

Human Allied Statistical Relational Artificial Intelligence

Sriraam Natarajan University of Texas-Dallas Research Faculty Dr. Gautam Kunapuli

Who we are!

Current Students (PhD)

Mayukh Das, Srijita Das, Devendra Dhami, Alex Hayes, Navdeep Kaur, Nandini Ramanan, Kaushik Roy, Harsha Kokel, Changbing Li, Yuqiao Chen...

Alumni (PhD) Phillip Odom , Shuo Yang, Tushar Khot

Key Collaborators

Kristian Kersting, Jude Shavlik, David Page, Dan Roth, Jana Doppa, Ron Parr, William Cohen, Kay Connelly, Clinical collaborators

Funding agencies

DARPA (CwC, DEFT & Machine Reading), NSF (SCH), AFRL, ARO (YIP, STIR), AFOSR (SBIR), NIH (R01), Indiana (Precision Medicine), XEROX PARC, Amazon, Intel and TURVO Inc.



Human Allied Al



Can we build systems that can seamlessly interact with, learn from and collaborate with human expert?



ARO YIP 13 DARPA DEFT AFOSR STTR 14 NIH R01 15 PARC Faculty Award ARO STIR 17

Human Allied Al



Kunapuli et al ICDM 13 Odom et al AAAI 15, AAMAS 16, ECML 16, ILP 16, AIME 15 Yang ICDM 14, ECML 13 Macleod CHASE 16

. . .

Apprenticeship Learning



DARPA CwC

Human Allied Al

Das SDM 16, HMCL WS 17, AAMAS 18 KBS (under review)



Our Prior Work towards HAAI

- 1. Tractable modeling of Multi-modal data
- 2. Efficient Reasoning in large domains
- 3. Faithful Temporal Modeling
- 4. Effective Sequential decision-making
- 5. Scalable learning with complex relational data
- 6. Explicit modeling of noise and uncertainty

Bottom line: Take your data spreadsheet ...

Features





Deep Learning

and many more ...

Heterogenous data networks abound

[Lu, Krishna, Bernstein, Fei-Fei "Visual Relationship Detection" CVPR 2016]



Most data in the world is stored in relational databases

Mining Electronic Health Records is an opportunity to save our lifes and a lot of money.

Unfortunately, they are noisy and interconnected



Actually, most data in the world stored in relational databases

De Raedt, Kersting, Natarajan, Poole, Statistical Relational Artificial Intelligence: Logic, Probability, and Computation. Morgan and Claypool Publishers, ISBN: 9781627058414, 2016.

Key Observation

Two trends that drive modern data science

- 1. Race to deeply understand data
- 2. Data in a large number of formats

We need a crossover of Machine Learning and Probabilistic/ Statistical Databases



[Getoor, Taskar MIT Press '07; De Raedt, Frasconi, Kersting, Muggleton, LNCS'08; Domingos, Lowd Morgan&Claypool '09; Natarajan, Kersting, Khot, Shavlik Springer Brief'15; Russell CACM 58(7): 88-97 '15]



Statistical Relational Artificial Intelligence Logic, Probability, and Computation

ristian Kersting iraam Nataraja

AI and ML: State-of-the-Art

Learning

Decision trees, Optimization, SVMs, ...

Logic

Resolution, WalkSat, Prolog, description logics, ...

Probability

Bayesian networks, Markov networks, Gaussian Processes...

Logic + Learning

Inductive Logic Programming (ILP)

Learning + Probability

EM, Dynamic Programming, Active Learning, ...



Logic + Probability

Nillson, Halpern, Bacchus, KBMC, ICL, ...



Example - Relational Probability Trees

To predict heartAttack(Person)



Outline

- State-of-the-art Structure Learning
- Effective Learning
- Advice Taking
- Actively Seeking Advice

Learning

- Parameter Learning Where do the numbers come from
- Structure Learning neither logic program nor models are fixed
- Evidence
 - Parital assignments of values to variables {burglary = false, earthquake = true, alarm = ?, johncalls = ?, marycalls = true}



Up next inside Learning

- Parameter learning graphical models
- Parameter learning StaR AI models
- Structure learning logical models & graphical models
- Structure learning StaRAI models

