Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Learning to Play Minigolf

(Mohammad Khansari, Klas Kronander and Aude Billard)

Lukas Sommerhalder
18.01.2024


Photo: Khansari et al.

Agenda

1

Introduction

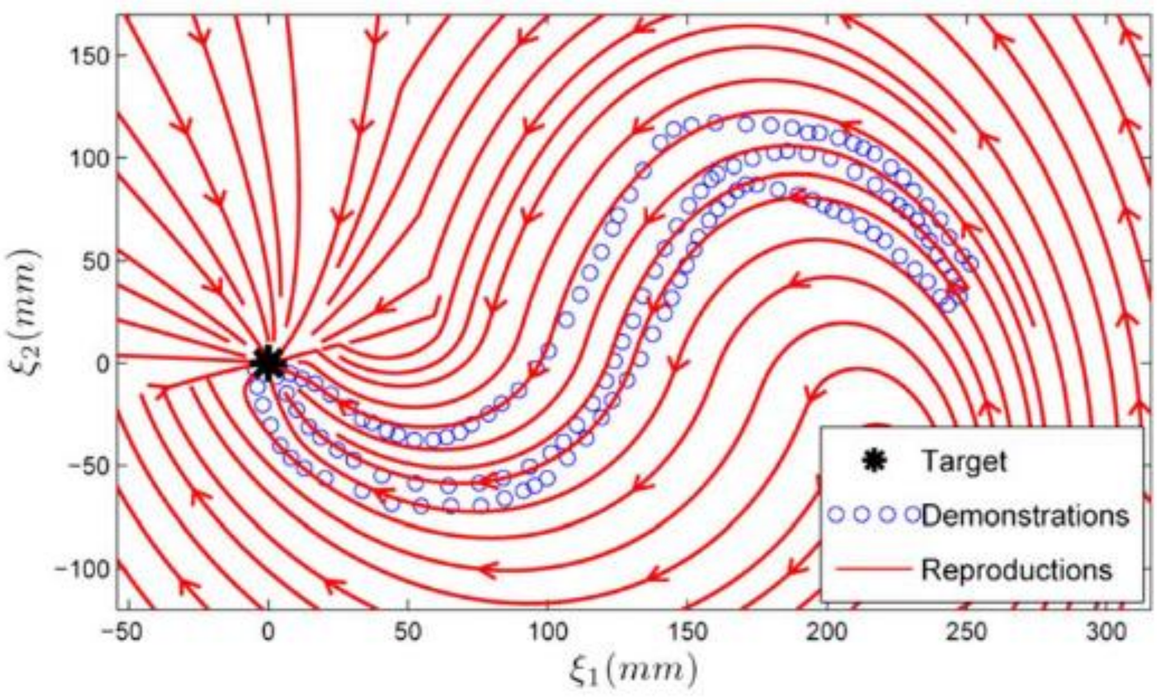

Picture: Khansari et al.

Objectives

- Introduction to:
    - Gaussian Mixture Models (GMM)
    - Gaussian Mixture Regression (GMR)
- Lyapunov Stability
- Stable Estimator of Dynamical Systems (SEDS)
- Usage of these concepts for a Minigolf-Robot

Capabilities of the whole model

- Hit the ball and put it in
- Reproduction of demonstrated hitting motions
- Estimate a successfull speed and direction
- Rotation and scaling of the hitting motion
- Robust against perturbations
    - Initial golf club position
    - Iinitial ball position
    - Deviations during the execution of the shot

Video: Teaching robot how to swing a golf club

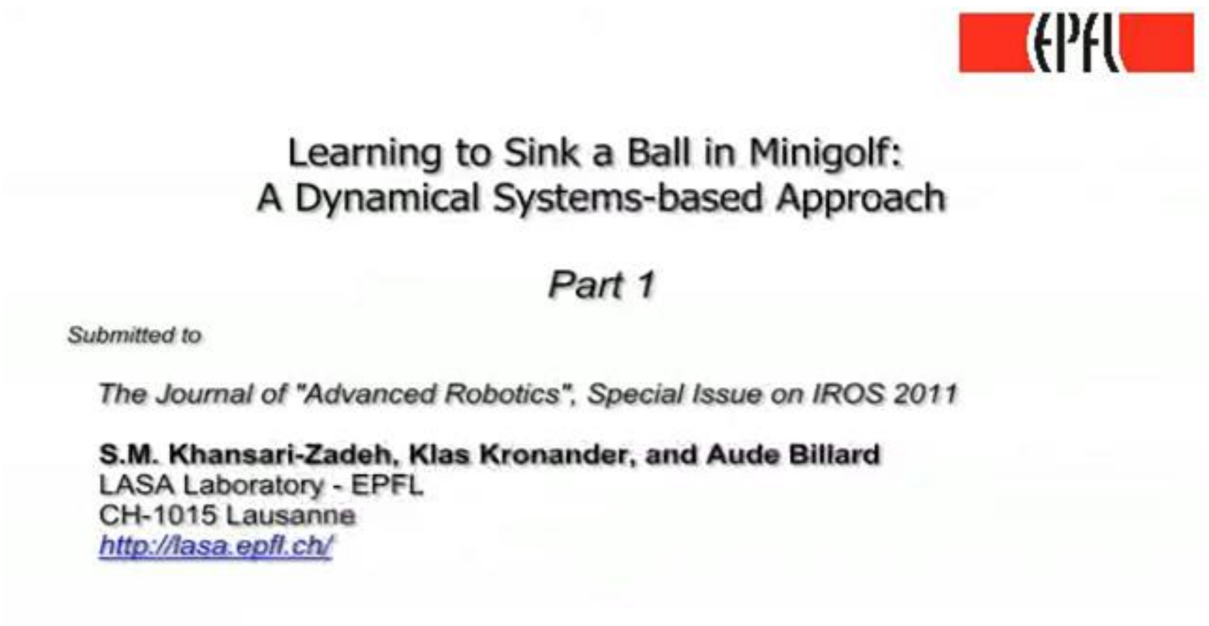

www.youtube.com/watch?v=hHq7QmuxTIw

2

SEDS

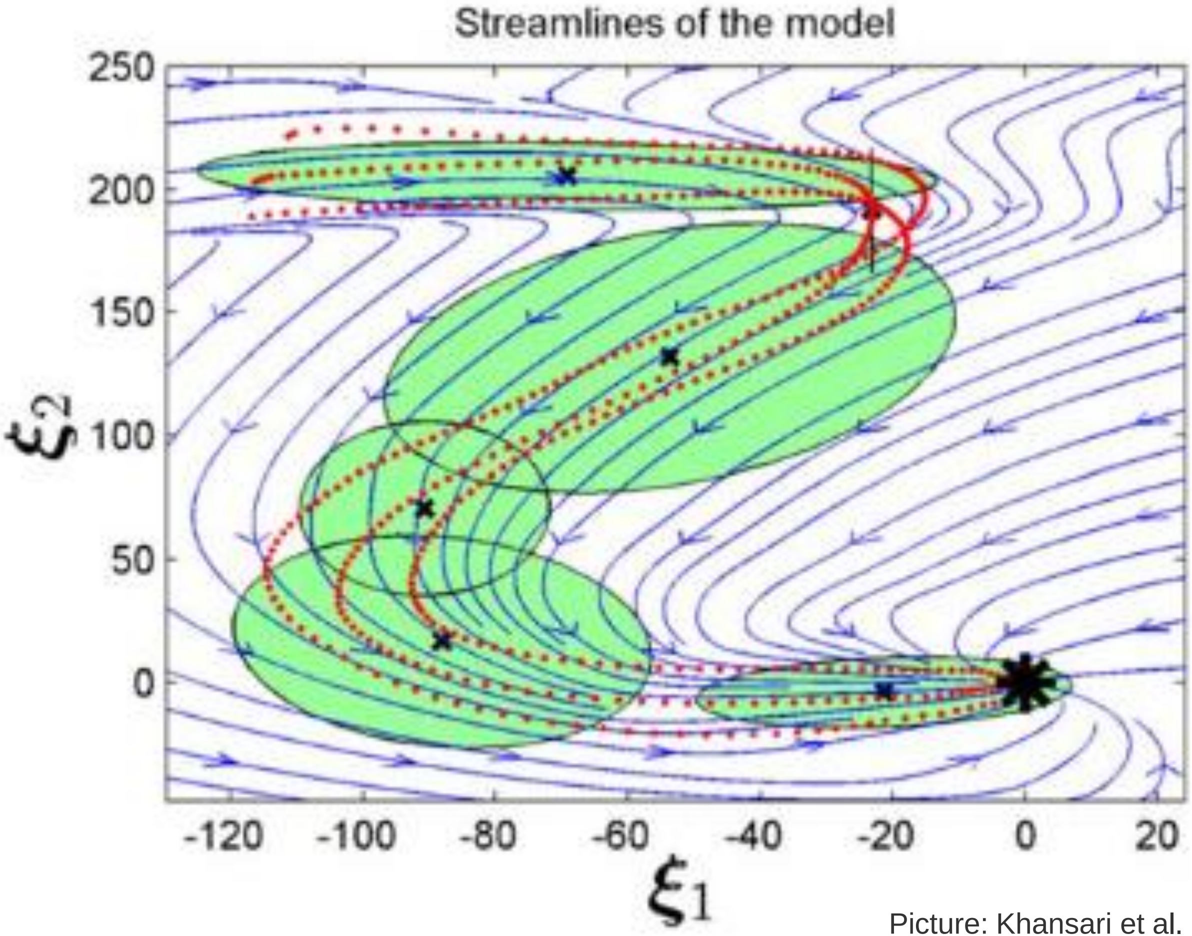

Picture: Khansari et al.

