Statistical Model Checking and Rare Events

Paolo Zuliani

Joint work with Edmund M. Clarke
Computer Science Department, CMU

Probabilistic Verification

- Verification of stochastic system models via statistical model checking

- Temporal logic specification:
  - "the amount of p53 exceeds 10^5 within 20 minutes"

- If Φ = "p53 exceeds 10^5 within 20 minutes"

Probability (Φ) = ?

Equivalently

- A biased coin (Bernoulli random variable):
  - Prob (Heads) = p Prob (Tails) = $1-p$
  - p is unknown
- Question: What is p?
- A solution: flip the coin a number of times, collect the outcomes, and use statistical estimation

Statistical Model Checking

Key idea

(Haakan Younes, 2001)

- System behavior w.r.t. property Φ can be modeled by a Bernoulli random variable of parameter p:
    - System satisfies Φ with (unknown) probability p
- Question: What is p?
- Draw a sample of system simulations and use:
    - Statistical estimation: returns "p in interval (a,b)" with high probability

Statistical Model Checking

- Statistical Model Checking is a Monte Carlo method
- Problems arise when p is very small (rare event)
- The number of simulations (coin flips) needed to estimate p accurately grows too large
- Need to deal with this …

Rare events

- Estimate Prob($X \geqslant t$) = p_t, when p_t is *small* (say 10^{-9})

Rare events

- Estimate Prob(X$\geqslant t$) = p_t, when p_t is *small* (say 10^{-9})
- Standard (Crude) Monte Carlo: generate K i.i.d. samples of X; return the estimator e_K

$$e_K = \frac{1}{K}\sum_{i=1}^{K} I(X_i \geqslant t) = \frac{k_t}{K}$$

- Prob ($e_K \rightarrow p_t$) = 1 for $K \rightarrow \infty$ (strong law LN)

Rare events

- E[e_K] = p_t
- Var[e_K] = $\dfrac{p_t(1-p_t)}{K}$

Rare events

- $E[e_K] = p_t$
- $Var[e_K] = \dfrac{p_t(1-p_t)}{K}$
- By the Central Limit Theorem (CLT), the distribution of e_K converges to a normal distribution with:
  - mean p_t
  - variance $\dfrac{p_t(1-p_t)}{K}$
- Relative Error (RE) = $\dfrac{\sqrt{\mathrm{var}[e_K]}}{\mathrm{E}[e_K]} = \dfrac{\sqrt{p_t(1-p_t)}}{p_t\sqrt{K}}$

Rare events

- RE = $\dfrac{\sqrt{p_t(1-p_t)}}{p_t\sqrt{K}}$

- Fix K, then RE is unbounded as $p_t \to 0$

- More accuracy $\to$ more samples

- Want confidence interval of relative accuracy δ and coverage probability c, i.e., estimate e_K must satisfy:

$$\text{Prob}(|\, e_K - p_t \,| < \delta \cdot p_t) \geq c$$

- How many samples do we need?

Rare events

- From the CLT, a 99% (approximate) confidence interval of relative accuracy δ needs about

$$K \approx \frac{1-p_t}{p_t \delta^2} \text{ samples}$$

Thus, $\text{Prob}(|\, e_K - p_t \,| < \delta p_t) \approx 0.99$

Rare events

- From the CLT, a 99% (approximate) confidence interval of relative accuracy δ needs about

$$K \approx \frac{1-p_t}{p_t \delta^2} \text{ samples}$$

Thus, Prob$(|\, e_K - p_t \,| < \delta p_t) \approx 0.99$

- Examples:
  - $p_t = 10^{-9}$ and $\delta = 10^{-2}$ (*ie*, 1% relative accuracy) we need about 10^{13} samples!!
  - Bayesian estimation requires about 6×10^6 samples with $p_t = 10^{-4}$ and $\delta = 10^{-1}$

A solution

- Importance Sampling (1940s)
- A variance-reduction technique
- Can result in dramatic reduction in sample size

