

Microsoft



Microsoft® Research

Faculty Summit 2012

Riviera Maya, Mexico | May 23-25 | In partnership with CONACYT



Probabilistic Graphical Models: Applications in Biomedicine






L. Enrique Sucar, INAOE
Puebla, México

May 2012



What do you see?



What we see depends on our previous knowledge (model) of the world and the information (data) from the images → Bayesian framework



Outline

- Probabilistic Graphical Models
- Bayesian Networks
 - Endoscopy assistant
- Temporal Bayesian Networks
 - Predicting HIV mutations
- Markov Decision Processes
 - User adaptation for rehabilitation
- Conclusions




Bayesian Models

- In the Bayesian approach we combine our previous knowledge (*priors*) with the evidence (*likelihood*) to arrive to conclusions (*posterior*):

$$P (H | E) \propto P (H) P (E | H)$$

- However, if we apply it in naive way its complexity grows exponentially on the size (number of variables) of the model
- Probabilistic graphical models take advantage of the independence relations among the variables in a domain to develop more efficient representations as well as inference and learning techniques



Probabilistic Graphical Models


- Given a set of (discrete) random variables,

$$X = X_1, X_2, \dots, X_N$$

- The joint probability distribution,

$$P(X_1, X_2, \dots, X_N)$$

- specifies the probability for each combination of values (the joint space). From it, we can obtain the probability of a variable(s) (marginal), and of a variable(s) given the other variables (conditional)



Probabilistic Graphical Models

- A Probabilistic Graphical Model is a compact representation of a joint probability distribution, from which we can obtain marginal and conditional probabilities
- It has several advantages over a “flat” representation:
 - It is generally much more compact (space)
 - It is generally much more efficient (time)
 - It is easier to understand and communicate
 - It is easier to build (from experts) or learn (from data)



Probabilistic Graphical Models


- A graphical model is specified by two aspects:
 - A Graph, $G(V,E)$, that defines the structure of the model
 - A set of local functions, $f(\mathbf{Y}_i)$, that defines the parameters (probabilities), where \mathbf{Y}_i is a subset of \mathbf{X}
- The joint probability is defined by the product of the local functions:

$$P(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N) = \prod_{i=1}^n f(\mathbf{Y}_i)$$



Probabilistic Graphical Models

- This representation in terms of a graph and a set of local functions (called potentials) is the basis for *inference* and *learning* in PGMs
 - **Inference:** obtain the marginal or conditional probabilities of any subset of variables Z given any other subset Y
 - **Learning:** given a set of data values for X (that can be incomplete) estimate the structure (graph) and parameters (local function) of the model




Probabilistic Graphical Models


- We can classify graphical models according to 3 dimensions:
 - **Directed vs. Undirected**
 - **Static vs. Dynamic**
 - **Probability vs. Decision**

Probabilistic Graphical Models

- Directed




- Undirected




Probabilistic Graphical Models

- Static




- Dynamic




Probabilistic Graphical Models

- Only random variables



- Considers decisions and utilities





Types of PGMs

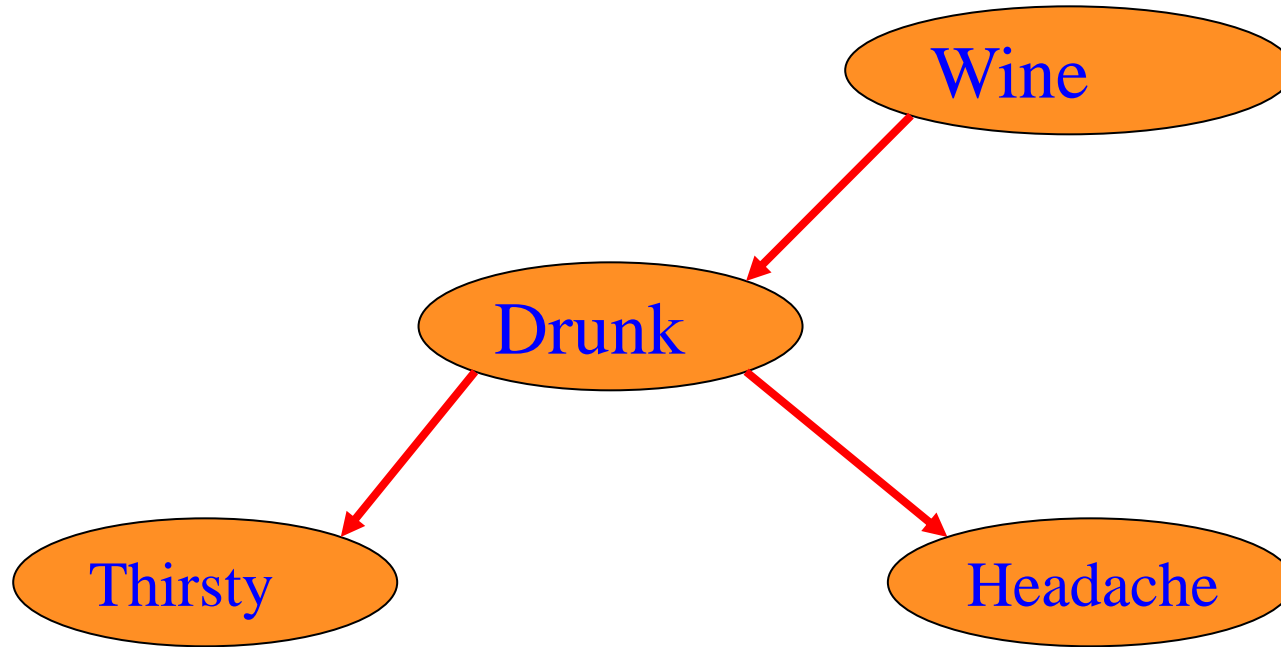
- There are different classes of PGMs:
 - Bayesian classifiers
 - Bayesian networks
 - Hidden Markov models
 - Dynamic Bayesian networks
 - Temporal Bayesian networks
 - Markov Random Fields
 - Influence diagrams
 - Markov decision processes



Bayesian Networks

- Bayesian networks (BN) are a graphical representation of dependencies between a set of random variables. A Bayesian net is a Directed Acyclic Graph (DAG) in which:
 - **Node: Propositional variable.**
 - **Arcs: Probabilistic dependencies.**
- An arc between two variables represents a direct dependency, usually interpreted as a *causal* relation.

An example of a Bayesian Network



Represents (in a compact way) the joint probability distribution:

$$P(W,D,T,H) = P(W) P(D|W) P(T|D) P(H|D)$$



Structure

- The topology of the network represents the dependencies (and independencies) between the variables
- Conditional independence relations between variables or sets of variables are obtained by a criteria called *D-separation*

