



Kernel Machines and Reinforcement Learning, Workshop of the 23rd ICML, Pittsburgh

Equi-gradient TD Learning TD(λ) on adaptative Basis Functions Network

Manuel Loth,
Manuel Davy, Rémi Coulom, Philippe Preux

 **INRIA**
FUTURS *Sequel* Lille, France

June 29th 2006

Introduction

- Kernel methods
new *Regression method* \sim kernel method
(regularized sample-based linear approximation)

- Reinforcement learning

Introduction

- Kernel methods
new *Regression method* \sim kernel method
(regularized sample-based linear approximation)

equi-gradient descent

- Reinforcement learning

Introduction

- Kernel methods
new *Regression method* \sim kernel method
(regularized sample-based linear approximation)

equi-gradient descent

- Reinforcement learning
Scheme for using it in *TD(λ)*

Introduction

- Kernel methods

new *Regression method* \sim kernel method
(regularized sample-based linear approximation)

equi-gradient descent

- Reinforcement learning

Scheme for using it in *TD(λ)*

equi-gradient TD(λ)

Regression

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathbb{R}$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(\mathbf{x}_1), \dots, \hat{f}(\mathbf{x}_n))^\top$$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(\mathbf{x}_1), \dots, \hat{f}(\mathbf{x}_n))^\top$$

minimize $\mathcal{L}(\mathbf{y}, \hat{\mathbf{y}})$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(\mathbf{x}_1), \dots, \hat{f}(\mathbf{x}_n))^\top$$

$$\text{minimize } \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \left| \hat{f} \right|.$$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(\mathbf{x}_1), \dots, \hat{f}(\mathbf{x}_n))^\top$$

$$\text{minimize } \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \left| \hat{f} \right|.$$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(\mathbf{x}_1), \dots, \hat{f}(\mathbf{x}_n))^\top$$

$$\text{minimize } \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \|\hat{f}\|$$

kernel methods:

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(x_1), \dots, \hat{f}(x_n))^\top$$

$$\text{minimize } \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \left| \hat{f} \right|.$$

kernel methods:

- $k : \mathcal{X}^2 \rightarrow \mathbb{R} \sim$ similarity

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(x_1), \dots, \hat{f}(x_n))^\top$$

$$\text{minimize } \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \|\hat{f}\|$$

kernel methods:

- $k : \mathcal{X}^2 \rightarrow \mathbb{R} \sim$ similarity
- 1 basis function / learning point: $\phi_i(\mathbf{x}) = k(x_i, \mathbf{x})$

Regression

- unknown $f : \mathcal{X} \rightarrow \mathbb{R}$
- samples $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathbb{R}$
- **accurate** and **simple** approximation \hat{f}

$$\mathbf{y} = (y_1, \dots, y_n)^\top$$

$$\hat{\mathbf{y}} = (\hat{f}(x_1), \dots, \hat{f}(x_n))^\top$$

$$\text{minimize } \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \|\hat{f}\|.$$

kernel methods:

- $k : \mathcal{X}^2 \rightarrow \mathbb{R} \sim$ similarity
- 1 basis function / learning point: $\phi_i(\mathbf{x}) = k(\mathbf{x}_i, \mathbf{x})$
- $\hat{f}(\mathbf{x}) = \sum_i w_i \phi_i(\mathbf{x})$

regularized sample-based linear approximators

regularized sample-based linear approximators

- samples

regularized sample-based linear approximators

- samples \rightsquigarrow basis functions $\phi_1, \dots, \phi_m : \mathcal{X} \rightarrow \mathbb{R}$

regularized sample-based linear approximators

- samples \rightsquigarrow basis functions $\phi_1, \dots, \phi_m : \mathcal{X} \rightarrow \mathbb{R}$
- $\hat{f}(\mathbf{x}) = \sum_i w_i \phi_i(\mathbf{x})$

regularized sample-based linear approximators

- samples \rightsquigarrow basis functions $\phi_1, \dots, \phi_m : \mathcal{X} \rightarrow \mathbb{R}$
- $\hat{f}(x) = \sum_i w_i \phi_i(x)$