Maximum Likelihood Estimation

- MLE is the fraction of positive counts over total counts
- Example: Bayes net parameters estimate each CPD using MLE
- For each P(X=x|Pa(X) = y), count the number of examples of x and y

$$P(X = x | Pa(X) = y) = \frac{N(X = x, Pa(X) = y)}{N(Pa(X) = y)}$$

EM Idea

- If data is complete, ML parameter estimation is easy: simple counting (1 iteration)
- But what if there are missing values, i.e., we are facing incomplete data?
 - 1. Complete data (Imputation)
 - most probable?, average?, ... value
 - 2. Count
 - 3. Iterate

EM Idea: complete the data



EM Idea: complete the data





- Parameter learning graphical models
- Parameter learning StaR AI models
- Structure learning logical models & graphical models
- Structure learning StaRAI models

Relational Parameter Estimation

	<u>E</u>	Backg	ground	
Model(1)		m(ann.dorothy).		
pc(brian)=b,		f(brian dorothy)		
bt(ann)	bt(ann)=a.		ily fred)	
bt(Mo	del(2)		(fred)	
bt(bt(c	bt(cecily)=ab,		y,irea),	
bt(h	enry)=a		bob),	
bt(f	bt(fred)=?,		Model(3)	
bt(k	bt(kim)=a,		pc(rex)=b,	
bt(b	bt(bob)=b		bt(doro)=a,	
	_		bt(brian)=?	





Relational Parameter Estimation

		B	ack	ground	
Model(1)		m(ann.dorothy).			
pc(brian)=b,		f(brian dorothy)			
bt(ann)=a.		.(,	,	ilu frod)	
bt(Model(2)			iny,ired),	
ht(bt(cecily)=ab,		,	y,fred),	
ыų	bt(henry)=	/)=a,		bob),	
	bt(fred)=?	?, <u>Mo</u>		del(3)	
	bt(kim)=a.		pc(rex)=b,		
	ht(hoh)=h		bt(doro)=a.		
	51(505) 5		ht(brian)=2		
			DI(I		





So, apply "standard" EM



Solution 1: Aggregators



Problem: Does not take into account the interaction between Rain and Temp

Solution 2: Combining Rules



- The 3 distributions are combined into one final distribution
- Gradient-descent and EM Methods exist

MLN Weight Learning

• Parameter tying: Groundings of same clause



- It is #P-complete to count the number of true groundings
- Generative learning: Pseudo-likelihood
- Discriminative learning: Cond. likelihood

Slide based on Domingos' talk

Generative Learning

- Function to optimize: $f(w) = \sum_{i} w_{i}n_{i}(x) \log Z$ • Gradient: $Z = \sum_{x} \exp\left(\sum_{i} w_{i}n_{i}(x')\right)$
 - $\frac{\partial}{\partial w_j} f(w) = n_j(x) \frac{1}{Z} \sum_{x'} \exp\left(\sum_i w_i n_i(x')\right) n_j(x')$ $= n_j(x) \sum_{x'} P(x') n_j(x')$ $= n_j(x) E[n_j(x)]$

Counts in training data

Weighted sum over all possible worlds No evidence, just sets of constants Very hard to approximate

Slide based on Domingos' talk

Pseudo-likelihood

$$PL(x) = \prod_{l} P(X_{l} = x_{l} | MB(x_{l}))$$
$$\log PL(x) = \sum_{l} \log P(X_{l} = x_{l} | MB(x_{l}))$$

$$P(X_{l} = x_{l} | MB(x_{l})) = \frac{P(x)}{P(x_{[X_{l}=0]}) + P(x_{[X_{l}=1]})}$$

=
$$\frac{1/Z \exp(\Sigma w_{i} n_{i}(x))}{1/Z \exp(\Sigma w_{i} n_{i}(x_{[X_{l}=0]})) + 1/Z \exp(\Sigma w_{i} n_{i}(x_{[X_{l}=1]}))}$$

$$\frac{\partial}{\partial w_j} \log PL(x) = \sum_l n_j(x) - P(X_l = 0 \mid MB(X_l)) n_j(x_{[X_l = 0]}) - P(X_l = 1 \mid MB(X_l)) n_j(x_{[X_l = 1]})$$
$$= \sum_l n_j(x) - E_{x'_l}[n_j(x_{[X_l = x'_l]})]$$
Slide based on Domingos' talk

Pseudo-likelihood

$$PL(x) = \prod_{l} P(X_{l} = x_{l} \mid MB(x_{l}))$$

While effective, still hard to count in many data sets

• Approximate counting techniques exist (Sarkhel et al. AAAI 2016, Das et al. SDM 2016)