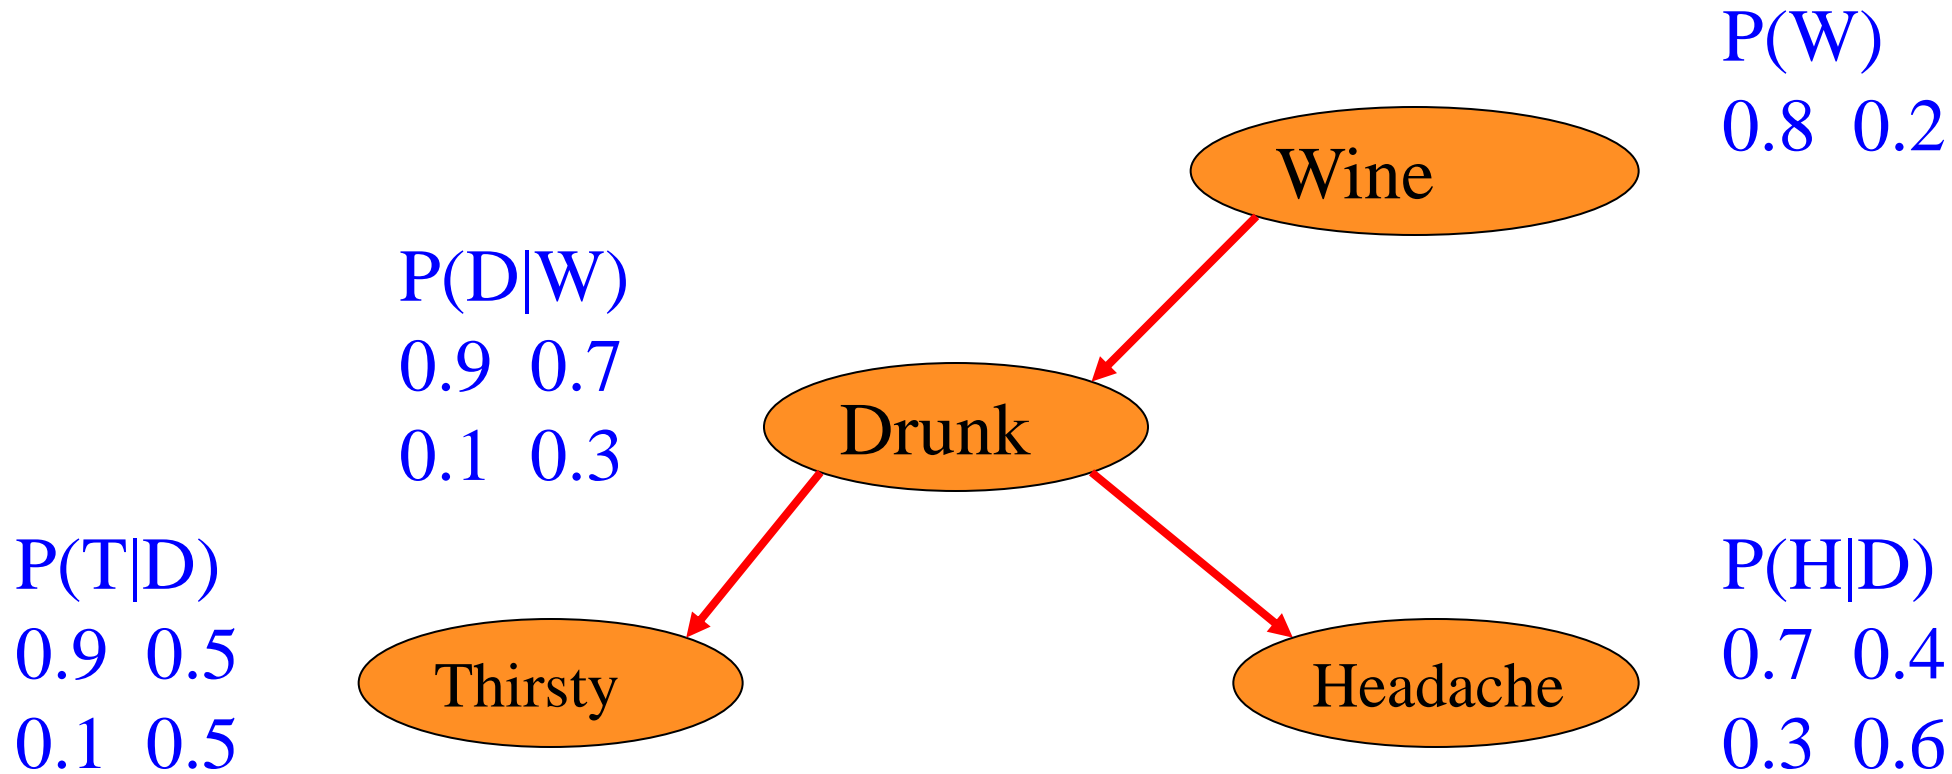


Parameters


Conditional probabilities of each node given its parents.

- **Root nodes:** vector of prior probabilities
- **Other nodes:** matrix of conditional probabilities

Parameters for the example



Inference



Given certain evidence, E , estimate the posterior probability of the other variables, H, C



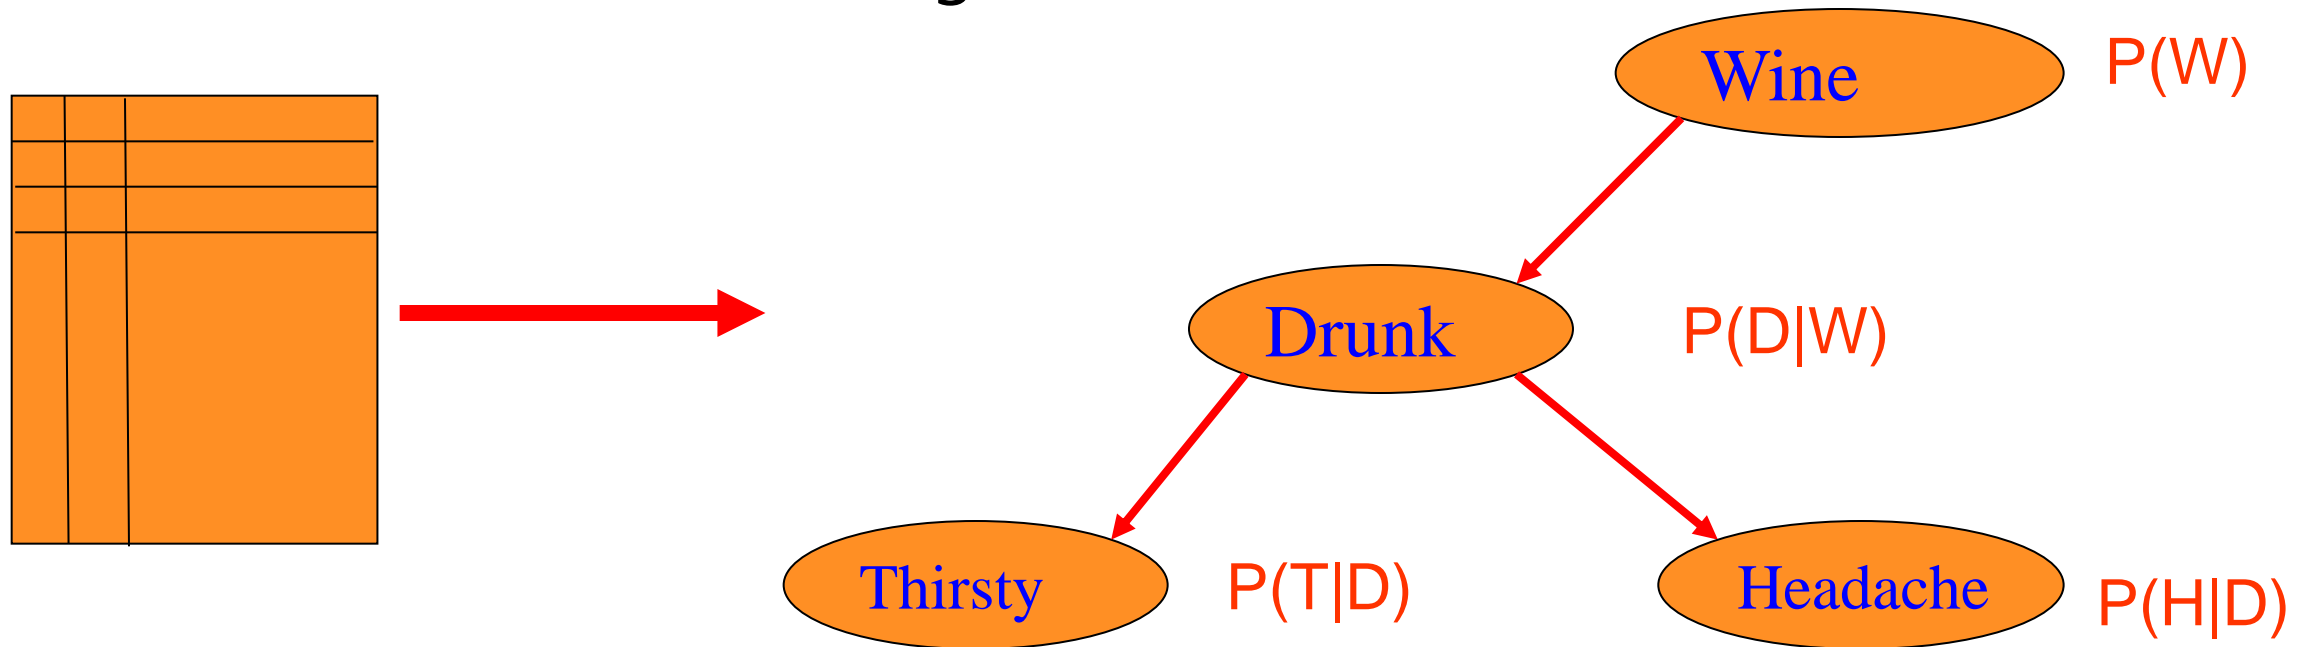
Inference

There are several inference algorithms:

- Variable elimination
- Message passing (Pearl's algorithm)
- Junction Tree
- Stochastic simulation
- ...
- In the worst case it an NP-Hard problem, however given a sparse graph the state of the art algorithms are very efficient

Learning


- Learning in Bayesian networks can be divided into two aspects:
 - Structure Learning
 - Parameter Learning




Structure Learning

Two general schemes:

Search and score



Independence tests









Structural Improvement

- Learning techniques require a large amount of data to obtain *good* models; an alternative is to **combine expert knowledge and data**
- We propose a method that starts from a subjective structure (given by an expert) and then improves it with data
- Assuming a tree structure, the conditional independence of child nodes given its parent are verified; if they are not independent there are 3 alternatives:
 - **Node elimination**
 - **Node combination**
 - **Node insertion**

Structural improvement






Endoscopy


- Endoscopy is a tool for direct observation of the human digestive system
- Navigating an endoscope is difficult due to the variability and dynamics of the human colon
- Thus, it is desirable to build a **semi-automatic system that can assist an endoscopist**
- The main challenge is to **recognize the "objects" in endoscopy images** which can be confused, such as "*lumen*" & "*diverticula*"
- The low-level vision algorithms can fail so we propose a **Bayesian network that combines the information and arrives to final decisions**

Colon Image





Low level features – dark region






Low level features –
shape from shading
(pq histogram)



(a) Colon Image



(b) Depth map (needle diagram)