$$\Phi = \begin{pmatrix} \phi_1(x_1) & \dots & \phi_m(x_1) \\ \vdots & \ddots & \vdots \\ \phi_1(x_n) & \dots & \phi_m(x_n) \end{pmatrix}, \quad \mathbf{w} = (w_1, \dots, w_m)^T$$

regularized sample-based linear approximators

- samples \rightsquigarrow basis functions $\phi_1, \dots, \phi_m : \mathcal{X} \rightarrow \mathbb{R}$
- $\hat{f}(x) = \sum_i w_i \phi_i(x)$

$$\Phi = \begin{pmatrix} \phi_1(x_1) & \dots & \phi_m(x_1) \\ \vdots & \ddots & \vdots \\ \phi_1(x_n) & \dots & \phi_m(x_n) \end{pmatrix}, \quad \mathbf{w} = (w_1, \dots, w_m)^\top$$

$$\hat{\mathbf{y}} = \Phi \mathbf{w}$$

regularized sample-based linear approximators

- samples \rightsquigarrow basis functions $\phi_1, \dots, \phi_m : \mathcal{X} \rightarrow \mathbb{R}$
- $\hat{f}(x) = \sum_i w_i \phi_i(x)$

$$\Phi = \begin{pmatrix} \phi_1(x_1) & \dots & \phi_m(x_1) \\ \vdots & \ddots & \vdots \\ \phi_1(x_n) & \dots & \phi_m(x_n) \end{pmatrix}, \quad \mathbf{w} = (w_1, \dots, w_m)^\top$$

$$\hat{\mathbf{y}} = \Phi \mathbf{w}$$

$$\mathcal{L}(\mathbf{y}, \Phi \mathbf{w}) + \lambda |\Phi \mathbf{w}|$$

LASSO

LASSO

$$\mathcal{L}(\mathbf{y}, \Phi \mathbf{w}) + \lambda |\Phi \mathbf{w}|$$

LASSO

$$(\mathbf{y} - \Phi\mathbf{w})^2 + \lambda |\Phi\mathbf{w}|$$

LASSO

$$(\mathbf{y} - \Phi\mathbf{w})^2 + \lambda \sum_i |w_i| \quad (\text{LASSO})$$

LASSO

$$(\mathbf{y} - \Phi\mathbf{w})^2 + \lambda \sum_i |w_i| \quad (\text{LASSO})$$

- Basis Pursuit [Chen 95](#)

LASSO

$$(\mathbf{y} - \Phi\mathbf{w})^2 + \lambda \sum_i |w_i| \quad (\text{LASSO})$$

- Basis Pursuit [Chen 95](#)
- Adaptive Ridge Regression [Granvalet 98](#)

LASSO

$$(\mathbf{y} - \Phi\mathbf{w})^2 + \lambda \sum_i |w_i| \quad (\text{LASSO})$$

- Basis Pursuit [Chen 95](#)
- Adaptive Ridge Regression [Granvalet 98](#)
- Matching Pursuit [Mallat & Zhang 93](#)

LASSO

$$(\mathbf{y} - \Phi\mathbf{w})^2 + \lambda \sum_i |w_i| \quad (\text{LASSO})$$

- Basis Pursuit [Chen 95](#)
- Adaptive Ridge Regression [Granvalet 98](#)
- Matching Pursuit [Mallat & Zhang 93](#)
- iterative gradient descent [Osborne *et al.* 00](#)

Equi-gradient descent

Least-Angle Regression Stagewise/l_aSSO

Equi-gradient descent

Least-Angle Regression Stagewise/l_aSSO

- Efron 2002 variable selection

Equi-gradient descent

Least-Angle Regression Stagewise/laSSO

- Efron 2002 variable selection
- Guigue 2005 kernelization

Equi-gradient descent

Least-Angle Regression Stagewise/laSSO

- Efron 2002 variable selection
- Guigue 2005 kernelization
- Generalization, simplification → *Equi-gradient descent*