Importance Sampling

- The fundamental Importance Sampling identity

$$
\begin{aligned}
p_t &= E[I(X \geqslant t)] \\
&= \int I(x \geqslant t) f(x)\, dx \\
&= \int I(x \geqslant t) \frac{f(x)}{f_*(x)} f_*(x)\, dx \\
&= \int I(x \geqslant t) W(x) f_*(x)\, dx \\
&= E_*[I(X \geqslant t) W(X)]
\end{aligned}
$$

f is the density of *X*

Importance Sampling

- The fundamental Importance Sampling identity

$$
\begin{aligned}
p_t &= E[I(X \geqslant t)] \\
&= \int I(x \geqslant t) f(x)\, dx \\
&= \int I(x \geqslant t) \boxed{\frac{f(x)}{f_*(x)}} f_*(x)\, dx \\
&= \int I(x \geqslant t) W(x) f_*(x)\, dx \\
&= E_*[I(X \geqslant t) W(X)]
\end{aligned}
$$

likelihood ratio

f is the density of *X*

Importance Sampling

- Estimate p_t= E[$X \geqslant t$] = Prob($X \geqslant t$)
- A sample X_1,… X_K iid as f
- The crude Monte Carlo estimator is

$$\hat{p}_t = \frac{1}{K} \sum_{i=1}^{K} I(X_i \geqslant t) = \frac{k_t}{K}, \qquad X_i \sim f$$

Importance Sampling

- Estimate $p_t = E[X \geqslant t] = \text{Prob}(X \geqslant t)$
- A sample $X_1, \ldots X_K$ iid as f
- The crude Monte Carlo estimator is

$$\hat{p}_t = \frac{1}{K} \sum_{i=1}^{K} I(X_i \geqslant t) = \frac{k_t}{K},$$

$$\boxed{X_i \sim f}$$

sampling from f

Importance Sampling

- Define a biasing density f_*
- Compute the IS estimator

$$\hat{p}_t = \frac{1}{K} \sum_{i=1}^{K} I(X_i \geqslant t) W(X_i), \qquad X_i \sim f_*$$

where $W(x) = \dfrac{f(x)}{f_*(x)}$ is the likelihood ratio

Importance Sampling

- Define a biasing density f_*

- Compute the IS estimator

$$\hat{p}_t = \frac{1}{K}\sum_{i=1}^{K} I(X_i \geqslant t) W(X_i), \qquad \boxed{X_i \sim f_*}$$

sampling from f_* !

where $W(x) = \dfrac{f(x)}{f_*(x)}$ is the likelihood ratio

Importance Sampling

- Need to choose a "good" biasing density (low variance)
- Optimal density: $$f_*(x) = \frac{I(x \geq t) f(x)}{p_t}$$

$$e_K = \frac{1}{K}\sum_{i=1}^{K} I(X \geq t) W(X) = \frac{1}{K}\sum_{i=1}^{K} I(X \geq t)\frac{f(X)}{f_*(X)}$$

$$= p_t \frac{1}{K}\sum_{i=1}^{K} I(X \geq t)\frac{f(X)}{I(X \geq t) f(X)} = p_t$$

- Zero variance! (But …)

Importance Sampling

- Need to choose a "good" biasing density (low variance)
- Optimal density: $$f_*(x) = \frac{I(x \geq t) f(x)}{p_t}$$

unknown

$$e_K = \frac{1}{K}\sum_{i=1}^{K} I(X \geq t) W(X) = \frac{1}{K}\sum_{i=1}^{K} I(X \geq t)\frac{f(X)}{f_*(X)}$$

$$= p_t \frac{1}{K}\sum_{i=1}^{K} I(X \geq t)\frac{f(X)}{I(X \geq t) f(X)} = p_t$$

- Zero variance! (But …)

Cross-Entropy Method
(R. Rubinstein)

- Suppose the density of X in a family of densities $\{f(\cdot\,;v)\}$
  - the "nominal" f is $f(x;u)$

- **Key idea**: choose a parameter v such that the *distance* between f_* and $f(\cdot\,;v)$ is minimal

- The Kullback-Leibler divergence (cross-entropy) is a measure of "distance" between two densities

- First used for rare event simulation by Rubinstein (1997)

Cross-Entropy Method

- The KL divergence (cross-entropy) of densities *g*, *h* is

$$D(g,h) = E_g\left[\ln \frac{g(X)}{h(X)}\right] = \int g(x)\ln g(x)dx - \int g(x)\ln h(x)dx$$

