Optimization of a sequential decision problem in prenatal ultrasound

E. Le Pennec
R. Besson - S. Allassonnière

ÉCOLE POLYTECHNIQUE

Journée COSMOS - 07/11/2019

Outline

Birth and Rare Diseases

A Few Numbers

- 780.000 births/year in France, 5 millions births/year in Europe
- 3 to 4% are affected by at least one congenital abnomalies
- Rare diseases: 3 millions patients in France, 30 millions in Europe.

Medical Setting

Prenatal Ultrasound Diagnosis

- **France:** three compulsory ultrasound tests during pregnancy.
- Some classical measures (e.g. Down syndrome).
- **No strict examination protocol.**

Necker Hospital Obstetrician

- Rare disease expertise.
- Among world largest medical database.
- Will to **systematize** their knowledge.

Proposed Tool

Ultrasound as a Sequential Process

- Ultrasound exam seen as a **sequence of measures**.
- **Goals:**
    - **Reduce the time** required to obtain a diagnosis
    - **Avoid to miss** a rare disease.

Diagnosis Assistance Tool

- **Propose** the next measure to make.
- **Show** the current most probable diseases.
- **Easy to use GUI** implemented in **R**!

What's inside this tool?

Charade

Ultrasound Diagnostic

Reset Page

CHARADE
direct au diagnostic

Choix d'une anomalie

ventricular septal defect (0.0361%)

not present | present

Historique

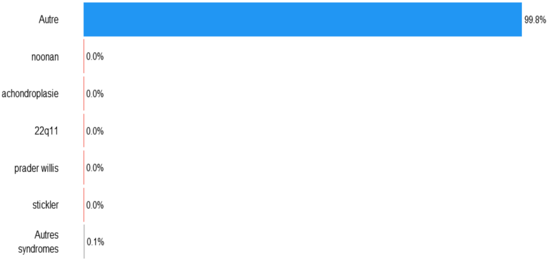

Outline

Data at Hands

id disease	id symptom	probability of symptom knowing the disease
16	29	0.39
16	136	0.67
16	149	0.50
16	176	0.16
16	181	0.50
16	231	0.75

- Rare diseases: very few cases even in the world largest DB!

Excel Type Dataset

- Expert database build from OrphaData (E. Spaggiari).
- 81 diseases, 202 symptoms (signs visible with ultrasound):
  - Disease probability: $P[D = d_j]$
  - Symptom probability given each disease: $P[S_i = k \mid D = d_j]$.

- Database will be enriched from the future exams.

Our Goals

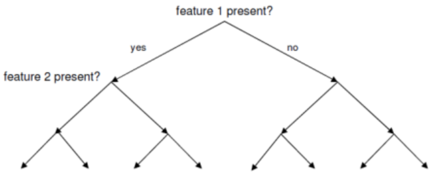

Medical Goals

- Guide a (non rare disease expert) sonographer to assess as fast as possible potential diseases.
- Propose her/him the next symptom to check.

Technical Goals

- Build a *good* decision tree (a *good* policy).
- Develop a GUI that can be easily used.

Markov Decision Process

State, Action and Policy

- State: $\mathbb{S} = \{P, A, U\}^{202}$ (presence,absence,not yet looked at) for each symptom.
- Action: $\mathbb{A} = \{1, \ldots, 202\}$ next symptom.
- Policy: $\pi : s \in \mathbb{S} \mapsto a \in \mathbb{A}$ next symptom given the state.

Probabilistic setting

- Natural Markovian modeling: $\mathcal{S}_{t+1}$ depends only on $\mathcal{S}_t$ and a_t!

Markovian Decision Process

- Any strategy π defines a law on $(\mathcal{S}_t)$ starting from $\mathcal{S}_0$.
- Let T be the stopping time before a diagnosis can be posed.
- We need to find $\pi^\star$ such that $\pi^\star(\mathcal{S}_0) = \operatorname{argmin}_\pi \mathbb{E}[T|\mathcal{S}_0]$!

Problems to be Solved

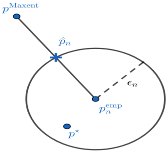

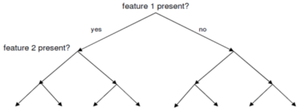

Environment Learning with Maximum Entropy Principle

- We have $P[S_i \mid D]$ but we need to know $P[S_{i_1}, \ldots, S_{i_K} \mid D]$.
- We need to take into account future exams.
- **Idea:** add some expert knowledge and maximize uncertainty, interpolate between the expert model and the data.