Equi-gradient descent

Equi-gradient descent

Given λ

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^T$

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^T$

$\mathbf{w}_a =$ **active** (non-zero) weights

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^T$

$\mathbf{w}_a =$ **active** (non-zero) weights

$\Phi_a =$ active basis functions

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^T$

$\mathbf{w}_a =$ **active** (non-zero) weights

$\Phi_a =$ active basis functions

$\hat{\mathbf{y}} = \Phi_a \mathbf{w}_a$

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^\top$

$\mathbf{w}_a = \text{active}$ (non-zero) weights

$\Phi_a = \text{active}$ basis functions

$\hat{\mathbf{y}} = \Phi_a \mathbf{w}_a$

$$\text{minimize } (\mathbf{y} - \Phi_a \mathbf{w}_a)^2 + \lambda \sum_i |w_i|$$

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^\top$

$\mathbf{w}_a = \text{active}$ (non-zero) weights

$\Phi_a = \text{active basis functions}$

$\hat{\mathbf{y}} = \Phi_a \mathbf{w}_a$

$$\text{minimize } (\mathbf{y} - \Phi_a \mathbf{w}_a)^2 + \lambda \sum_i |w_i|$$

$$\Phi_a^\top (\mathbf{y} - \Phi_a \mathbf{w}_a) = \lambda \text{sgn}_a$$

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^\top$

$\mathbf{w}_a =$ **active** (non-zero) weights

$\Phi_a =$ active basis functions

$\hat{\mathbf{y}} = \Phi_a \mathbf{w}_a$

$$\text{minimize } (\mathbf{y} - \Phi_a \mathbf{w}_a)^2 + \lambda \sum_i |w_i|$$

$$\Phi_a^\top (\mathbf{y} - \Phi_a \mathbf{w}_a) = \lambda \text{sgn}_a$$

$$\text{for all active } w_i, \quad \left| \frac{\partial \mathcal{L}}{\partial w_i} \right| = \lambda$$

Equi-gradient descent

Given λ

Given $\text{sgn}(\mathbf{w}) = (0, 0, -1, 0, 1, \dots)^\top$

$\mathbf{w}_a =$ **active** (non-zero) weights

$\Phi_a =$ active basis functions

$\hat{\mathbf{y}} = \Phi_a \mathbf{w}_a$

$$\text{minimize } (\mathbf{y} - \Phi_a \mathbf{w}_a)^2 + \lambda \sum_i |w_i|$$

$$\Phi_a^\top (\mathbf{y} - \Phi_a \mathbf{w}_a) = \lambda \text{sgn}_a$$

$$\text{for all active } w_i, \quad \left| \frac{\partial \mathcal{L}}{\partial w_i} \right| = \lambda$$

$$\text{for all inactive } w_i, \quad \left| \frac{\partial \mathcal{L}}{\partial w_i} \right| \leq \lambda$$

Regularization path


Regularization path

samples

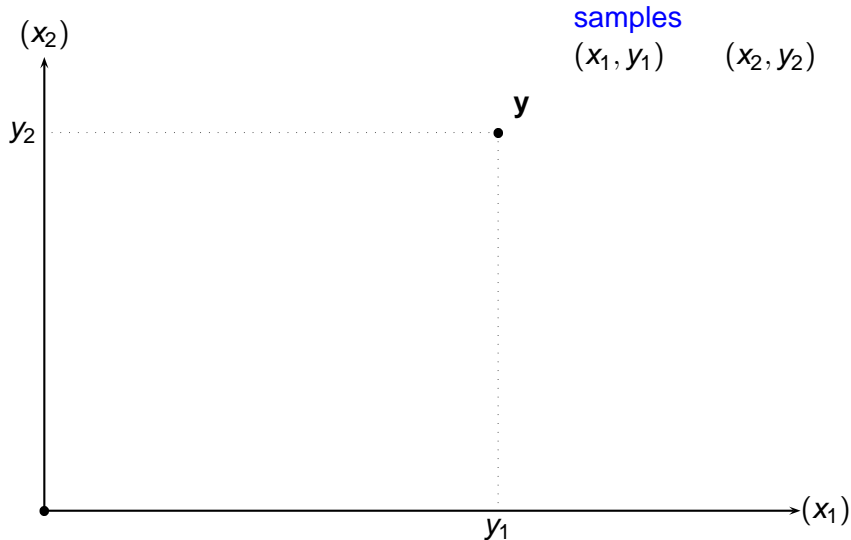
(x_1, y_1)