- $D(g,h) \geqslant 0$ (= 0 IFF $g = h$)
- $D(g,h) \neq D(h,g)$

Cross-Entropy Method

- The KL divergence (cross-entropy) of densities *g*, *h* is

$$D(g,h) = E_g\left[\ln \frac{g(X)}{h(X)}\right] = \int g(x)\ln g(x)dx - \int g(x)\ln h(x)dx$$

- $D(g,h) \geqslant 0$ (= 0 IFF $g = h$)
- $D(g,h) \neq D(h,g)$

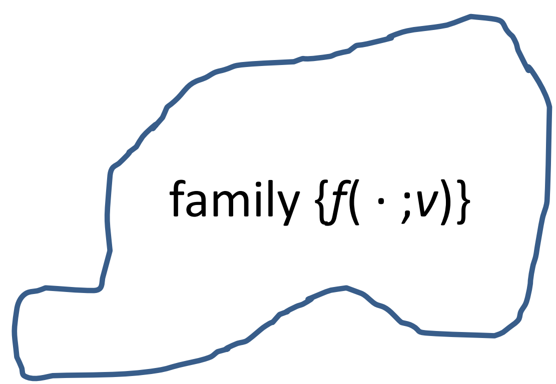

Cross-Entropy Method

- The KL divergence (cross-entropy) of densities *g*, *h* is

$$D(g,h) = E_g\left[\ln \frac{g(X)}{h(X)}\right] = \int g(x)\ln g(x)dx - \int g(x)\ln h(x)dx$$

- $D(g,h) \geqslant 0$ (= 0 IFF $g = h$)
- $D(g,h) \neq D(h,g)$

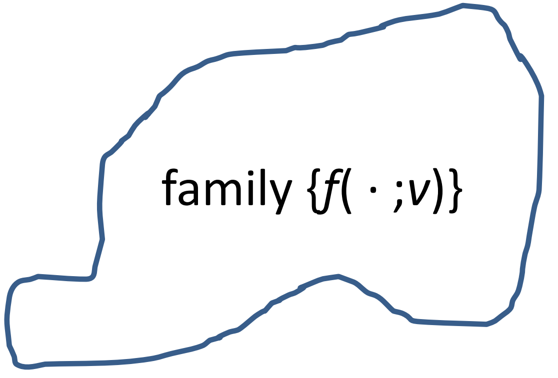

optimal density f_*

Cross-Entropy Method

- The KL divergence (cross-entropy) of densities g, h is

$$D(g,h) = E_g\left[\ln \frac{g(X)}{h(X)}\right] = \int g(x)\ln g(x)dx - \int g(x)\ln h(x)dx$$

- $D(g,h) \geqslant 0 \qquad (= 0 \quad \text{IFF} \quad g = h)$
- $D(g,h) \neq D(h,g)$