Diagnostic Strategy Optimization by Reinforcement Learning

- Find a policy that allows to detect the disease while minimizing the average duration.
- **Idea:** recast the problem as a planning issue and find the optimal strategy.

Outline

Uncertainty and Entropy

Environment Learning

- We have $P[S_i \mid D]$ but we need to know $P[S_{i_1}, ..., S_{i_K} \mid D]$.
- Idea: add some expert knowledge and maximize uncertainty.

Expert knowledge

- Some symptoms can not occur simultaneously...
- Need at least a certain number of symptoms to talk about a syndrome.

Uncertainty

- General idea: choose a solution that maximize the uncertainty while respecting the constraints (probability/impossibility).
- Uncertainty measured by entropy.

MaxEnt Principle

Environment Learning

- We have $P[S_i \mid D]$ but we need to know $P[S_{i_1}, ..., S_{i_K} \mid D]$.

- Naive idea: $P[S_{i_1}, ..., S_{i_K} \mid D] = P[S_{i_1} \mid D] \times ... \times P[S_{i_K} \mid D]$ (Conditional independence)

Data and Expert Knowledge

- Conditional probabilities: $P[S_i \mid D]$
- Medical constraints: $P[S_{i_k}, S_{i_{k'}} \mid D] = 0$...
- Mathematical constraints: P should be a probability...

MaxEnt Principle

- Maximize the entropy of the distribution $P[S_{i_1}, ..., S_{i_k} \mid D]$ under mathematical and medical constraints.

- Numerical scheme available.
- Interp. between maxent and maximum likelihood.

Interp. between expert model and data

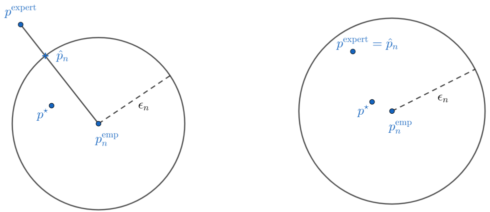

Concentration inequality for the empirical distribution

- With probability $1-\delta$, $\mathbb{KL}(p_n^{\text{emp}}||p^\star) \leq \frac{1}{n}\log\left(\frac{c(f(n)+1)}{\delta}\right)$ with $f(n) = o(n^{\frac{2^{K_1}-2}{2}})$.

- Rk: Pinsker ineq. yields a similar concentration ineq. for L^1.

Interpolation: barycentre between expert model and data

- $$\hat{p}^{\mathcal{L}}_{\epsilon_n} = \operatorname*{argmin}_{p\in\mathcal{C}/\mathcal{L}(p_n^{\text{emp}},p)\leq\epsilon_n} \mathcal{L}(p^{\text{expert}}, p)$$

  where $\epsilon_n := \epsilon_n^\delta = \text{argmin}_l\, \mathbb{P}[\mathcal{L}(p_n^{\text{emp}}, p^\star) \leq l] \geq 1-\delta$.

Some Theoretical Results

Barycentre for the L^1 norm

$$\hat{p}_n^1 = \operatorname*{argmin}_{p \in \mathcal{C} / \|p_n^{\text{emp}} - p\|_1 \leq \epsilon_n} \|p^{\text{expert}} - p\|_1$$

Theorem

- $\exists \alpha_n \in [0,1]$ such that
$$\hat{p}_n^1 = \alpha_n p^{\text{expert}} + (1-\alpha_n) p_n^{\text{emp}}$$
where $\alpha_n = \dfrac{\epsilon_n}{\|p_n^{\text{emp}} - p^{\text{expert}}\|_1}$ if $\epsilon_n \leq \|p_n^{\text{emp}} - p^{\text{expert}}\|_1$ and $\alpha_n = 1$ otherwise.
- For all $n \in \mathbb{N}$, we have with probability at least $1-\delta$:
$$\|p^\star - \hat{p}_n^1\|_1 \leq 2 \min\{\epsilon_n, \|p^\star - p^{\text{expert}}\|_1\}$$

- Smoothed version of the choice between the two solutions.
- Similar result for KL divergence.