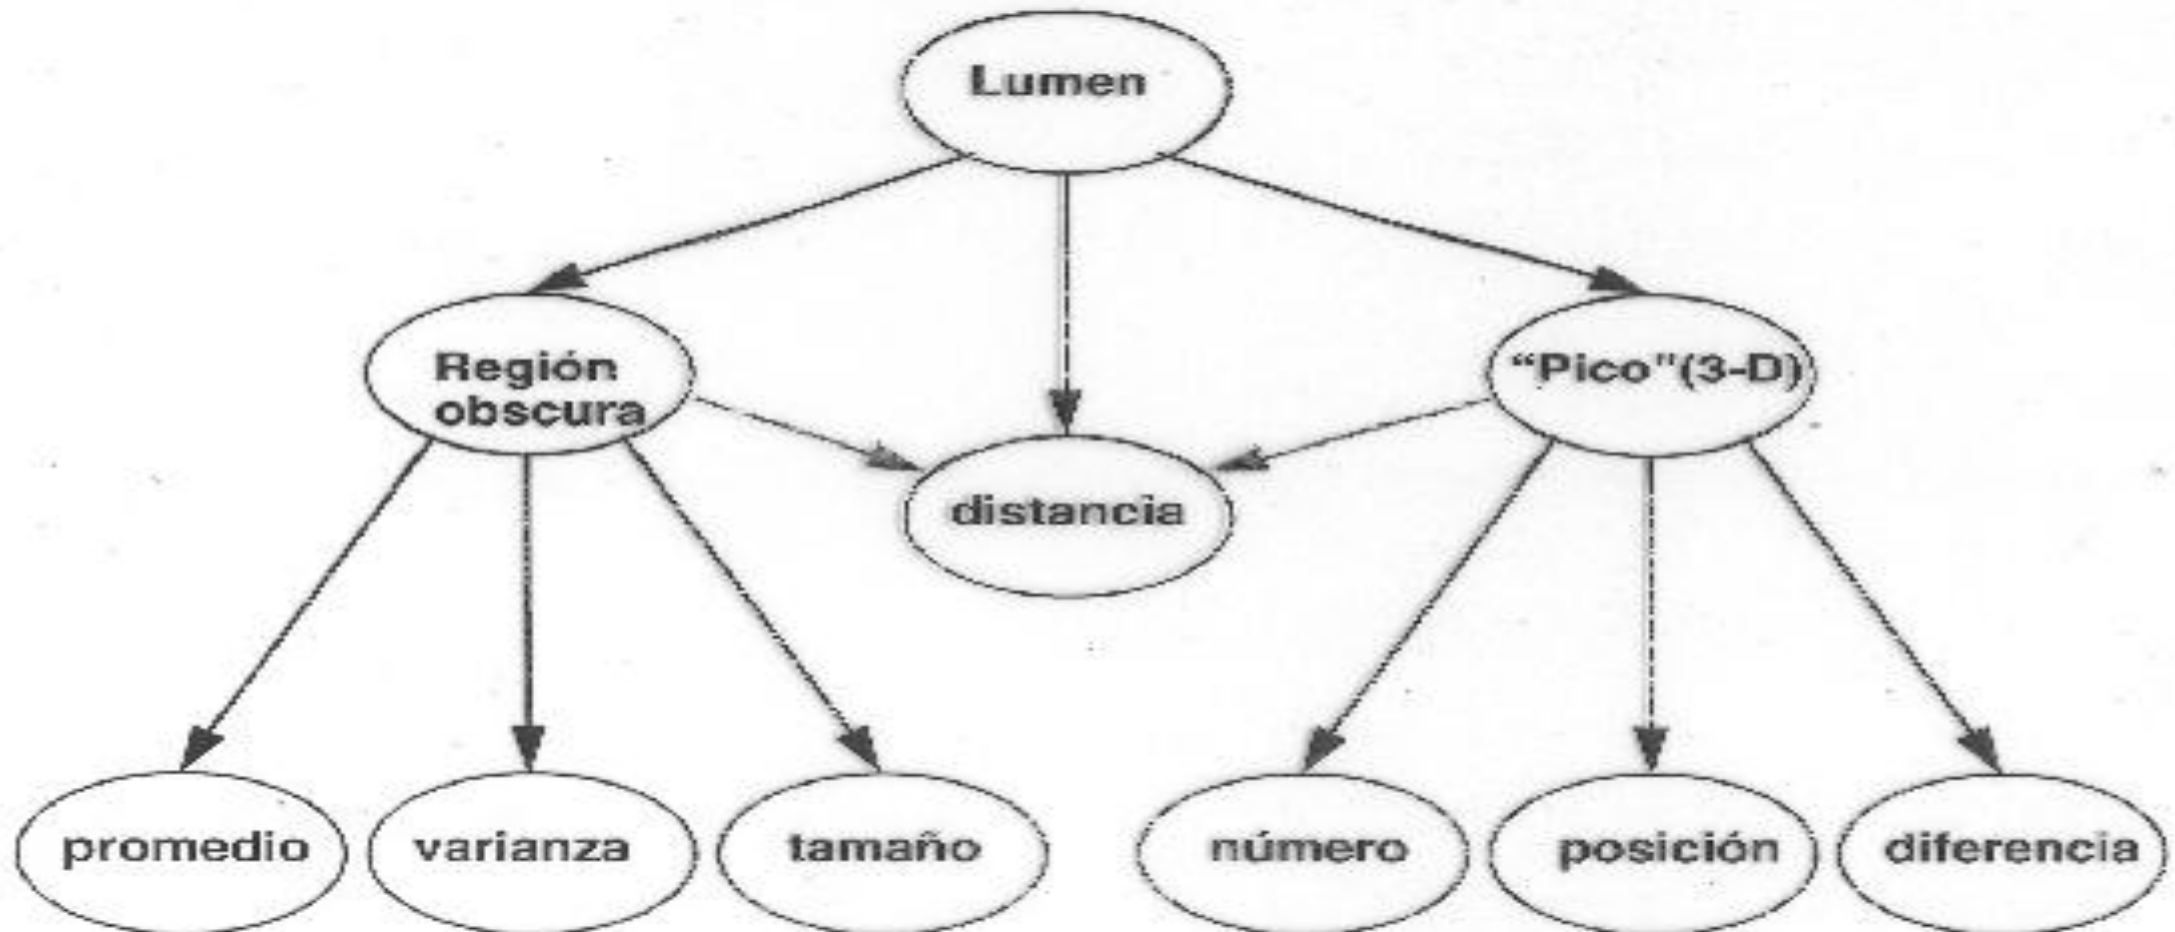




Model Construction

- The structure of the BN was built with the help of an expert endoscopist
- Later it was improved based on the structural improvement technique
- Parameters were learned from videos of *real* colonoscopy sessions

BN for endoscopy (partial)