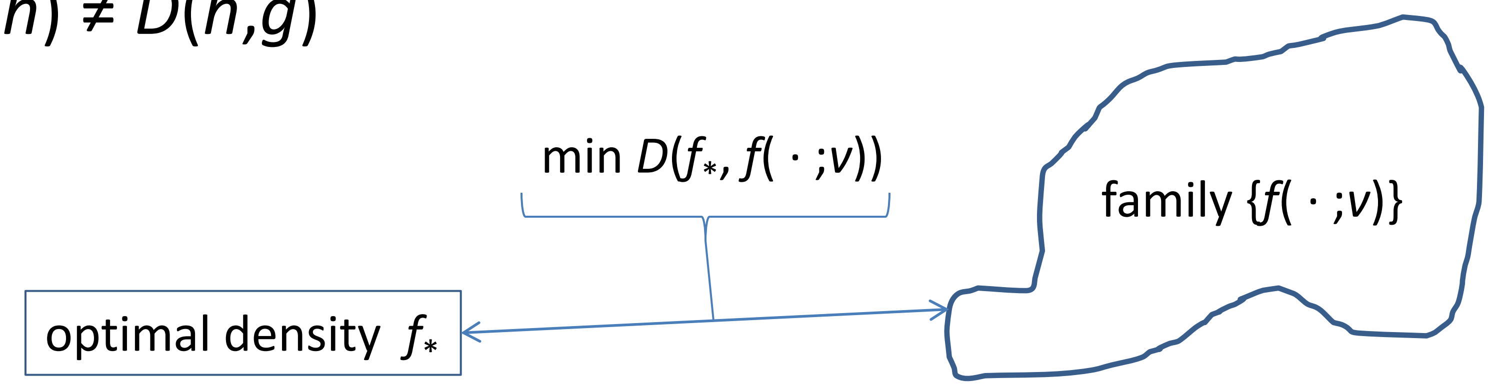

Cross-Entropy Method

- The Cross-Entropy Method has two basic steps

Cross-Entropy Method

- The Cross-Entropy Method has two basic steps

1. find $v_* = \underset{v}{\arg\min}\, D(f_*(\cdot), f(\cdot; v))$

Cross-Entropy Method

- The Cross-Entropy Method has two basic steps

1. find $v_* = \arg\min_{v} D(f_*(\cdot), f(\cdot; v))$

2. run importance sampling with biasing density $f(\cdot\,; v_*)$

Cross-Entropy Method

- The Cross-Entropy Method has two basic steps

1. find $v_* = \arg\min_{v} D(f_*(\cdot), f(\cdot; v))$

2. run importance sampling with biasing density $f(\cdot\,; v_*)$

Cross-Entropy Method

- The Cross-Entropy Method has two basic steps

1. find $v_* = \underset{v}{\arg\min}\, D(f_*(\cdot), f(\cdot; v))$
2. run importance sampling with biasing density $f(\cdot\,; v_*)$

- Step 2 is “easy”
- Step 1 is **not so easy**

Cross-Entropy Method

- Step 1:

$v_* =$

Cross-Entropy Method

- Step 1:

$$\boldsymbol{v}_* = \arg\min_{v} E_{f_*}\left[\ln\frac{f_*(X)}{f(X;v)}\right] = \arg\min_{v} \int f_*(x)\ln f_*(x)dx - \int f_*(x)\ln f(x;v)dx$$

Cross-Entropy Method

- Step 1:

$$v_* = \arg\min_{v} E_{f_*}\left[\ln \frac{f_*(X)}{f(X;v)}\right] = \arg\min_{v} \overbrace{\int f_*(x)\ln f_*(x)dx - \int f_*(x)\ln f(x;v)dx}^{\text{always} \geqslant 0}$$

Cross-Entropy Method

- Step 1:

$$v_* = \arg\min_{v} E_{f_*}\left[\ln \frac{f_*(X)}{f(X;v)}\right] = \arg\min_{v} \underbrace{\int f_*(x)\ln f_*(x)dx - \int f_*(x)\ln f(x;v)dx}^{\text{always} \geqslant 0}$$

$$= \arg\max_{v} \int f_*(x)\ln f(x;v)dx$$

Cross-Entropy Method

- Step 1:

$$v_* = \arg\min_{v} E_{f_*}\left[\ln \frac{f_*(X)}{f(X;v)}\right] = \arg\min_{v} \overbrace{\int f_*(x)\ln f_*(x)dx - \int f_*(x)\ln f(x;v)dx}^{\text{always} \geqslant 0}$$

$$= \arg\max_{v} \int f_*(x)\ln f(x;v)dx = \arg\max_{v} \int \frac{I(x \geq t)f(x;u)}{p_t}\ln f(x;v)dx$$

Cross-Entropy Method

- Step 1:

$$v_* = \underset{v}{\arg\min}\, E_{f_*}\left[\ln \frac{f_*(X)}{f(X;v)}\right] = \underset{v}{\arg\min}\, \overbrace{\int f_*(x) \ln f_*(x)dx - \int f_*(x) \ln f(x;v)dx}^{\text{always} \geqslant 0}$$

$$= \underset{v}{\arg\max} \int f_*(x) \ln f(x;v)dx = \underset{v}{\arg\max} \int \frac{I(x \geq t) f(x;u)}{p_t} \ln f(x;v)dx$$

$$= \underset{v}{\arg\max} \int I(x \geq t) f(x;u) \ln f(x;v)dx$$

Cross-Entropy Method

- Step 1:

$$v_* = \arg\min_{v} E_{f_*}\left[\ln \frac{f_*(X)}{f(X;v)}\right] = \arg\min_{v} \underbrace{\int f_*(x)\ln f_*(x)dx - \int f_*(x)\ln f(x;v)dx}_{\text{always} \geqslant 0}$$

$$= \arg\max_{v} \int f_*(x)\ln f(x;v)dx = \arg\max_{v} \int \frac{I(x \geq t)f(x;u)}{p_t}\ln f(x;v)dx$$

$$= \arg\max_{v} \int I(x \geq t)f(x;u)\ln f(x;v)dx = \arg\max_{v} E_u[I(X \geq t)\ln f(X;v)]$$

Cross-Entropy Method

- For certain families $\{f(\cdot\,;v)\}$ (*eg*, one-dim exponential) the problem

$$v_* = \underset{v}{\arg\max}\, E_u[I(X \geq t) \ln f(X;v)]$$

can be solved analytically:

$$v_* = \frac{E_u[I(X \geq t)X]}{E_u[I(X \geq t)]}$$

Cross-Entropy Method

- For certain families $\{f(\cdot\,; v)\}$ (*eg*, one-dim exponential) the problem

$$v_* = \underset{v}{\arg\max}\, E_u[I(X \geq t) \ln f(X; v)]$$

can be solved analytically:

$$v_* = \frac{E_u[I(X \geq t)X]}{E_u[I(X \geq t)]}$$

Cross-Entropy Method

- In practice: get $X_1, \ldots, X_K$ samples iid as $f(\cdot\,;u)$ and compute the approximation

$$\bar{v}_* = \frac{\sum_{i=1}^{K} [I(X_i \geq t) X_i]}{\sum_{i=1}^{K} [I(X_i \geq t)]}$$

Cross-Entropy Method

- In practice: get $X_1, \ldots, X_K$ samples iid as $f(\cdot\,;u)$ and compute the approximation

$$\bar{v}_* = \frac{\sum_{i=1}^{K} [I(X_i \geq t) X_i]}{\sum_{i=1}^{K} [I(X_i \geq t)]}$$

In general, one would have to (numerically) solve the problem

$$\frac{1}{K} \sum_{i=1}^{K} I(X_i \geq t) \nabla \ln f(X_i; v) = 0$$

Cross-Entropy Method

- Problem: If $\{X \geqslant t\}$ is a rare event, then this fails

$$\bar{v}_* = \frac{\sum_{i=1}^{K} [I(X_i \geq t) X_i]}{\sum_{i=1}^{K} [I(X_i \geq t)]}$$

Cross-Entropy Method

- Problem: If $\{X \geqslant t\}$ is a rare event, then this fails

$$\bar{v}_* = \frac{\sum_{i=1}^{K} [I(X_i \geq t) X_i]}{\sum_{i=1}^{K} [I(X_i \geq t)]}$$

- Most terms in both sums will be zero!

Cross-Entropy with Rare Events

- Rubinstein gave an algorithm that computes v_* adaptively
- Estimate Prob $\{S(X) \geqslant t\}$ [$S(X)$ = sample performance]

Cross-Entropy with Rare Events

- Rubinstein gave an algorithm that computes v_* adaptively
- Estimate Prob $\{S(X) \geqslant t\}$ [$S(X)$ = sample performance]
- **Idea**: compute v_* for a non-rare event $\{S(X) \geqslant t'\}$, and iterate until t' converges to t

Cross-Entropy with Rare Events

- Rubinstein gave an algorithm that computes v_* adaptively
- Estimate Prob $\{S(X) \geqslant t\}$ [$S(X)$ = sample performance]
- **Idea**: compute v_* for a non-rare event $\{S(X) \geqslant t'\}$, and iterate until t' converges to t
- Fix, say $0.01 < \rho < 0.1$; get $X_1, \ldots X_K$ samples iid as $f(\cdot\,;u)$

Cross-Entropy with Rare Events

- Rubinstein gave an algorithm that computes v_* adaptively
- Estimate Prob $\{S(X) \geqslant t\}$ [$S(X)$ = sample performance]
- **Idea**: compute v_* for a non-rare event $\{S(X) \geqslant t'\}$, and iterate until t' converges to t
- Fix, say $0.01 < \rho < 0.1$; get $X_1, \ldots X_K$ samples iid as $f(\cdot\,;u)$
- Compute $t' = (1\text{-}\rho)$ sample quantile of the $S(X_j)$

  then Prob $\{S(X) \geqslant t'\} \geqslant \rho$ (approx.)