Introduction SEDS

- Dynamical System (DS)

$$\dot{\boldsymbol{\xi}} = \boldsymbol{f}(\boldsymbol{\xi})$$

Multidimensional Kinematic Variable:
e.g. End-effector position/orientation,
joint angles

- Challenge
  - Finding a model of a globally asymptotically stable DS
  - With few demonstrations

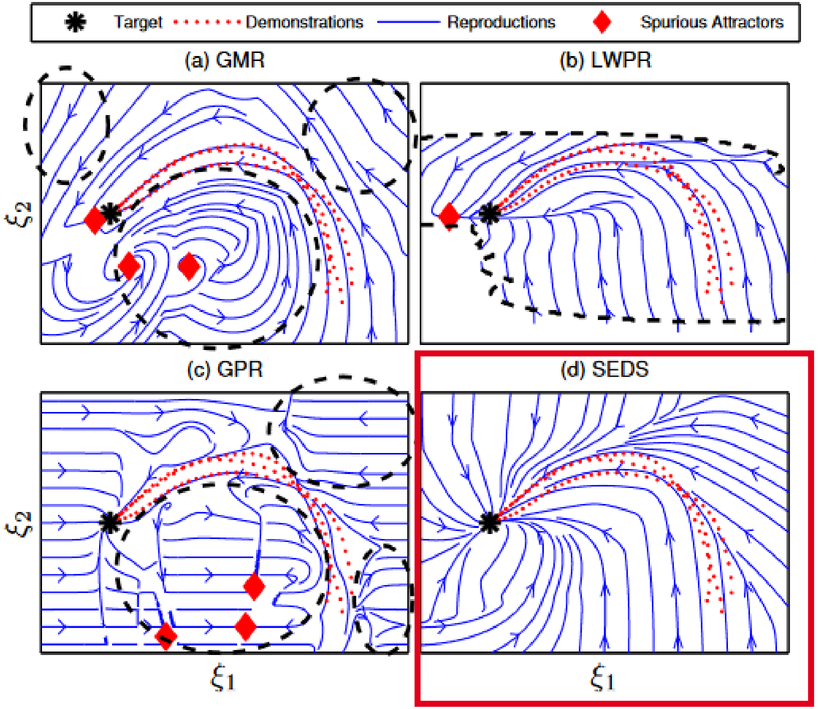

Picture: Khansari et al.

Demonstration Data

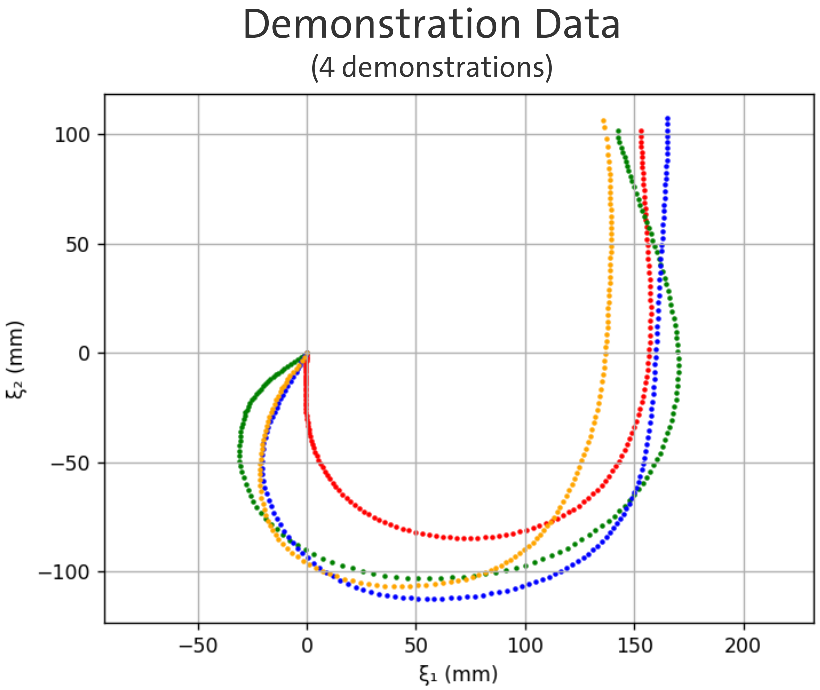

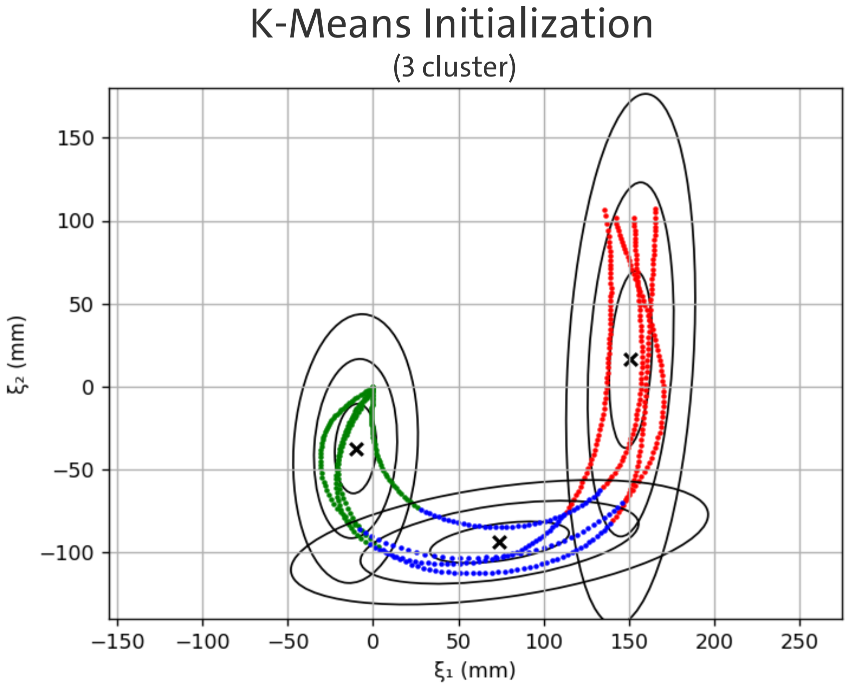

Each point describes its coordinates and velocities.

Multivariate Gaussian Distribution I

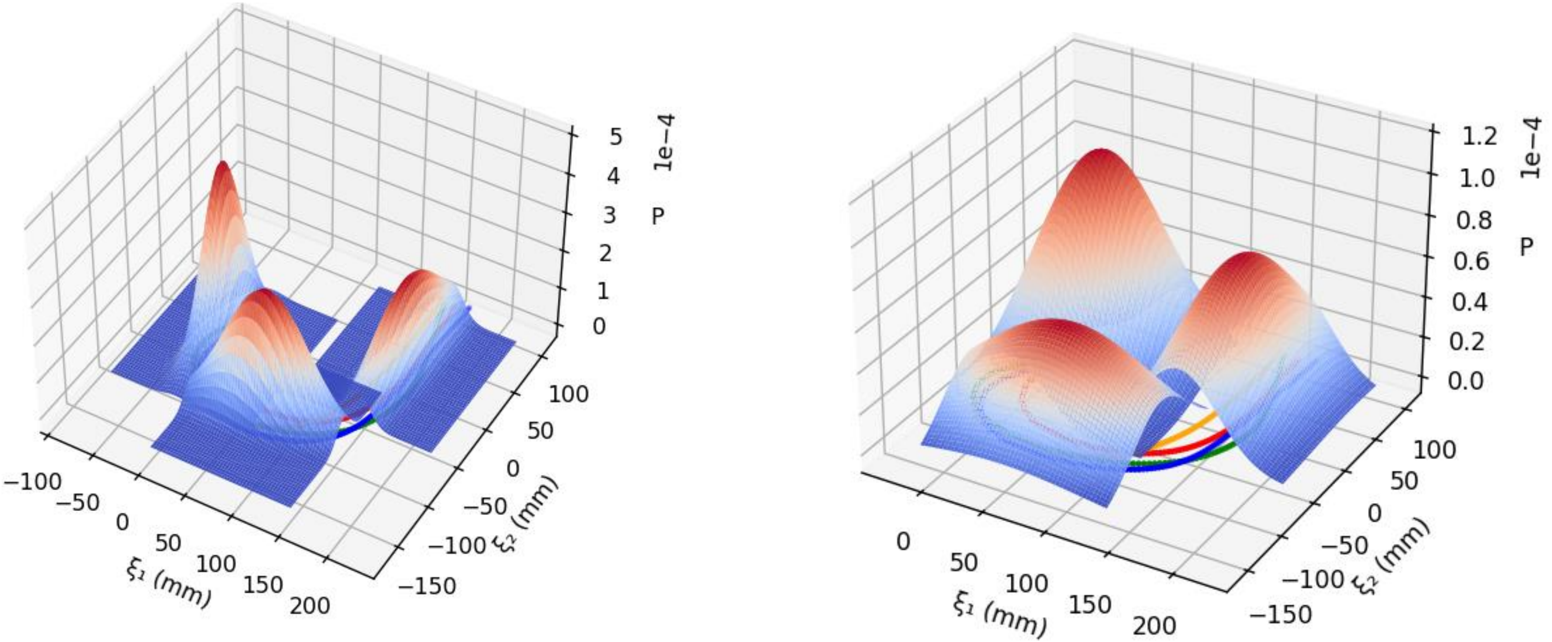

Examples of Multivariate Gaussians
K-Means Initialization (left), SEDS (right)

Multivariate Gaussian Distribution II

Probability Density Function (PDF):

$$\mathcal{N}(\xi^{t,n}, \dot{\xi}^{t,n}; \theta^k) = \frac{1}{\sqrt{(2\pi)^{2d}|\Sigma_\xi^k|}} e^{-\frac{1}{2}([\xi^{t,n}, \dot{\xi}^{t,n}] - \mu^k)^T (\Sigma^k)^{-1} ([\xi^{t,n}, \dot{\xi}^{t,n}] - \mu^k)} \quad \begin{cases} \forall n \in 1..N \\ t \in 0..T^n \end{cases}$$