(x_2, y_2)


Regularization path



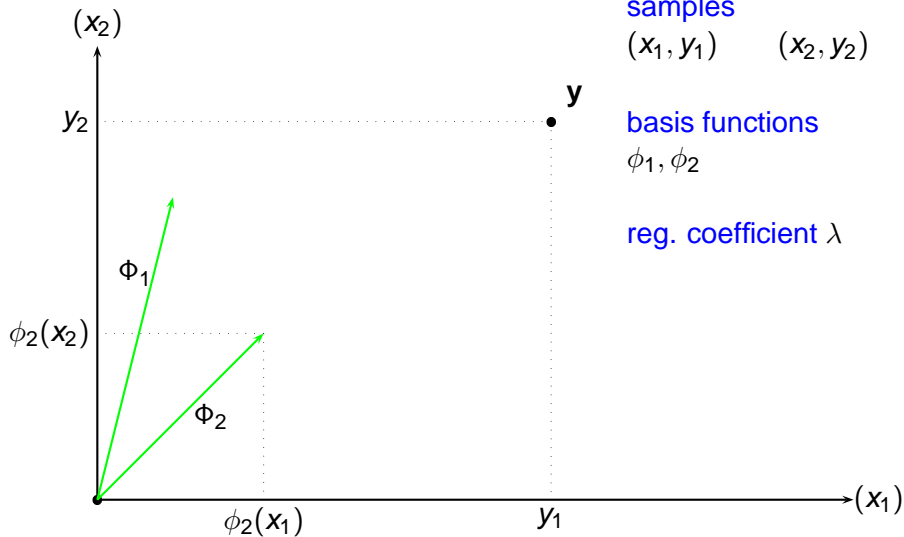
Regularization path




Regularization path




Regularization path



Regularization path



Regularization path



samples

(x_1, y_1) (x_2, y_2)

basis functions


ϕ_1, ϕ_2

reg. coefficient λ

$\hat{f}_\lambda = w_1 \phi_1(x) + w_2 \phi_2(x)$

$\hat{y}_\lambda = w_1 \Phi_1 + w_2 \Phi_2$

Regularization path



samples

(x_1, y_1) (x_2, y_2)

basis functions

ϕ_1, ϕ_2

reg. coefficient λ

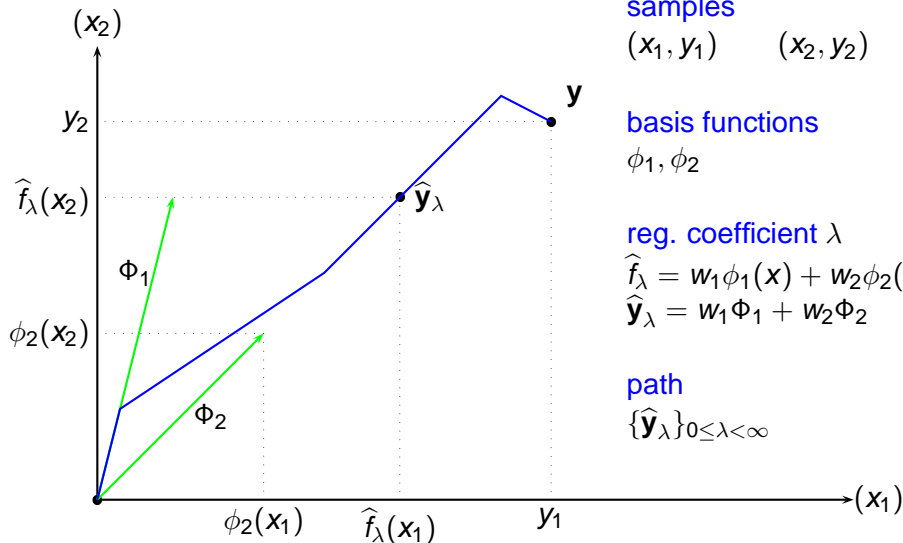
$\hat{f}_\lambda = w_1\phi_1(x) + w_2\phi_2(x)$

