



Andreas Holzinger

185.A83 Machine Learning for Health Informatics
2018S, VU, 2.0 h, 3.0 ECTS
Lecture 06 - Module 04 – Week 20 - 15.05.2018



Probabilistic Graphical Models Part 2: From Bayesian Networks to Probabilistic Topic Models

a.holzinger@hci-kdd.org

<http://hci-kdd.org/machine-learning-for-health-informatics-course>



Holzinger Group, hci-kdd.org

1

Machine Learning Health 06

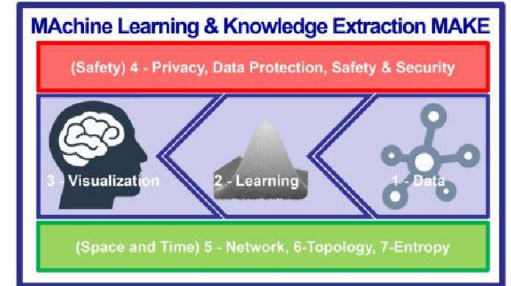


Science is to test crazy ideas –
Engineering is to put these ideas into Business
Lucky Students ☺

Holzinger Group, hci-kdd.org

2

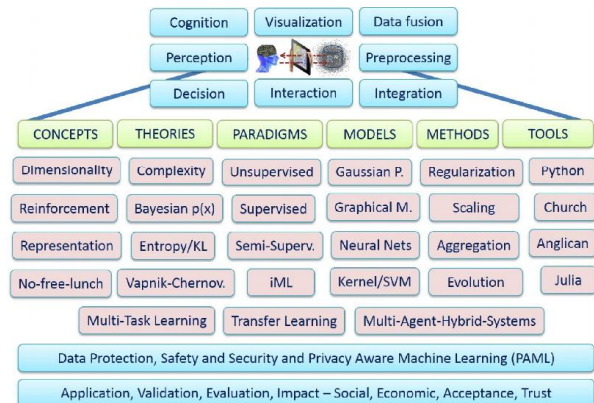
Machine Learning Health 06



Holzinger Group, hci-kdd.org

3

Machine Learning Health 06



Holzinger, A. 2016. Machine Learning for Health Informatics. In: LNCS 9605, pp. 1-24, doi:10.1007/978-3-319-50478-0_1.

Holzinger Group, hci-kdd.org

4

Machine Learning Health 06

- 00 Reflection
- 01 Probabilistic Decision Making
- 02 Probabilistic Programming Part II
- 03 Probabilistic Topic Models
- 04 Knowledge Representation in Net Medicine
- 05 ML on Graphs Examples
- 06 Digression: Similarity
- 07 Graph Measures
- 08 Point Clouds from Natural Images



Holzinger Group, hci-kdd.org

5

Machine Learning Health 06



Holzinger Group, hci-kdd.org

6

Machine Learning Health 06

04

- 1) learn from prior data
- 2) extract knowledge
- 2) generalize,
 - i.e. guessing where a probability mass function concentrates
- 4) fight the curse of dimensionality
- 5) disentangle underlying explanatory factors of data, i.e.
- 6) understand the data in the context of an application domain



Holzinger Group, hci-kdd.org

7

Machine Learning Health 06

$$\mathbb{E}[f] = \int f(z)p(z)dz$$

$$\hat{f} = \frac{1}{L} \sum_{l=1}^L f(z^{(l)})$$



Compute $a_i := \sum_j f_{ij}x_j$
Draw u from $\text{Uniform}(0, 1)$
If $u < 1/(1 + e^{-2a_i})$
 $x_i := +1$
Else
 $x_i := -1$

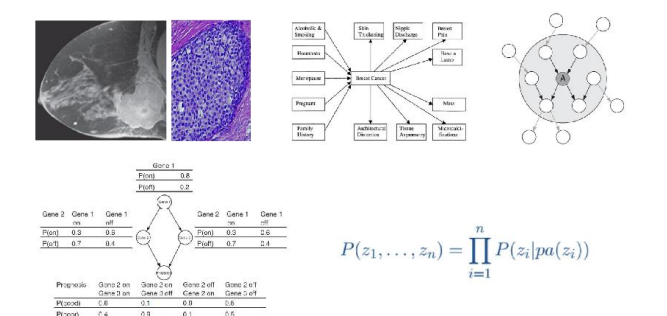


Propp, J. G. & Wilson, D. B.
1996. Exact sampling with
coupled Markov chains and
applications to statistical
mechanics. Random
structures and Algorithms,
9, (1-2), 223-252.

Holzinger Group, hci-kdd.org

8

Machine Learning Health 06



Gevaert, O., Smet, F. D., Timmerman, D., Moreau, Y. & Moor, B.
D. (2006) Predicting the prognosis of breast cancer by
integrating clinical and microarray data with Bayesian networks.
Bioinformatics, 22, 14, 184-190.

Holzinger Group, hci-kdd.org

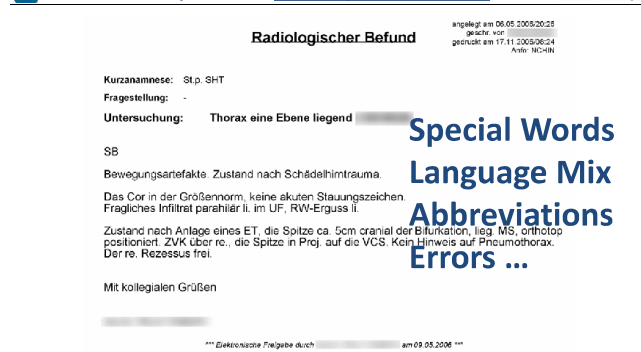
9

Machine Learning Health 06

- For certain cases it is tractable if:
 - Just one variable is unobserved
 - We have singly connected graphs (no undirected loops -> belief propagation)
 - Assigning probability to fully observed set of variables
- Possibility: Monte Carlo Methods (generate many samples according to the Bayes Net distribution and then count the results)
- Otherwise: approximate solutions, NOTE:

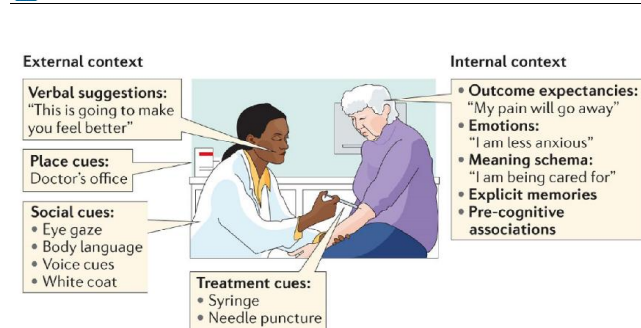
Sometimes it is better to have an approximate solution to a complex problem – than a perfect solution to a simplified problem

- HWI =
 - Harnwegsinfekt
 - Hinterwandinfarkt
 - Hinterwandischämie
 - Hakenwurminfektion
 - Halswirbelimmobilisation
 - Hip Waist Index
 - Height-Width Index
 - Heart-Work Index
 - Hemodynamically weighted imaging
 - High Water Intake
 - Hot water irrigation
 - Hepatic weight index
 - Häufig wechselnder Intimpartner
- Leitung = Nervenleitung, Abteilungsleitung, Stromleitung, Wasserleitung, Harnleitung, Ableitung, Vereinsleitung ☺...




Holzinger, A., Geierhofer, R. & Errath, M. 2007. Semantische Informationsextraktion in medizinischen Informationssystemen. *Informatik Spektrum*, 30, (2), 69-78.

- Intelligence?
 - Hundreds of controversial definitions – very hard to define;
 - For us: ability to solve problems, to make decisions and to acquire and apply knowledge and skills
- Learning?
 - Different definitions – relatively hard to define
 - basically acquisition of knowledge through prior experience
- Problem Solving?
 - Process of finding solutions to complex issues
- Reasoning?
 - ability of our mind to think and understand things
- Sense Making?
 - Process of giving meaning to experience
- Causality?
 - Relationship between cause and effect
- Decision Making?
 - Process of “de-ciding” (“ent-scheiden”) between alternative options



Nature Reviews | Neuroscience

Wager, T. D. & Atlas, L. Y. 2015. The neuroscience of placebo effects: connecting context, learning and health. *Nat Rev Neurosci*, 16, (7), 403-418, doi:10.1038/nrn3976



01 Probabilistic Decision Making

Laplace, P-S. 1781. Mémoire sur les probabilités. *Mémoires de l'Académie Royale des sciences de Paris*, 1778, 227-332.

- Type 1 Decisions:** related to the **diagnosis**, i.e. computers are used to assist in diagnosing a disease on the basis of the individual patient data. Questions include:
 - What is the probability that this patient has a myocardial infarction on the basis of given data (patient history, ECG, ...)?
 - What is the probability that this patient has acute appendices, given the signs and symptoms concerning abdominal pain?
- Type 2 Decisions:** related to **therapy**, i.e. computers are used to select the best therapy on the basis of clinical evidence, e.g.:
 - What is the best therapy for patients of age x and risks y, if an obstruction of z % is seen in the left coronary artery?
 - What amount of insulin should be prescribed for a patient during the next 5 days, given the blood sugar levels and the amount of insulin taken during the recent weeks?

Harold C. Sox, Michael C. Higgins & Douglas K. Owens 1988. Medical decision making, Second Edition, Chichester, Wiley.

Decision Making under Uncertainty HCI-KDD

$\mathbb{E}[f] = \int p(x)f(x) dx$

$\mathbb{E}[f] \simeq \frac{1}{N} \sum_{n=1}^N f(x_n)$

Holzinger Group, hci-kdd.org 19 Machine Learning Health 06

Remember: Expected Utility Theory $E(U|d)$ HCI-KDD

For a single decision variable an agent can select $D = d$ for any $d \in \text{dom}(D)$.
The expected utility of decision $D = d$ is

$E(U | d) = \sum_{x_1, \dots, x_n} P(x_1, \dots, x_n | d) U(x_1, \dots, x_n, d)$

An optimal single decision is the decision $D = d_{\max}$ whose expected utility is maximal:

$d_{\max} = \arg \max_{d \in \text{dom}(D)} E(U | d)$

Von Neumann, J. & Morgenstern, O. 1947. Theory of games and economic behavior, Princeton university press.