2d components (d coordinates + d time derivatives)

Parameters:

$$\theta^k = \{\pi^k, \mu^k, \Sigma^k\} \quad , \quad \mu^k = \begin{pmatrix} \mu_\xi^k \\ \mu_{\dot{\xi}}^k \end{pmatrix} \quad , \quad \Sigma^k = \begin{pmatrix} \Sigma_\xi^k & \Sigma_{\xi\dot{\xi}}^k \\ \Sigma_{\dot{\xi}\xi}^k & \Sigma_{\dot{\xi}}^k \end{pmatrix}$$

Weight parameter for the GMM (not used at this point)

Gaussian Mixture Model

Weight of cluster k:

$$\mathcal{P}(k) = \pi^k$$

Conditional PDF (k-th Cluster PDF):

$$\mathcal{P}(\xi^{t,n}, \dot{\xi}^{t,n} | k) = \mathcal{N}(\xi^{t,n}, \dot{\xi}^{t,n}; \mu^k, \Sigma^k)$$

Probability Density Function of the GMM:

$$\mathcal{P}(\xi^{t,n}, \dot{\xi}^{t,n}; \theta) = \sum_{k=1}^{K} \mathcal{P}(k)\mathcal{P}(\xi^{t,n}, \dot{\xi}^{t,n} | k) \quad \begin{cases} \forall n \in 1 \dots N \\ t \in 0 \dots T^n \end{cases}$$

Gaussian Mixture Regression I

Gaussian Mixture Regression (GMR):

$$\dot{\xi} = \sum_{k=1}^{K} \frac{\mathcal{P}(k)\mathcal{P}(\xi|k)}{\sum_{i=1}^{K} \mathcal{P}(i)\mathcal{P}(\xi|i)} \left(\mu_{\dot{\xi}}^{k} + \Sigma_{\dot{\xi}\xi}^{k} (\Sigma_{\xi}^{k})^{-1} (\xi - \mu_{\xi}^{k})\right)$$

Simplification through substitution:

$$\dot{\xi} = \hat{f}(\xi) = \sum_{k=1}^{K} \underbrace{h^k(\xi)}_{\text{non-linear}} \underbrace{(A^k \xi + b^k)}_{\text{linear}} \quad \begin{cases} A^k = \Sigma_{\dot{\xi}\xi}^{k} (\Sigma_{\xi}^{k})^{-1} \\ b^k = \mu_{\dot{\xi}}^{k} - A^k \mu_{\xi}^{k} \\ h^k(\xi) = \frac{\mathcal{P}(k)\mathcal{P}(\xi|k)}{\sum_{i=1}^{K} \mathcal{P}(i)\mathcal{P}(\xi|i)} \end{cases}$$

non-linear

linear

Gaussian Mixture Regression II

Visualization of a 1-D Model with three gaussians

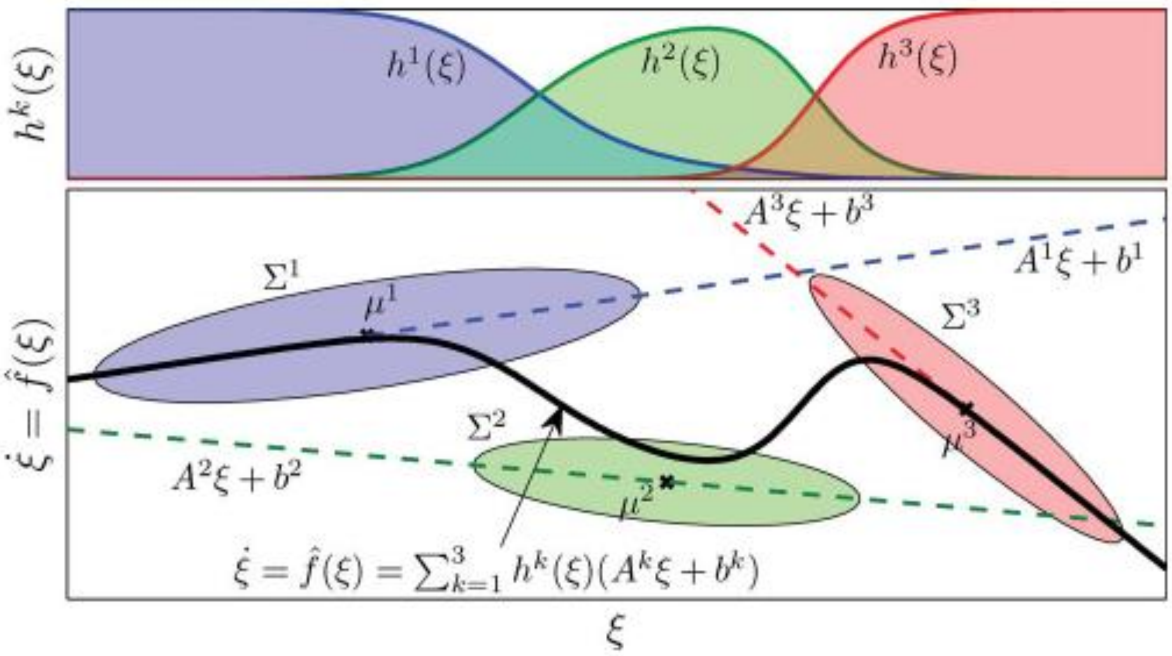

Picture: Khansari et al.

Finding Parameters for the GMM

- The usual algorithm (Expectation-Maximization) should not be used because:
  - Doesn‘t ensure globally asymptotically stability.
  - Minimizing the log likelihood might not be optimal.
- Solution
  - Adding constraints to ensure stability.
  - Allow various goals for optimization:
    - Mean Square Error
    - Log-Likelihood
    - Direction Deviation

Lyapunov Stability Theorem

Lyapunov Function: $V(\xi) \,:\, \mathbb{R}^d \rightarrow \mathbb{R}$

$$V(\xi) > 0 \qquad \forall \xi \in \mathbb{R}^d, \qquad \xi \neq \xi^*$$
$$\dot{V}(\xi) < 0 \qquad \forall \xi \in \mathbb{R}^d, \qquad \xi \neq \xi^*$$
$$V(\xi^*) = 0, \qquad \dot{V}(\xi^*) = 0.$$

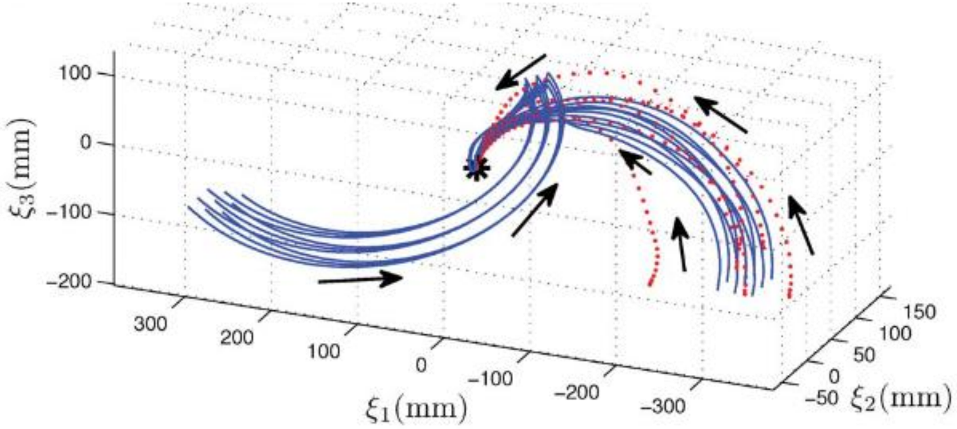

Picture: Khansari et al.

Lyapunov Function

$$V(\xi) = \frac{1}{2}(\xi - \xi^*)^T(\xi - \xi^*)$$

$$\dot{V}(\xi) = \frac{dV}{dt} = \frac{dV}{d\xi}\frac{d\xi}{dt}$$

$$= \frac{1}{2}\frac{d}{d\xi}\left((\xi - \xi^*)^T(\xi - \xi^*)\right)\dot{\xi}$$

$$= (\xi - \xi^*)^T\dot{\xi}$$

is a Lyapunov Function with:

$$\dot{\xi} = \hat{f}(\xi) = \sum_{k=1}^{K} h^k(\xi)(A^k\xi + b^k)$$

and constraints:

$$b^k = -A^k\xi^* \qquad \forall k = 1 \dots K$$

$$A^k + (A^k)^T \prec 0$$

negative definite

Optimization Goals

$$\min_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \underbrace{-\frac{1}{T}\sum_{n=1}^{N}\sum_{t=0}^{T^n} \log \mathcal{P}([\boldsymbol{\xi}^{t,n}; \dot{\boldsymbol{\xi}}^{t,n}] | \boldsymbol{\theta})}_{\text{Maximizing the likelihood}}$$

$$\min_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \underbrace{\frac{1}{2T}\sum_{n=1}^{N}\sum_{t=0}^{T^n} \|\hat{\dot{\xi}}^{t,n} - \dot{\xi}^{t,n}\|^2}_{\text{Minimizing the MSE}}$$

$$\min_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \underbrace{-\frac{1}{T}\sum_{n=1}^{N}\sum_{t=0}^{T^n} \frac{(\dot{\boldsymbol{\xi}}^{t,n})^T \boldsymbol{f}(\boldsymbol{\xi}^{t,n}; \boldsymbol{\theta})}{\|\dot{\boldsymbol{\xi}}^{t,n}\| \|\boldsymbol{f}(\boldsymbol{\xi}^{t,n}; \boldsymbol{\theta})\|}}_{\text{Minimizing the angle between demonstrations and estimation}}$$