 $(X_{[X_i=1]})$

$$\frac{\partial}{\partial w_j} \log PL(x) = \sum_l n_j(x) - P(X_l = 0 \mid MB(X_l)) n_j(x_{[X_l = 0]}) - P(X_l = 1 \mid MB(X_l)) n_j(x_{[X_l = 1]})$$
$$= \sum_l n_j(x) - E_{x'_l}[n_j(x_{[X_l = x'_l]})]$$
Slide based on Domingos' talk

- Parameter learning graphical models
- Parameter learning StaR AI models
- Structure learning logical models & graphical models
- Structure learning StaRAI models

Probabilistic Graphical Models




Inductive Logic Programming = Machine Learning + Logic Programming The Problem Specification [Muggleton, De Raedt JLP96]

- Given:
 - *Examples:* first-order atomic formulas (atoms), each labeled positive or negative.
 - Background knowledge: definite clause (if-then rules) theory.
 - Language bias: constraints on the form of interesting new rules (clauses).

ILP Specification (Continued)

- Find:
 - A hypothesis h that meets the language constraints and that, when conjoined with B, implies (lets us prove) all of the positive examples but none of the negative examples.
- To handle real-world issues such as noise, we often relax the requirements, so that h need only entail significantly more positive examples than negative examples.

A Common Approach

- Use a greedy covering algorithm.
 - Repeat while some positive examples remain uncovered (not entailed):
 - Find a *good clause* (one that covers as many positive examples as possible but no/few negatives).
 - Add that clause to the current theory, and remove the positive examples that it covers.
- ILP algorithms use this approach but vary in their method for finding a *good clause*.

- Parameter learning graphical models
- Parameter learning StaR AI models
- Structure learning logical models & graphical models
- Structure learning StaRAI models

Vanilla SRL Approach [De Raedt, Kersting ALT04]



• Traverses the hypotheses space a la ILP

. . .

• Replaces ILP's 0-1 covers relation by a "smooth", probabilistic one [0,1]

$$\operatorname{cover}(e, H, B) = P(e|H, B)$$

 $\operatorname{cover}(E, H, B) = \prod_{e \in E} \operatorname{cover}(e, H, B)$

Structure learning methods for MLNs

• Top-down approach:

- GSL[Kok & Domingos, 2005], DSL[Biba et al., 2008]
- Start from unit clauses and search for new clauses
- Bottom-up approach:
 - BUSL [Mihalkova & Mooney, 2007], Hypergraph Lifting [Kok & Domingos, 2009], Structural Motifs [Kok & Domingos, 2010]
 - Use data to generate candidate clauses
- Max-Margin Approach:
 - Discriminative learning [Huynh & Mooney, 2008]
 - Effectively learns horn clauses
 - Uses regularization to force parameters to zero
 - Later extended to online setting

Learning via Hypergraph Lifting

[Kok & Domingos, ICML'09]



- Relational DB can be viewed as hypergraph
 - Nodes ' Constants
 - Hyperedges ´True ground atoms

Learning via Hypergraph Lifting

[Kok & Domingos, ICML'09]



Learning via Hypergraph Lifting

[Kok & Domingos, ICML'09]







Clause Creation



Clause Creation

Advises(p, s) and Teaches(p, c) and TAs(s, c)
Advises(p, s) V not Teaches(p, c) V notTAs(s, c)
Advises(p, s) V Teaches(p, c) V notTAs(s, c)

Outline (Higher Level)

- State-of-the-art Structure Learning
- Effective Learning
- Advice Taking
- Actively Seeking Advice



Functional Gradient Boosting

Key Insight: Learn multiple weak models rather than a single complex model

- First explored in the context of StaRAI by Kersting & Driessens



Friedman et al 2001, Dietterich et al. 2004, Kersting & Driessens 08, Natarajan et al. MLJ 2012, Springer Brief '15

Relational Dependency Network

- Cyclic directed graphs
- Approximated as product of conditional distributions





Sum all gradients to get final ψ

$$\psi_m = \psi_0 + \Delta_1 + \dots + \Delta_m$$

Markov Logic Networks

Richardson & Domingos '05

• Weighted logic





Normalization term sums over all world states

$$PLL(\mathbf{X} = \mathbf{x}) = \sum_{x_i \in \mathbf{x}} \log P(x_i | \mathbf{MB}(x_i))$$
• Learning approaches maximize oseudo-
logli
Key Insight: View MLNs as sets of RDNs