$\hat{y}_\lambda = w_1\Phi_1 + w_2\Phi_2$

path

$\{\hat{y}_\lambda\}_{0 \leq \lambda < \infty}$

Regularization path



samples

(x_1, y_1) (x_2, y_2)

basis functions

ϕ_1, ϕ_2

reg. coefficient λ


$\hat{f}_\lambda = w_1\phi_1(x) + w_2\phi_2(x)$

$\hat{y}_\lambda = w_1\phi_1 + w_2\phi_2$


path

$\{\hat{y}_\lambda\}_{0 \leq \lambda < \infty}$


Regularization path




Regularization path



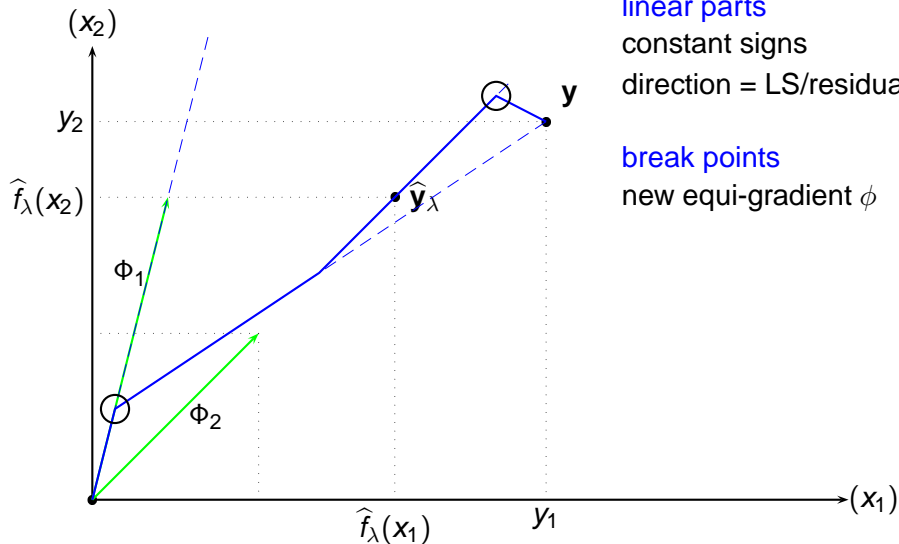
Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path