Stability Constraints

$$b^k = -A^k \xi^*$$
$$A^k + (A^k)^T \prec 0$$
$$\Sigma^k \succ 0$$
$$0 < \pi^k \leq 1$$
$$\textstyle\sum_{k=1}^{K} \pi^k = 1$$

$$\forall k \in 1 \dots K$$

Constraints of the Gaussians

Example Trajectory Reproduction

Streamlines of the model

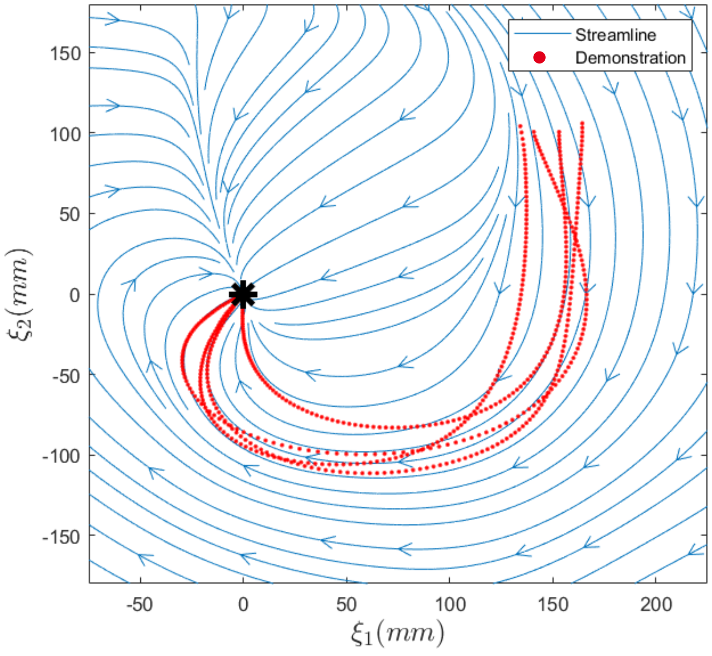

Trajectory reproduction

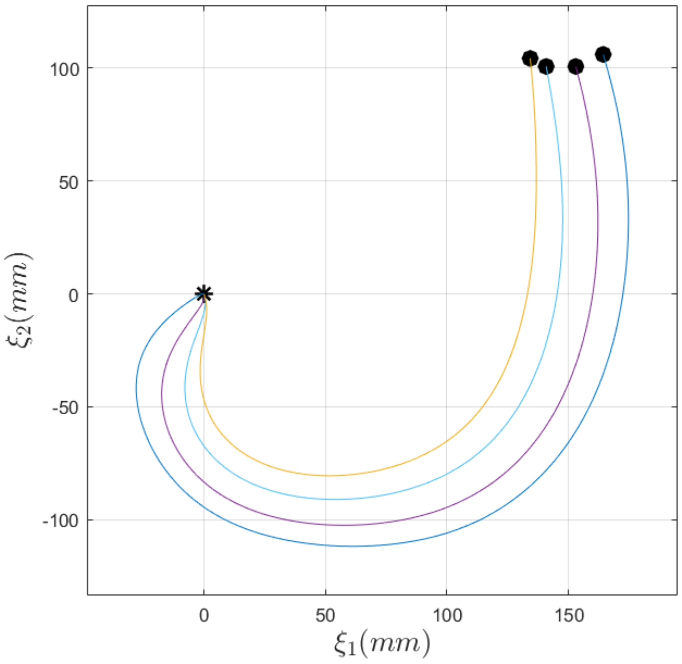

Example Velocity Reproduction

Demonstration/Reproduction

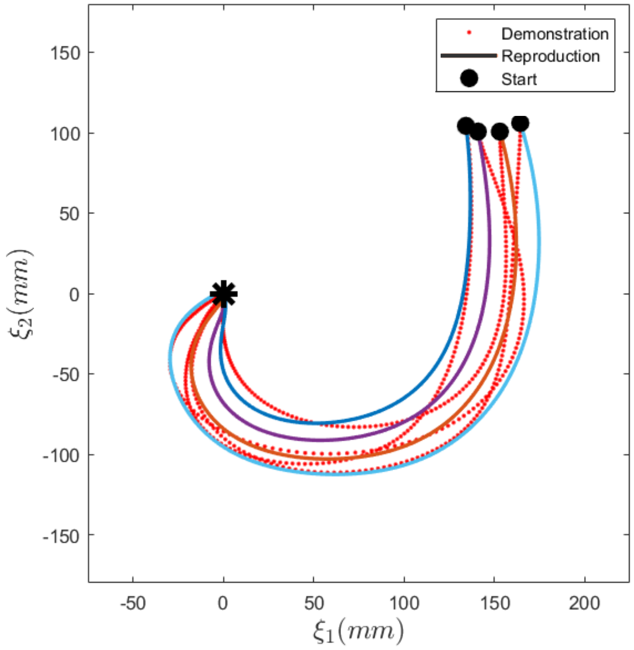

Demonstration/Velocity Predictions

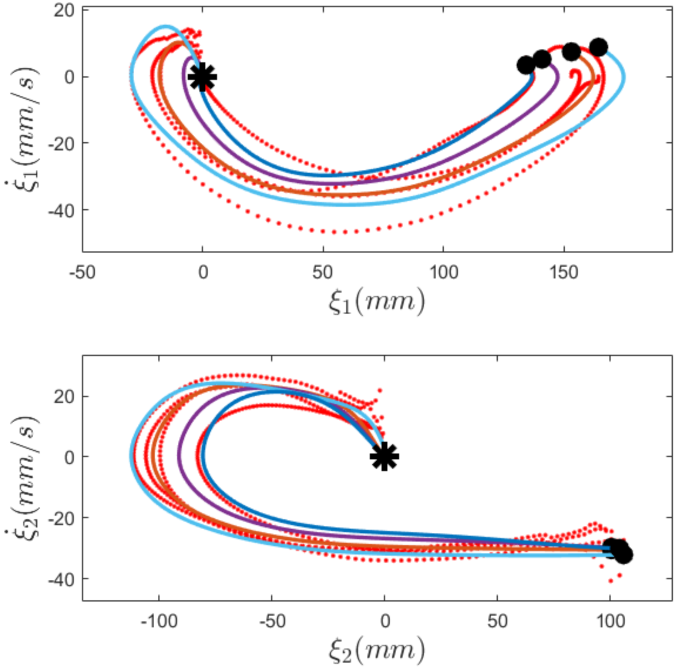

Robot Control Loop

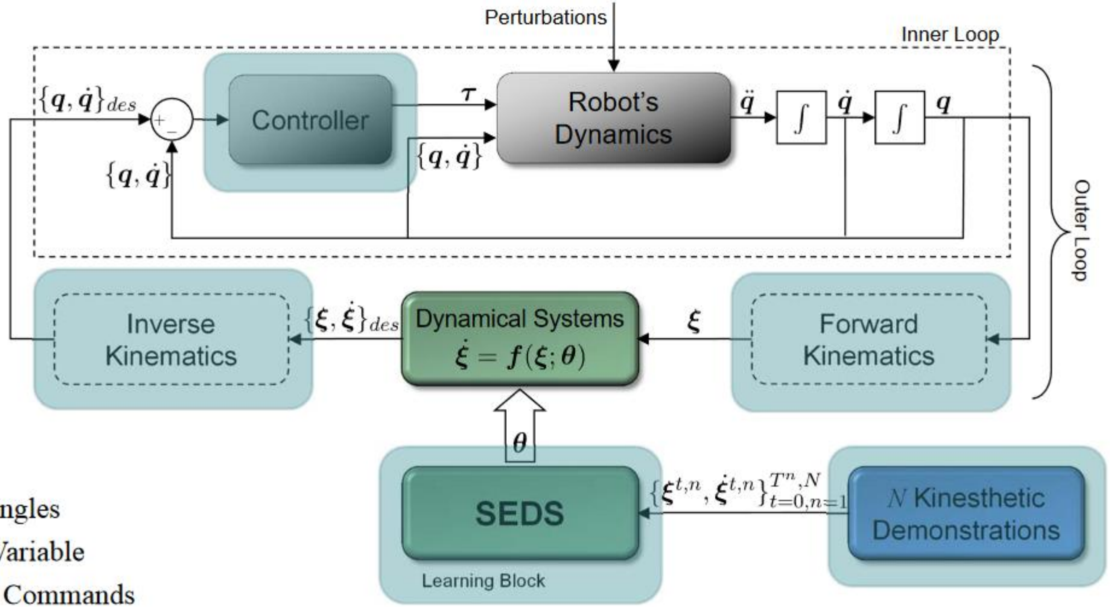

q : Joint Angles

ξ : State Variable

τ : Motor Commands

Picture: Khansari et al.