Functional gradient for MLNs

 $_{i}))$

RDN



Regression tree uses aggregators (e.g. Exists) MLN

Learning optimizes a product of conditional distributions x_i)

Probability of x Each conditional distribution not learned independently

Ψ(X)

Regression tree scales output by the number of groundings (Shavlik & Natarajan '09)

 $-n_j(x_i=0;\mathbf{MB}(x_i))$

+1

MLN from Trees



 $w_1: p(X), q(X, Y) \to target(X)$ $w_2: p(X), \neg(\exists Y, q(X, Y)) \to target(X)$ $w_3: \neg p(X) \to target(X)$

Learning Clauses

- Same as squared error for trees
- Force weight on false branches (w₃, w₂) to be 0
- Hence no existential vars needed

$$w_1 : p(X), q(X, Y) \rightarrow target(X)$$

Similar algorithm for learning Relational Logistic Regression (Ramanan et al., KR 18 - to appear)

Break!!!



RFGB-EM



Relational RL (under review)



What can be learned?



Natarajan et al (2010, 2011, 2012, 2013) Khot et al (2011, 2013), Yang et al (2016)

What can be learned?

Natarajan et al (2010, 2011, 2012, 2013) Khot et al (2011, 2013), Yang et al (2016), Hadiji et al (2015), Yang et al (2017)



Impressive Results on Standard Domains



Predicting the	Algo	Likelihood	AUC-ROC	AUC-PR	Time
advisor for a	Boosting	0.810	0.961	0.930	9s
student	RPT	0.805	0.894	0.863	1s
	MLN	0.730	0.535	0.621	93 hrs
Movie Cit Recommendation	Image: With State of the s			ning from	
Unstructured Web Text Web	- 10 (- 115 - 11 M	Scale o G facts desc k drug-dise M facts on 1	f Learning cribing the ease intera NLP tasks	Structure recommendations	endations

Natarajan et al. MLJ'12, Khot et al. ICDM '11, Natarajan et al. IJCAI '11, Natarajan et al. IAAI '13 Weiss et al. IAAI '12 AI Magazine '12, Natarajan et al. IJMLC '13, Khot et al. MLJ' 14

Several Real Applications



Weiss et al (2012,2013). Natarajan et al (2013,2012, 2014, 2015), Shivram et al (2014), Picado et al (2014) Soni et al (2016), Viswanathan et al (2016), Odom et al (2014,2015a, 2015b), Yang et al (KBS 2017)

Scaling Boosting

Efficient counting using Databases



1. Encode the data into relational database

- 2.Reformulate counting as join queries
- 3.Reuse the join query results for different clauses with common predicates and in subsequent induction steps
- 4.Use modes to restrict the search process

DB Boost Results



Approximate Counting via Graph DB

(Das et al., under review)



Compute summary statistics of the motifs on the resulting hypergraph Approximate the exact counts with these statistics



Related work by Das '16, Venugopal et al, '15 (MLNs)

Try it yourself

• https://starling.utdallas.edu/software/boostsrl/

Tutorial

<u>https://starling.utdallas.edu/software/boostsrl/wiki/</u>

BoostSRL: "Boosting for S ×		Θ – □
$ \rightarrow$ \mathcal{C} \bigtriangleup \triangleq Secure https://starling.utda	Illas.edu/software/boostsrl/	\$
As with the standard gradient-boo sequence of regression models. Th models (i.e. regression models that and the output are essentially first- predicates.	sting approach, our approach turns the e key difference to the standard appro coperate on relational data). We assum order regression trees where the inner	e model-learning problem to learning a aches is that we learn relational regression ie the data to be in predicate-logic format nodes contain conjunctions of logical
Latest Release	License	Wiki
tag v1.0.0	license GPL-3.0	Documentation
Getting Started		
Prerequisites		
• Java (tested with openjdk 1.8.0	_ 144)	
Installation		
• Download stable jar file.		
 Download stable source with g 	git.	
git clone -b master https://gith	nub.com/boost-starai/BoostSRL.git	
 Nightly builds with git. 		
git clone -b development https:,	//github.com/boost-starai/BoostSRL.git	
Basic Usage		
[hayesall@hawk Datasets]\$ ls -R (Cora/:	Cora/	



BoostSRL assumes that data are contained in files with data structured in predicate-logic format.

