Geometrically Coupled Monte Carlo Sampling

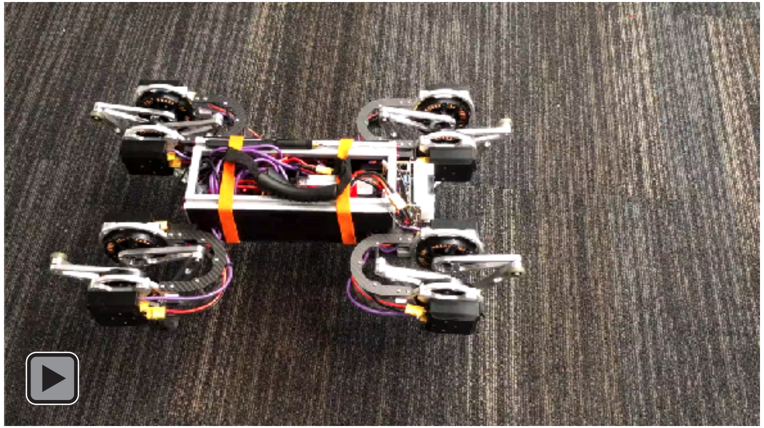

Mark Rowland

Krzysztof Choromanski

François Chalus

Aldo Pacchiano

Tamas Sarlos

Richard E. Turner

Adrian Weller

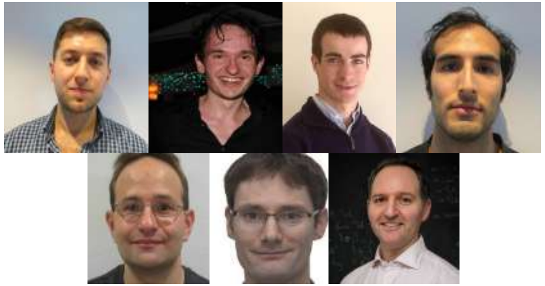

Geometrically Coupled Monte Carlo Sampling

Central goal: $\mathbb{E}_{X\sim\mu}\left[f(X)\right]$

Unbiased Monte Carlo estimation: $\frac{1}{m}\sum_{i=1}^{m} f(X_i)\,,\quad (X_i)_{i=1}^{m}\sim\mu$

Can we do better than i.i.d.?

Key contribution: *K*-optimality. Optimise the objective below over the joint distribution of $X_{1:m}$

$$\mathbb{E}_{f\sim \mathrm{GP}(0,K)}\left[\mathbb{E}_{X_{1:m}}\left[\left(\frac{1}{m}\sum_{i=1}^{m} f(X_i) - \mathbb{E}_{X\sim\eta}\left[f(X)\right]\right)^2\right]\right]$$

This leads to a multi-marginal transport problem, which is often analytically solvable.

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g} \sim \mu}[F_\theta(\mathbf{g})]$$

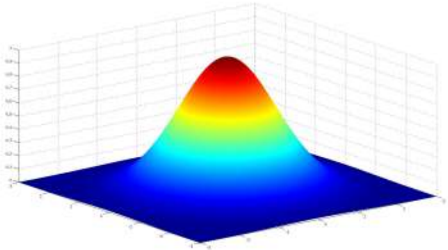

isotropic distribution

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g}\sim\mu}[F_\theta(\mathbf{g})]$$

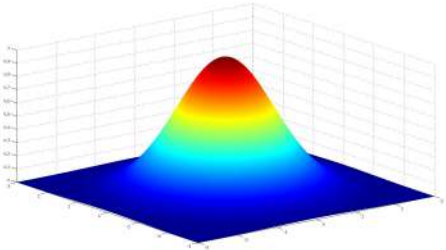

isotropic distribution

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g} \sim \mu}[F_\theta(\mathbf{g})]$$