Regularization path



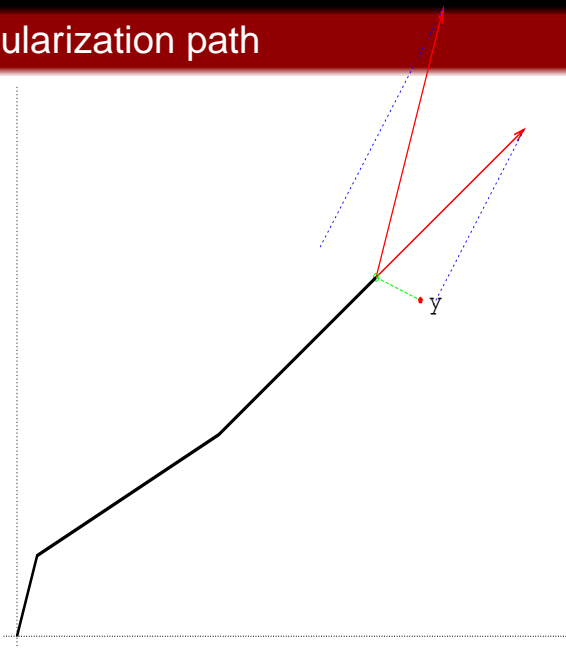
Regularization path




Regularization path




Regularization path



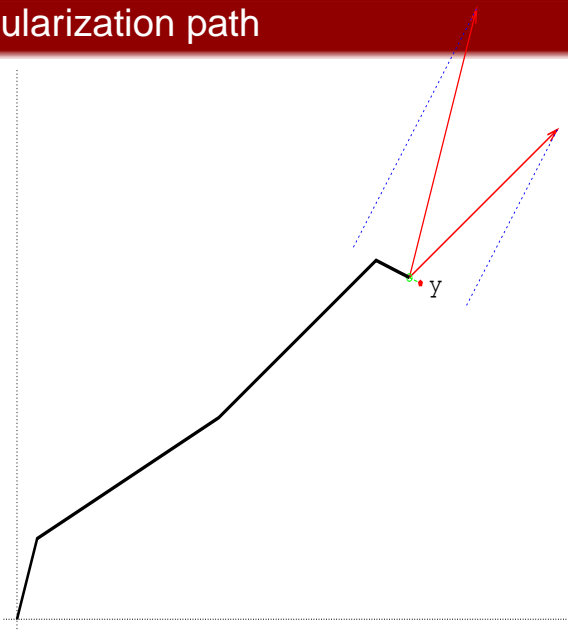
Regularization path




Regularization path



Regularization path



Regularization path



Complexity

Each step (activation/deactivation) is $O(|\mathcal{D}| + |\mathcal{A}|^2)$

- linear in dictionary size
- quadratic in number of active basis functions


Demo

Reinforcement Learning

- TD(λ)
- continuous state space
- discrete time
- updates after each trajectory


Radial Basis Functions Network

$$\hat{V}(\mathbf{x}) = \sum w_i \phi_i(\mathbf{x})$$




Radial Basis Functions Network

$$\hat{V}(\mathbf{x}) = \sum w_i \phi_i(\mathbf{x})$$




Update

$$\hat{V}(\mathbf{x}) = \sum (w_i + \Delta w_i) \phi_i(\mathbf{x})$$




Update

$$\hat{V}(\mathbf{x}) = \sum w_i \phi_i(\mathbf{x}) + \sum \Delta w_i \phi_i(\mathbf{x})$$




Independant regression

temporal differences




Gradient descent on fixed basis network

temporal differences



Equi-gradient descent on extended network

temporal differences



Temporal regularization

Temporal regularization

Continuously add basis functions?!

Temporal regularization

Continuously add basis functions?!

- Put a preference on existing basis functions

Temporal regularization

Continuously add basis functions?!

- Put a preference on existing basis functions
- Remove zero-weighted basis functions

Temporal regularization

Continuously add basis functions?!

- Put a preference on existing basis functions
 $\phi_i \leftarrow \rho\phi_i$
- Remove zero-weighted basis functions

Temporal regularization

Continuously add basis functions?!


- Put a preference on existing basis functions
$$\phi_i \leftarrow \rho \phi_i \Rightarrow w_i \leftarrow \frac{1}{\rho} w_i$$
- Remove zero-weighted basis functions

Preliminary simple experiments


Preliminary simple experiments

- Inverted pendulum
- Gaussian basis functions on normalized state space
- Updates after each episode
- Stopping EG descents at $|\hat{\mathbf{y}}|^2 = 70\%|\mathbf{y}|^2$


Experiments: preference




Experiments: multi-kernels




Experiments: symmetry



Experiments: comparison



Experiments: comparison



Conclusion

Conclusion

Summary:

Conclusion

Summary:

- **efficient** and **easy** way to select basis functions in TD(λ)

Conclusion

Summary:

- **efficient** and **easy** way to select basis functions in TD(λ)
- **robust**, no unintuitive parameters

Conclusion

Summary:

- **efficient** and **easy** way to select basis functions in TD(λ)
- **robust**, no unintuitive parameters

Perspectives:

Conclusion

Summary:

- **efficient** and **easy** way to select basis functions in TD(λ)
- **robust**, no unintuitive parameters

Perspectives:

- experiments on other problems

Conclusion

Summary:

- **efficient** and **easy** way to select basis functions in TD(λ)
- **robust**, no unintuitive parameters

Perspectives:

- experiments on other problems
- automatically build basis function dictionary based on topology, TD variance, ... (wavelets, low-dimensional projections, ...)

Take home message

- Feature selection is easy!
- Basic use of it in RL \rightarrow
efficient & easy-to-tune TD(λ)