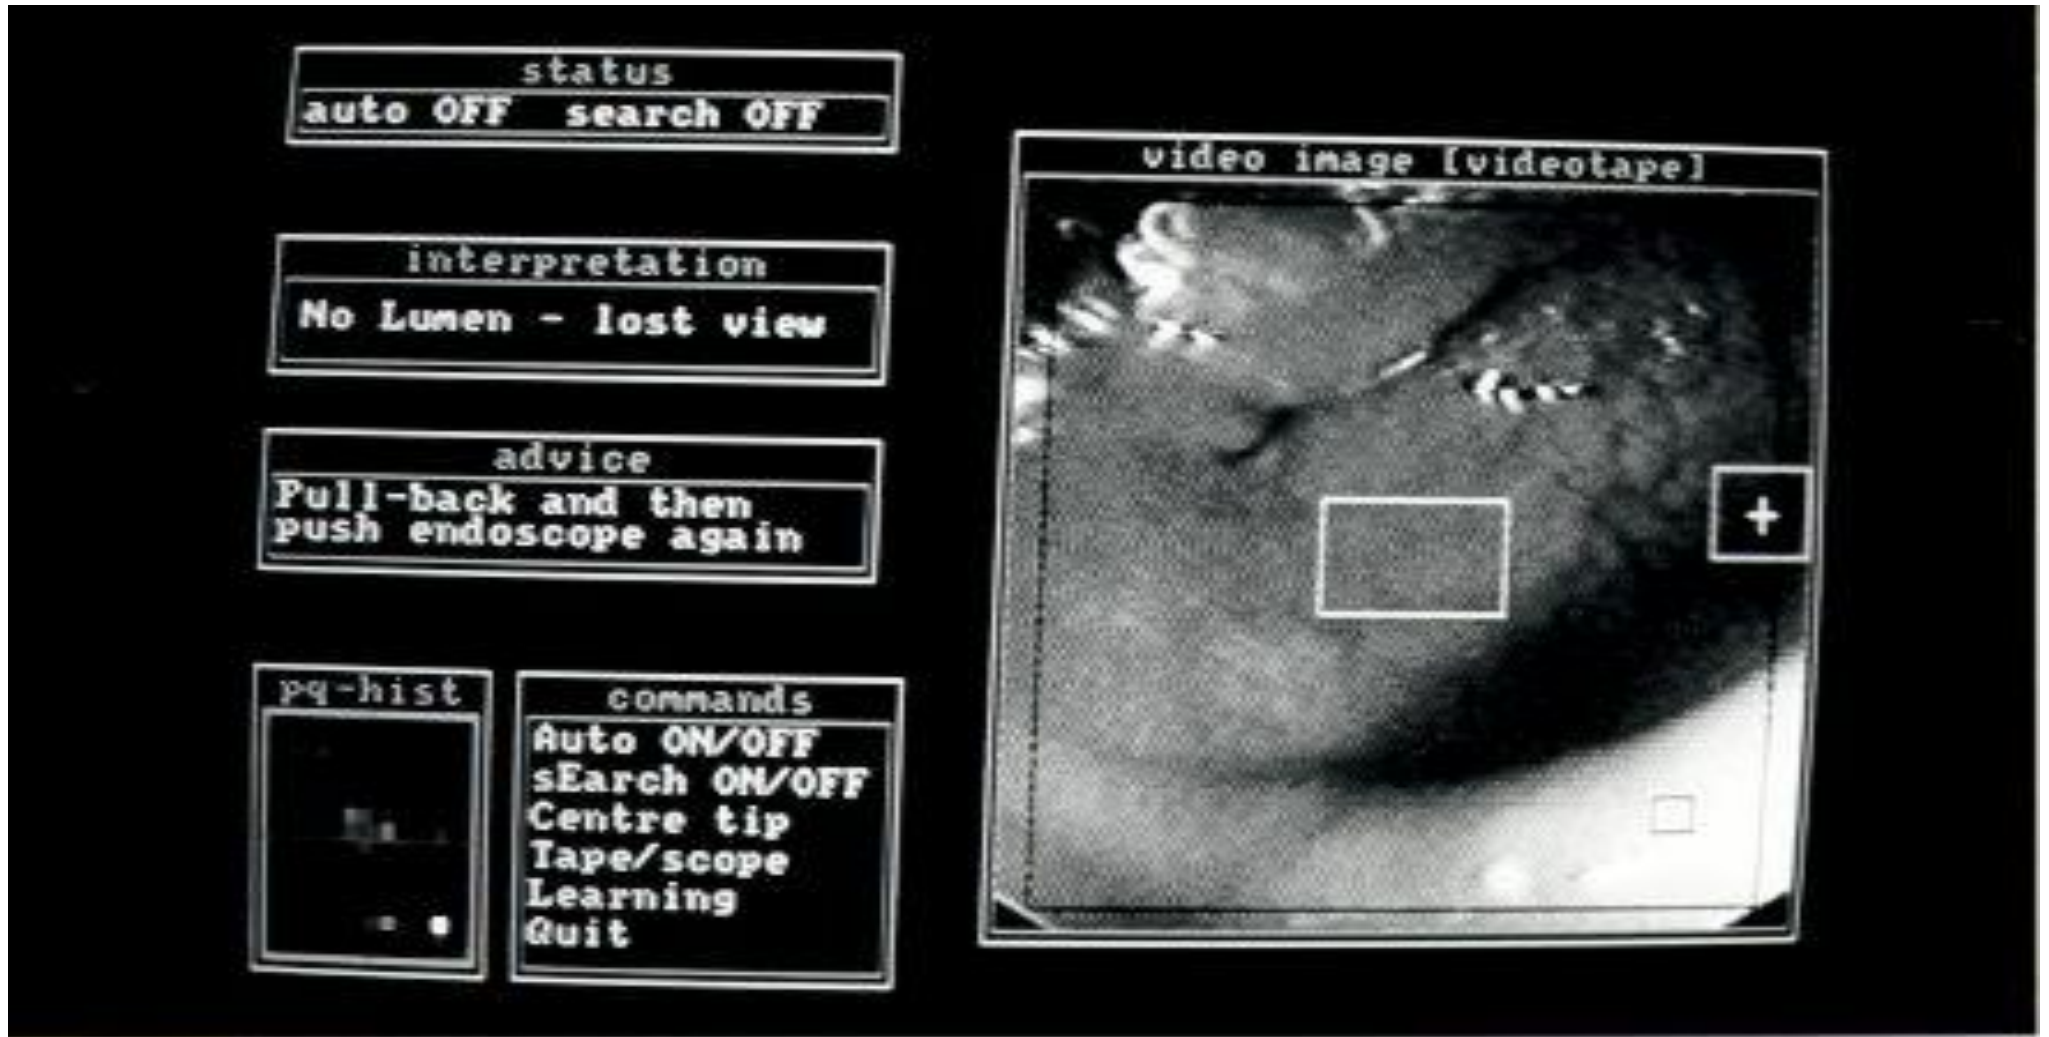
Semi-automatic Endoscope



Endoscope navigation system: example 1



Endoscope navigation system: example 2

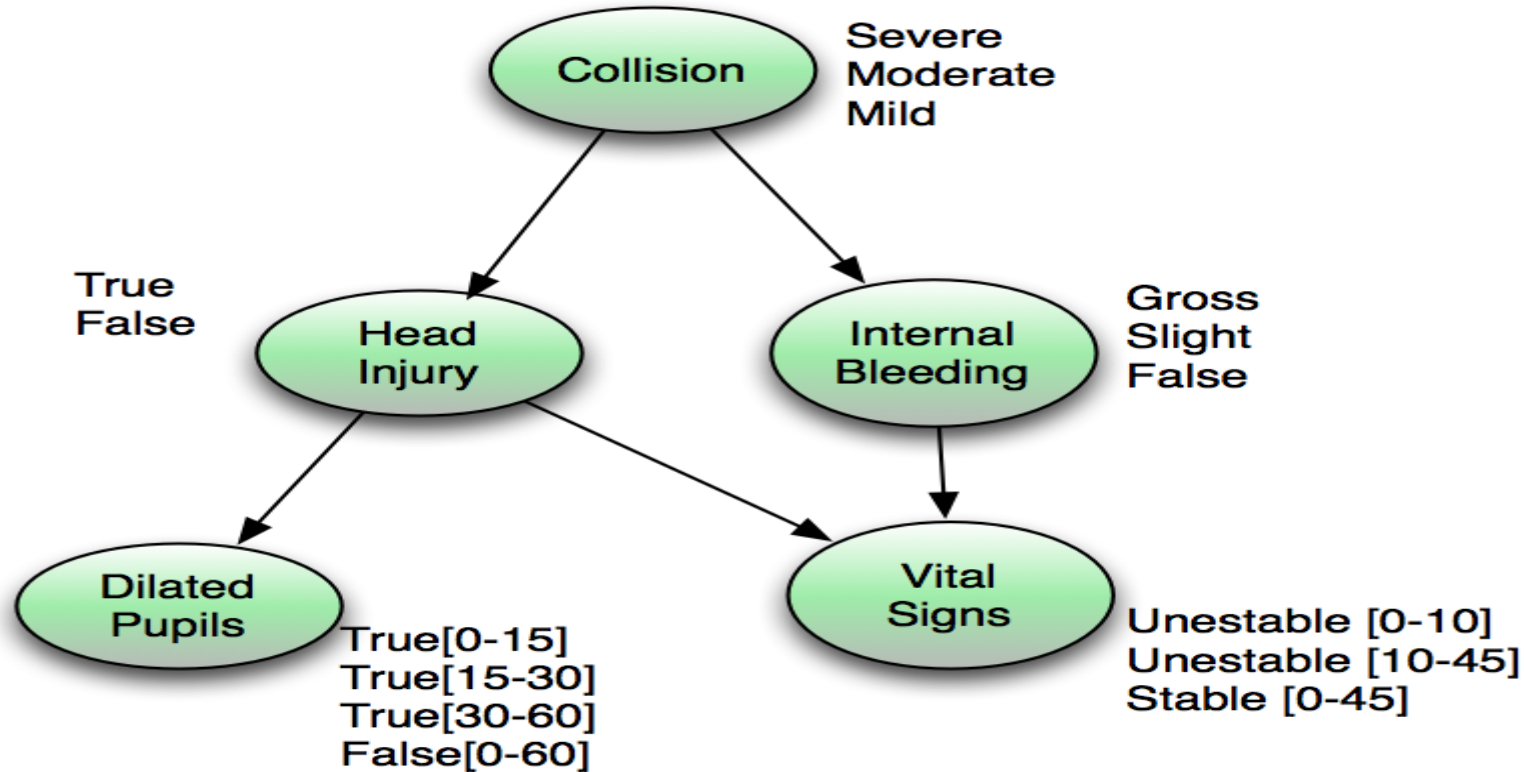




Temporal Nodes Bayesian Networks (TNBN)

- An alternative to Dynamic Bayesian Networks to **model dynamic processes with uncertainty**
- **Temporal information is within the nodes** in the model, which represent the time of occurrence of certain events
- The links represent temporal-causal relation
- Adequate for applications in which there are **few state changes in the temporal range**

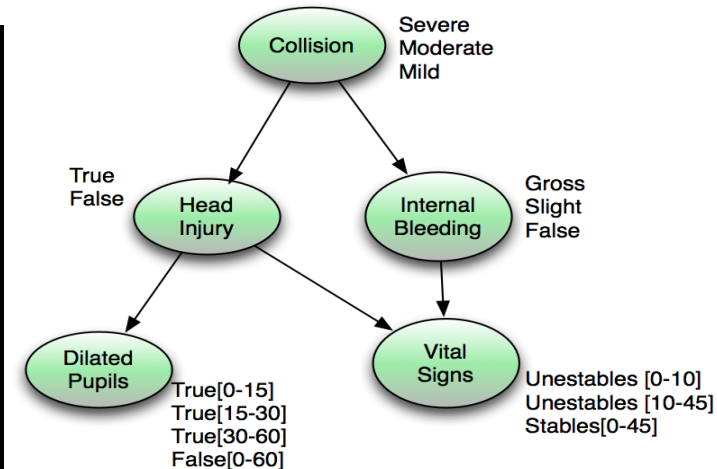
Example



Learning a TNBN

- Obtains the structure, the intervals and the parameters of the TNBN from data.

Collision	Head Inj.	Internal Ble.	Pupils Dil.	Vital Signs
severe	yes	gross	15	10
severe	yes	slight	25	20
mild	no	false	25	----
mild	no	false	21	----
moderate	yes	slight	----	20
moderate	yes	false	----	----
mild	no	false	----	----





Learning a TNBN

1. **Initial discretization of the temporal variables**, using an Equal-Width discretization (EWD).
 - Obtains an initial approximation of the intervals for all the Temporal Nodes (TN).
2. **Standard BN structural learning**, using the K2 learning algorithm
 - Obtains an initial structure
3. **Refines the intervals for each TN** by means of clustering using a Gaussian mixture model (GMM).