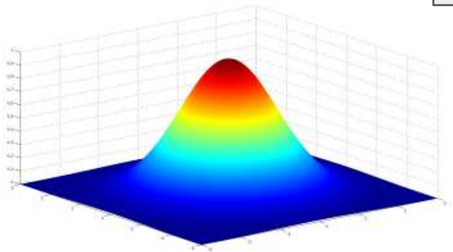

isotropic distribution

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g}\sim\mu}[F_\theta(\mathbf{g})]$$

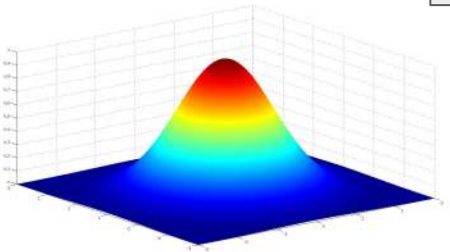

isotropic distribution

$$\widehat{J}(\theta) = \sum_{i=1}^{2k} \frac{1}{2k}[F_\theta(\mathbf{g}^i) + F_\theta(-\mathbf{g}^i)]$$

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g} \sim \mu}[F_\theta(\mathbf{g})]$$

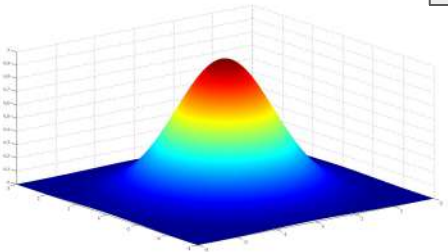

isotropic distribution

$$\widehat{J}(\theta) = \sum_{i=1}^{2k} \frac{1}{2k}[F_\theta(\mathbf{g}^i) + F_\theta(-\mathbf{g}^i)]$$

antithetic pair

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g}\sim\mu}[F_\theta(\mathbf{g})]$$

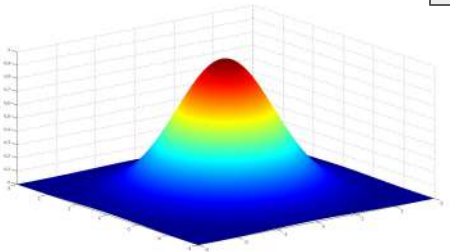

isotropic distribution

$$\widehat{J}(\theta) = \sum_{i=1}^{2k} \frac{1}{2k}[F_\theta(\mathbf{g}^i) + F_\theta(-\mathbf{g}^i)]$$

antithetic pair

Typical approach to Monte Carlo Sampling:

- Independent Antithetic Pairs
- Coupled Samples of Equal Lengths

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g} \sim \mu}[F_\theta(\mathbf{g})]$$

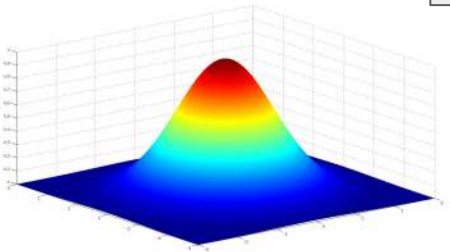

isotropic distribution

$$\widehat{J}(\theta) = \sum_{i=1}^{2k} \frac{1}{2k} \underbrace{[F_\theta(\mathbf{g}^i) + F_\theta(-\mathbf{g}^i)]}_{\text{antithetic pair}}$$

Typical a... h to M... rlo Sampling:

- **Independ... c Pairs**
- **Coupled ... Equal Lengths**

GCMC in Robotics - Policy Search - An Overview

$$J(\theta) = \mathbb{E}_{\mathbf{g}\sim\mu}[F_\theta(\mathbf{g})]$$

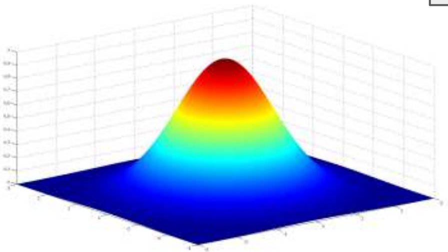

isotropic distribution

$$\widehat{J}(\theta) = \sum_{i=1}^{2k} \frac{1}{2k}\underbrace{[F_\theta(\mathbf{g}^i) + F_\theta(-\mathbf{g}^i)]}_{\text{antithetic pair}}$$

antithetic pair

Typical a h to M rlo Sampling:

- Independ c Pairs
- Coupled qual Lengths

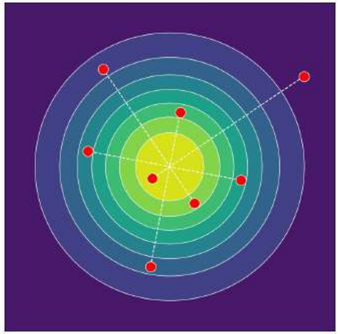

GCMC:

- orthogonal directions of different antithetic pairs
- correlated unequal lengths within a pair
- variance reduction

GCMC for Policy Search - Details

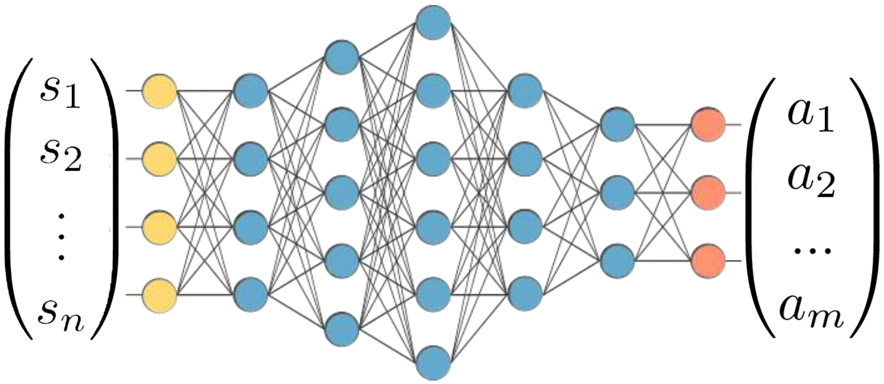

$$\pi : S \rightarrow A$$

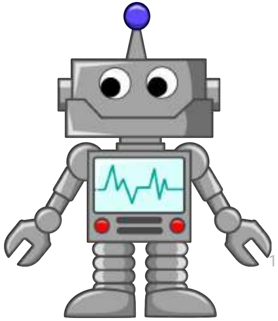

GCMC for Policy Search

$$F : \mathbb{R}^d \rightarrow \mathbb{R}$$

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_L \end{pmatrix}$$

$$F$$

$$R_{\text{TOTAL}}$$

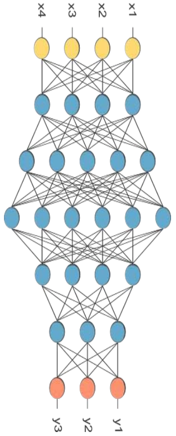

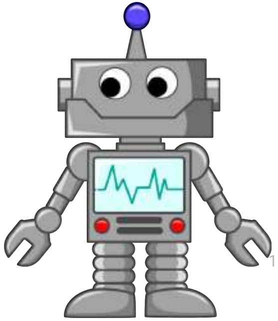

GCMC for Policy Search

$$F : \mathbb{R}^d \rightarrow \mathbb{R}$$

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_L \end{pmatrix}$$

R_{TOTAL}

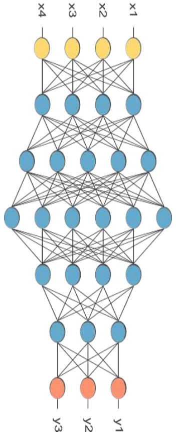

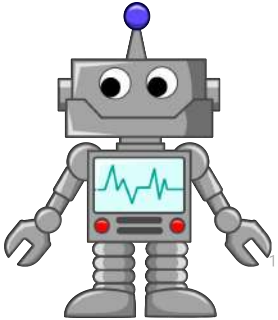

GCMC for Policy Search

$$F : \mathbb{R}^d \rightarrow \mathbb{R}$$

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_L \end{pmatrix}$$

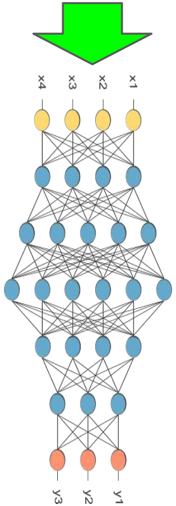

$$R_{\mathrm{TOTAL}}$$

Towards smooth relaxations

$$\max_{\mu \in \mathcal{P}(\mathbb{R}^d)} \mathbb{E}_{\theta \sim \mu}[F(\theta)]$$

Gaussian smoothings

$$\max_{\theta \in \mathbb{R}^d} J(\theta) = \mathbb{E}_{\phi \sim N(\theta, \sigma^2 I)} [F(\phi)]$$

GCMC for Policy Search

$$F : \mathbb{R}^d \rightarrow \mathbb{R}$$