Holzinger Group, hci-kdd.org 20 Machine Learning Health 06

Bayesian Data Analysis Third Edition <http://www.stat.columbia.edu/~gelman/book/> HCI-KDD

Andrew Gelman, John B. Carlini, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin

https://github.com/avehtari/BDA_py_demos
<http://www.stat.columbia.edu/~gelman/book/data/>

Holzinger Group, hci-kdd.org 21 Machine Learning Health 06

Human learning vs. Machine Learning HCI-KDD

- Example 1: Inverse Probability
- Example 2: Diagnosis
- Example 3: Language understanding

$p(h|d) \propto p(D|\theta) * p(h)$

$P(\text{words}|\text{sounds}) \propto P(\text{sounds}|\text{words}) * P(\text{words})$

Recognize speech
Wreck a nice beach

- Learning ensures that new observations (d) match our previous hypotheses (h)

Holzinger Group, hci-kdd.org 22 Machine Learning Health 06

Cognition as probabilistic inference HCI-KDD

- Visual perception, language understanding, motor learning, associative learning, categorization, concept learning, reasoning, causal inference, ...
- Learning concepts from (few!) examples
- Learning and applying intuitive theories (balancing complexity vs. fit optimality)

Holzinger Group, hci-kdd.org 23 Machine Learning Health 06

Modeling basic cognitive capacities as intuitive Bayes HCI-KDD

- Similarity
- Representativeness and evidential support
- Causal judgement
- Coincidences and causal discovery
- Diagnostic inference
- Predicting the future

Tenenbaum, J. B., Griffiths, T. L. & Kemp, C. 2006. Theory-based Bayesian models of inductive learning and reasoning. Trends in cognitive sciences, 10, (7), 309-318.

Holzinger Group, hci-kdd.org 24 Machine Learning Health 06

Human brains as probabilistic reasoning machines HCI-KDD

LTM: Prior knowledge \mathcal{H}

STM: $p(\theta|D) = \frac{p(D|\theta) * p(\theta)}{p(D)}$

Uncertain world

Holzinger Group, hci-kdd.org 25 Machine Learning Health 06

Similar as our RL-Agent seeks to maximize rewards HCI-KDD

Intelligent behavior arises from the actions of an individual seeking to **maximize its received reward signals in a complex and changing world**

Agent
Representation
Learning algorithm
Action selection policy

Environment

State $x^{(t)}$
Reward $r^{(t)}$
Action $a^{(t)}$

Sutton, R. S. & Barto, A. G. 1998. Reinforcement learning: An introduction, Cambridge MIT press

Holzinger Group, hci-kdd.org 26 Machine Learning Health 06

De-cision (Ent-scheidung) between alternatives HCI-KDD

$\{a, b, c\}$
→ decision that is best for worst case
Non-deterministic model
~ Adversarial search

$\{a(p_a), b(p_b), c(p_c)\}$
→ decision that maximizes expected utility value
Probabilistic model

Holzinger Group, hci-kdd.org 27 Machine Learning Health 06

02 Probabilistic Programming

Learning representations (θ, h) from observed data

Observed data: 

Training data: $\mathcal{D} = x_{1:n} = \{x_1, x_2, \dots, x_n\} \quad x, y \in \mathcal{H}$

Feature Parameter: θ or hypothesis $h \in \mathcal{H}$

Prior belief \approx prior probability of hypothesis h : $p(\theta) \quad p(h)$

Likelihood $\approx p(x)$ of the data that h is true $p(\mathcal{D}|\theta) \quad p(d|h)$

Data evidence \approx marginal $p(x)$ that $h = \text{true}$ $p(\mathcal{D}) \quad \sum_{h \in \mathcal{H}} p(d|h) * p(h)$

Posterior $\approx p(x)$ of h after seen ("learn") data $d \quad p(\theta|\mathcal{D}) \quad p(h|d)$

$$\text{posterior} = \frac{\text{likelihood} * \text{prior}}{\text{evidence}} \quad p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta) * p(\theta)}{p(\mathcal{D})}$$

$$p(h|d) = \frac{p(d|h) * p(h)}{\sum_{h \in \mathcal{H}} p(d|h) p(h)}$$

Probabilistic Programming Languages

Image credit to Frank Wood (2016)


Graphical Models: BUGS, STAN, Factorio, Infer.NET

PL	AI	ML	STATS
2010 Figaro HANSAI		ProbLog, BLOG, Church, Infer.NET, Factorio	LibBi, STAN, JAGS
2000 IBAL	Prism	KMP	WinBUGS
1990			BUGS

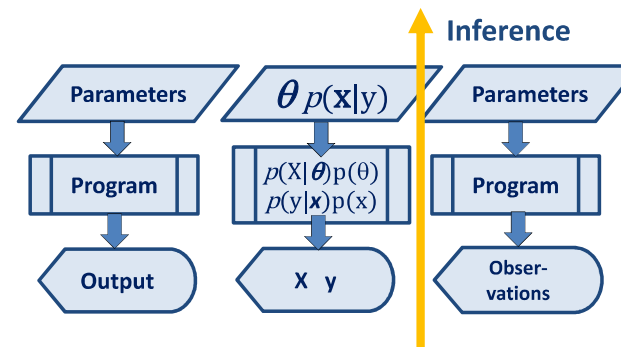
- 1940s: connecting wires to represent 0s and 1s
- 1950s: assemblers, FORTRAN, COBOL, LISP
- 1960s: ALGOL, BCPL (\rightarrow B \rightarrow C), SIMULA
- 1970s: Prolog, FP, ML, Miranda
- 1980s: Eiffel, C++
- 1990s: Haskell, Java, Python

- Dan ROY: Probabilistic Programming Wiki <http://www.probablistic-programming.org/wiki/Home>
- Frank WOOD, many tutorials, slides, code and papers <http://www.robots.ox.ac.uk/~fwood/teaching/index.html>
- Avi PFEFFER 2016. Practical probabilistic programming, Shelter Island (NY), Manning <https://www.manning.com/books/practical-probabilistic-programming>

Look also for work of:
Andrew GORDON
Noah GOODMAN
Josh TENENBAUM
John WINN
Rob ZINKOV
Vikash MANSINGHA
David WINGATE



Probabilistic Programming Concept



Frank Wood, Jan-Willem Van De Meent & Vikash Mansinghka. A New Approach to Probabilistic Programming Inference. AISTATS 2014, Reykjavik, JMLR, 1024-1032

Some selected PPLs

- <https://github.com/pymc-devs/pymc>
- <http://infern.net/azurewebsites.net/>
- <http://mc-stan.org/>
- <https://github.com/p2t2/figaro>
- <https://sites.google.com/site/bloginference/>
- <http://projects.csail.mit.edu/church/wiki/Church>
- <http://factorio.cs.umass.edu/>
- <http://www.openbugs.net/w/FrontPage>
- <http://mcmc-jags.sourceforge.net/>

- Take patient information, e.g., observations, symptoms, test results, -omics data, etc. etc.
- Reach conclusions, and **predict** into the future, e.g. how likely will the patient be ...
- Prior = belief before making a particular observation
- Posterior = belief after making the observation and is the prior for the next observation – intrinsically incremental

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \quad p(x_i|y_j) = \frac{p(y_j|x_i)p(x_i)}{\sum p(x_i, y_j)p(x_i)}$$

Comparison

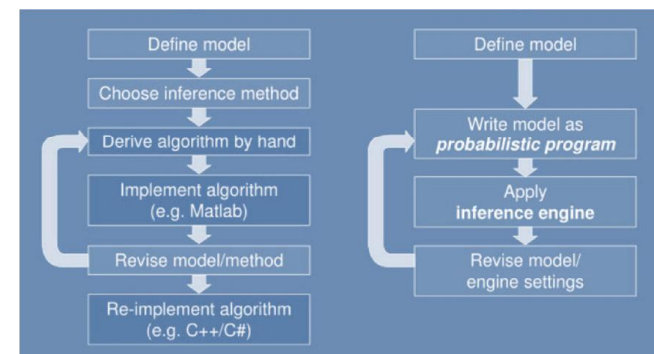


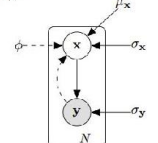
Image credit to John WINN (2010)

Try out WebPPL ("web-people") <http://dippl.org>

```

1 var obs = loadJson("data.json");
2 var guideNet = nn.mlp(1, [(nOut: 3, activation: nn.sigmoid), (nOut: 2)], 'guideNet');
3 var model = function() {
4   var mu_x = nn.mlp(1, [(nOut: 1)], 'mu_x');
5   var sigma_x = softmax(modelParam((name: 'sigma_x')));
6   var sigma_y = softmax(modelParam((name: 'sigma_y')));
7   var latent = sample(Gaussian(mu_x, sigma_x));
8   var nInout = vector(2);
9   var nOutput = nEval(guideNet, nInout);
10  var x = sample(Gaussian(mu_x, sigma_x));
11  guide: Gaussian(mu: 1, get(nOutput, 0));
12  sigma: softmax(1, get(nOutput, 1));
13  });
14  observe(Gaussian((mu: x, sigma: sigma_x)), y);
15  return (x, x);
16  });
17  return latent;
18  });

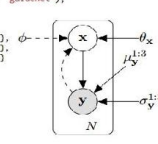
```




```

1 var obs = loadJson("data.json");
2 var nComp = 3;
3 var guideNet = nn.mlp(1, [(nOut: 3, activation: nn.sigmoid), (nOut: nComp-1)], 'guideNet');
4 var model = function() {
5   var theta_x = softmax(modelParam((name: 'theta_x')));
6   var theta_y = softmax(modelParam((name: 'theta_y')));
7   var mu_x = softmax(modelParam((name: 'mu_x')));
8   var mu_y = softmax(modelParam((name: 'mu_y')));
9   var latent = sample(Gaussian(mu_x, sigma_x));
10  var nInout = vector(2);
11  var nOutput = nEval(guideNet, nInout);
12  var x = sample(Gaussian(mu_x, sigma_x));
13  guide: Discrete(softmax(theta_x), (psi: theta_x));
14  sigma: softmax(1, get(nOutput, 1));
15  });
16  observe(Gaussian((mu: x, sigma: sigma_x)), y);
17  return (x, x);
18  });
19  return latent;
20  });

```



Daniel Ritchie, Paul Horsfall & Noah D Goodman 2016. Deep Amortized Inference for Probabilistic Programs. arXiv:1610.05735.