Video: Introduction to SEDS

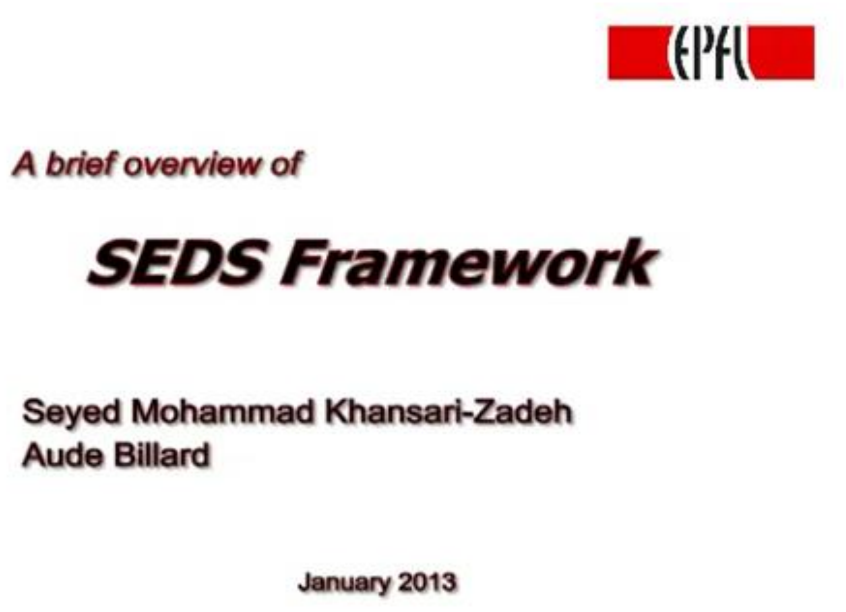

https://www.youtube.com/watch?v=qc5_as8qxBI

3

Hitting Motion

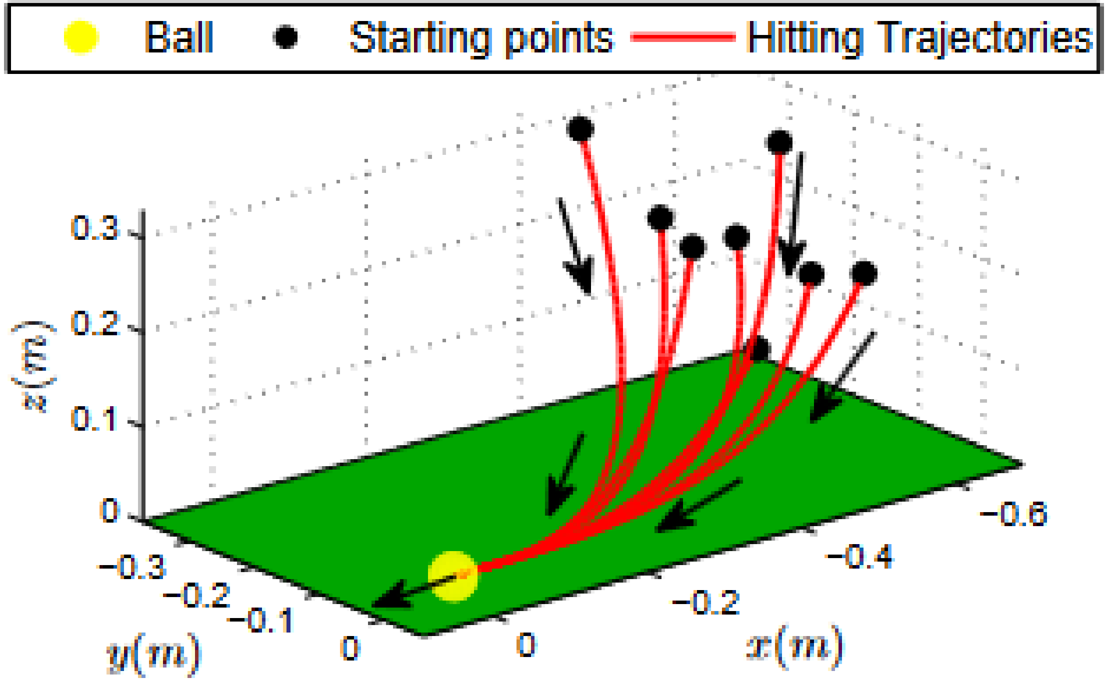

Picture: Khansari et al.

Getting a Target Field from SEDS I

Picture: Khansari et al.

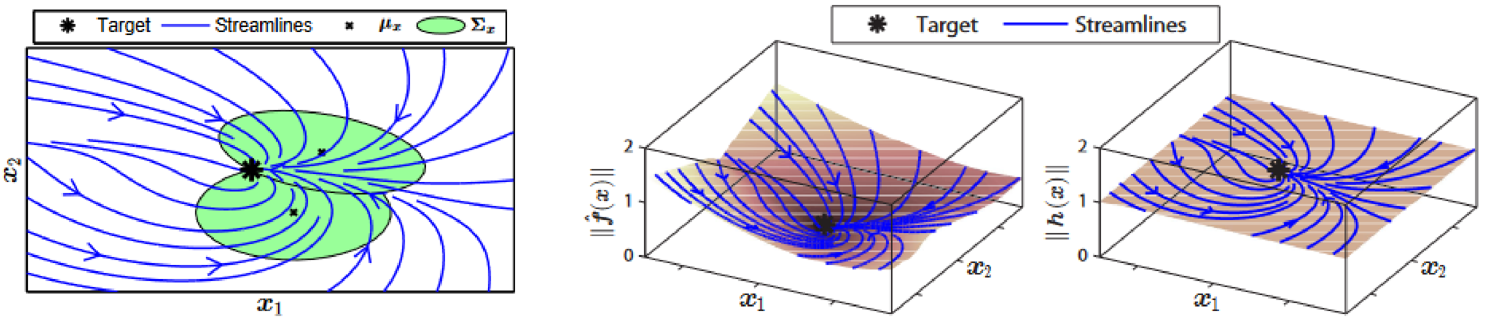

SEDS

Normalized Streamlines

Getting a Target Field from SEDS II

- Getting a normalized field of motion to reach the target with a non-zero velocity:

$$\boldsymbol{h}(\boldsymbol{x};\boldsymbol{\theta}) = \frac{\hat{\boldsymbol{f}}(\boldsymbol{x};\boldsymbol{\theta})}{\|\hat{\boldsymbol{f}}(\boldsymbol{x};\boldsymbol{\theta})\|} \qquad \forall \boldsymbol{x} \in \mathbb{R}^3 \backslash \boldsymbol{x}^*$$

Target Position

$$\boldsymbol{h}(\boldsymbol{x}^*;\boldsymbol{\theta}) = \lim_{\boldsymbol{x}\to\boldsymbol{x}^*} \boldsymbol{h}(\boldsymbol{x};\boldsymbol{\theta})$$

- The vector field $\boldsymbol{h}(\boldsymbol{x};\boldsymbol{\theta})$ conserves the convergence of the SEDS flow but induces a flow of constant speed.

Modified SEDS

- Modification of the SEDS for trajectories with non-zero velocities at the target point:

$$\dot{x} = f_h(x) = v(x)h(x)$$

Strength Factor

Target Field/
velocity vector with constant speed

Modified SEDS Optimization Problem

$$\min_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = -\sum_{n=1}^{N}\sum_{t=0}^{T^n} \omega^{t,n} \frac{(\dot{\boldsymbol{x}}^{t,n})^T \dot{\boldsymbol{x}}^{t,n}(\boldsymbol{\theta})}{\|\dot{\boldsymbol{x}}^{t,n}\|\|\dot{\boldsymbol{x}}^{t,n}(\boldsymbol{\theta})\|}$$

$$\begin{cases} \boldsymbol{\mu}_{\dot{\boldsymbol{x}}}^k + \boldsymbol{\Sigma}_{\dot{\boldsymbol{x}}\boldsymbol{x}}^k (\boldsymbol{\Sigma}_{\boldsymbol{x}}^k)^{-1}(\boldsymbol{x}^* - \boldsymbol{\mu}_{\boldsymbol{x}}^k) = 0 \\ \boldsymbol{\Sigma}_{\dot{\boldsymbol{x}}\boldsymbol{x}}^k (\boldsymbol{\Sigma}_{\boldsymbol{x}}^k)^{-1} + (\boldsymbol{\Sigma}_{\boldsymbol{x}}^k)^{-1}(\boldsymbol{\Sigma}_{\dot{\boldsymbol{x}}\boldsymbol{x}}^k)^T \prec 0 \\ -\boldsymbol{\Sigma}^k \prec 0 \\ 0 < \pi^k \leq 1 \\ \sum_{k=1}^{K} \pi^k = 1 \end{cases} \quad \forall k \in 1..K$$

$\boldsymbol{\theta}^k = \{\pi^k, \boldsymbol{\mu}^k, \boldsymbol{\Sigma}^k\}$ and $\boldsymbol{\theta} = \{\boldsymbol{\theta}^1..\boldsymbol{\theta}^K\}$

The optimization problem minimizes the angle between the demonstrations ($\dot{\boldsymbol{x}}^{t,n}$) and estimations ($\dot{\boldsymbol{x}}^{t,n}(\boldsymbol{\theta}) = \hat{\boldsymbol{f}}(\boldsymbol{x}^{t,n};\boldsymbol{\theta})$) as before but with weights ($\omega^{t,n}$) :

$$\omega^{t,n} = \frac{t}{T^n}(\omega_u - \omega_l) + \omega_l$$

- ω_l - weight of the first data point
- ω_u - weight of the last data point

The direction at the target point matters the most in Minigolf.

Target Field

The learned GMR parameters can now be used to estimate the target direction at varying positions:

$$\hat{\boldsymbol{f}}(\boldsymbol{x};\boldsymbol{\theta}) = \sum_{k=1}^{K} h^k(\boldsymbol{x};\boldsymbol{\theta})(\boldsymbol{\mu}^k_{\dot{\boldsymbol{x}}} + \boldsymbol{\Sigma}^k_{\dot{\boldsymbol{x}}\boldsymbol{x}}(\boldsymbol{\Sigma}^k_{\boldsymbol{x}})^{-1}(\boldsymbol{x} - \boldsymbol{\mu}^k_{\boldsymbol{x}})) \quad \text{with} \quad h^k(\boldsymbol{x};\boldsymbol{\theta}) = \frac{\pi^k \mathcal{N}(\boldsymbol{x};\boldsymbol{\theta}^k)}{\sum_{i=1}^{K} \pi^i \mathcal{N}(\boldsymbol{x};\boldsymbol{\theta}^i)}$$

$$\boldsymbol{h}(\boldsymbol{x};\boldsymbol{\theta}) = \frac{\hat{\boldsymbol{f}}(\boldsymbol{x};\boldsymbol{\theta})}{\|\hat{\boldsymbol{f}}(\boldsymbol{x};\boldsymbol{\theta})\|} \qquad \forall \boldsymbol{x} \in \mathbb{R}^3 \backslash \boldsymbol{x}^*$$

Normalized Streamlines

Strength Factor

The strength factor $v(\boldsymbol{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ is a positive scalar and defines the intensity / velocity of a motion which the robot should follow.