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_L \end{pmatrix}$$

R_{TOTAL}

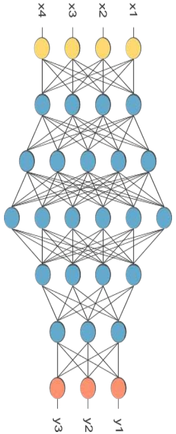

Towards smooth relaxations

$$\max_{\mu \in \mathcal{P}(\mathbb{R}^d)} \mathbb{E}_{\theta \sim \mu}[F(\theta)]$$

Gaussian smoothing gradient

$$\nabla J(\theta) = \frac{1}{\sigma} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_d)} [F(\theta + \sigma \mathbf{g}) \mathbf{g}]$$

GCMC for Policy Search

$$F : \mathbb{R}^d \rightarrow \mathbb{R}$$

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_L \end{pmatrix}$$

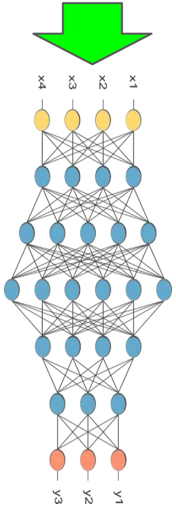

R_{TOTAL}

Towards smooth relaxations

$$\max_{\mu \in \mathcal{P}(\mathbb{R}^d)} \mathbb{E}_{\theta \sim \mu}[F(\theta)]$$

Gaussian smoothing gradient

$$\nabla J(\theta) = \frac{1}{\sigma} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_d)}[F(\theta + \sigma \mathbf{g})\mathbf{g}]$$

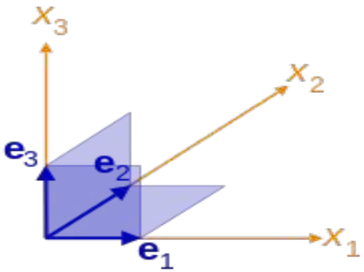

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 1 \end{pmatrix}$$

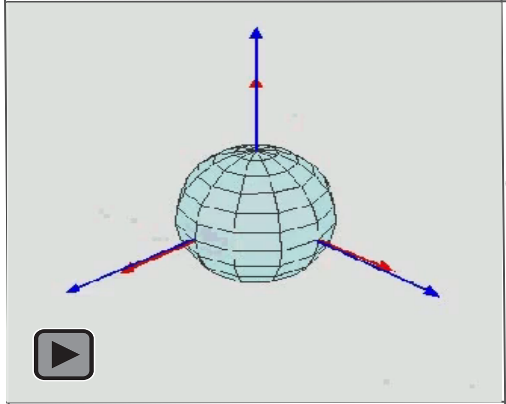

$$\mathbf{G}_{\mathrm{ort}} = \begin{pmatrix} g^{\mathrm{ort}}_{1,1} & g^{\mathrm{ort}}_{1,2} & \dots & g^{\mathrm{ort}}_{1,n} \\ g^{\mathrm{ort}}_{2,1} & g^{\mathrm{ort}}_{2,2} & \dots & g^{\mathrm{ort}}_{2,n} \\ \dots & \dots & \dots & \dots \\ g^{\mathrm{ort}}_{m,1} & g^{\mathrm{ort}}_{m,n} & \dots & g^{\mathrm{ort}}_{m,n} \end{pmatrix}$$