Diederik P Kingma & Max Welling 2013. Auto-encoding variational Bayes. arXiv:1312.6114 (1983 citations as of 13.05.2018 07:00)

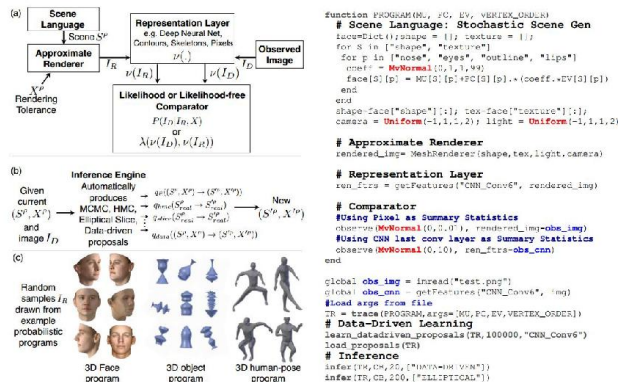
Algorithm 1 Minibatch version of the Auto-encoding VB (AEVB) algorithm. Either of the two SGVB estimators in section 2.3 can be used. We use settings $M = 100$ and $L = 1$ in experiments.

```

θ, φ ← Initialize parameters
repeat
  XM ← Random minibatch of M datapoints (drawn from full dataset)
  ε ← Random samples from noise distribution p(ε)
  g ← ∇θ, φ  $\tilde{\mathcal{L}}^M(\theta, \phi; X^M, \epsilon)$  (Gradients of minibatch estimator (8))
  θ, φ ← Update parameters using gradients g (e.g. SGD or Adagrad (DHS10))
until convergence of parameters (θ, φ)
return θ, φ
    
```

x	y
program source code	program output
scene description	image
policy and world	rewards
cognitive process	behavior
simulation	constraint

Image credit to Frank Wood (2016)



Kulkarni, Kohli, Tenenbaum & Mansinghka. Picture: A probabilistic programming language for scene perception. Proceedings of the IEEE conference on computer vision and pattern recognition, 2015. 4390-4399.

Deep Probabilistic Programming Languages: A Qualitative Study

Guillaume Baudart IBM Research
guillaume.baudart@ibm.com
Martin Hirzel IBM Research
hirzel@us.ibm.com
Louis Mandel IBM Research
lmandel@us.ibm.com

ABSTRACT

Deep probabilistic programming languages try to combine the advantages of deep learning with those of probabilistic programming languages. If successful, this would be a big step forward in machine learning and programming languages. Unfortunately, as of now, this new crop of languages is hard to use and understand. This paper addresses this problem directly by explaining deep probabilistic programming languages and indirectly by characterizing their current strengths and weaknesses.

CIC CONCEPTS

• Theory of computation → Probabilistic computation;
• Computing methodologies → Neural networks;
• Software and its engineering → Domain specific languages;

KEYWORDS

DL, PPL, DSL

1 INTRODUCTION

A deep probabilistic programming language (PPL) is a language for specifying both deep neural networks and probabilistic models. In other words, a deep PPL draws upon programming languages,

These frameworks provide automatic differentiation (users need not manually calculate gradients for gradient descent), GPU support (to efficiently execute vectorized computations), and Python-based embedded domain-specific languages [18].

Deep PPLs, which have emerged just recently [28–32], aim to combine the benefits of PPLs and DL. Ideally, programs in deep PPLs would directly represent uncertainty, yield explainable models, and require only a small amount of training data; be easy to write in a well-designed programming language, and match the breakthrough accuracy and fast training times of DL. Realizing all of these promises would yield tremendous advantages. Unfortunately, this is hard to achieve. Some of the strengths of PPLs and DL are seemingly at odds, such as explainability vs. automated feature engineering, or learning from small data vs. optimizing for large data. Furthermore, the barrier to entry for work in deep PPLs is high, since it requires non-trivial background in fields as diverse as statistics, programming languages, and deep learning. To tackle this problem, this paper characterizes deep PPLs, thus lowering the barrier to entry, providing a programming-language perspective early when it can make a difference, and shining a light on gaps that the community should try to address.

This paper uses the Stan PPL as a representative of the state of the art in regular (not deep) PPLs [9]. Stan is a main-stream, mature,

Approximate Bayesian Image Interpretation using Generative Probabilistic Graphics Programs

Vikash K. Mansinghka^{1,2}, Tejas D. Kulkarni^{1,2}, Yura N. Perov^{1,2,3}, and Joshua B. Tenenbaum^{1,2}

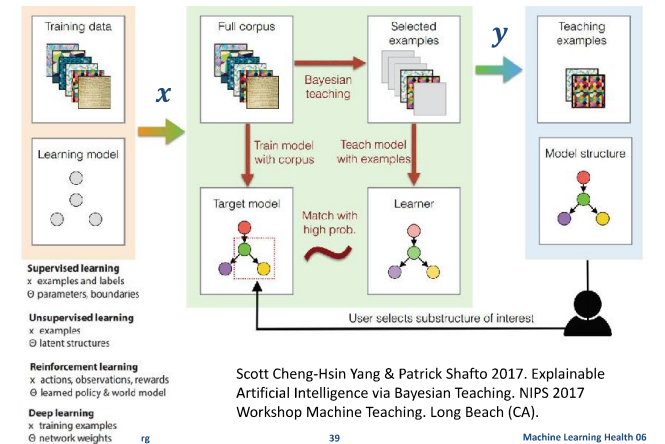
¹Computer Science and Artificial Intelligence Laboratory, MIT

²Department of Brain and Cognitive Sciences, MIT

³Institute of Mathematics and Computer Science, Siberian Federal University

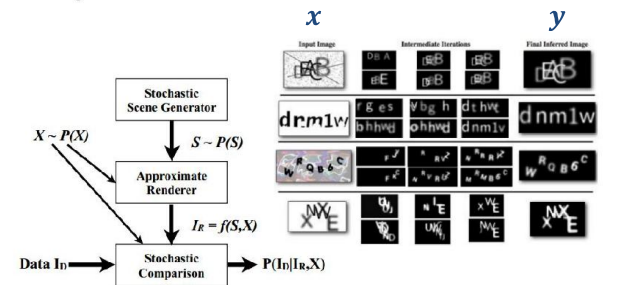
Abstract

The idea of computer vision as the Bayesian inverse problem to computer graphics has a long history and an appealing elegance, but it has proved difficult to directly implement. Instead, most vision tasks are approached via complex bottom-up processing pipelines. Here we show that it is possible to write short, simple probabilistic graphics programs that define flexible generative models and to automatically invert them to interpret real-world images. Generative probabilistic graphics programs (GPGPs) consist of a stochastic scene generator, a renderer based on graphics software, a stochastic likelihood model linking the renderer's output and the data, and latent variables that adjust the fidelity of the renderer and the tolerance of the likelihood. Representations and algorithms from computer graphics are used as the deterministic backbone for highly approximate and stochastic generative models. This formulation combines probabilistic programming, computer graphics, and approximate Bayesian computation, and depends only on general-purpose, automatic inference techniques. We describe two applications: reading sequences of degraded and adversarially obscured characters, and inferring 3D road models from vehicle-mounted camera images. Each of the probabilistic graphics programs we present relies on under 20 lines of probabilistic code, and

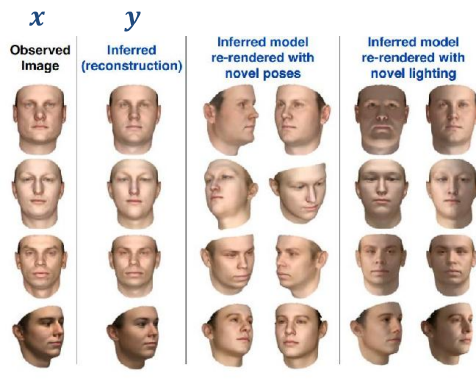


Scott Cheng-Hsin Yang & Patrick Shafto 2017. Explainable Artificial Intelligence via Bayesian Teaching. NIPS 2017 Workshop Machine Teaching. Long Beach (CA).

$$P(S|I_D) \propto \int P(S)P(X)\delta_{f(S,X)}(I_R)P(I_D|I_R,X)dX$$


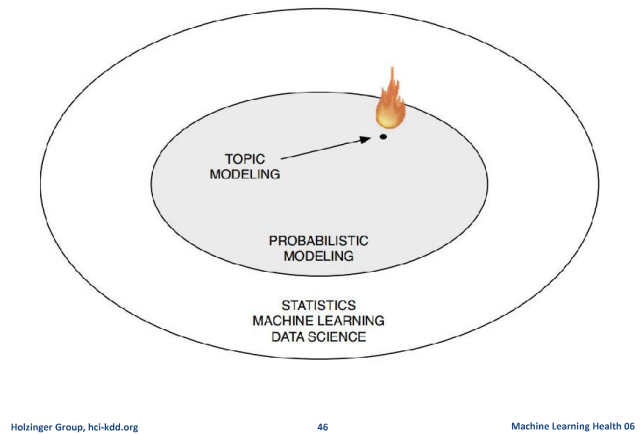


Vikash K. Mansinghka, Tejas D. Kulkarni, Yura N. Perov & Josh Tenenbaum. Approximate Bayesian image interpretation using generative probabilistic graphics programs. In: Burges, Christopher J. C., Bottou, Leon, Welling, Max, Ghahramani, Zoubin & Weinberger, Kilian Q., eds. Advances in Neural Information Processing Systems, 2013 Lake Tahoe. NIPS, 1520-1528.



Kulkarni, Kohli, Tenenbaum & Mansinghka. Picture: A probabilistic programming language for scene perception. Proceedings of the IEEE conference on computer vision and pattern recognition, 2015. 4390-4399.






Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C. & Byers, A. H. (2011) *Big data: The next frontier for innovation, competition, and productivity*. Washington (DC), McKinsey Global Institute.