An estimate of the strength factor can be learned from demonstrations through various regression techniques like GMR with regard to $v(\boldsymbol{x}) > 0$.

E.g. GMR: $$v(\boldsymbol{x}) = \sum_{k=1}^{K_{SF}} h_{SF}^k(\boldsymbol{x}) \left(\boldsymbol{\mu}_{SF,v}^k + \boldsymbol{\Sigma}_{SF,v\boldsymbol{x}}^k (\boldsymbol{\Sigma}_{SF,\boldsymbol{x}}^k)^{-1} (\boldsymbol{x} - \boldsymbol{\mu}_{SF,\boldsymbol{x}}^k)\right)$$

Control of Hitting Direction

- Default hitting speed and direction are given through the demonstrations.
- To change the hitting direction and hitting speed, proceed as follows:

$$\dot{\boldsymbol{x}} = \kappa\, \boldsymbol{R}_\alpha\, \boldsymbol{f}_h(R_\alpha^T \boldsymbol{x};\boldsymbol{\theta}) = \kappa\, \boldsymbol{R}_\alpha\, v(\boldsymbol{R}_\alpha^T \boldsymbol{x})\, \boldsymbol{h}(\boldsymbol{R}_\alpha^T \boldsymbol{x};\boldsymbol{\theta})$$

3. Define the hitting speed with gain κ.

1. Transform the input to the desired reference frame.

2. Transform the result back to the desired hitting direction

4

Hitting Parameters

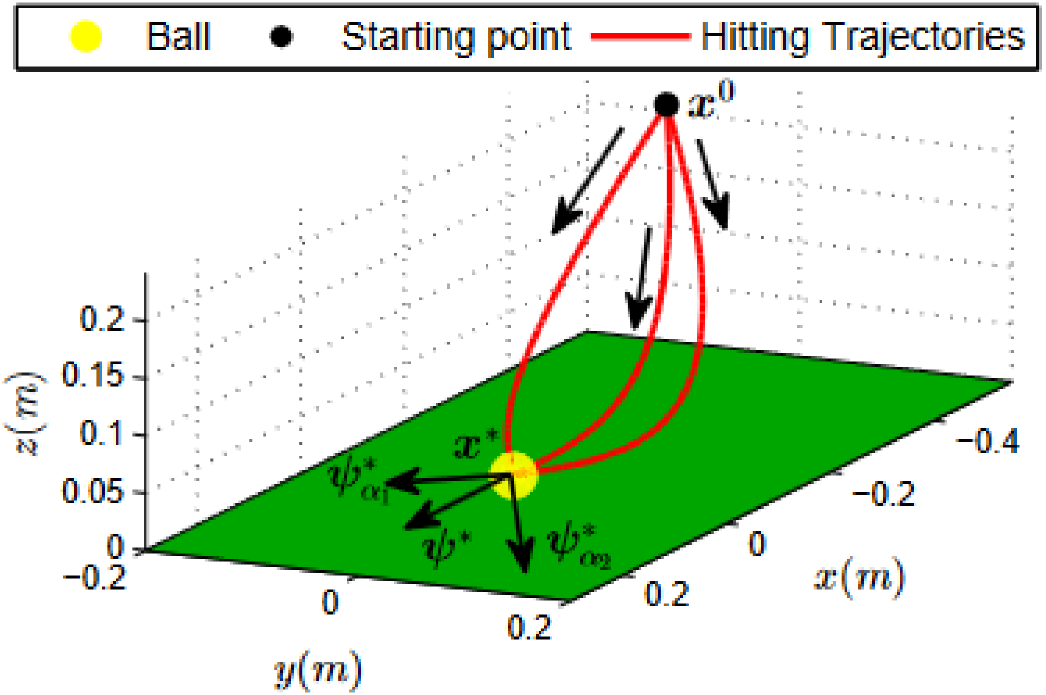

Picture: Khansari et al.

Hitting Parameters

- The goal is to find good hitting parameters for varying situations.
- Hitting parameters:
  - Hitting speed
  - Hitting direction
- Situation:
  - Ball position
  - Goal position

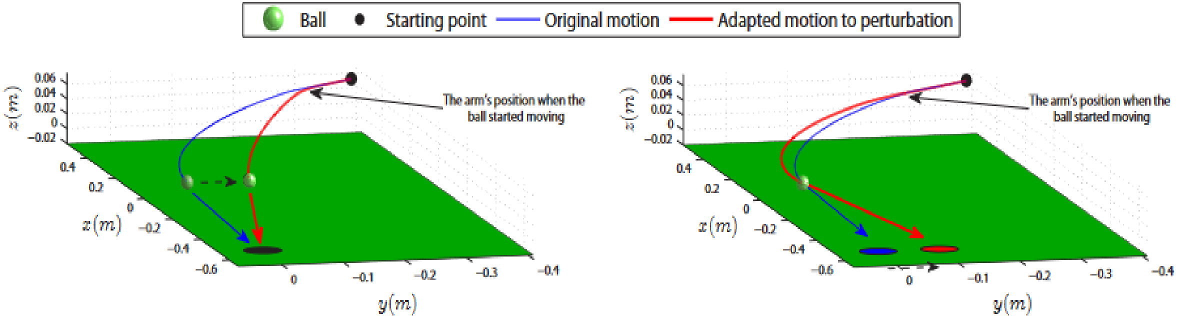

Picture: Khansari et al.

Training Data

$$\{S,\alpha,\kappa\} = \{s^m, g_\alpha(s^m) + \epsilon_\alpha, g_\kappa(s^m) + \epsilon_\kappa\}_{m=1}^{M}$$

Training Data
- Situation (S)
- Angle (α)
- Gain (κ)

Function that maps the situation to a successful angle

Function that maps the situation to a successful speed

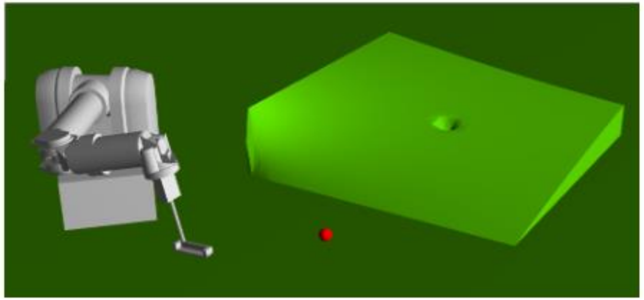

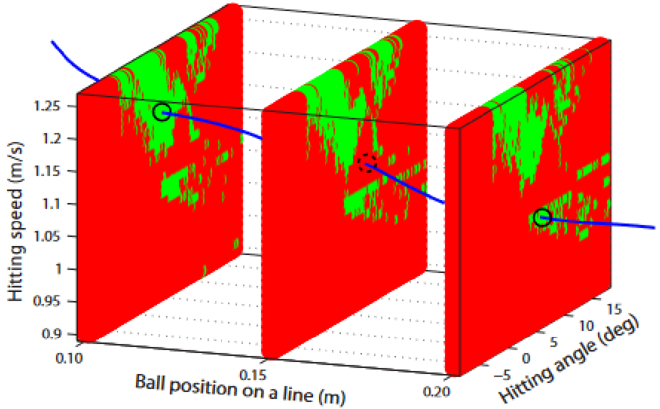

Pictures: Khansari et al.

Hitting Parameters Prediction with GMR

1. The parameters can be optimized to maximize the likelihood of the training set (e.g. Expectation-Maximization Algorithm).

2. Finding hitting parameters for unseen inputs with GMR:

$$\hat{g}(s^*) = \sum_{k=1}^{K_{HP}} h_{HP}^k(\boldsymbol{s}) \left(\boldsymbol{\mu}_{HP,\alpha\kappa}^k + \boldsymbol{\Sigma}_{HP,\alpha\kappa\boldsymbol{s}}^k (\boldsymbol{\Sigma}_{HP,\boldsymbol{s}}^k)^{-1} (\boldsymbol{s} - \boldsymbol{\mu}_{HP,\boldsymbol{s}}^k)\right)$$

Prediction of successful hitting parameters

Non-linear part (as on previous slides)

Comparison of GMR and Gaussian Process Regression

Situation	Method	Attempts	Successful	Ratio
Flat field	GMR	10	9	90%
	GPR	10	10	100%
Multihill field	GMR	10	8	80%
	GPR	10	8	80%
Arctan field (sim)	GMR	30	24	80%
	GPR	30	28	93%

Adopted from Khansari et al.

5

Minigolf Workflow

Photo: Khansari et al.

Stages

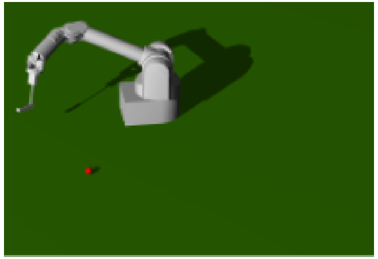

(a) Rest position

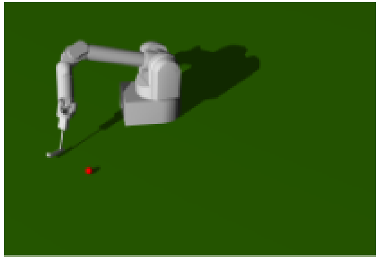

(b) Hitting motion

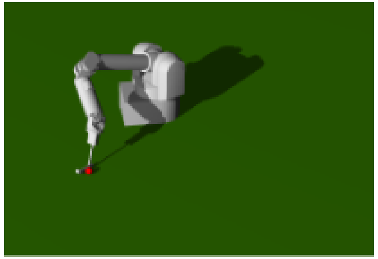

(c) Hit!