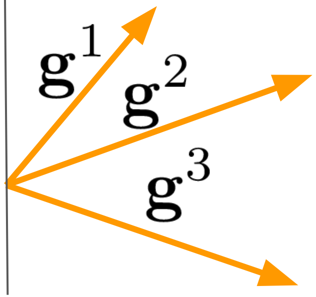

$$\mathbf{G} = \begin{pmatrix} g_{1,1} & g_{1,2} & \dots & g_{1,n} \\ g_{2,1} & g_{2,2} & \dots & g_{2,n} \\ \dots & \dots & \dots & \dots \\ g_{m,1} & g_{m,n} & \dots & g_{m,n} \end{pmatrix}$$

Coupled antithetic pairs for Monte Carlo gradient estimation

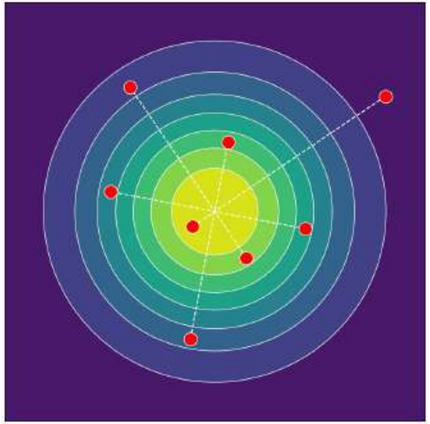

Baseline gradient estimator with antithetic pairs (Salimans et al. 2017):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma} \sum_{i=1}^{N} (F(\theta + \sigma\epsilon_i)\epsilon_i - F(\theta - \sigma\epsilon_i)\epsilon_i)$$

Coupled antithetic pairs for Monte Carlo gradient estimation

Baseline gradient estimator with antithetic pairs (Salimans et al. 2017):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma}\sum_{i=1}^{N}(F(\theta+\sigma\epsilon_i)\epsilon_i - F(\theta-\sigma\epsilon_i)\epsilon_i)$$

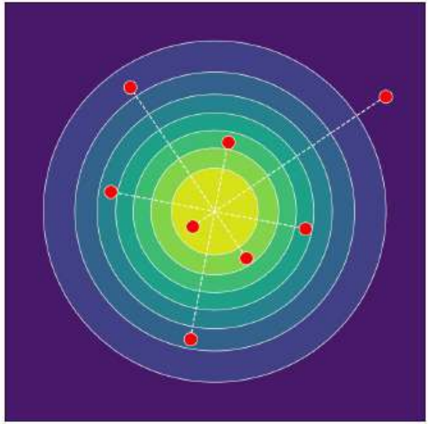

Coupled antithetic pairs for Monte Carlo gradient estimation

Baseline gradient estimator with antithetic pairs (Salimans et al. 2017):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma}\sum_{i=1}^{N}(F(\theta+\sigma\epsilon_i)\epsilon_i - F(\theta-\sigma\epsilon_i)\epsilon_i)$$

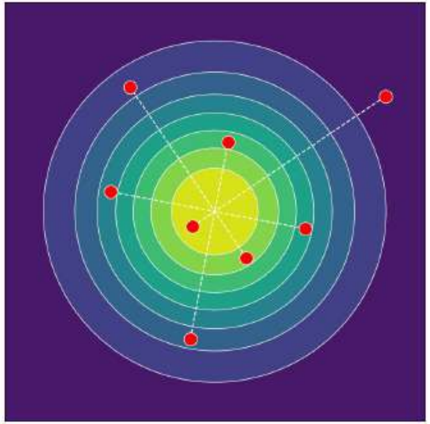

Antithetic inverse lengths coupling estimator (Rowland, Choromanski et al. 2018):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma}\sum_{i=1}^{N}(F(\theta+\sigma\epsilon_i)\epsilon_i - F(\theta-\sigma\epsilon_i^{'})\epsilon_i^{'})$$