Cross-Entropy with Rare Events

- Rubinstein gave an algorithm that computes v_* adaptively
- Estimate Prob $\{S(X) \geqslant t\}$ [$S(X)$ = sample performance]
- **Idea**: compute v_* for a non-rare event $\{S(X) \geqslant t'\}$, and iterate until t' converges to t
- Fix, say $0.01 < \rho < 0.1$; get $X_1, \ldots X_K$ samples iid as $f(\cdot\,;u)$
- Compute t' = $(1\text{-}\rho)$ sample quantile of the $S(X_j)$

  then Prob $\{S(X) \geqslant t'\} \geqslant \rho$ (approx.)
- Now compute v_* "as usual". Iterate until $t'=t$

Cross-Entropy with Rare Events

- Does NOT work with statistical model checking
- Problem: sample quantile computation
- Order the sample performances

$$S_{(1)} \leqslant \ldots \leqslant S_{(i)} \leqslant \ldots \leqslant S_{(K)}$$

$\boxed{(1 - \rho)K}$

- In statistical model checking, sample performances are either 0 (property false) or 1 (property true)

Cross-Entropy with Rare Events

- However …

$$v_* = \frac{E_u[I(X \geq t)X]}{E_u[I(X \geq t)]} = \frac{E_w[I(X \geq t)W(X;u,w)X]}{E_w[I(X \geq t)W(X;u,w)]}$$