HIV


- HIV among **fastest evolving organisms**
- The HIV evolves (among other pressures) in response to antiretroviral therapy
- Although mutations conferring drug resistance are mostly known, the **dynamics of the appearance chain of mutations** remains poorly understood
- We use TNBN for modeling the relationships between antiretroviral drugs and HIV mutations, in order to analyze **temporal occurrence of specific mutations in HIV that may lead to drug resistance.**



Mutational Networks

- Mutational networks are “drug-associated mutational pathways in the protease gene, revealing the co-occurrence of mutations and its temporal relationships”
- If we could predict the most likely evolution of the virus in any host, then it would be plausible to select an appropriate antiretroviral regimen that prevents the appearance of mutations, effectively increasing HIV control.

Antiretrovirals



Antiretroviral therapy (ART) generally consists of well-defined combinations of three or four ARV drugs in order to reduce the possibility of development of drug resistance mutations.

<http://us.viramune.com/consumer/hiv-treatment>

Experiments

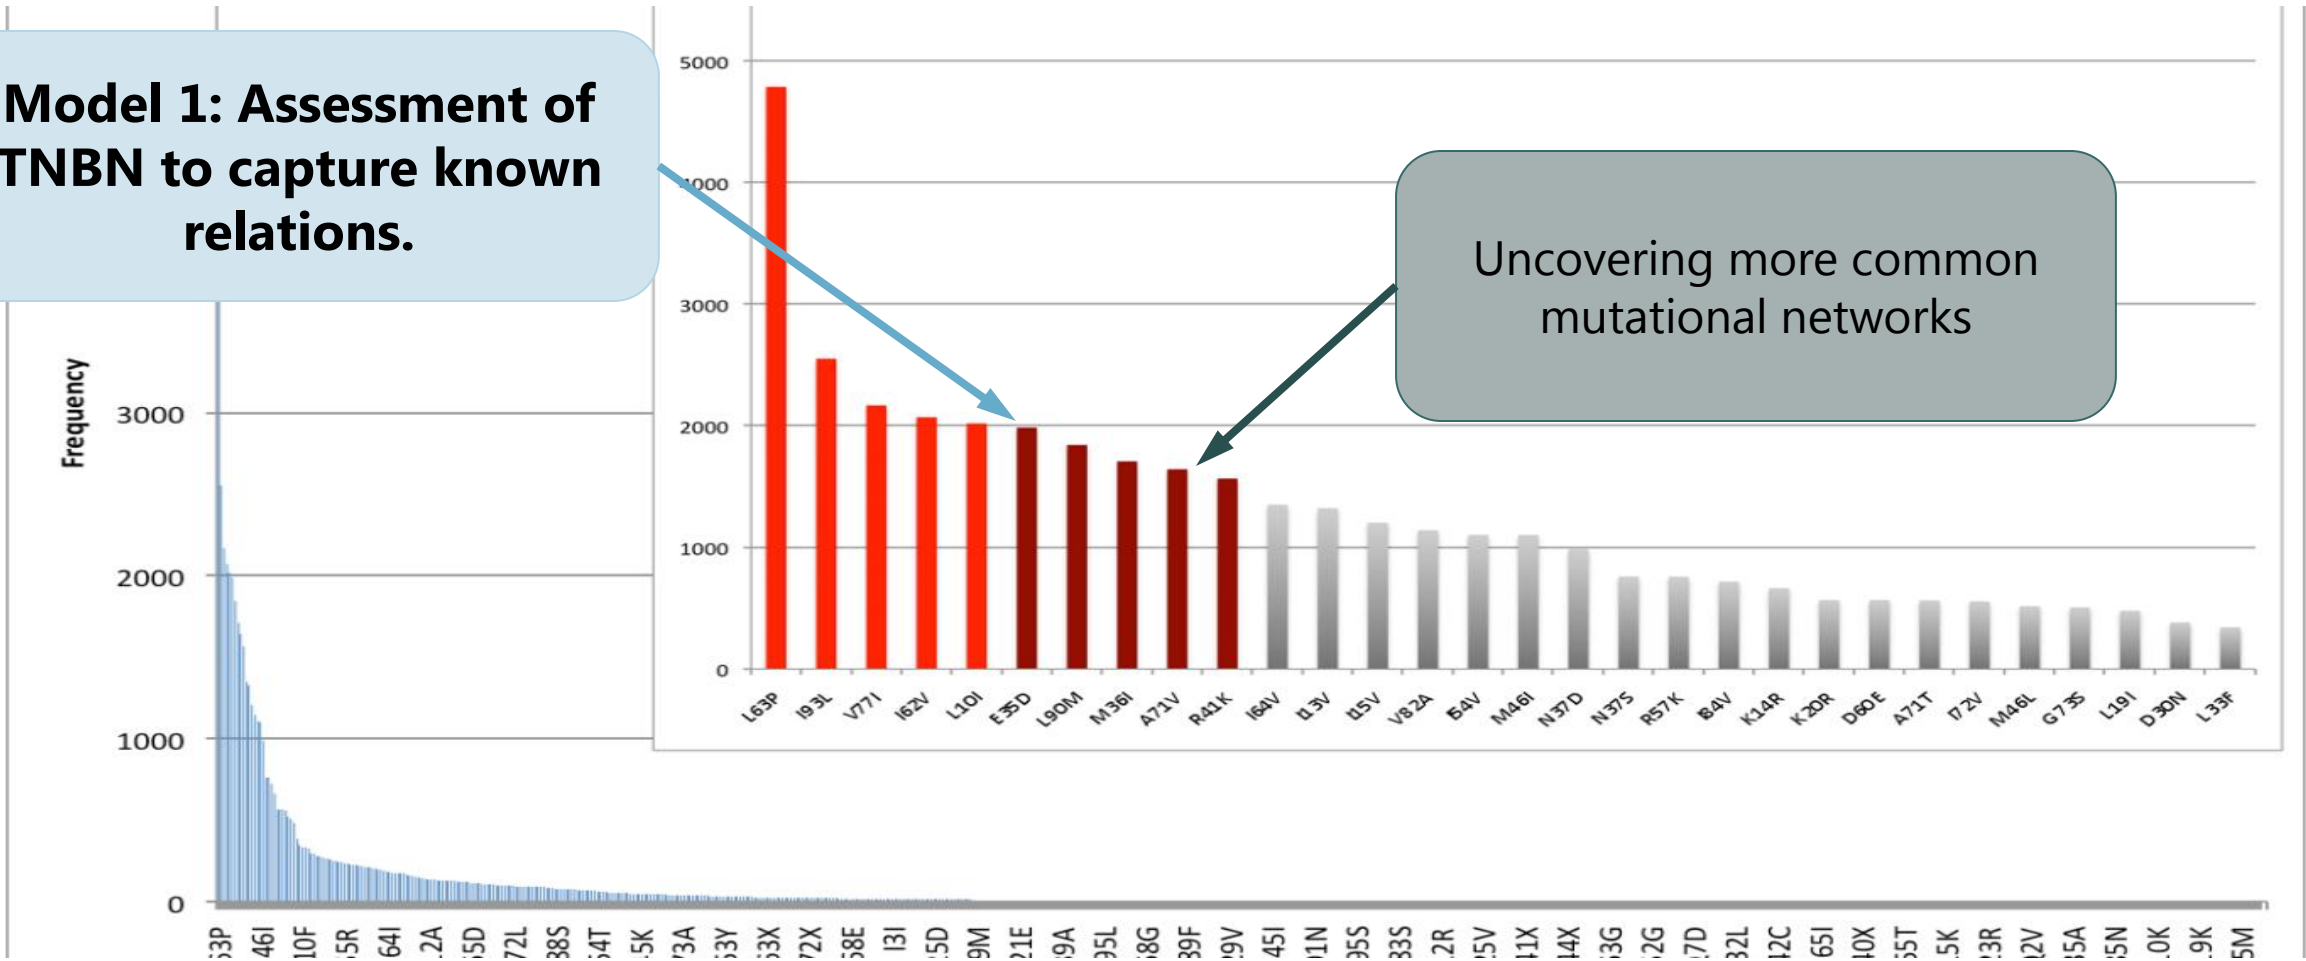
- Data and processing
 - HIV Stanford database (HIVDB) - HIV Drug Resistance Database
 - 2373 patients with subtype B was retrieved
 - Data retrieved contains a history consisting of a variable number of studies.

Patient	Initial Treatment	List of Mutations	Weeks
P1	LPV, FPV, RTV	L63P, L10I, V77I, I62V	15 30 10
P2	NFV, RTV, SQV	L10I V77I	25 45

Defining target mutations

Model 1: Assessment of TNBN to capture known relations.