Positive Examples:



Negative Examples:

Outline

- State-of-the-art Structure Learning
- Effective Learning
- Advice Taking in SRL
- Actively Seeking Advice

Humans can do more (or less)!



• SRL uses FOL as underlying representation

- Used for designing classical AI systems

 SRL research is primarily data-driven or fully expert-driven





Advice for SRL



Goal: Expand Knowledge-Based Systems to SRL





Quantity of Training Data

Fung et al. 2002, Towell and Shavlik 1994, Kunapuli et al. 2013


Yang & Natarajan ECML '13, Yang et al. ICDM '14



Powerful framework that can incorporate different kinds of advice

Odom et al. AAAI '15

Types of Advice

Privileged Information





Question ^[1]	Score
What is your age? (1 point)	
What is the time to the nearest hour? (1 point)	
Give the patient an address, and ask him or her to repeat it at the end of the test. (1 point)	
e g. 42 West Street	
What is the year? (1 point)	i
What is the name of the hospital or number of the residence where the patient is situated? (1 point)	
Can the patient recognize two persons (the doctor, nurse, home help, etc.)? (1 point)	
What is your date of birth? (day and month sufficient) (1 point)	
In what year did World War 1 begin? (1 point)	
(other dates can be used, with a preference for dates some time in the past.)	
Name the present monarch/dictator/prime minister/president. (1 point)	1
(Alternatively, the question "When did you come to [this country]? " has been suggested)	
Count backwards from 20 down to 1. (1 point)	1

Wearable Sensors



Deployment/ Test

м	Venu	0	Ind Dataset Ave.	and include 18 Mar	100 Care 840	ten 1	_	Welco	omel Steve Rathschild	_		Looput	Help		_	-		
0	+ DHR	÷	Dashboard D	ate Compl	ete History *	Encounter N	, Provider no	ne* Filter	Print Report: Fin	ancial Analysis	*			1			ì	
	Patient Dashboard									Dashboard Column(s) 2 •					1	10.000		1
ł.									land.					R.		100		
			encounters	·			~	108			^		189	100 m 10 m	and the		7	
			Dete Dector Name Diagnosis				Status	Vital Dute HT WT Entered BT DR/20/2011 233AM 8 to 83 km Ritech Selvu					100 100	Y	200			
			01/17/2012	Steve				Doen	Add Vitab					18 6	-	1000		
			11:09:4M	Rothschild					Notes			^		(R) 📖	-	1.1		
			11/09/2011 6:529M	Rothschild	Ald		Open		And Andrew Andrew Andrew				10 10 10 10		S 40			
			09/28/2011	Steve	Abdominal	oain. left lower quadra	rt .	Doen	11/00/2011 6-5204	Gran Pathoohild	Standard I	in the second		10.00		16		
			102AM	Rothschild	8				09/28/2011 1-02AM	Steve Rothschild	Followap	In-Progress		10.100	1000	11		
			248AM	Rothschild	wall initial e	pisode of care	eveela	Dpen	09/28/2011 1:02AM	Steve Rothschild	Standard 2	In-Progress		100	108 C	Y		
			07/26/2011	Steve	Acute myo	ardial infarction of an	erolateral	Doen	07/25/2011 248AM	Steve Rothschild	Chiro Exam	In-Progress		100	1000	Sec. C.		
			2.52AM	Rethochild	wall initial of	pisode of care			07/26/2011 248AM	Steve Rothschild	Consultation Form 2	In-Progress		1	100 C	and the second		
		1	-				_	_	07/26/2011 2:48AM	Steve Rothschild	Standard 2	In-Progress		100	A . A	per la	l	
	Physical Examination of		1985 *						12					100	1 11	1 N		ł
	- Prescription -								Add Notes						and the Y	100		
	I have been criptica	Alerjes ^					Problems	Poblems				SH -		all and				
	Prescription History		Allergy Type		Allergy Name		Reaction		Problems		Date			AL .	1	N E		1
			Drug		penicillin (obsei	eta)	hives		AA anyloid nephrop	sthy	08/30/2011 7-254M			131		S-3-01		
			Drug A-Cof I			Drug A-Cof DH well								101	2012	Acres in		
			Add Allergy						Orders					181	10.00	3600		
			Diagnosis					~	Contract Data		famile fame			1	14	a she		
			ICD D	escription	Status	Condition	Out		10/11/2011 12/00/44	Office Viel	Ed Back Encoded 10 Mar	ter.				1000		
	Registration		Variant Creut	afeldt-Jakob		Improved,	09/27/29	11	10/07/2011 647AM	Tripluceride	Los Press Polaster av mins	10			2.2.2	A DOWN	i	2
	Appointment	٠	disease			Metastatic, sore	8.02PM		Add Orders	. appende							1	



What is this Advice?