where $W(x;u,w) = \dfrac{f(x;u)}{f(x;w)}$ for an arbitrary parameter *w*

<u>Work in progress</u>

Example: Fuel Control System

The Stateflow/Simulink model

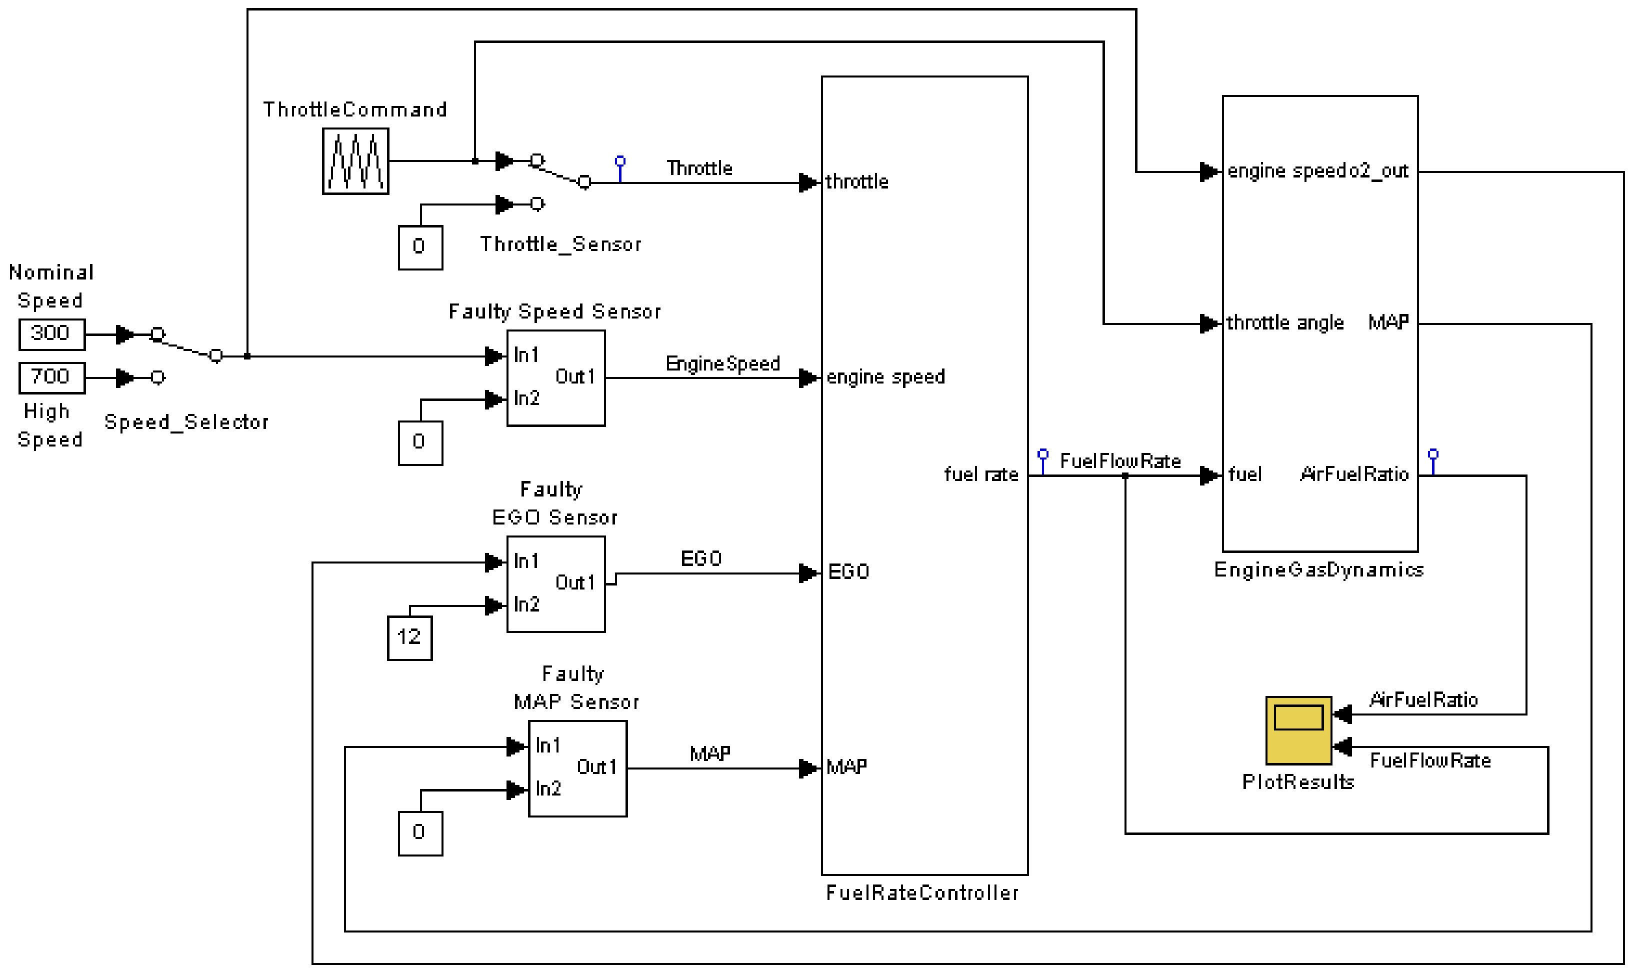

Verification

- We want to estimate the probability that

$$\mathcal{M}, \textit{FaultRate} \models \boldsymbol{F^{100}}\ \boldsymbol{G^{1}}(\textit{FuelFlowRate} = 0)$$

- *"It is the case that within 100 seconds, FuelFlowRate is zero for 1 second"*

- *FaultRate = 1/3600s* (same value for the three sensors)

Importance Sampling

- Ran cross-entropy method to estimate optimal biasing density with *FaultRates* = {1/7, 1/8, 1/9}
- Used 100 samples for this, and obtained
  - $NewRates_{*}$ = {1/2.007, 1/1.0113, 1/1.7277}
- Run importance sampling with 1,000 samples and $NewRates_{*}$
  - Probability estimate 9.1855×10^{-15}

Conclusions

- Need to be able to deal with rare events in statistical model checking
- The Cross-Entropy method is an interesting, semi-automatic technique
- Research: adaptive technique for stat. model checking
- [Further benefit: cross-entropy method also applies to optimization, *eg*, finding policies for MDPs]

The End

Questions?