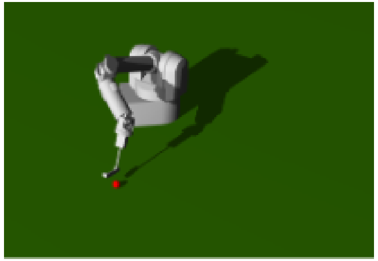

(d) Braking

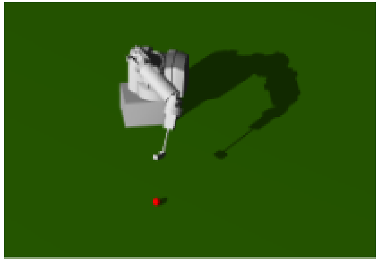

(e) Braking

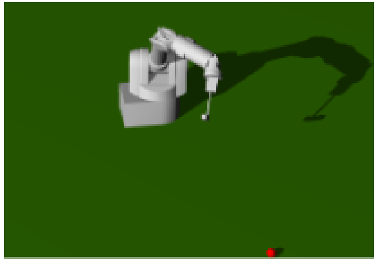

(f) Idle

Picture: Khansari et al.

Training

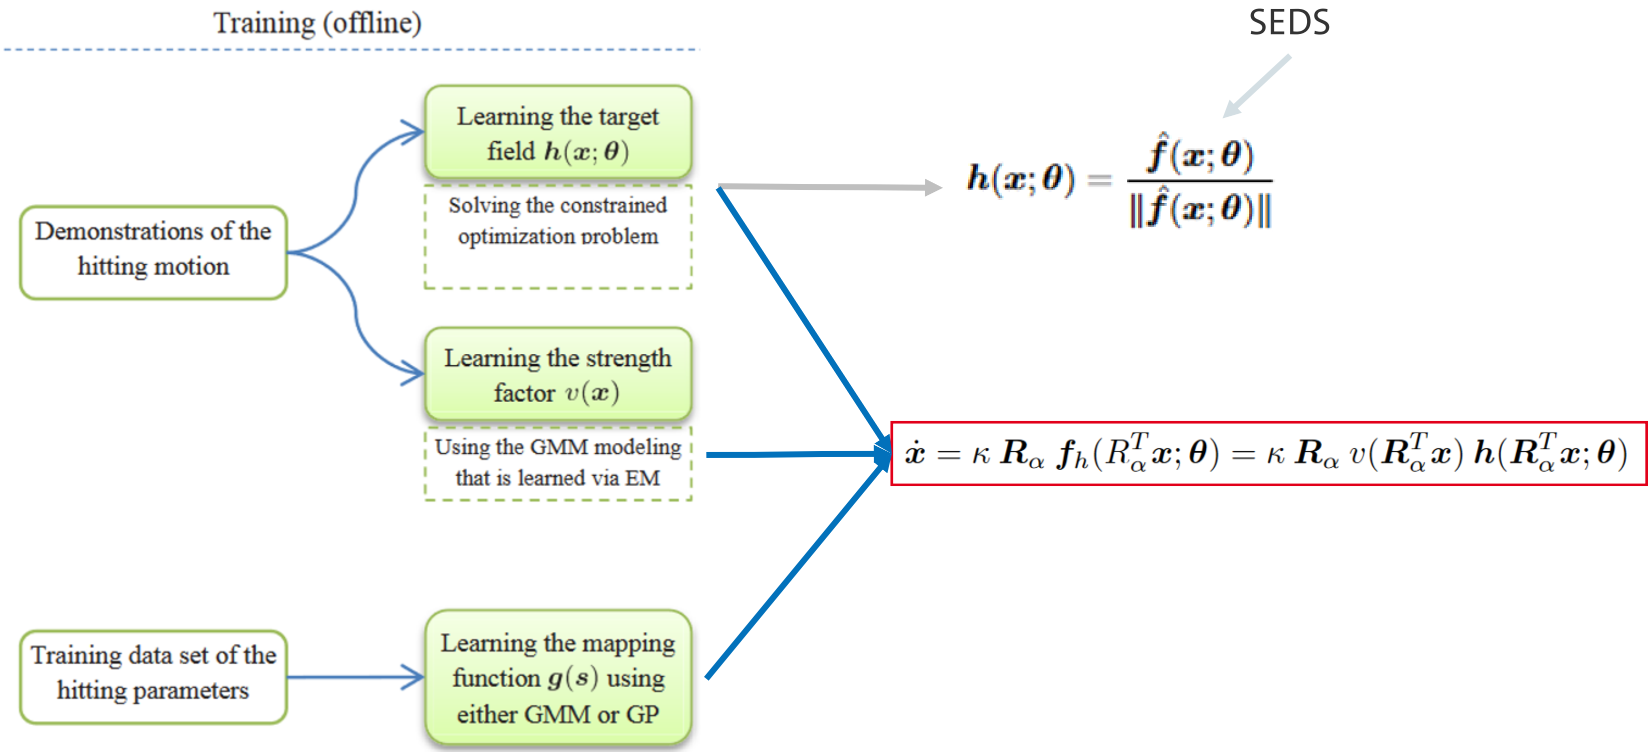

Picture: Khansari et al.

Execution

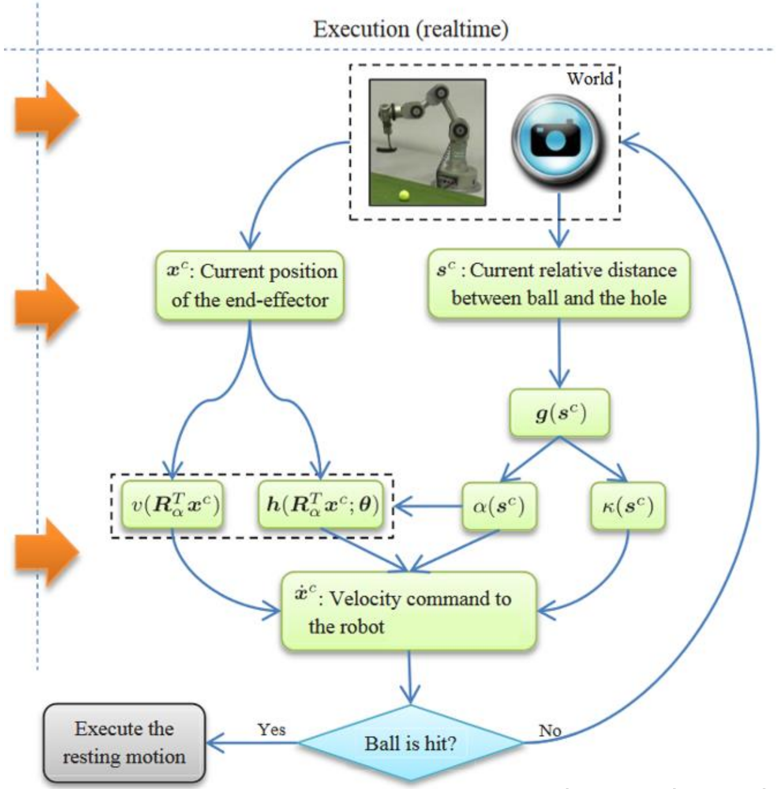

Picture: Khansari et al.

Workflow

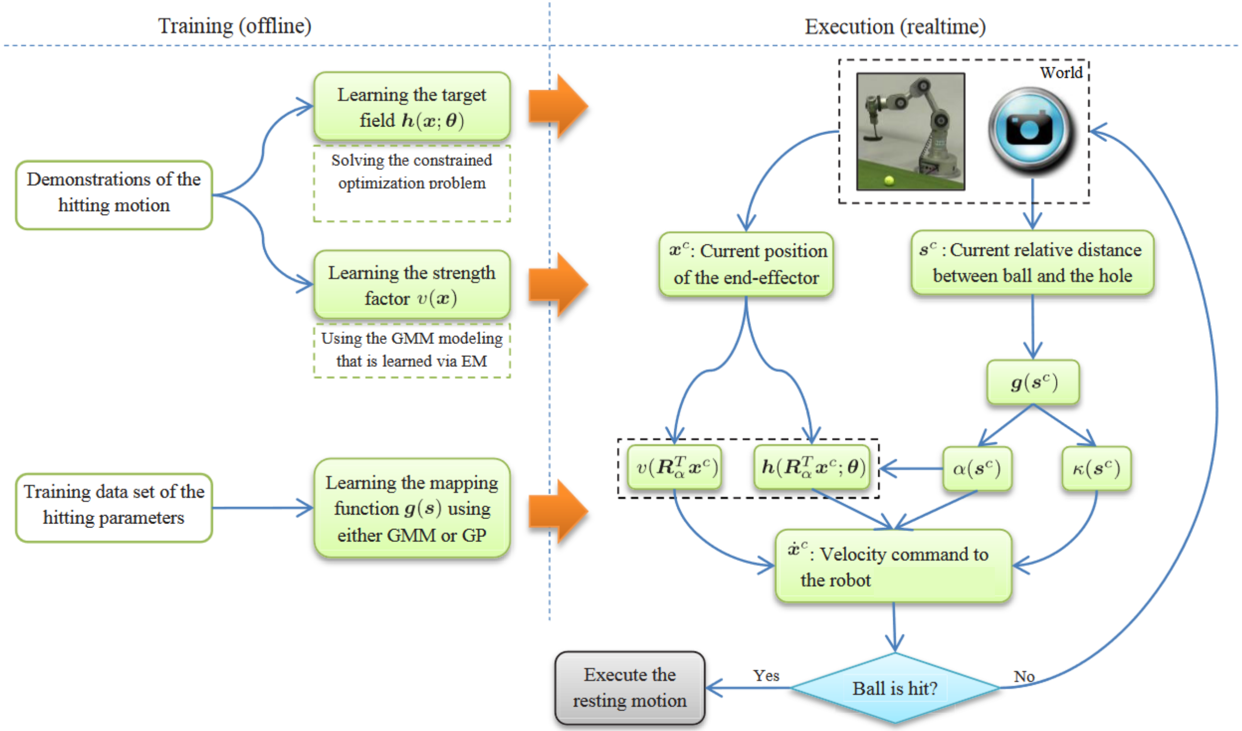

Picture: Khansari et al.