Advice for RFGB: Update Gradients

Idea: Incorporate advice into the gradients



Odom et al. AAAI '15, Yang et al. ICDM '14

Learning Framework

$$MLL(\mathbf{x}, \mathbf{y}) = \sum_{x_i \in \mathbf{x}} \log \frac{\exp(\varphi(x_i; y_i))}{\sum_{y'} \exp(\varphi(x_i; y') + cost(y_i, y', \varphi))}$$

Label Preferences:

$$cost(x_i, \varphi) = -\lambda \times \varphi(x_i) \times [n_t(x_i) - n_f(x_i)]$$

Advice $\uparrow P(label_i)$

Advice $\downarrow P(label_i)$

Gradients:

$$\Delta(x_i) = I(y_i = 1) - P(y_i = 1; \psi) + \lambda [n_t(x_i) - n_f(x_i)]$$

(Aside) Linear Program Formulation for Sequential Decision-Making



Kunapuli et al. (2013)

Sample Results Relational Classification





Sample Results Cost-sensitive Learning





Advice for Adverse Drug Events NLP

- "If a drug and an effect are present in a proposed ADE and a sentence contains both the drug and effect, the ADE is true"
- "If a drug and an effect are present in a proposed ADE, and a sentence contains both the drug and effect, and the sentence contains the pattern 'effect after drug', the ADE is true"
- "If a drug and an effect are present in a proposed ADE, and a sentence contains both the drug and effect, and the sentence contains the pattern 'drug-induced effect', the ADE is true"

Experiments: Adverse Drug Events



Outline

- State-of-the-art Structure Learning
- Effective Learning
- Advice Taking in SRL
- Actively Seeking Advice

Knowledge-Based Learning



Fung et al. 2002



Active vs Passive Learning



Active Learning

- Learn initial model from training data m_i
- Generate prediction over data $P_{m_i}(y_i|x_i)$
- Calculate uncertainty $H(P_{m_i}(y_i|x_i))$
- Select example(s) $\operatorname{argmax}_{x_i} H(P_{m_i}(y_i|x_i))$

Relational Active Advice Seeking

- Learn initial model from training data
- Generate prediction over data $P_{m_i}(y_i|x_i)$
- Calculate uncertainty $H(P_{m_i}(y_i|x_i))$
- Select example(s) $\operatorname{argmax} H(P_{m_i}(y_i|x_i))$

Select **advice clause** with the highest uncertainty

Odom & Natarajan (2016)

Relational Query Generation



Predict which states vote democrat?

Relational Query Generation



Predict which states vote democrat?

Relational Query Generation



Queries:

Region(x,West),Gov(x,y),Party(y,Dem)
PrimaryIndustry(x, Manufacturing)

Relational Active Advice Seeking

Query: P ^ Q

0.8

0.6

0.2

- Query Conjunction of literals
 - Defines set of examples
 - Learned by fitting regression trees on current uncertainty
- Expert Response Selects Preferred/Avoided labels



(Aside) Sequential Decision-Making

Goal: Select states (c_j) according to their cluster quality (f) and the total uncertainty of all states $(S = c_j \cup c_k)$

$$\arg\max_{\boldsymbol{c}_j} f(|\boldsymbol{c}_j|) [U_c(\boldsymbol{c}_j) + U_c(\boldsymbol{c}_k)]$$

We simplify this intractable global optimization problem by:

- 1. Clustering the states to define a finite set of queries
- 2. Considering an uniform cluster quality when selecting a query



Odom & Natarajan (2016)

(Aside) Preference Guided Planning



- Guides search progressively
- n-step Roll-out
- Evaluate each **option** $m_i : S(m_i)$ where $m_i \in M_\tau$

