III Simulation-based Inference (SBI)

setting: X: observables, i.e. variables we know or can measure

Y hidden properties / states: variables we would like to know, but cannot measure

assumptions

① hidden variables are more fundamental, e.g. the Y cause the X

② we have a scientific theory how the X arise from the Y "forward process"

③ theory is implemented as an algorithm $\hat{=}$ computer simulation

$\Rightarrow$ can do "in-silico experiments" (as opposed to in-vivo and in-vitro)

three types of variables in a simulation:

- Y: inputs, X: outputs, optionally: η: random numbers for non-deterministic simulation
- deterministic $X = \phi(Y)$ simulation prog. $\Rightarrow p^S(X \mid Y) = \delta(X - \phi(Y))$
- non-deterministic $X = \phi(Y, \eta)$ "noise outsourcing" $\Rightarrow p^S(X \mid Y) = \phi_{\#}\, p^S(Y, \eta)$

simulation paradigm: - If we knew Y, we could predict X.

- Since we don't know Y, try multiple Y and generate alternative scenarios.

"worst case", "best case", "typical case"

SBI should improve upon this state-of-the-art.

two important special cases (structured Ys)

- mixed effects model

$$Y = \begin{cases} Y_G : & \text{hidden global properties, same for all members of a group/population} \\ Y_{Li} : & \text{hidden local properties, different for every individual/instance} \end{cases}$$

hierarchical model: • $Y_G \sim p^S(Y_G)$, for $i = 1, \dots, N$: $Y_{Li} \sim p^S(Y_L \mid Y_G)$

$X_i \sim p^S(X \mid Y_G, Y_{Li})$

task: find $p^S(Y \mid X)$

$\Rightarrow X_i \not\perp X_{i'}$, but $X_i \perp X_{i'} \mid Y_G$ (don't use the naive i.i.d. assumption)

- dynamical systems (time-dependent behavior)

$$Y = \begin{cases} Y_G & \text{: global hidden properties, do not change over time (in a given system)} \\ S_t & \text{: hidden state at time } t \end{cases}$$

- $Y_G \sim p^S(Y_G)$, $S_0 \sim p^S(S_0 \mid Y_G)$ initial state
- for $t = 1, \dots, T$: $S_t \sim p^S(S \mid Y_G, S_{<t})$

  $X_t \sim p^S(X \mid Y_G, S_{\leq t})$

- if t is continuous: differential equation, stochastic differential equation

  t discrete: hidden Markov model

  $X_t \not\perp X_{t'}$, but $X_t \perp X_{t'} \mid S_t$ if Markov property fulfilled

Main tasks of SBI:

1. surrogate modeling: train a model $p(X|Y)$ that emulates simulation $X = \phi(Y, \eta)$
   - speed-up: $X = \phi(Y, \eta)$ is often slow $\Rightarrow$ replace with a fast generative model such that $p(X|Y) \approx p^S(X|Y)$
   - forward inference: often, $X = \phi(Y, \eta)$ only defines $p^S(X|Y) = \phi_\# p^S(Y, \eta)$ implicitly, but doesn't allow to calculate $p^S(X = x \mid Y)$
     "likelihood-free inference, implicit likelihood"
     approximate true likelihood by $p(X|Y) \approx p^S(X|Y)$
2. inverse inference: "run the simulation backwards $Y = \phi^{-1}(X)$"
   usually intractable (no analytic solution) and/or ill-posed (no unique solution, non-invertible due to information loss $Y \rightarrow X$)
   - classical solution: pick a single solution ("best") using constraints & regularization
   - SBI solution: probabilistic treatment via Bayes rule: $p^S(Y|X) = p^S(Y)\, p^S(X|Y) / p^S(X)$

define equivalence classes $F(X) = \{ Y : \exists \eta \text{ with } \phi(Y,\eta) = X \}$

all Ys that could have produces given X

posterior: $p^s(Y/X)$ assigns a "plausibility" to every $Y \in F(X)$

$F(X) = \phi^{-1}(X)$ often an infinite set

$\Rightarrow$ focus on high-prob. members

problems: - if likelihood $p^s(X|Y)$ is only implicitly defined, Bayes rule cannot be calculated

- even if $p^s(X/Y)$ (or a surrogate $p(X/Y)$ is known, Bayes rule is usually intractable

$\Rightarrow$ learn generative model for posterior $p(Y/X) \approx p^s(Y|X)$

model misspecification & outlier detection

- simulation is <u>not</u> reality $\underbrace{p^s(Y) \cdot p^s(X|Y)}_{\text{simulation}} \approx \underbrace{p^*(Y)\, p^*(X|Y)}_{\text{truth in real life}}$

$\Rightarrow$ use SBI to detect if $p^s(X,Y) \neq p^*(X,Y)$

- is an observed outcome $X^{obs} \sim p^*(X,Y)$ compatible with $p^s(X,Y)$ ...

if not $\Rightarrow$ simulation is unrealistic, "simulation gap"

- is a set of outcomes $\{X_i^{obs}\}_{i=1}^N$ compatible with $p^s(X, Y)$?

(4) model comparison and selection:

if we have competing theories $X = \phi^{(l)}(Y^{(l)}, \eta^{(l)})$

$\Rightarrow$ determine which l describes X^{obs} or $\{X_i^{obs}\}_{i=1}^N$ best, if any?

(5) digital twins: in a mixed effects setting, given $\{X_i^{obs}\}_{i=1}^N$

determine Y_G and Y_{Li} accurately enough to predict

$$X_i^{future} = \phi(Y_G, Y_{Li}, \eta) \quad \text{accurately}$$

classical: base treatment decisions mainly on Y_G ("treatment guidelines")

(after appropriate stratification of population into subgroups)

desired: "precision medicin": use Y_G and Y_{Li}

(6) experimental design and active learning: How should we measure $\{X_i^{obs}\}_{i=1}^N$

to learn as much as possible about Y with given experimental budget