Table I. Effect of Performance Change on Space Density

Type of indexing	(a) Effect of performance improvement on space density				(b) Effect of performance deterioration on space density			
	Cluster organization A (155 clusters, 2:1 overlap)		Cluster organization B (83 clusters, 1:3 overlap)		Cluster organization A (155 clusters, 2:1 overlap)		Cluster organization B (83 clusters, 1:3 overlap)	
	Standard term frequency weights (f_{ij})	Term frequency with inverse doc. freq. ($f_{ij} \cdot ID_{F_i}$)	Standard term frequency weights (f_{ij})	Term frequency with inverse doc. freq. ($f_{ij} \cdot ID_{F_i}$)	Standard term frequency weights (f_{ij})	Term frequency with inverse doc. freq. ($f_{ij} \cdot ID_{F_i}$)	Standard term frequency weights (f_{ij})	Term frequency with inverse doc. freq. ($f_{ij} \cdot ID_{F_i}$)
Recall-precision output*	—	+14%	—	+14%	—	−10.1%	—	−10.1%
Average similarity between documents and corresponding cluster centroids (x)	.712	.668 (−.044)	.650	.589 (−.061)	.712	.741 (+.029)	.650	.696 (+.046)
Average similarity between cluster centroids and main centroid	.500	.454 (−.046)	.537	.482 (−.055)	.500	.555 (+.055)	.537	.574 (+.037)
Average similarity between pairs of cluster centroids (y)	.273	.209 (−.064)	.315	.252 (−.063)	.273	.329 (+.056)	.315	.362 (+.047)
Ratio y/x	.273/.712 = .383	.209/.668 = .311 (−19%)	.315/.650 = .485	.252/.589 = .428 (−12%)	.273/.712 = .383	.329/.741 = .444 (+16%)	.315/.650 = .485	.362/.696 = .520 (+7%)

* From [2].

Gerard M. Salton, Andrew Wong & Chungshu S. Yang 1975. Vector-Space Model for automatic indexing. Communications of the ACM, 18, (11), 613-620, doi:10.1145/361219.361220.



Information Retrieval and Language Processing
C.A. Montgomery
Editor
A Vector Space Model for Automatic Indexing
G. Salton, A. Wong
and C. S. Yang
Cornell University

In a document retrieval, or other pattern matching environment where stored entities (documents) are compared with each other or with incoming patterns (search requests), it appears that the best indexing (property) space is one where each entity lies as far away from the others as possible; in these circumstances the value of an indexing system may be expressible as a function of the density of the object space; in particular, retrieval performance may correlate inversely with space density. An approach based on space density computations is used to choose an optimum indexing vocabulary for a collection of documents. Typical evaluation results are shown, demonstrating the usefulness of the model.

Key Words and Phrases: automatic information retrieval, automatic indexing, content analysis, document space

CR Categories: 3.71, 3.73, 3.74, 3.75

1. Document Space Configurations



Consider a document space consisting of documents D_i , each identified by one or more index terms T_j , the terms may be weighted according to their importance, or unweighted with weights restricted to 0 and 1.² A typical three-dimensional index space is shown in Figure 1, where each item is identified by up to three distinct terms. The three-dimensional example may be extended to i dimensions when i different index terms are present. In that case, each document D_i is represented by a i -dimensional vector $D_i = (d_{i1}, d_{i2}, \dots, d_{in})$, d_{ij} representing the weight of the j th term.

Given the index vectors for two documents, it is possible to compute a similarity coefficient between them, (D_1, D_2) , which reflects the degree of similarity in the corresponding terms and term weights. Such a similarity measure might be the inner product of the two vectors, or alternatively an inverse function of the angle between the corresponding vector pairs, when the term assignment for two vectors is identical, the angle will be zero, producing a maximum similarity measure.

Instead of identifying each document by a complete vector originating at the 0-point in the coordinate system, the relative distance between the vectors is preserved by normalizing all vector lengths to one, and considering the projection of the vectors onto the envelope of the space represented by the unit sphere. In that case, each document may be depicted by a single

- $D = \langle d_1, d_2, \dots, d_n \rangle$
- $d_i = t_1, t_2, \dots, t_k$
- $w_{i,j} = \begin{cases} 1, & t_i \in d_j \\ 0, & t_i \notin d_j \end{cases} \Rightarrow d_j = (0, 1, 1, 0, 1, \dots, 1)^T$
- $w_{i,j} = \begin{cases} (1 + \log f_{i,j}) * \log \frac{N}{n_i}, & \text{if } f_{i,j} > 0 \\ 0 & \text{otherwise} \end{cases}$

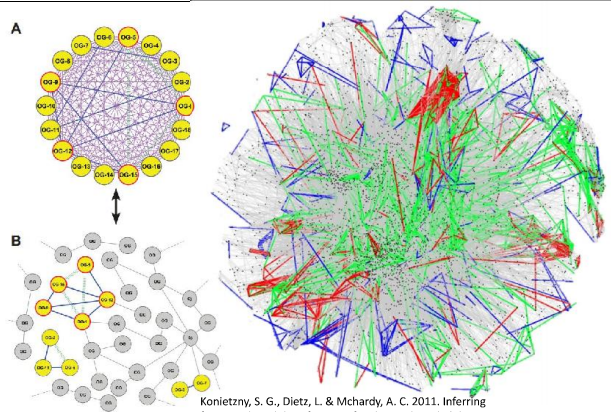
										Bild Senden	PDF	Fall	Diagnostik
	Datum	Zeit	OE	Kurzbez	VMA	Status	Klass.	erg. Enc.	Column	Fall			
<hr/>													
↳	2010018065	17.11.1953		MIKARDIO MK Kardioamb									
↳	Leistungen (KAL, RAD, Therap)	ambulanter Fall 13.01.2010											
↳	EKG (12 Ableitungen)	13.01.2010	08:00	MIKARDIO MK Kardioamb	DUSLTMO	OK				2010018066			
↳	RPR-Intervall-Untersuchung	13.01.2010	08:00	MIKARDIO MK Kardioamb	DUSLTMO	OK				2010018066			
↳	Schrittmacherkontrolle	13.01.2010	08:00	MIKARDIO MK Kardioamb	DUSLTMO	OK				2010018066			
↳	2010022197	ambulanter Fall 04.01.2010		CKTTRANS CK Transp									
↳	Diagnosen Desam (3)												
↳	2003414995	stationärer Fall 20.12.2009		MEDANGIO Med Angio									
↳	Diagnosen Desam (14)												
↳	Leistungen (KAL, RAD, Therap)												
↳	Becken- u. Beckenbodenuntersuchung	22.12.2009	16:36	KINRADIO KIN Raum B	STABELM	OK	PTA			2003414995			
↳	Leisten- u. Leistenbanduntersuchung	21.12.2009	08:30	MIKARDIO MK Kardioamb	SPARANOR	OK				2003414995			
↳	Ernährungsuntersuchung	21.12.2009	08:30	MIKARDIO MK Kardioamb	SPARANOR	OK				2003414995			
↳	2003453621	stationärer Fall 17.11.2009		CKTOMC CK TCMC									
↳	Diagnosen Desam (12)												
↳	Leistungen (MEL) (3)												
↳	Physiolo. i. R. i. stat. Aufnahm	23.11.2009	08:05	CKPHYSIO CK Physio	BEITWALT	OK				2003453621			
↳	Organische, Bildwandlungs	17.11.2009	08:12	CKTROP CK Trop	SCHWIMCH	OK		Organische, Bildwandlungs		2003453621			
↳	Leistungen (KAL, RAD, Therap)												
↳	Becken- u. Beckenbodenuntersuchung	22.12.2009	16:36	KINRADIO KIN Raum B	STABELM	OK	PTA			2003453621			
↳	Leisten- u. Leistenbanduntersuchung	21.12.2009	08:30	MIKARDIO MK Kardioamb	SPARANOR	OK				2003453621			
↳	Ernährungsuntersuchung	21.12.2009	08:30	MIKARDIO MK Kardioamb	SPARANOR	OK				2003453621			
↳	2003453621	ambulanter Fall 16.09.2009		MIKARDIO MK Kardioamb									
↳	Diagnosen Desam (8)												
↳	Leistungen (KAL, RAD, Therap)												
↳	Blutdruck- Langzeit (24 Std)	17.09.2009	10:59	MIKARDIO MK Kardioamb	PULDERM	OK				2003278723			
↳	Blutdruck- Langzeit (24 Std)	16.09.2009	12:02	MIKARDIO MK Kardioamb	PULDERM	OK	an			2003278723			
↳	2008197548	stationärer Fall 21.04.2006		CKOMG CK OMG									
↳	Diagnosen Desam (5)												
↳	Leistungen (KAL, RAD, Therap)												
↳	1-feldokumenbilddiag.	25.04.2006	00:49	MIKARDIO MK Kardioamb	FITTHIE	OK				2009107546			
↳	EKG (12 Ableitungen)	29.04.2006	08:49	MIKARDIO MK Kardioamb	FITTHIE	OK				2009107546			
↳	RPR-Intervall-Untersuchung	29.04.2006	08:49	MIKARDIO MK Kardioamb	FITTHIE	OK				2009107546			
↳	Schrittmacherkontrolle	29.04.2006	08:49	MIKARDIO MK Kardioamb	FITTHIE	OK				2009107546			
↳	EKG (12 Ableitungen)	23.04.2006	10:43	MIKARDIO MK Kardioamb	KOBENOR	OK				2009107546			
↳	RPR-Intervall-Untersuchung	23.04.2006	10:43	MIKARDIO MK Kardioamb	KOBENOR	OK				2009107546			
↳	Konsult.	21.04.2006	10:22	KNKONS N K.N.Konsult	LANNMICH	OK				2009107546			




Given the parameters α and β , the joint distribution of a topic mixture θ , a set of N topics z , and a set of N words w is given by:

$$p(\theta, z, w | \alpha, \beta) = p(\theta | \alpha) \prod_{n=1}^N p(z_n | \theta) p(w_n | z_n, \beta)$$

Blei, D. M., Ng, A. Y. & Jordan, M. I. 2003. Latent dirichlet allocation. The Journal of machine Learning research, 3, 993-1022.