Coupled antithetic pairs for Monte Carlo gradient estimation

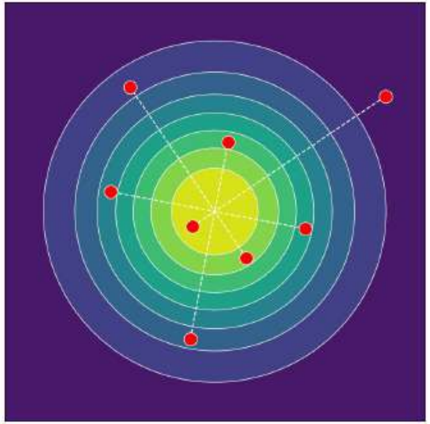

Baseline gradient estimator with antithetic pairs (Salimans et al. 2017):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma}\sum_{i=1}^{N}(F(\theta+\sigma\epsilon_i)\epsilon_i - F(\theta-\sigma\epsilon_i)\epsilon_i)$$

Antithetic inverse lengths coupling estimator (Rowland, Choromanski et al. 2018):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma}\sum_{i=1}^{N}(F(\theta+\sigma\epsilon_i)\epsilon_i - F(\theta-\sigma\epsilon_i^{'})\epsilon_i^{'})$$

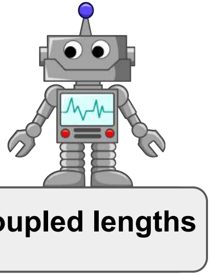

Coupled antithetic pairs for Monte Carlo gradient estimation

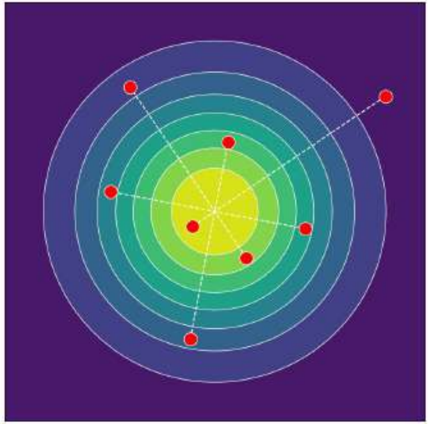

Baseline gradient estimator with antithetic pairs (Salimans et al. 2017):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma} \sum_{i=1}^{N} (F(\theta + \sigma\epsilon_i)\epsilon_i - F(\theta - \sigma\epsilon_i)\epsilon_i)$$

Antithetic inverse lengths coupling estimator (Rowland, Choromanski et al. 2018):

$$\widehat{\nabla}_N J(\theta) = \frac{1}{2N\sigma} \sum_{i=1}^{N} (F(\theta + \sigma\epsilon_i)\epsilon_i - F(\theta - \sigma\epsilon_i^{\prime})\epsilon_i^{\prime})$$

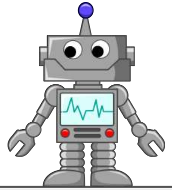

coupled lengths

$$\epsilon_i \sim \mathcal{N}(0, \mathbf{I}_d) \quad \Big| \quad \epsilon_i^{\prime} = \epsilon_i \frac{F^{-1}_{\chi(d)}(1 - F_{\chi(d)}(\|\epsilon_i\|_2))}{\|\epsilon_i\|_2}$$

Experimental results: Minitaur Learning How to Walk with antithetic coupled samples + linear policies