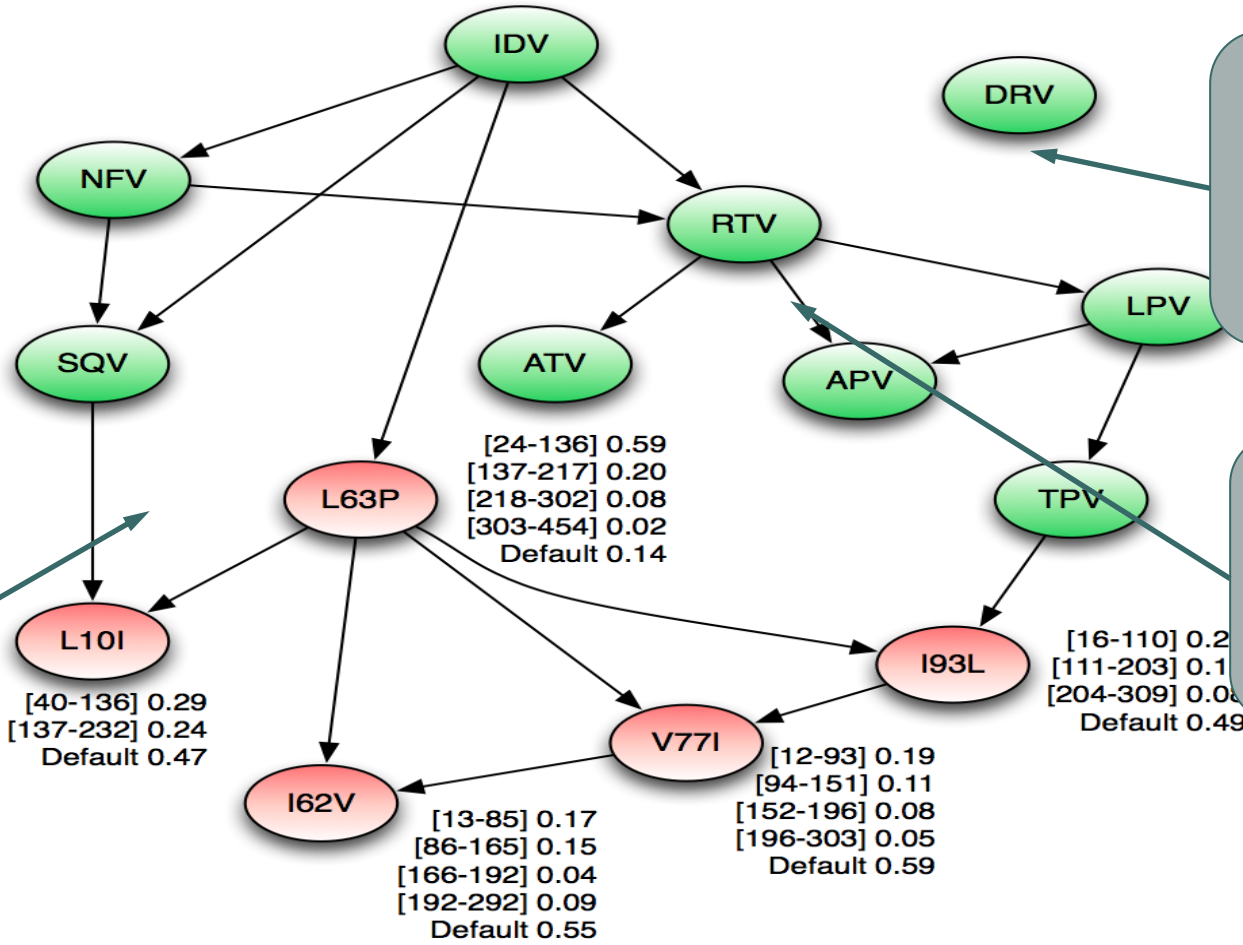
Uncovering more common mutational networks



Results: Model 1

Known relationships are appropriately captured

Link between SQV and L10I



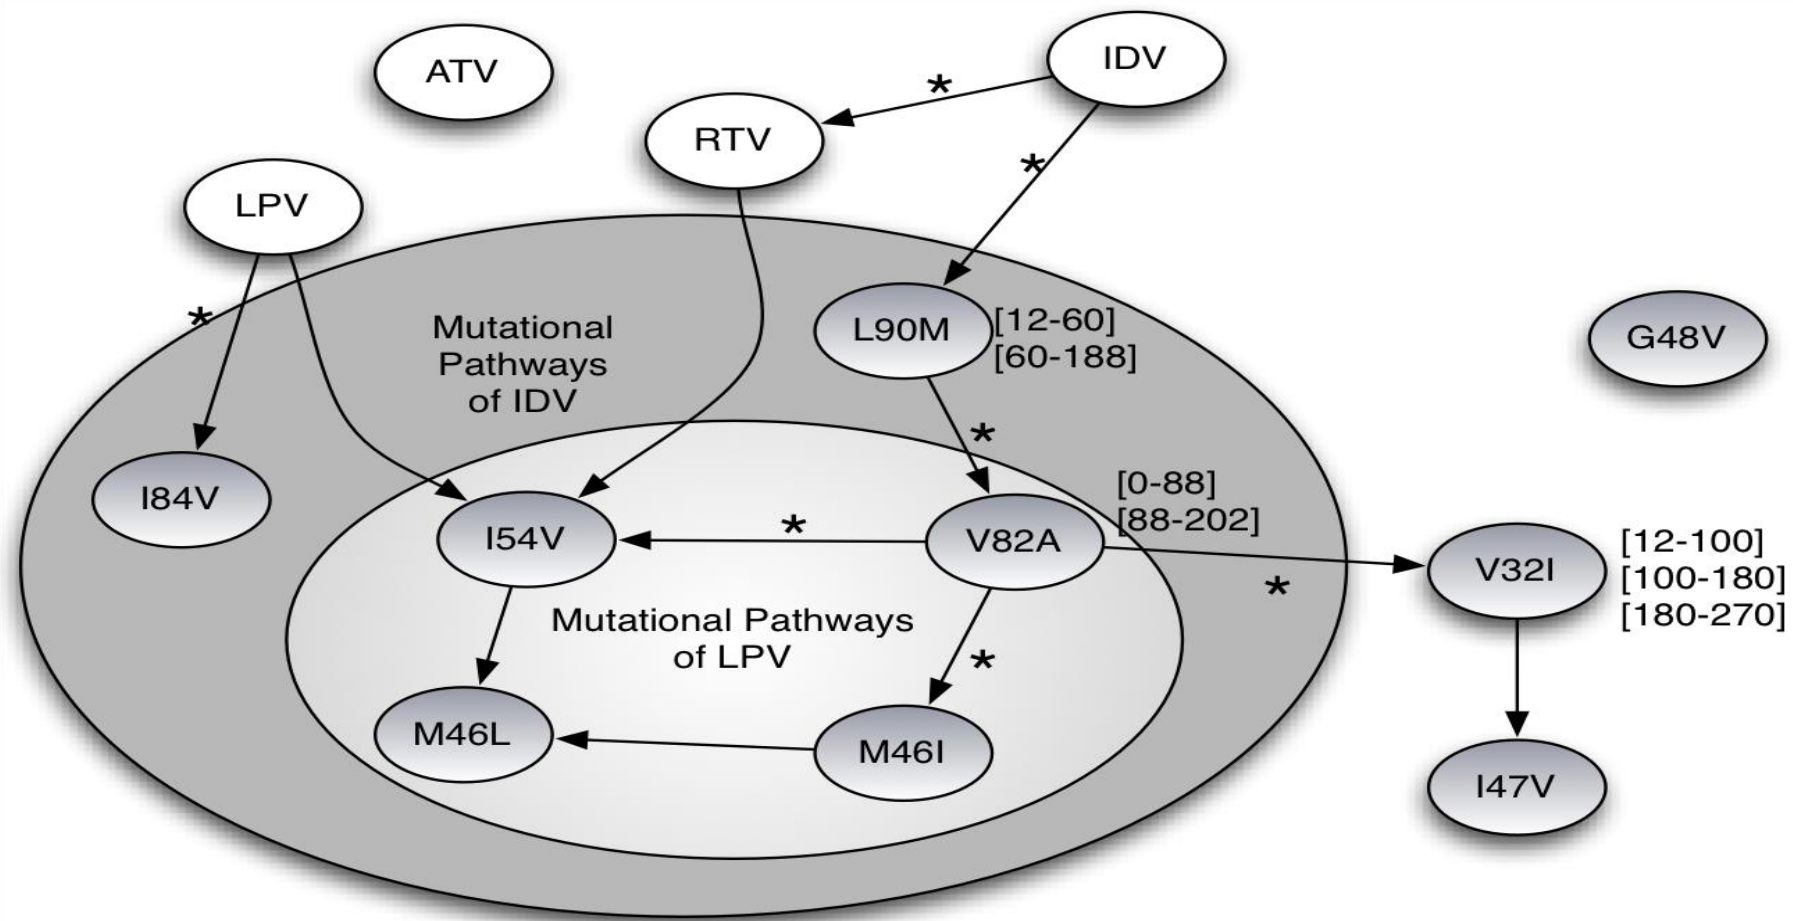
DRV isolation: New drug is hardly ever given as a first treatment

Role of RTV is often used in conjunction with other drug (boosting)



Model 2

- Use expert's information to select a subset of mutations and drugs of special importance.
- We used the Major HIV Drug Resistance Mutations and four drugs highly used in the past and nowadays.



The model was able to capture some mutational pathways already known (obtained by clinical experimentation).