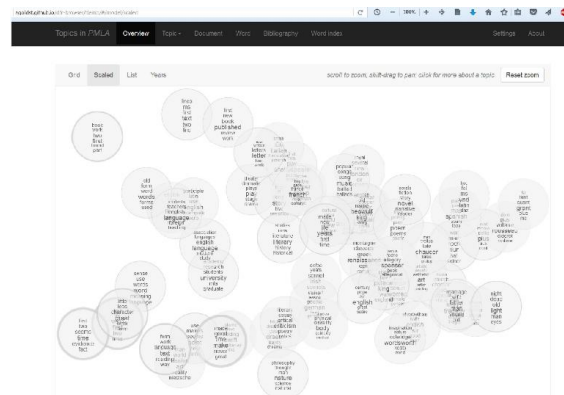


Konietzny, S. G., Dietz, L. & Mchardy, A. C. 2011. Inferring functional modules of protein families with probabilistic topic models. BMC bioinformatics, 12, (1), 1.



$$\cos(\phi) = \frac{q \cdot d_j}{\|q\| \|d_j\|}$$

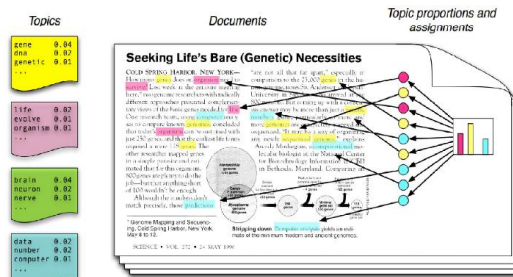
Salton, G., Wong, A. & Yang, C. S. 1975. Vector-Space Model for automatic indexing. Communications of the ACM, 18, (11), 613-620.




<http://agoldst.github.io/dfr-browser/demo/#/model/scaled>

Goal: to get insight in unknown document collections

See a nice demo <http://agoldst.github.io/dfr-browser/demo/#/model/grid>

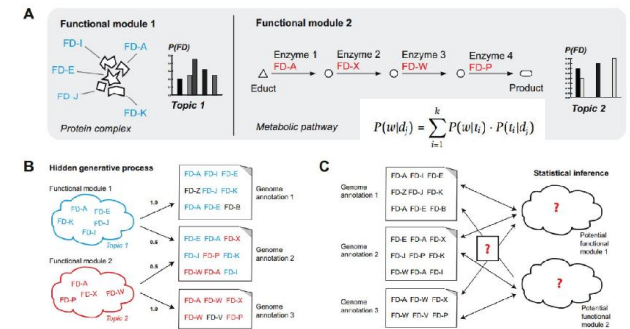


Each doc is a random mix of corpus-wide topics and each word is drawn from one of these topics

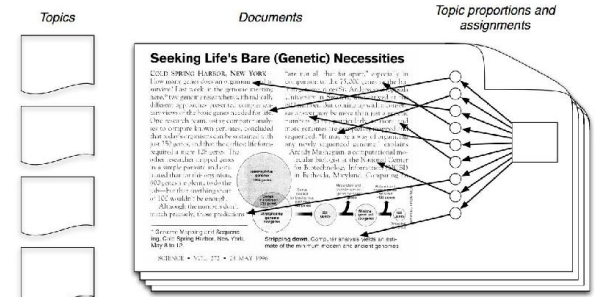


- Documents = categorical distributions over a large space of predefined vocabulary
- Topics = categorical distributions
- Generative model = each document can be seen as a convex combination of the topic distributions

Teh, Y. W., Jordan, M. I., Beal, M. J. & Blei, D. M. 2006. Hierarchical dirichlet processes. Journal of the american statistical association, 101, (476), 1566-1581.



Konietzny, S. G., Dietz, L. & Mchardy, A. C. 2011. Inferring functional modules of protein families with probabilistic topic models. BMC bioinformatics, 12, (1), 1.

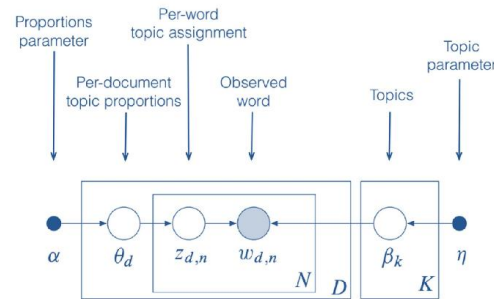


We only observe the docs – the other structure is hidden; then we compute the posterior $p(t, p, \alpha | docs)$


human	evolution	disease	computer
genome	evolutionary	host	models
dna	species	bacteria	information
genetic	organisms	diseases	data
genes	life	resistance	computers
sequence	origin	bacterial	system
gene	biology	new	network
molecular	groups	strains	systems
sequencing	phylogenetic	control	model
map	living	infectious	parallel
information	diversity	malaria	methods
genetics	group	parasite	networks
mapping	new	parasites	software
project	two	united	new
sequences	common	tuberculosis	simulations

Columns sort
probability
word given to

D. Blei




- Encodes assumptions on data with a factorization of the joint
- Connects assumptions to algorithms for computing with data
- Defines the posterior (through the joint)

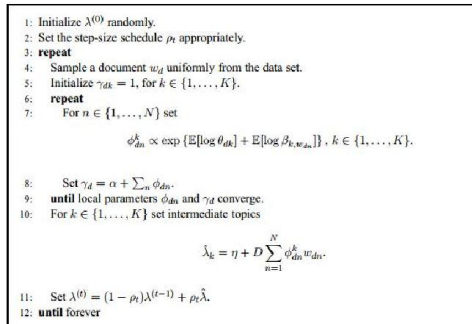


$$p(\beta, \theta, z | \mathbf{w}) = \frac{p(\beta, \theta, z, \mathbf{w})}{\int_{\beta} \int_{\theta} \sum_z p(\beta, \theta, z, \mathbf{w})}$$


We can't compute the denominator, the marginal $p(\mathbf{w})$, therefore we use **approximate inference**;
However, this do not scale well ...



Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. 2013. Stochastic variational inference. The Journal of Machine Learning Research, 14, (1), 1303-1347.




Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. 2013. Stochastic variational inference. The Journal of Machine Learning Research, 14, (1), 1303-1347.




1. Sample a document
2. Estimate the local variational parameters using the current topics
3. Form intermediate topics from those local parameters
4. Update topics as a weighted average of intermediate and current topics

Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. 2013. Stochastic variational inference. The Journal of Machine Learning Research, 14, (1), 1303-1347.



Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. 2013. Stochastic variational inference. The Journal of Machine Learning Research, 14, (1), 1303-1347.



Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. 2013. Stochastic variational inference. The Journal of Machine Learning Research, 14, (1), 1303-1347.

- Flexible and expressive components for building models
- Scalable and generic inference algorithms
- Easy to use software to stretch probabilistic modeling into the health domain
- Topic models are only one approach towards detection of topics in text collections
- More general: Identify re-occurring patterns in data collections generally ...
- Much open work for you in the future ☺

- Particular topic models
 - Stanford topic model toolbox <http://nlp.stanford.edu/software/tmt>
 - Topic modeling at Princeton <http://www.cs.princeton.edu/~blei/topicmodeling.html>
 - MALLET (Java) <http://mallet.cs.umass.edu>
 - Network topic models: Bayes-stack <https://github.com/bgamari/bayes-stack>
 - Gensim (Python) <http://radimrehurek.com/gensim/>
 - R package for Topic models. <http://epub.wu.ac.at/3987/>
- Frameworks for generative models
 - Variational inference: Infer.net <http://research.microsoft.com/infernet/>
 - Gibbs sampling: OpenBUGS <http://openbugs.net/>






Image credit to Eric Xing, Carnegie Mellon University, Pittsburgh




Holzinger, A., Ofner, B., Dehmer, M.: Multi-touch Graph-Based Interaction for Knowledge Discovery on Mobile Devices: State-of-the-Art and Future Challenges. In: LNCS 8401, pp. 241–254, (2014)

Dehmer, M., Emmert-Streib, F., Pickl, S. & Holzinger, A. (eds.) 2016. Big Data of Complex Networks, Boca Raton, London, New York: CRC Press Taylor & Francis Group.



04 Knowledge Representation in Network Medicine



Nature Reviews | Molecular Cell Biology

Image description find here: http://www.nature.com/nrm/journal/v6/n2/fig_tab/nrm1570_F1.html

- Nodes
 - drugs
 - clinical guidelines
 - patient conditions (indication, contraindication)
 - pharmacological groups
 - tables and calculations of medical scores
 - algorithms and other medical documents
- Edges: 3 crucial types of relations inducing medical relevance between two active substances
 - pharmacological groups
 - indications
 - contra-indications




Image credit to Anna Goldenberg, Toronto






Image credit to Anna Goldenberg, Toronto






Holzinger Group, hci-kde Learning Health 06

- Holzinger, A., Ofner, B., & Dehmer, M. (2014). Multi-touch Graph-Based Interaction for Knowledge Discovery on Mobile Devices: State-of-the-Art and Future Challenges. LNCS 8401 (pp. 241–254). Berlin, Heidelberg: Springer.





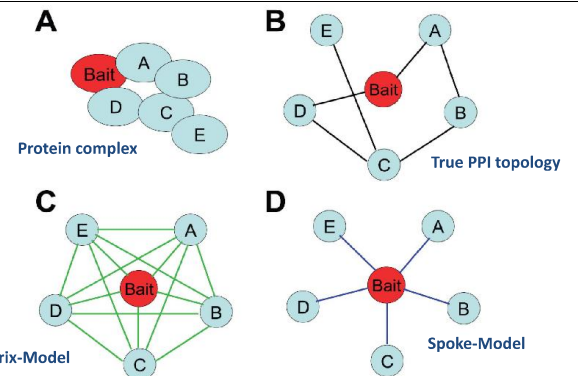

Transcriptional regulatory network with two components:
TF = transcription factor
TG = target genes
(TF regulates the transcription of TG)

Protein-Protein interaction network


Metabolic network
(constructed considering the reactants, chemical reactions and enzymes)

Costa, L. F., Rodrigues, F. A. & Cristino, A. S. (2008) Complex networks: the key to systems biology. *Genetics and Molecular Biology*, 31, 3, 591-601.

Salgado, H., Santos-Zavaleta, A., Gama-Castro, S., Peralta-Gil, M., Peñalosa-Spinola, M. I., Martínez-Antonio, A., Karp, P. D. & Collado-Vides, J. (2006). The comprehensive updated regulatory network of *Escherichia coli* K-12. *BMC bioinformatics*, 7, (1), 5.