Some Numerical Results

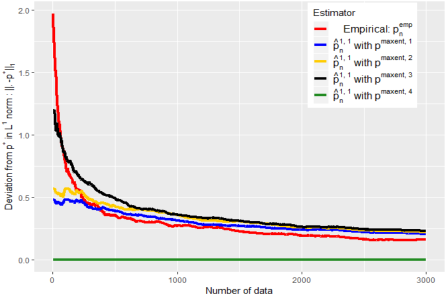

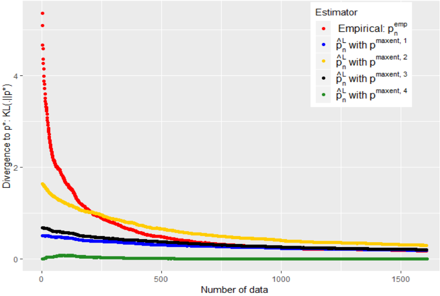

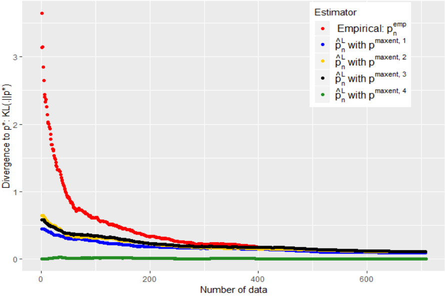

Low Sensitivity to δ

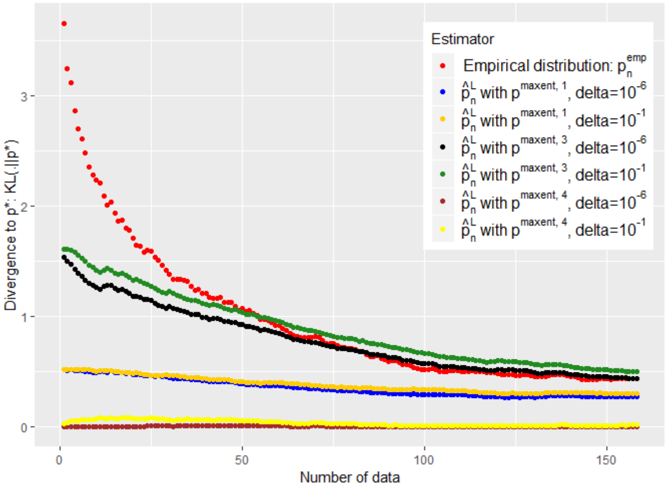

Outline

Diagnostic Strategy Optimization

Diagnostic Strategy Optimization.

- Find a policy that allows to detect the disease while minimizing the average duration.

Measure of Performance

- Number of questions before being able to diagnose a disease.

Alternative Formulations

- Trade-off: cost of misdiagnosis/cost of medical tests to perform.
- Reach the lowest uncertainty under fixed budget constraint (time, money).

Non Adversarial Game

- The disease and symptoms do not change during the exam.
- Strategy: given what has been seen, what is the next symptom to look at?

Stochastic Shortest Path

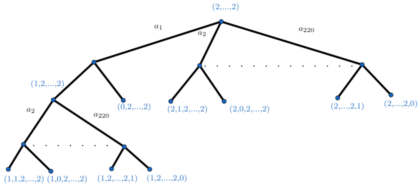

Stochastic Shortest Path

- T is a stooping time at some *final states*
- How to minimize the expectation of T?

Final States

- Entropy based criterion: $H(D \mid \mathcal{S}) \leq \epsilon$

MDP

- Rewards: $\forall \mathcal{S}_t, a_t,\ r(\mathcal{S}_t, a_t) = -1$

How to Solve It?

Dynamic Programming?

- Most natural approach in a MDP setting!
- **Tool:** Fixed point algorithm.
- **Issue:** we have a simulator rather than the MDP transition proba...

Tabular Reinforcement Learning?

- Most natural approach in a simulator setting!
- **Tool:** Stochastic approximation of the fixed point algorithm.
- **Issue:** Very large state space (3^{202})...

Approximate Reinforcement Learning!

- Only remaining direction!
- **Tool:** Functional and stochastic approximation!
- **Issue:** Danger Zone!

Numerical Experiments

Naive approach: Breiman CART

- Greedy policy that optimize the expectation of next step entropy.

Baseline: Actor-critic with REINFORCE

- Linearly parametrized policy using next step entropy expectation and other simple features

Deep Q-Learning

- Q-Learning with Neural Networks.