LPV : M46I/L, I54V/T/A/S and V82T/F/S (Kempf et al., 2001) ,
IDV: V82A/T/F/S/M, M46I/L, I54V/T/A, I84V and L90M (Bélec et al., 2000; Descamps et al., 2005)



Markov decision processes (MDPs)

- Ideal framework for [planning under uncertainty](#).
- Main features:
 - Considers the uncertainty in the actions
 - Considers the utility of the plan
 - It allows to obtain optimal solutions
 - Considers uncertainty in the observations (POMDP)



MDP

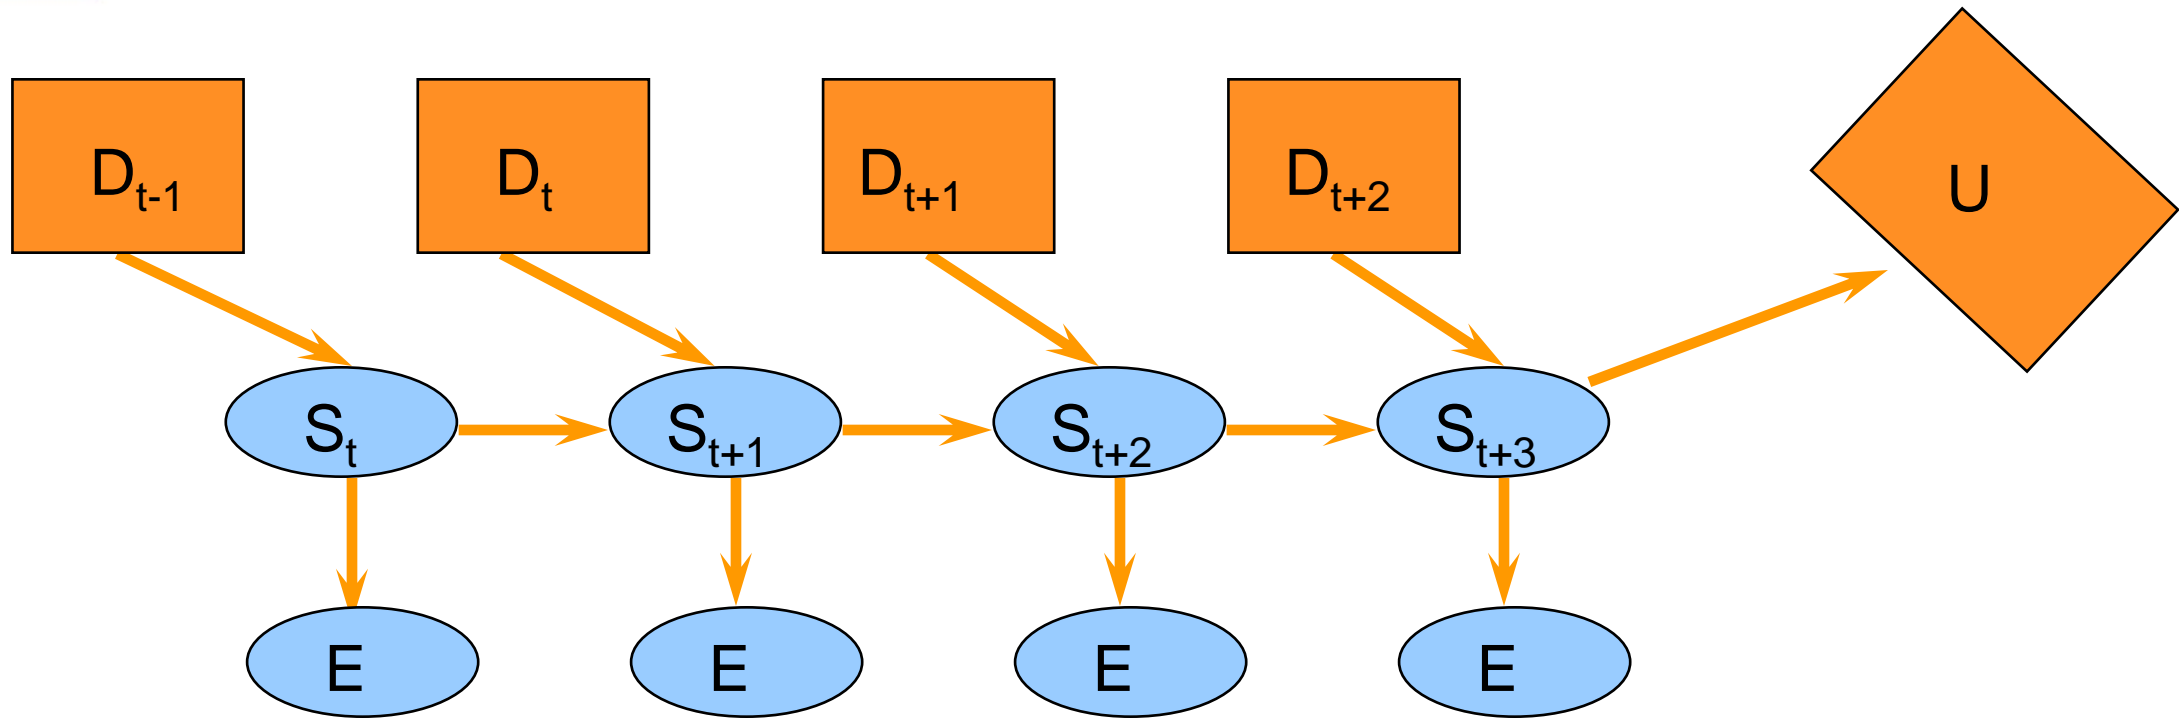
- Formally, a discrete MDP is defined by:
 - A finite set of states, S
 - A finite set of actions, A
 - A transition model, $P(s' | s, a)$
 - A reward function for each state-action, $r(s, a)$



POMDP

- Besides the MDP model, a POMDP has:
 - An observation probability distribution, $P(O|S)$
 - An initial probability distribution, $P(S)$

A POMDP as a Dynamic Decision Network





Basic solution techniques


- There are two main classes of algorithms:
 - **Dynamic programming techniques**: consider a known model (transition and reward functions) which is solved to obtain the optimal policy
 - Montecarlo and **reinforcement learning**: the model is not known, so the optimal policy is obtained by exploring the environment

Factored MDPs

The state is decomposed in a set of factors or state variable:

$$X = \{x_1, x_2, x_3, x_4, x_5\}$$

So the transition function is represented as a two-stage DBN per action



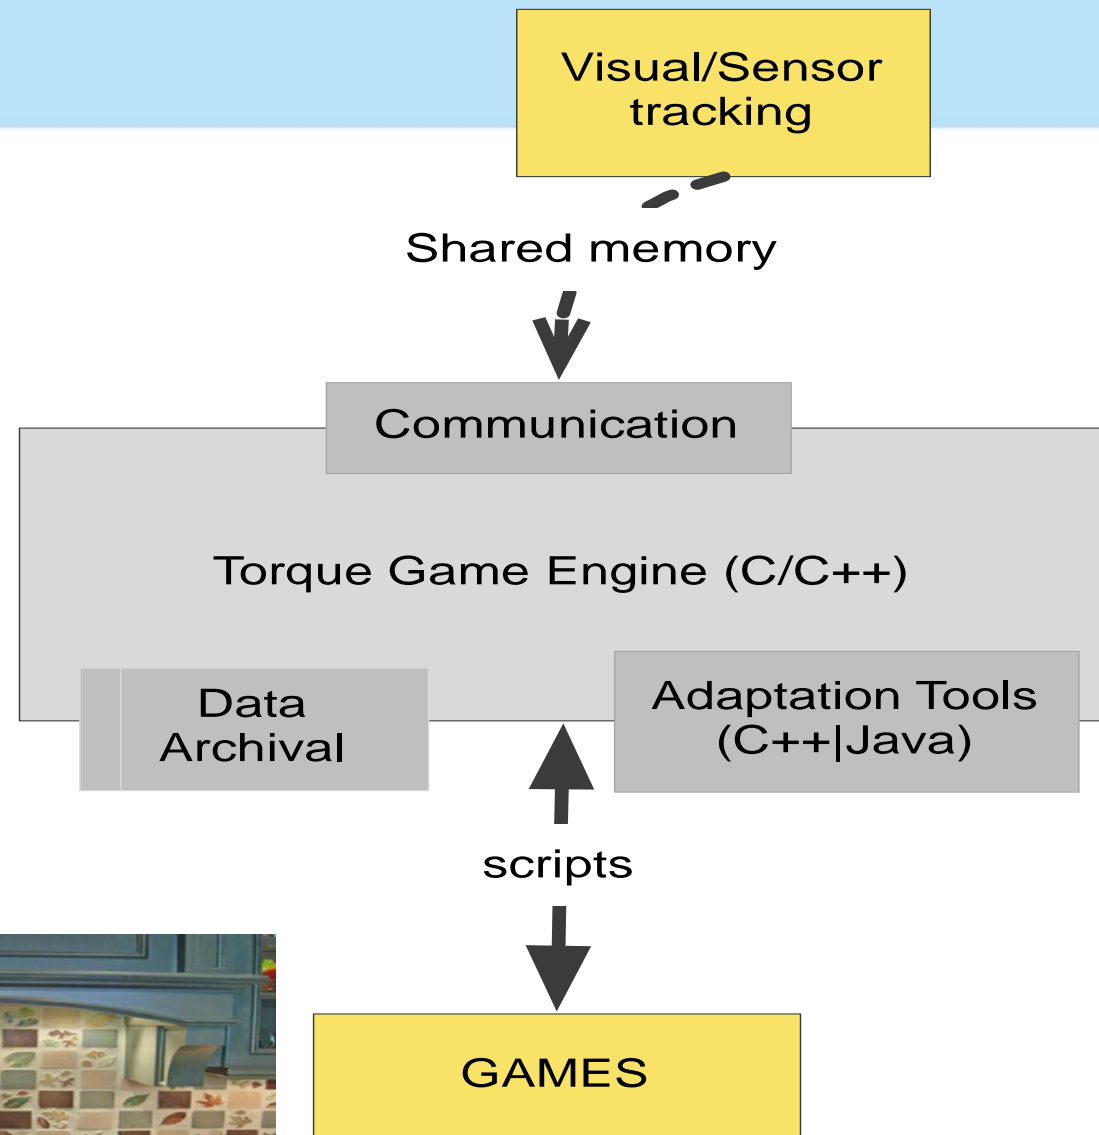



Gesture Therapy

- Many people suffer strokes (15 million worldwide per year)
- 80% lose arm and hand movement skills
- Physical and occupational therapy can help, but:
 - Expensive (requires a therapist)
 - Usually not enough
 - Patients loose motivation
- Robotic systems are too expensive for use at home or small clinics
- Develop low-cost technology that allows stroke patients to practice intensive movement training at home without the need of an always present therapist