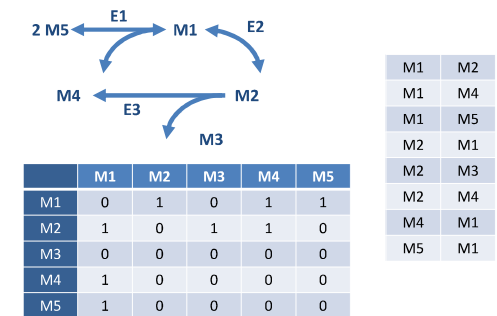
Wang, Z. & Zhang, J. Z. (2007) In search of the biological significance of modular structures in protein networks. *PLoS Computational Biology*, 3, 6, 1011-1021.



Boyen, P., Van Dyck, D., Neven, F., van Ham, R. C. H. J. & van Dijk, A. (2011) SLIDER: A Generic Metaheuristic for the Discovery of Correlated Motifs in Protein-Protein Interaction Networks. *Computational Biology and Bioinformatics*, IEEE/ACM Transactions on, 8, 5, 1344-1357.

Input: PPI-network $G = (V, E, \lambda)$, $\ell, d \in \mathbb{N}$, $d < \ell$
Output: $\{X^*, Y^*\}$ best correlated motif pair found in G
1: $\{X^*, Y^*\} \leftarrow \text{randomMotifPair}()$
2: $\text{maxsup} \leftarrow f(\{X^*, Y^*\}, G)$
3: $\text{sup} \leftarrow -\infty$
4: **while** $\text{maxsup} > \text{sup}$ **do**
5: $\{X, Y\} \leftarrow \{X^*, Y^*\}$
6: $\text{sup} \leftarrow \text{maxsup}$
7: **for all** $\{X', Y'\} \in N(\{X, Y\})$ **do**
8: **if** $f(\{X', Y'\}, G) > \text{maxsup}$ **then**
9: $\{X^*, Y^*\} \leftarrow \{X', Y'\}$
10: $\text{maxsup} \leftarrow f(\{X', Y'\}, G)$

Boyen et al. (2011)



Matrix contains many sparse elements - In this case it is computationally more efficient to represent the graph as an adjacency list



Hodgman, C. T., French, A. & Westhead, D. R. (2010) *Bioinformatics. Second Edition*. New York, Taylor & Francis.

Schmid, A. K., Reiss, D. J., Pan, M., Koide, T. & Baliga, N. S. (2009) A single transcription factor regulates evolutionarily diverse but functionally linked metabolic pathways in response to nutrient availability. *Molecular Systems Biology*, 5, 1-9.


http://www.nature.com/msb/journal/v5/n1/fig_tab/msb200940_F6.html

Electronic patient records remain a unexplored, but potentially rich data source for example to discover correlations between diseases.

Roque, F. S., Jensen, P. B., Schmock, H., Dalgaard, M., Andreatta, M., Hansen, T., Søebye, K., Bredkjær, S., Juul, A., Werge, T., Jensen, L. J. & Brunak, S. (2011) Using Electronic Patient Records to Discover Disease Correlations and Stratify Patient Cohorts. *PLoS Computational Biology*, 7, 8, e1002141.



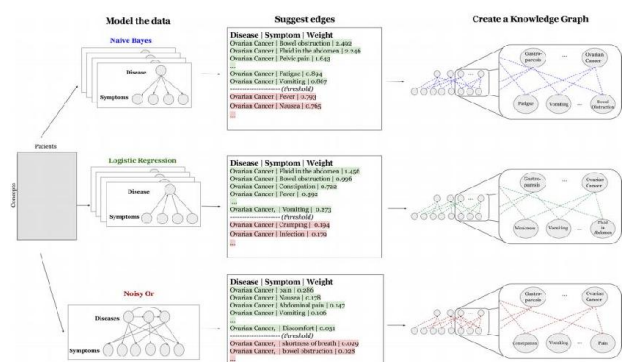
Roque, F.S. et al (2011) Using Electronic Patient Records to Discover Disease Correlations and Stratify Patient Cohorts. *PLoS Comput Biol*, 7, 8, e1002141.



The two main forms of lymphoma are Hodgkin lymphoma and non-Hodgkin lymphoma (NHL). Lymphoma occurs when cells of the immune system called lymphocytes, a type of white blood cell, grow and multiply uncontrollably. Cancerous lymphocytes can travel to many parts of the body, including the lymph nodes, spleen, bone marrow, blood, or other organs, and form a mass called a tumor. The body has two main types of lymphocytes that can develop into lymphomas: B-lymphocytes (B-cells) and T-lymphocytes (T-cells).

www.lymphoma.org

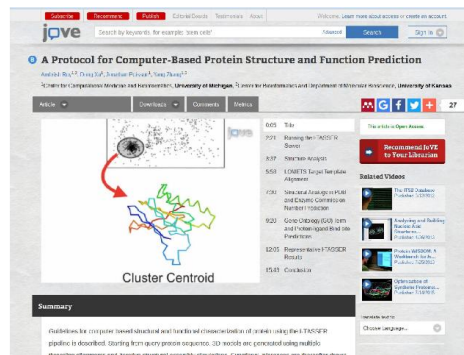

http://imagebank.hematology.org/



- Homology modeling is a knowledge-based prediction of protein structures.
- In homology modeling a protein sequence with an unknown structure (the target) is aligned with one or more protein sequences with known structures (the templates).
- The method is based on the principle that homologue proteins have similar structures.
- Homology modeling will be extremely important to personalized and molecular medicine in the future.**

- Discover unexplored interactions in PPI-networks and gene regulatory networks
- Learn the structure
- Reconstruct the structure

Dittrich, M. T., Klau, G. W., Rosenwald, A., Dandekar, T. & Müller, T. 2008. Identifying functional modules in protein-protein interaction networks: an integrated exact approach. Bioinformatics, 24, (13), 1223-1231.



05 Machine Learning on Graphs Relevant for Health Informatics

SCIENTIFIC REPORTS

OPEN

Learning a Health Knowledge Graph from Electronic Medical Records

Maya Rotmensch¹, Yoni Halpern², Abdulhakim Tilmat³, Steven Horig^{1,4} & David Sontag^{5,6}

Received: 3 March 2017
Accepted: 1 June 2017
Published online: 20 July 2017

Demand for clinical decision support systems in medicine and self-diagnostic symptom checkers has substantially increased in recent years. Existing platforms rely on knowledge bases manually compiled through a labor-intensive process or automatically derived using simple pairwise statistics. This study explored an automated process to learn high quality knowledge bases linking diseases and symptoms directly from electronic medical records. Medical concepts were extracted from 273,174 de-identified patient records and maximum likelihood estimation of three probabilistic models was used to automatically construct knowledge graphs. Logistic regression, naive Bayes classifier and a Bayesian network using noisy OR gates. A graph of disease-symptom relationships was elicited from the learned parameters and the constructed knowledge graphs were evaluated and validated, with permission, against Google's manually-constructed knowledge graph and against expert physician opinions. Our study shows that direct and automated construction of high quality health knowledge graphs from medical records using rudimentary concept extraction is feasible. The noisy OR model produces a high quality knowledge graph reaching precision of 0.85 for a recall of 0.6 in the clinical evaluation. Noisy OR significantly outperforms all tested models across evaluation frameworks ($p < 0.001$).

Nodes: proteins
Links: physical interactions (binding)

Puzzling pattern:

Hubs tend to link to small degree nodes.


Why is this puzzling?

In a random network, the probability that a node with degree k links to a node with degree k' is:

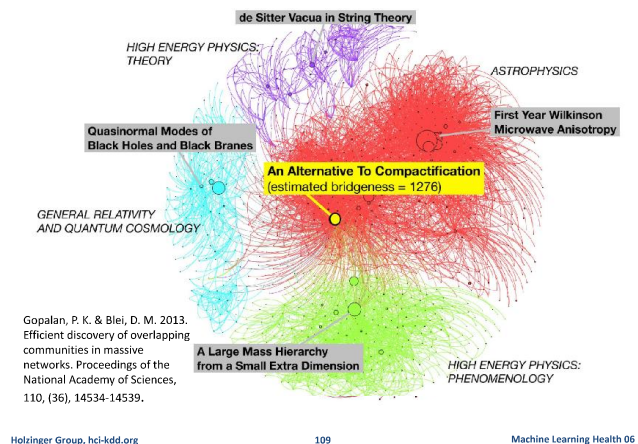
$$p_{kk'} = \frac{kk'}{2L}$$

$$k=50, k'=13, N=1,458, L=1746$$

$$p_{50,13} = 0.15 \quad p_{13,50} = 0.0004$$




Jeong, H., Mason, S. P., Barabasi, A. L. & Oltvai, Z. N. 2001. Lethality and centrality in protein networks. Nature, 411, (6833), 41-42.



- Similar Property Principle: Molecules having similar structures should have similar activities.
- Structure-based representations: Compare molecules by comparing substructures, e.g.
 - Sets as vectors: Measure similarity by the cosine distance
 - Sets as sets: Measure similarity by the Jaccard distance
 - Sets as points: Measure similarity by Euclidean distance
- Problems: Dimensionality, Non-Euclidean cases

- Describe the clinical decision making process!
- Which type of graph is particularly useful for inference and learning?
- What is the key challenge in the application of graphical models for health informatics?
- What was Judea Pearl (1988) discussing in his paper, for which he received the Turing award?
- What main difficulties arise during breast cancer prognosis?
- What can be done to increase the robustness of prognostic cancer tests?
- Inference in Bayes Nets is NP-complete, but there are certain cases where it is tractable, which ones?

- A) Discovery of unexplored interactions
- B) Learning and Predicting the structure
- C) Reconstructing the structure
- Which joint probability distributions does a graphical model represent?
- How can we learn the parameters and structure of a graphical model?