Video: Robot playing mini golf on challenging fields

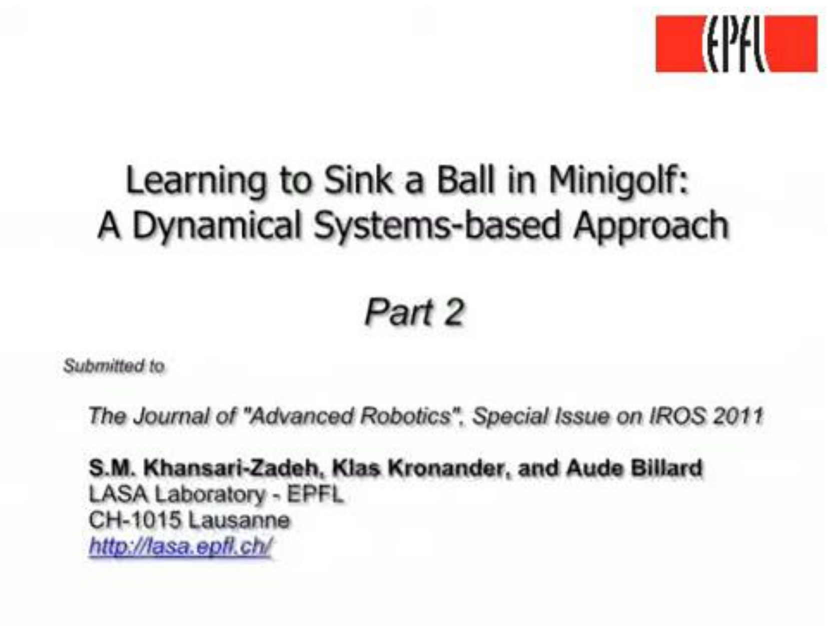

https://www.youtube.com/watch?v=agGZ8itP830

6

Conclusion

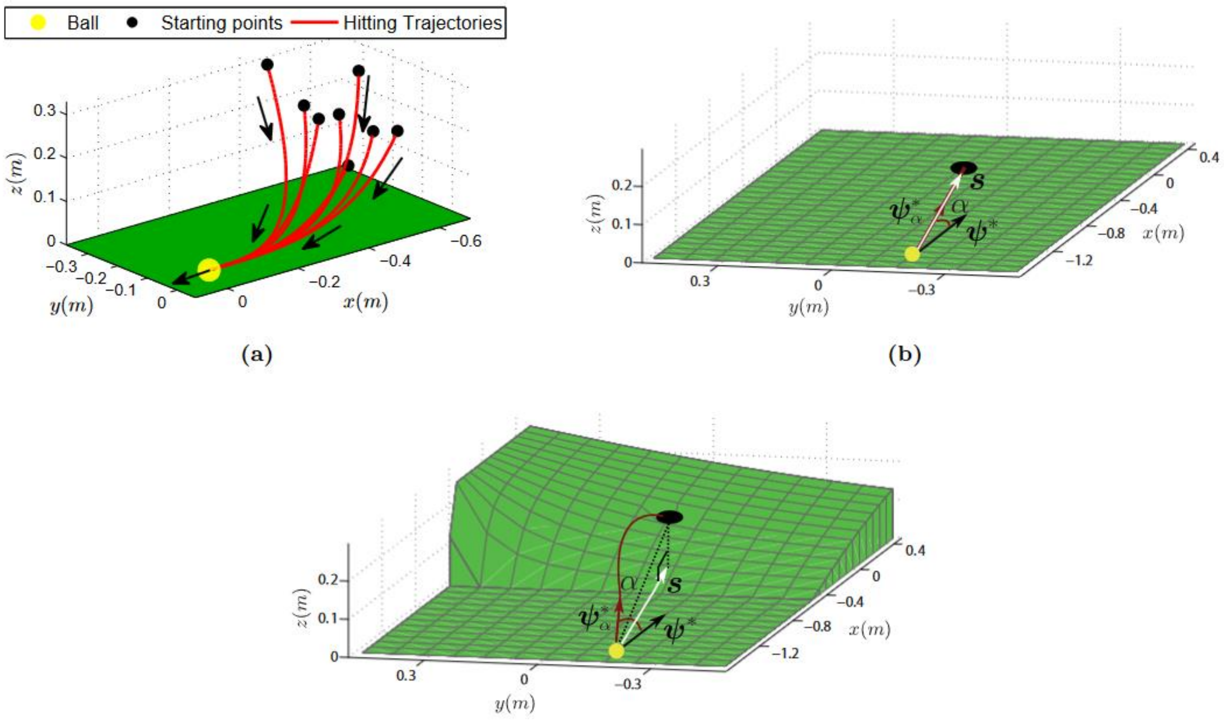

Picture: Khansari et al.

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Conclusion

- The task to learn Mini Golf can be separated into two subtasks:
    1. To learn how to hit the ball.
    2. To learn successful hitting angles and hitting speeds.
- The modified SEDS is a powerful tool for the first task:
    - Changes to the start position are easy to implement.
    - The effects of disturbances when swinging the club are dampened.

7

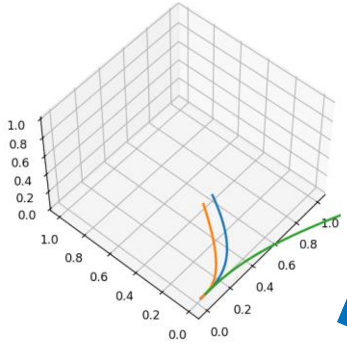

SEDS

Integration into the Minigolf Project

Picture: www.mybotshop.de

Current State

Simulated demonstrations

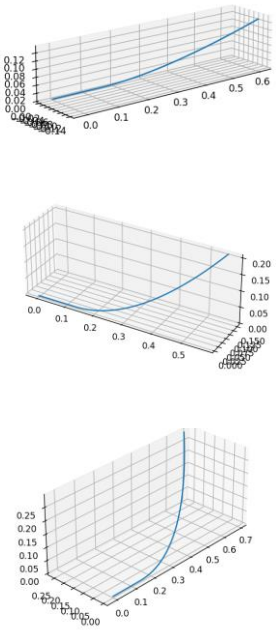

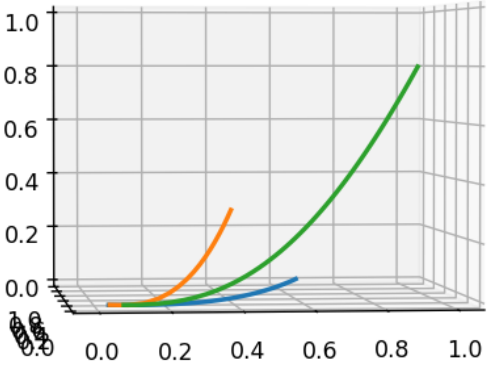

Simulation of the trajectory with the SEDS Model

Todo

- Building a workflow to achieve the suggested velocities from the SEDS on the robot.
  - Inverse Kinematics
  - Collision Detection
- Optional (might be necessary to achieve high velocities):
  - Building a model to estimate good start positions
  - Building a model to estimate a good acceleration profile

8

Discussion

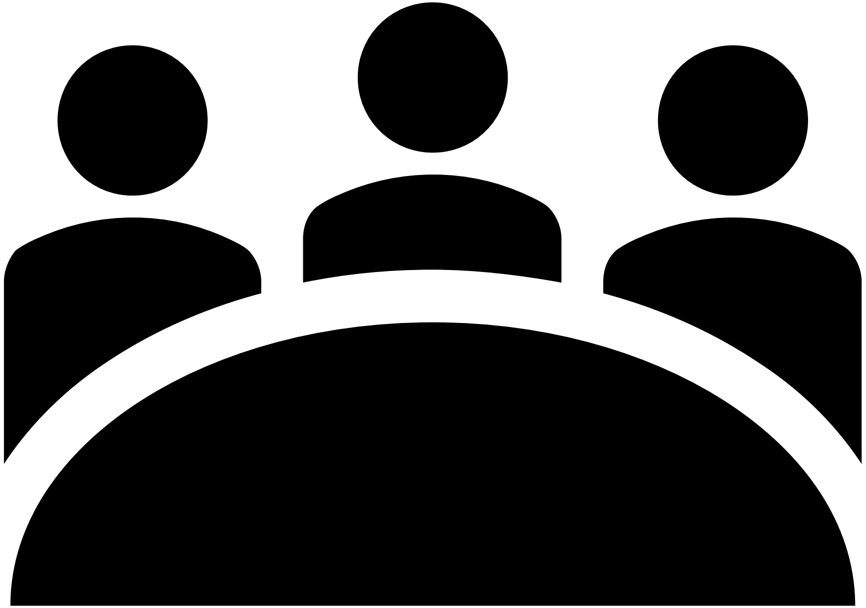

References

- S. M. Khansari-Zadeh and A. Billard, Learning Stable Nonlinear Dynamical Systems With Gaussian Mixture Models, IEEE Transactions on Robotics, vol. 27, no. 5, pp. 943-957, 2011, DOI: 10.1109/TRO.2011.2159412.
- S.M. Khansari-Zadeh, K. Kronander & A. Billard, Learning to Play Minigolf: A Dynamical System-Based Approach, Advanced Robotics, vol. 26, no. 17, pp. 1967-1993, 2012, DOI: 10.1080/01691864.2012.728692.
- K. Kronander, M. S. M. Khansari-Zadeh and A. Billard, Learning to control planar hitting motions in a minigolf-like task, 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, San Francisco, pp. 710-717, 2011, DOI: 10.1109/IROS.2011.6094402.