- Nothing specific for the first two approaches...
- Having a way to *rank* the actions turns out to be of tremendeous importance for the physicians.

Very High Dimension Case!

Issues

- DQN is unstable with TD in our setting (too slow to backpropagate the rewards?)
- Much better results using MC!
- Still hard to optimize everything from the beginning!

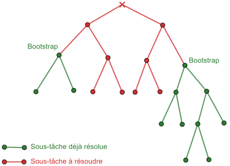

Dimension Reduction Trick

- State space partitioning to solve several smaller sub-problems.

State Partition and MC

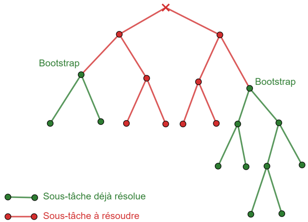

State Partitions

- Partition obtained by solving the problem starting from an abnomalies and falling back to previously computed strategy as soon as one reach a common state.
- Similar to a *n*-step bootstrapping!
- Works well with MC as *n* is not too large.

- Task ordering has to be specified!

Subtask Dimension

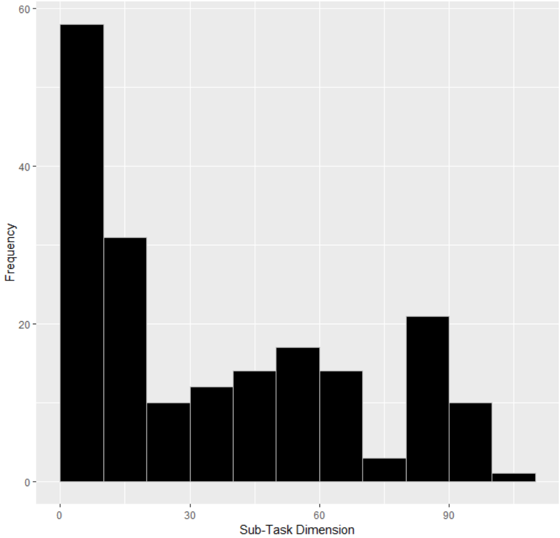

Optimal Decision Tree for a Small Subtask

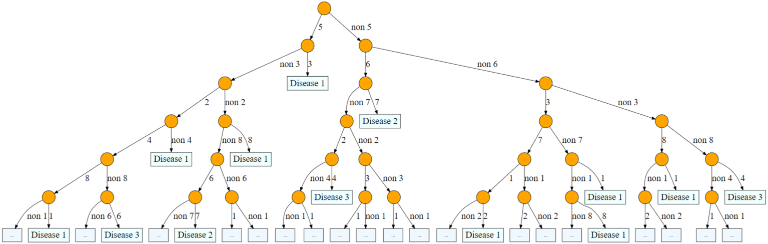

Optimal Policy for Small Subsets?