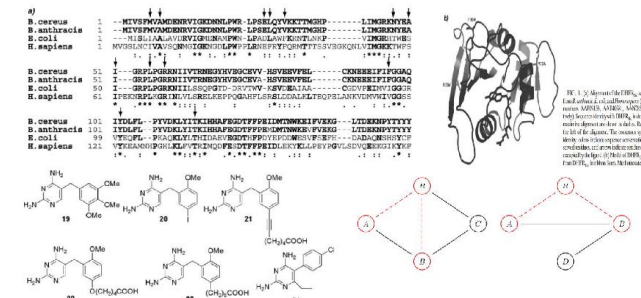


The chemical space

- 10^{10} possible small organic molecules
- 10^{22} stars in the observable universe



- Why do we want to apply ML to graphs?
- Describe typical ML tasks on the example of blood cancer cells!
- If you have a set of points – which similarity measures are useful?
- Why is graph comparison in the medical domain useful?
- Why is the Gromov-Hausdorff distance useful?
- What is the central goal of a generative probabilistic model?
- Describe the LDA-model and its application for topic modelling!



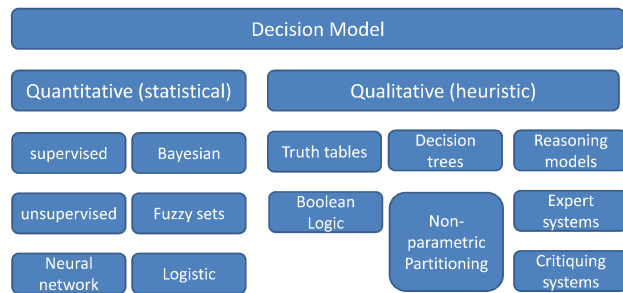
How similar are two graphs? How similar is their structure? How similar are their node and edge labels?

Questions

- Briefly describe the stochastic variational inference algorithms!
- What is the principle of a bandit?
- How does a multi-armed bandit (MAB) work?
- In which ways can a MAB represent knowledge?
- What is the main problem of a clinical trial – and maybe the main problem in clinical medicine?
- Why are rare diseases both important and relevant? Describe an example disease!
- What is the big problem in clinical trials for rare diseases?
- What did Richard Bellman (1956) describe with dynamic programming?
- Why are graph bandits a hot topic for ML research?

Thank you!

- 1=this is a factor graph of an undirected graph – we have seen this in protein networks (refer to slide Nr. 70 in lecture 5). Factor graph is bipartite and has two types of nodes: Variables, which can be either evidence variables (when we know its value) or query variables (when the value is unknown and we want to predict the value); and factors, which define the relationship between variables in the graph. Each factor can be connected to many variables and comes with a factor function to define the relationship between these variables. For example, if a factor node is connected to two variables nodes A and B, a possible factor function could be $\text{imply}(A,B)$, meaning that if the random variable A takes value 1, then so must the random variable B. Each factor function has a weight associated with it, which describes how much influence the factor has on its variables in relative terms. For more information please consult: <http://deepdive.stanford.edu/inference>
- 2= this is the decomposition of a tree, rooted at nodes into subtrees
- 3= an example for machine translation, Image credit to Kevin Gimpel, Carnegie Mellon University
- 4= the famous expectation-utility theory according to von Neumann and Morgenstern (1954): a decision-maker faced with risky (probabilistic) outcomes of different choices will behave as if he is maximizing the expected value of some function defined over the potential outcomes at some specified point in the future.
- 5= MYCIN –expert system that used early AI (rule-based) to identify bacteria causing severe infections, such as bacteremia and meningitis, and to recommend antibiotics, with the dosage adjusted for patient's body weight – the name derived from the antibiotics themselves, as many antibiotics have the suffix "-mycin".
- 6= metabolic and physical processes that determine the physiological and biochemical properties of a cell. These networks comprise the chemical reactions of metabolism, the metabolic pathways, as well as the regulatory interactions that guide these reactions.
- 7= With the sequencing of complete genomes, it is now possible to reconstruct the network of biochemical reactions in many organisms, from bacteria to human. Several of these networks are available online, e.g. Kyoto Encyclopedia of Genes and Genomes (KEGG), EcoCyc, BioCyc etc. Metabolic networks are powerful tools for studying and modelling metabolism.



Bemmel, J. H. v. & Musen, M. A. (1997) *Handbook of Medical Informatics*. Heidelberg, Springer.

h_1 = The identity of ORGANISM-1 is streptococcus
 h_2 = PATIENT-1 is febrile
 h_3 = The name of PATIENT-1 is John Jones

$CF[h_1, E] = .8$: There is strongly suggestive evidence (.8) that the identity of ORGANISM-1 is streptococcus
 $CF[h_2, E] = -.3$: There is weakly suggestive evidence (.3) that PATIENT-1 is not febrile
 $CF[h_3, E] = +1$: It is definite (1) that the name of PATIENT-1 is John Jones

Shortliffe, E. H. & Buchanan, B. G. (1984) *Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project*. Addison-Wesley.

Appendix

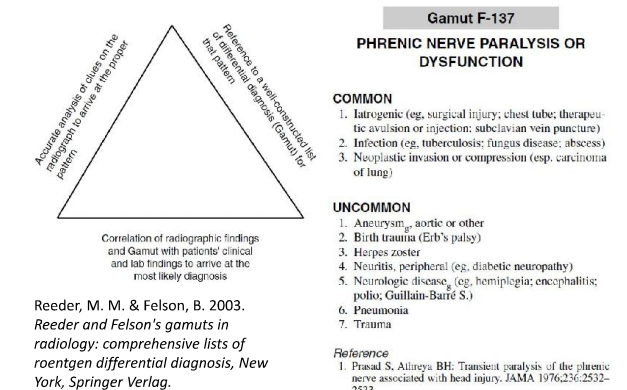
- The information available to humans is often imperfect – imprecise - uncertain.
- This is especially in the medical domain the case.
- An **human agent** can cope with deficiencies.
- Classical logic permits only **exact reasoning**:
- IF A is true THEN A is non-false and
IF B is false THEN B is non-true
- Most real-world problems do not provide this exact information, mostly it is inexact, incomplete, uncertain and/or **un-measurable!**

<https://www.youtube.com/watch?v=IVGWMOCKNWA> ("real nurse triage")



1) Reasoning under Uncertainty

- MYCIN is a rule-based Expert System, which is used for therapy planning for patients with bacterial infections
- Goal oriented strategy ("Rückwärtsverkettung")
- To every rule and every entry a certainty factor (CF) is assigned, which is between 0 and 1
- Two measures are derived:
 - MB: measure of belief
 - MD: measure of disbelief
- Certainty factor – CF of an element is calculated by:
 $CF[h] = MB[h] - MD[h]$
 - CF is positive, if more evidence is given for a hypothesis, otherwise CF is negative
 - $CF[h] = +1 \rightarrow h$ is 100 % true
 - $CF[h] = -1 \rightarrow h$ is 100% false



Reeder, M. M. & Felson, B. 2003. *Reeder and Felson's gamuts in radiology: comprehensive lists of roentgen differential diagnosis*. New York, Springer Verlag.

Example - Gamuts in Radiology

Reeder and Felson's GAMUTS IN RADIOLOGY

GAMUT G-25
EROSIVE GASTRITIS

COMMON

1. Acute gastritis (eg. alcohol abuse)
2. Crohn's disease
3. Drugs (eg. aspirin, NSAID, steroids)
4. Helicobacter pylori infection
5. Idiopathic
6. [Normal antral gastritis]
7. Peptic ulcer, hyperacidity

UNCOMMON

1. Corrosive gastritis
2. Cytosporidium gastritis
3. [Lymphoma]
4. Cryptocystic infection (eg. candidiasis [moniliasis], herpes simplex, cytomegalovirus)
5. Postoperative gastritis
6. Radiation therapy
7. Zollinger-Ellison S. multiple endocrine neoplasia (MEN) S.

* Superficial erosions or aphthoid ulcerations seen especially with cobble contrast technique.

[] This condition does not actually cause the gamut(s) imaging finding, but can produce imaging changes that simulate it.

Reeder, M. M. & Felson, B. (2003) *Reeder and Felson's gamuts in radiology: comprehensive lists of roentgen differential diagnosis*. New York, Springer Verlag.

<http://rfs.acr.org/gamuts/data/G-25.htm>

Future Outlook

The future is in integrative ML, i.e. combining relational databases, ontologies and logic with probabilistic reasoning models and statistical learning – and algorithms that have good scalability

W Smokes(x) ∧ Friends(x,y) ⇒ Smokes(y)

Big data
Big models

Learns a model over 900,030,000 random variables

Van Den Broeck, G., Taghipour, N., Meert, W., Davis, J. & De Raedt, L. Lifted probabilistic inference by first-order knowledge compilation. Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Three, 2011. AAAI Press, 2178-2185.

Probabilistic Program	Graphical Model
Variables	Variable nodes
Functions/operators	Factor nodes/edges
Fixed size loops/arrays	Plates
If statements	Gates (Minka & Winn)
Variable sized loops, Complex indexing, jagged arrays, mutation, recursion, objects/properties...	No common equivalent

Reasoning under uncertainty

- Take patient information, e.g., observations, symptoms, test results, -omics data, etc. etc.
- Reach conclusions, and predict into the future, e.g. how likely will the patient be re-admitted?
- Prior = belief before making a particular observation
- Posterior = belief after making the observation and is the prior for the next observation – intrinsically incremental

$$p(x_i|y_j) = \frac{p(y_j|x_i)p(x_i)}{\sum p(x_i, y_j)p(x_i)}$$

Quiz

- Diagram of a simple network with nodes 1, 2, 3, 4, 5, 6.
- Diagram of a more complex network with nodes 1, 2, 3, 4, 5, 6, 7, 8, 9, 10.
- Diagram of a network with nodes 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20.
- Equation: $E(U | d) = \sum_{x_1, \dots, x_n} P(x_1, \dots, x_n | d) U(x_1, \dots, x_n, d)$
- Text: h_1 = the identity of ORGANISM-1 is streptococcus, h_2 = PATIENT-1 is febrile, h_3 = The name of PATIENT-1 is John Jones.
- Diagram of a network with nodes A, B, C, D, E, F, G, H, I, J, K, L, M, N, O, P, Q, R, S, T, U, V, W, X, Y, Z.

Medical Example

Sequence Outcome

GGTGGGATTCGCTGCTTGGAGGAGGCT	Y
GGGCTTGGGAGGCTTTCAGTTCGAGCT	X
CTAGAGTGGGAGGAGGAGGCTTTCAGTTCAGCT	Y
CTAGAGGAGGCTTTCAGTTCAGTTCAGCT	Y
TTAGAGGAGGCTTTCAGTTCAGTTCAGCT	X
...	...