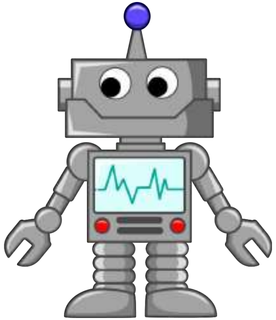

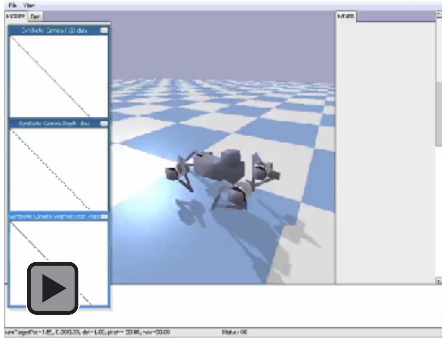

N=8

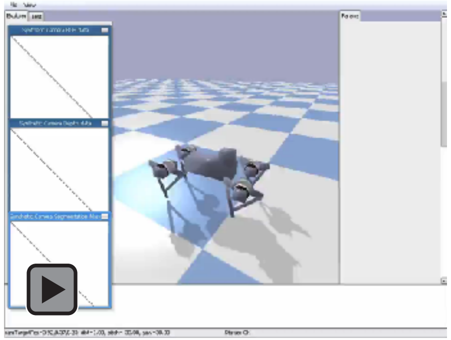

N=16

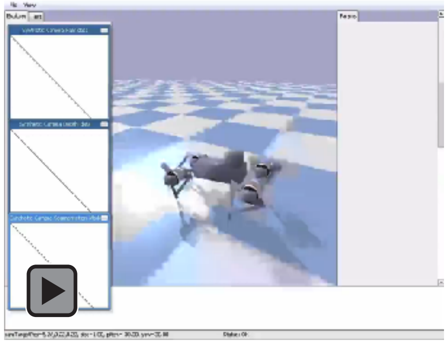

N=48

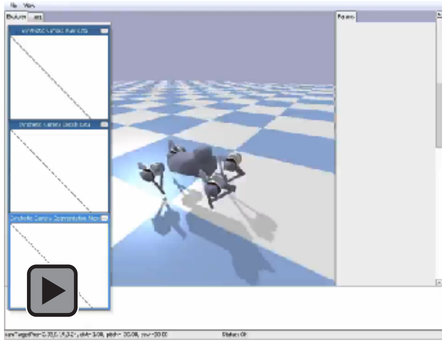

N=54

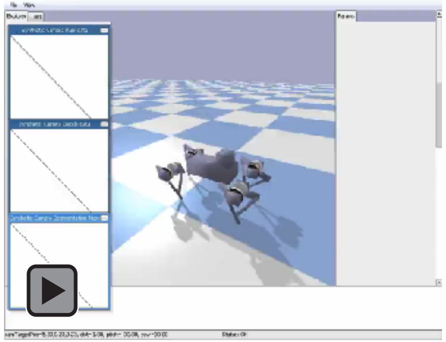

N=64

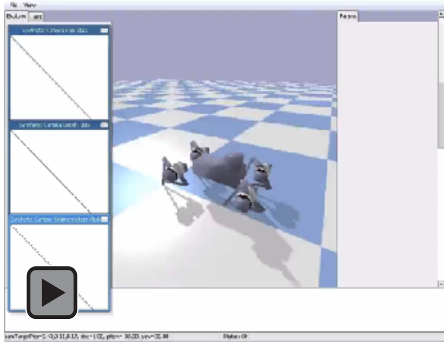

N=96

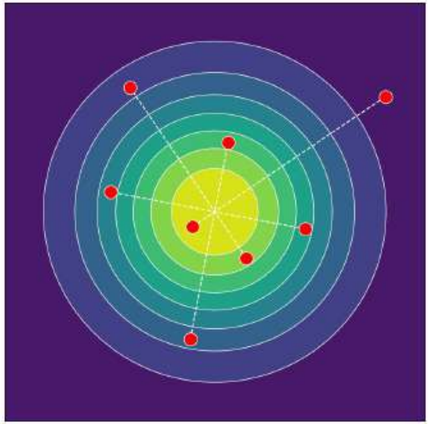

Thank you !!!