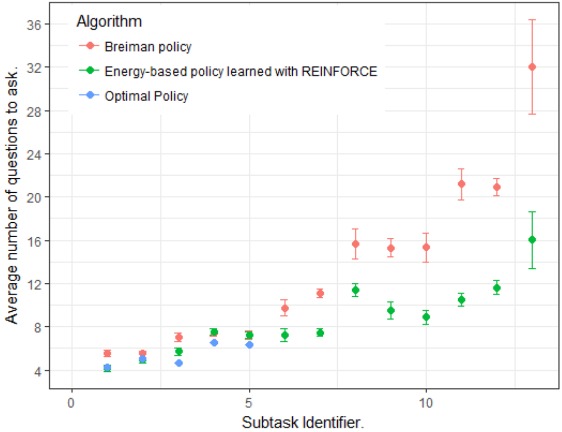

DQN vs REINFORCE vs Breiman

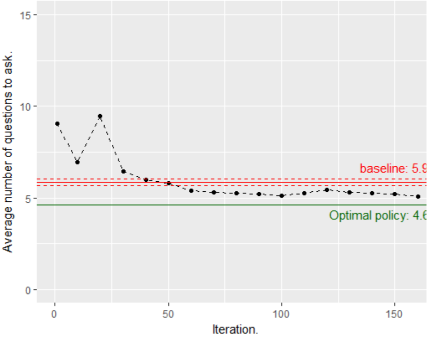

Task Dimension: 10

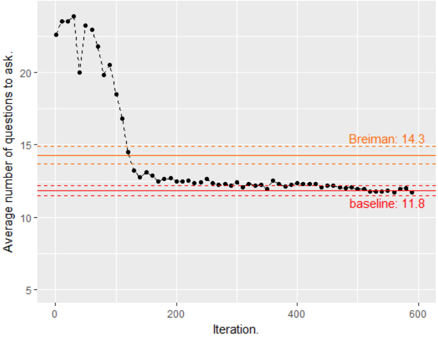

Task Dimension: 26

DQNs vs REINFORCE vs Breiman

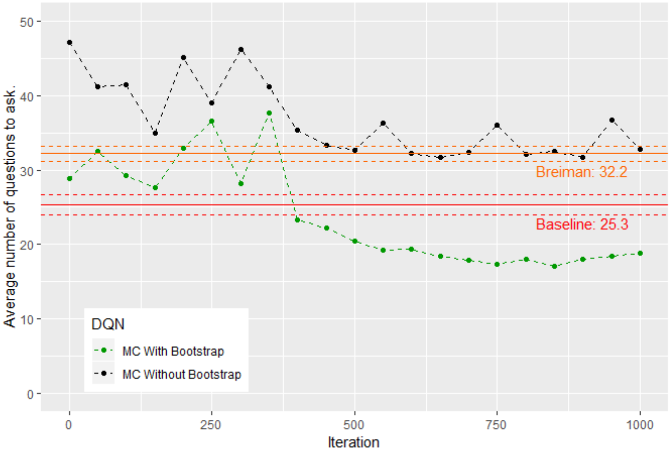

Task Dimension: 70

DQNs vs REINFORCE vs Breiman

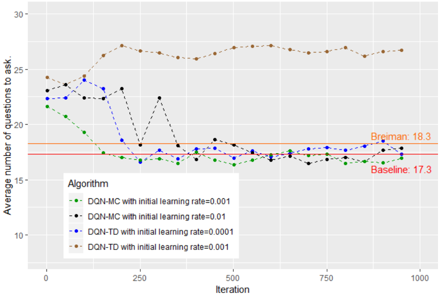

Task Dimension: 29

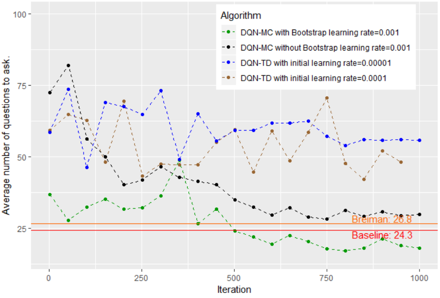

Task Dimension: 104

Take Away Message

Medical Goals

- Help obstetricians by improving/systematizing ultrasonic diagnostic **(MDP modeling)**
- Guide a (non rare disease expert) sonographer to assess as fast as possible potential diseases **(first prototype at Necker)**

Technical Goals

- Build an optimized decision tree:
  - Need to learn the environment **(MaxEnt and data assim.)**
  - Reinforcement learning **(Param. policy and MC vs Deep Q)**
- **Not yet** (theoretical) guarantees.

Take Away Message

- Reinforcement learning (or MDP) is an interesting tool.
- Formalization requires a true dialog between the mathematicans and the practicians.
- First prototype already tested by Necker.

The Team

Cliniciens

Statisticiens

Expert base
de données

Outline

References

R. Sutton and A. Barto.
Reinforcement Learning, an Introduction (2nd ed.)
MIT Press, 2018

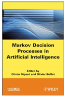

O. Sigaud and O. Buffet.
Markov Decision Processes in Artifical Intelligence.
Wiley, 2010

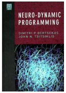

D. Bertsekas and J. Tsitsiklis.
Neuro-Dynamic Programming.
Athena Scientific, 1996

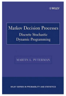

M. Puterman.
Markov Decision Processes. Discrete Stochastic Dynamic Programming.
Wiley, 2005

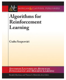

Cs. Szepesvári.
Algorithms for Reinforcement Learning.
Morgan & Claypool, 2010

- R. Besson, E. Le Pennec, S. Allassonnière, J. Stirnemann, E. Spaggiari, A. Neuraz, *A Model-Based Reinforcement Learning Approach for a Rare Disease Diagnostic Task*, arXiv preprint, 2018
- R. Besson, E. Le Pennec, S. Allassonnière, *Learning from both experts and data* arXiv preprint, 2019