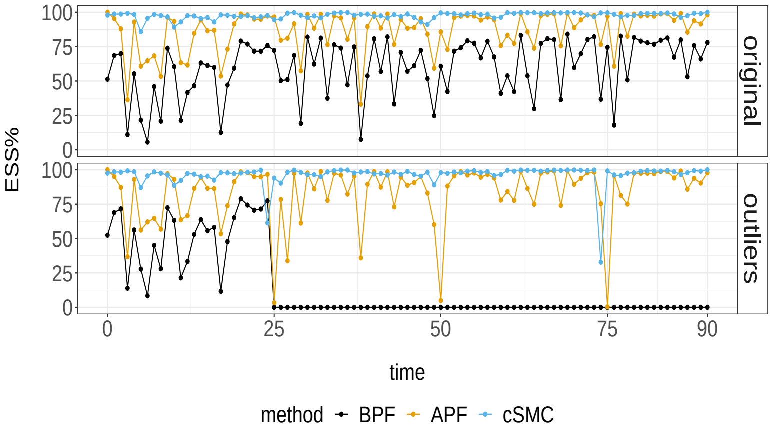

Bottom panel: replace observation at $t \in \{25, 50, 75\}$ by $2y_t$.

SMC comparison: variance and run-time

Variance of $\log \hat{p}_\theta(y_{0:T})$ at data generating parameter $\theta^\star$:

	BPF		APF		cSMC	
P	Var	Run-time	Var	Run-time	Var	Run-time
64	4.32	0.09	0.281	1.49	0.0696	1.46
128	2.39	0.17	0.154	2.95	0.0285	2.88
256	1.67	0.33	0.110	5.85	0.0164	5.72
512	0.88	0.63	0.056	11.72	0.0087	11.41
1024	0.55	1.25	0.026	23.46	0.0049	22.83
2048	0.31	2.49	0.011	47.48	0.0020	45.97

$N = 100$, $T = 90$, heterogeneous agents with $d = 2$.

SMC comparison: variance

Variance of $\log \hat{p}_\theta(y_{0:T})$ at $\beta_\lambda \neq \beta_\lambda^\star$:

	BPF	APF	cSMC
64	-	71.88	13.52
128	-	39.52	8.62
256	-	26.86	4.11
512	-	18.98	3.57
1024	-	13.38	2.05
2048	-	9.93	1.15

data generating parameters (DGP) is $\beta_\lambda^\star = (-1, 2)$ and non-DGP $\beta_\lambda = (-3, 0)$.
$N = 100$, $T = 90$, heterogeneous agents with $d = 2$.

Numerical illustration: log-likelihood

$$\log \hat{p}_{(\beta^1_\lambda, \beta^2_\lambda)}.$$

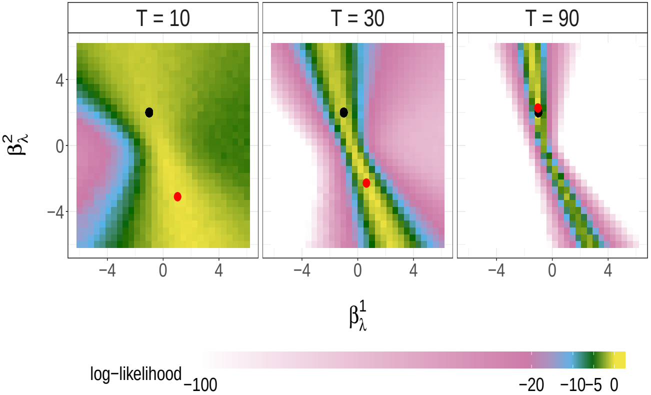

(black dot: data generating parameter, red dot: MLE.)

Numerical illustration: log-likelihood

$$\log \hat{p}_{(\beta_\lambda^2, \beta_\gamma^2)}.$$

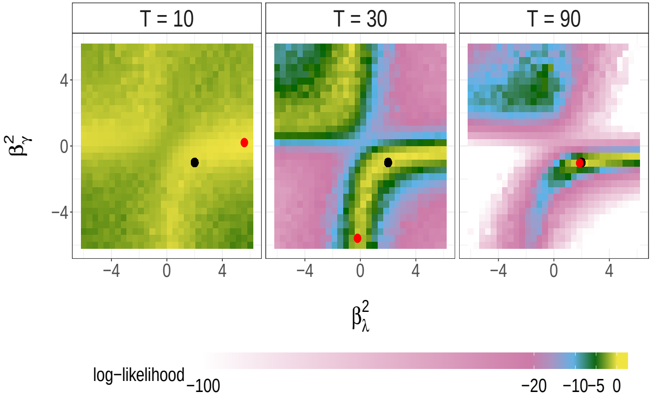

(black dot: data generating parameter, red dot: MLE.)