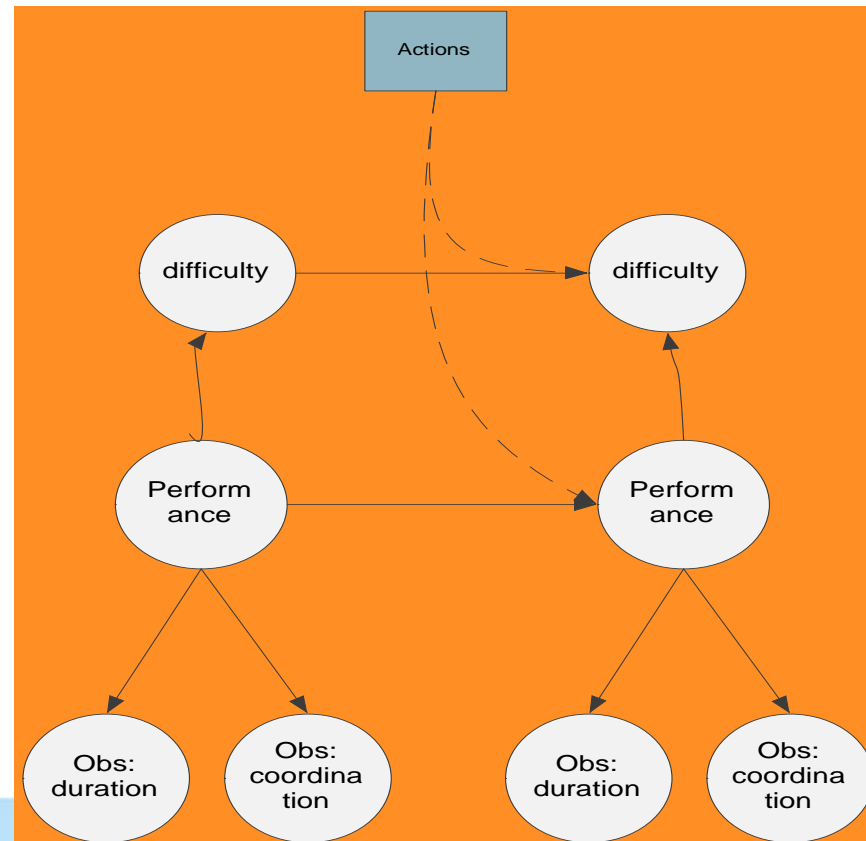
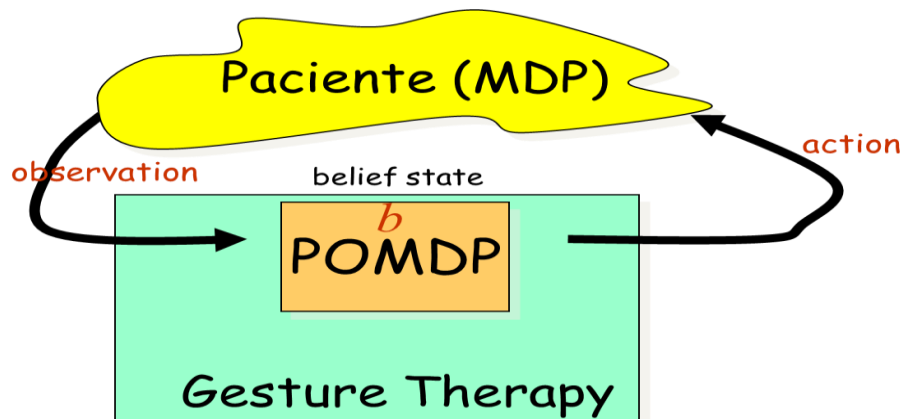
Gesture Therapy

- Simulated environment
- Monocular tracker
- Gripper
- Trunk compensation detection
- Adaptation to the patient




Adptation to the patient

The system estimates the patient “state” based on observing its performance in the game (speed, control) and decides the game difficulty accordingly according to the policy dictated by the POMDP



Evaluation

- Preliminary results with a “normal” person



Prototype of the system at the INNN rehabilitation unit



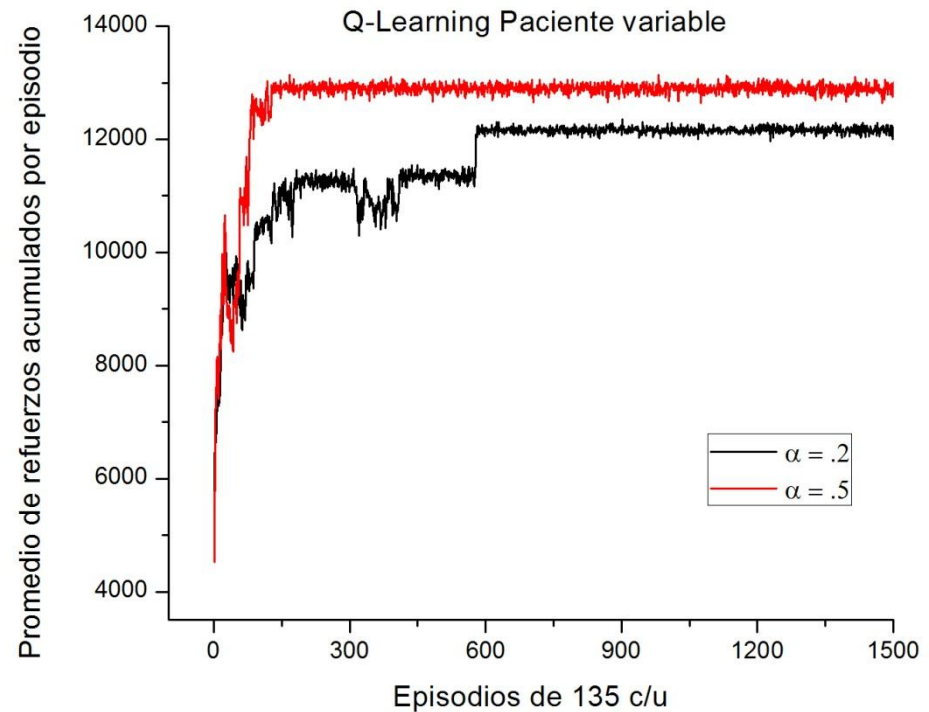


Policy adaptation

- The POMDP *model could be wrong* so the policy is not necessarily “optimal”
- Also, the *best policy could depend on the patient*
- We developed a policy adaptation algorithm based on RL+ reward shaping which *improves an initial policy based on the therapist feedback*

Initial results

- *Simulated therapist* – feedback based on the optimal policy





Conclusions

- The **Bayesian approach** combines prior knowledge (a priori probability) with evidence (likelihood) based on Bayes theorem
- **Graphical models** allow for an efficient and clear representation of probability distributions based on dependency & independency relations



Conclusions

- PGMs provide a set of techniques which can be applied to solve complex problems in biomedicine which require to model uncertainty, time and cost/utilities
- Biomedical applications pose interesting challenges that require novel developments in representation, inference and learning of PGMs



References

- L. E. Sucar, D.F. Gillies, D. A. Gillies, “Objective probabilities in expert systems”, *Artificial Intelligence*, 1993.
- L. E. Sucar, D. F. Gillies, “Probabilistic reasoning in high-level vision”, *Image and Vision Computing*, 1994.
- L. E. Sucar, R. Luis, R. Leder, J. Hernández, I. Sánchez, “GESTURE THERAPY: A Vision-Based System for Upper Extremity Stroke Rehabilitation”, *IEEE EMBC*, 2010
- P. Hernández, F. Orihuela, A. Rios, J. González, L.E. Sucar, “Unveiling HIV mutational networks associated to pharmacological selective pressure: a temporal Bayesian approach”, *AIME*, 2011.


Acknowledgements

- Collaborators:

- Eduardo Morales, Jesús González, Felipe Orihuela, Alma Rios, Pablo Hernández, Shender Ávila (INAOE, Mexico)
- Duncan Gillies (Imperial College, UK)
- Hesse Hoey (U. of Waterloo, Canada)
- Ron Leder (UNAM)
- David Reinkensmeyer (UC Irving, USA)

- Sponsors:

- INAOE
- CONACYT-Salud (Mexico)
- FONCICYT (Mexico-EU)
- RIC / Department of Education (USA)



Thanks!

Questions?

esucar@inaoep.mx
ccc.inaoep.mx/esucar