•
$$S(m_i) = Func\left(C_{m_i}, D_{m_i}, A_{m_i}^{(pref)}\right)$$

- $S(M_{\tau}) \rightarrow P(M_{\tau})$ (induce distribution)
- Uncertainty $U(M_{\tau}) = -\sum_{m \in M_{\tau}} p(m) \log(\frac{1}{p(m)})$
 - Query human based on uncertainty in decision

HMCL '17 Workshop at AAAI, AAMAS 18, KBS (under review)

Joint work with Jana Doppa, Rakibul Islam and Dan Roth

(Aside) HTN Planning Results

10 Domains ; Baselines : (1) Pre-Encoded/Upfront Preferences, (2) Random Queries & (3) No Preferences



AAMAS 2018



More Problems solved

Better solutions obtained

(Aside) Interfaces for Communication CwC – Joint work with Dan Roth, Jana Doppa, Julia Hockenmaier

<u>\$</u>		- 🗆 🗙
Select Game: 122	Simulation Count: 0	UCT-C: 0 Sample Count: 0 Start
K A	• •	Human and Planner Communication Panel
		<planner> Would you like to interact with the planner while solving (y/n)? <human> y <planner> Would you like to give any suggestions now, before we start? <human> no <planner> Okay, let me see if I can find a move that works <planner> Currently, I am trying to move the eight of diamonds blocking the three of diamonds. I'm not sure what to do now. Please suggest a move. <human> finish the ace of diamonds <planner> The ace of diamonds was finished.</planner></human></planner></planner></human></planner></human></planner>
Plan cost: 4.0		
free the king of diamonds finish the ace of hearts		
finish the ace of clubs finish the ace of diamonds		
		Send
·	/	

ER Diagrams for Preference Specification? (KCAP 2017)



Outline

- Effective Learning
- Advice Taking in SRL
- Actively Seeking Advice
- Wrap-Up

Precision Health Tasks

CARDIA EXAM COMPONE NTS—ALL YEARS Schedule of components in the core study, substudies, and ancillary studies by CARDIA exar

				Year	Exam'								Year	Exan			
	1985	1987	1990	1992	1995	2000	2005	2010		1985	1987	1990	1992	1995	2000	2005	2010
	0	2	5	7	10	15	20	25		0	2	5	7	10	15	20	25
CORE STUDY																	
BLOOD PRESSURE									Weight	х	х	x	x	х	х	x	х
Resting	x	x	х	х	x	x	x	х	S kinfolds	х	х	х	х	х	-	-	
S tanding		x	-	-	-	-	-		ChestCircumference	-	х	х	-	-	-	-	
Reactivity		х	-	-	-	-	-	-	Waist Circumference	х	х	х	х	х	х	х	х
CHEMIS TRIES									Hip Circumference	х	х	х	х	х	-	-	х
Genetic									Thigh Circumference	-	-	-	x	-	-	-	-
DNA Storage	-	-	x	-	x	x	x	х	Elbow Breadth	х	х	-	-	-	-	-	
Stored Cells for Cell Immortalization	-	-	-	-	-	x	-	-	S houlder Breadth		х						-
Plasma									Sitting Height		х	-	-			-	
Lipids	х	х	х	х	х	х	х	х	Toenails	-	х	-	-	-	-	-	
Lipoproteins	х	Х	Х	х	Х	х	х	-	EveColor	-	-	х		-	-	-	-
Apoproteins	х	х	-	-	-	-	-	-	Skin Reflectance	-	-	-	х	-	-	-	
CBC	х	-	-	-	-	-	-	-	MEDICAL HISTORY								
Lp(a)	-	-	х	-	-	-	-	-	Medical History	x	x	x	x	x	x	x	x
Fibringen	-	-	x	-	-	-	x		Illicit Drug Use	x	x	x	x	x	x	x	x
AnoE Phenotyne				х					Danth Cartificata		x	x	x	x	x	x	x
S tored Plasma		x	x	x	x	x	x	х	Montal Frants				x	x	x	x	x
C-Reactive Protein				х		х	х	х	Safety Onastionasia		x						
Interleukin-6							х		Interim Hospitalization		x	x	x	x	x	x	x
Serium									Chart Pric/Palaitations			x		x			
Cotinine	х	-	-	-		-	-	-	History of Lyng Deblam	x	x			x	x	x	
SMAC 12	x	-			-	-		-	Oral Contracention Winters					- X			
Fastion Insulin	x	-		x	x	x	x	х	Woman's Rangeluctive Haalth						x	x	x
Fasting Glucose	x	-		x	x	x	x	x	Slaar Habits						x	x	
Oral Gincore Tolerance Test					x		x	x	Tehase	x	x	x	x	x	x	x	x
Stored Serum	х	х	х	х	x	х	x	x	Alcohol	x	x	x	x	x	x	x	x
GGT	x				x			1.1	Weight History	x	x						x
Samm Crastinina	x	-		-	x	x	x	x	Social amongation	x	x	x	x	x	x	x	x
Uric Acid	x			-	x	x			FAMILY HISTORY OURSTIONNAIRE	~	~	~	~	~	~	~	~
Urine									Family History	x		x		x			x
Usinan: Continina					x	x	x	x	DEVELOAL ACTIVITY FITNESS				-	~	-	-	~
Albuminuria					x	x	x	x	7 Day Physical Activity	x							
ANTHRODOMETRY									Phone ind Activity Operation	÷	÷	÷	÷	÷	÷	÷	÷
Hainbe	x	x	x	x	x	x	x	x	Enystea Activity Questionnaire	÷	~	~	v	~	~	~	~
11Cigin.	~	~	~	~	~	~	~	**	Bracke Distriction 1981	~	-	v	~	v	-		-
1 mar of study indicates when original data col	laction or:	ment a		dies and	a contraction of the	-			Disecce Questioningite	-	-	~	~	- 0	-	-	
year or acousy assocates when original data col	uecuoti oci	correct, as	soay Of C	coung tray	, occur is	1121			nousenoid Unores	-	-		^	~	-		v
									secentary penavior Questionnaire	-	-	-	-	-	-	-	~