Summary

In the talk:

- agent-based models for transmission of infectious diseases,
- sequential Monte Carlo algorithms: APF, cSMC.

In the manuscripts:

- APF and cSMC for the SIR model,
- theoretical support of cSMC and BIF approximations,
- Bayesian parameter inference using particle Markov chain Monte Carlo.

Contributions

- Ju, N., Heng, J., and Jacob, P. E. (2021). **Sequential Monte Carlo algorithms for agent-based models of disease transmission**. *arXiv preprint*.

- Zhao, Q., Ju, N., Bacallado, S., and Shah, R. (2020). **BETS: The dangers of selection bias in early analyses of the coronavirus disease (COVID-19) pandemic**. *Annals of Applied Statistics*.

- Heng, J., Jacob. P.E., and Ju,N. (2020). **A simple Markov chain for independent Bernoulli variables conditioned on their sum**. *Under review*.

Discussions

- Model extensions: negative binomial noise; observations at regular intervals; observing difference in infection counts.
- Statistical properties: parameter identifiability and estimation consistency; choice of prior.
- Model comparison: equation-based vs. agent-based model, bias-variance trade-off.
- Network: unknown network; dynamic network.
- And more!

Concluding remarks

不积跬步 无以至千里
不积小流 无以成江海
– 荀子《劝学》
'Unless you collect
little steps,
you can never journey
a thousand miles;
Unless you gather
tiny streams,
you can never make
a river or a sea.'
– Xun Zi.

Acknowledgment

Pierre Jacob, Harvard

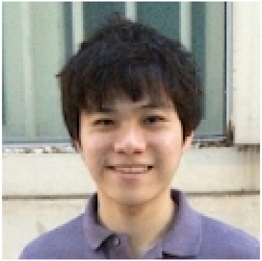

Jeremy Heng, ESSEC Singapore

Manuscripts:	agents: arXiv:2101.12156
	CondBern: arXiv:2012.03103 (*under review*)
	Covid: arXiv:2004.07743 (*AoAS*)
Code:	`https://github.com/nianqiaoju/agents`
Slides:	`https://nianqiaoju.github.io`

References I

Zhao, Q., Ju, N., Bacallado, S., and Shah, R. (2020). BETS: The dangers of selection bias in early analyses of the coronavirus disease (COVID-19) pandemic. *Annals of Applied Statistics*.

Heng, J., Jacob. P.E., and Ju,N. (2020). A simple Markov chain for independent Bernoulli variables conditioned on their sum. *arXiv preprint*.

Ju, N., Heng, J., and Jacob, P. E. (2021). Sequential Monte Carlo algorithms for agent-based models of disease transmission. *arXiv preprint*.

References II

Hoertel, N., Blachier, M., Blanco, C. et al.(2020). A stochastic agent-based model of the SARS-CoV-2 epidemic in France. *Nature Medicine*.

Chen, S.X. and Liu, J.S. (1997) Statistical applications of the Poisson-Binomial and conditional Bernoulli distributions. *Statistica Sinica*.

Barbour, A. D. and Ćekanavićius, V. (2002) Total variation asymptotics for sums of independent integer random variables. *Annals of Probability*.

Chen, R., Wang, X. and Liu, J. S. (2000). Adaptive joint detection and decoding in flat-fading channels via mixture Kalman filtering. *IEEE Trans. Inform. Theory*.

Pitt, M.K. and Shephard N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American statistical association*.

Johansen A.M. and Doucet A. (2008). A note on auxiliary particle filters. *Statistics & Probability Letter*.

Guarniero, P., Johansen, A.M., and Lee, A. (2017). The iterated auxiliary particle filter. *Journal of the American Statistical Association*.

Heng, J., Bishop, A.N., Deligiannidis, G. and Doucet, A. (2020). Controlled sequential Monte Carlo. *Annals Statistics*.

Bresler Y. (1986). Two-filter formulae for discrete-time non-linear Bayesian smoothing. *International Journal of Control*

Individual reproductive number

If all agents are the same $\lambda^n = \lambda, \gamma^n = \gamma$ for all $n \in [1 : N]$.

Basic reproductive number

$$R_0 = \lambda/\gamma.$$

If agents are heterogeneous, we define $R_0^n = \lambda^n/\gamma^n$.

Individual reproductive number