Cardiovascular events/procedures from EHR/Clinical Study (IAAI 12, IAAI13, AI Magazine 12, AIME 15a, AAAI 15, AAAI 16, AIME 15b, AIME 17a, AIME 17b)



Alzheimer's /diabetes prediction from FMRI (IJMLC 13, Neuroradiology, ICMLA 12Jour of Neurotrauma 15)



Adverse Drug events from EHR/Med abstracts AAAI 12, AIME 15, EMNLP (under preparation)

Theme	Questions	Answers							
	Age	Number entry							
	Gender	Male, Female, Other							
Demographic	Country of Residence	Text entry							
Information	Marital Status	{Married, Living with a Partner, Divorced, Separated, Widowed, Single, Other}							
	Employment	{Full time, Part time, Retired, Student, Disabled, Not Employed for Pay}							
	Education	{Less than grade 8, some high school, completed high school, technical/trade/vocational school, some college/university, completed college/university, some post-graduate education, completed post-graduate ed- ucation}							
1	Disease name	Text entry							
Disease Information	How many years has it been since you first started experiencing symptoms?	Number entry							
	How many years has it been since you were diag- nosed?	Number entry							
	How severely do your symptoms impact your life?	5 point scale from No impact to Extreme impact							
	How often do you use the internet?	{Several times a day, About once a day, Several times a week, Every few weeks, Less often, Never}							
	Do you own any of the following technologies? (Check all that apply)	{Desktop computer, Laptop computer, Cell phone, e- Reader, MP3 Player, Game console, Tablet}							
Technology	On your cell phone, do you have any applications that help you track or manage your health?	{Yes, No, I dont have a cellphone}							
Use	Do you ever use your cell phone to look up health or medical information?	$\{Y\!es,No,I\;dont\;have\;a\;cellphone\}$							

Predicting rare diseases/PPD from survey data (CHASE 16, CHASE 17)



Parkinsons' prediction from Study data (AIME 17)

Conclusions

- Al in the wild is more than a single table Graphs, different data types, relational DBs, ... are central to HAAI
- Symmetry/Relation aware AI is essential
- Sequential models and dynamic models that do not make simplistic assumptions are necessary
- Human is an ally in learning and AI system needs to efficiently use human knowledge and input
- Deployment: medicine, social science, traffic, journalism, ... and the Data Science Genome: Machines read and understand data science publications and help the user with their problem at hand

Next Steps

- Web-scale StaRAI
- More applications
 - NELL, relation extraction, more EHRs, imaging data
- Learning from Multiple Experts
- Efficient Inference for CTBNs
- "Truly" Hybrid Models
- Combine with efficient inference techniques
 - Approximate counting & Approximate inference
- Solve a *dream* problem
 - Relational POMDPs