1. Simple example: Nucleotide "A" may follow nucleotide "T" in the sequences more frequently for outcome X than for outcome Y.

$$P(A|T, X) > P(A|T, Y)$$

2. $P(\theta|D) = \frac{P(D|\theta) \cdot P(\theta)}{P(D)}$

3. $P(\theta|D) = \frac{P(D|\theta) \cdot P(\theta)}{P(D)}$

4. $P(\theta|D) = \frac{P(D|\theta) \cdot P(\theta)}{P(D)}$

5. $P(\theta|D) = \frac{P(D|\theta) \cdot P(\theta)}{P(D)}$

6. $P(\theta|D) = \frac{P(D|\theta) \cdot P(\theta)}{P(D)}$

Remember: 2 types of decisions (Diagnosis vs. Therapy)

- Type 1 Decisions:** related to the **diagnosis**, i.e. computers are used to assist in diagnosing a disease on the basis of the individual patient data. Questions include:
 - What is the probability that this patient has a myocardial infarction on the basis of given data (patient history, ECG, ...)?
 - What is the probability that this patient has acute appendices, given the signs and symptoms concerning abdominal pain?
- Type 2 Decisions:** related to **therapy**, i.e. computers are used to select the best therapy on the basis of clinical evidence, e.g.:
 - What is the best therapy for patients of age x and risks y , if an obstruction of more than z % is seen in the left coronary artery?
 - What amount of insulin should be prescribed for a patient during the next 5 days, given the blood sugar levels and the amount of insulin taken during the recent weeks?

Bemmel, J. H. V. & Musen, M. A. 1997. *Handbook of Medical Informatics*, Heidelberg, Springer.

Probabilistic-programming.org

- C → Probabilistic-C
- Scala → Figaro
- Scheme → Church
- Excel → Tabular
- Prolog → Problog
- Javascript → webPP
- Python → PyMC

PyMC Pythonic Markov chain Monte Carlo


05 Digression: What is similarity?



Image credit to Eamonn Keogh (2008)

Bronstein, A. M., Bronstein, M. M. & Kimmel, R. 2008. *Numerical geometry of non-rigid shapes*, New York, Springer.

<http://www.inf.usi.ch/bronstein/>



Correspondence quality = structure similarity (distortion)

Minimum possible correspondence distortion




Counts the number of "i-dimensional holes"
bi is the "i-th Betti number"




Emmy Noether (1882-1935)



$b_1=1$
 $b_2=0$




$b_1=0$
 $b_2=1$





$b_1=2$
 $b_2=1$

Betti numbers are computed as dimensions of Boolean vector spaces (E. Noether)


Zomorodian, A. & Carlsson, G. 2005. Computing Persistent Homology. *Discrete & Computational Geometry*, 33, (2), 249-274.



Bronstein, A. M., Bronstein, M. M. & Kimmel, R. 2008. *Numerical geometry of non-rigid shapes*, New York, Springer.



- Statement of Vin de Silva (2003), Pomona College:
- Let M be a topological or metric space, known as the *hidden parameter space*;
- let \mathbb{R}^d be a Euclidean space, the *observation space*,
- and let $f: M \rightarrow \mathbb{R}^d$ be a continuous embedding.
- Furthermore, let $X \subset M$ be a finite set of data points, perhaps the realization of a stochastic process, i.e., a family of random variables $\{X_i, i \in I\}$ defined on a probability space (Ω, \mathcal{F}, P) , and denote $Y = f(X) \subset \mathbb{R}^d$ the images of these points under the mapping f .
- We refer to X as *hidden data*, and Y as the *observed data*.
- M, f and X are unknown, but Y is - so can we identify M ?



Bronstein, A. M., Bronstein, M. M. & Kimmel, R. 2008. *Numerical geometry of non-rigid shapes*, New York, Springer.




Michail Gromov (1943-)

Gromov, M. (1984) Infinite groups as geometric objects.




Felix Hausdorff (1868-1942)



$$d_{GH}(X, Y) = \frac{1}{2} \min_C \max_{(x_i, y_i) \in C} |\delta_X(x_i) - \delta_Y(y_i)|$$

$$\forall x_i \exists y_i \text{ s.t. } (x_i, y_i) \in C \quad \forall y_i \exists x_i \text{ s.t. } (x_i, y_i) \in C$$

Discrete optimization over correspondences is NP hard !



- Mega Problem: To date none of our known methods, algorithms and tools scale to the massive amount and dimensionalities of data we are confronted in practice;
- we need much more research efforts towards making computational topology successful as a general method for data mining and knowledge discovery


Holzinger, A. 2014. On Topological Data Mining. In: *Lecture Notes in Computer Science*, LNCS 8401. Berlin Heidelberg: Springer, pp. 331-356, doi:10.1007/978-3-662-43968-5_19.

06 Review of basic concepts, metrics and measures

Baby Stuff: Computational Graph Representation


Adjacency (ə-ˈjā-sən(t)-sē) Matrix $A = (a_{jk})$

$$a_{jk} = \begin{cases} 1, & \text{if } \{j, k\} \in E \\ 0, & \text{otherwise} \end{cases}$$



$$a_{jk} = \begin{bmatrix} 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$$

Simple graph, symmetric, binary




$$w_{jk} = \begin{bmatrix} 0 & 0 & -3 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 3 & 0 \\ 3 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 1 & -5 \\ 1 & -2 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 5 & 0 \end{bmatrix}$$

Directed and weighted



For more information: Diestel, R. (2010) *Graph Theory, 4th Edition*. Berlin, Heidelberg, Springer.

Some Network Metrics (2/2)

- Centrality (d) = the level of "betweenness- centrality" of a node i ("hub-node in Slide 28);



- Nodal degree (e) = number of links connecting i to its neighbors: $k_i = \sum_j a_{ij}$

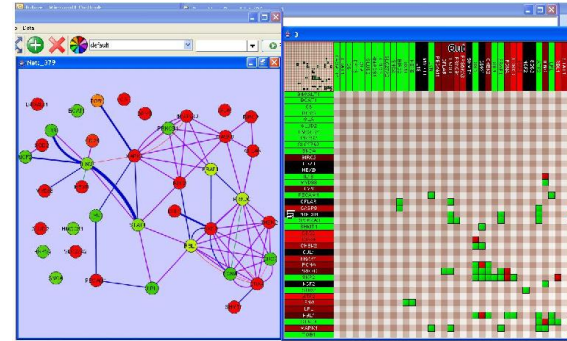


Modularity (f) = describes the possible formation of communities in the network, indicating how strong groups of nodes form relative isolated sub-networks within the full network (refer also to Slide 5-8).

Complex Biological Systems key concepts

- In order to understand complex biological systems, the three following key concepts need to be considered:
- (i) **emergence**, the discovery of links between elements of a system because the study of individual elements such as genes, proteins and metabolites is insufficient to explain the behavior of whole systems;
- (ii) **robustness**, biological systems maintain their main functions even under perturbations imposed by the environment; and
- (iii) **modularity**, vertices sharing similar functions are highly connected.
- Network theory can largely be applied for biomedical informatics, because many tools are already available



Example: Tool for Node-Link Visualization



Jean-Daniel Fekete http://wiki.cytoscape.org/InfoVis_Toolkit

Fekete, J.-D. The infovis toolkit. Information Visualization, INFOVIS 2004, 2004. IEEE, 167-174.


Network Topologies



Van Heuvel & Hulshoff (2010)

Network Basics on the Example of Bioinformatics

$G(V, E)$ Graph
 $V \dots$ vertex
 $E \dots$ edge $\{a, b\}$
 $a, b \in V; a \neq b$



Hodgman, C. T., French, A. & Westhead, D. R. (2010) *Bioinformatics. Second Edition*. New York, Taylor & Francis.

Some Network Metrics (1/2)


Order = total number of nodes n ; Size = total number of links (a):

$$\sum_i \sum_j a_{ij}$$

Clustering Coefficient (b) = the degree of concentration of the connections of the node's neighbors in a graph and gives a measure of local inhomogeneity of the link density:

$$C_i = \frac{2t_i}{k(k_i - 1)}$$

$$C = \frac{1}{n} \sum_i C_i$$

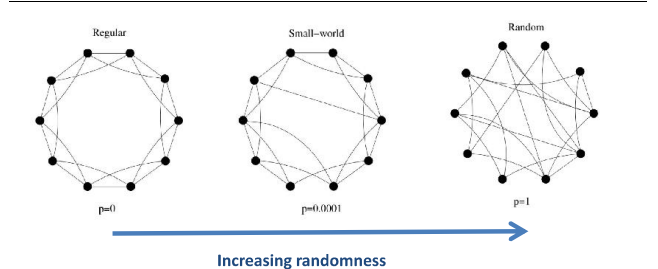


Path length (c) = is the arithmetical mean of all the distances:

$$l = \frac{1}{n(n-1)} \sum_{i \neq j} d_{ij}$$

Costa, L. F., Rodrigues, F. A., Travieso, G. & Boas, P. R. V. (2007) Characterization of complex networks: A survey of measurements. *Advances in Physics*, 56, 1, 167-242.


Small-World Networks



29.000 citations ...

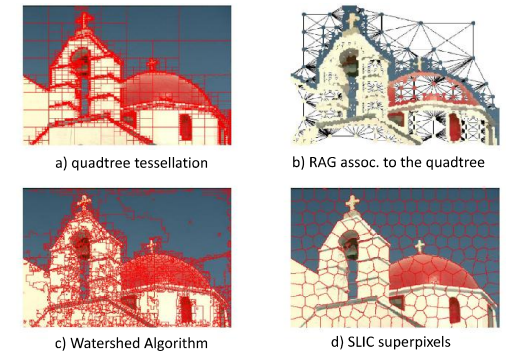
Watts, D. J. & Strogatz, S. (1998) Collective dynamics of small-world networks. *Nature*, 393, 6684, 440-442.

Milgram, S. 1967. The small world problem. *Psychology today*, 2, (1), 60-67.



Lézoray, O. & Grady, L. 2012. Graph theory concepts and definitions used in image processing and analysis. In: Lézoray, O. & Grady, L. (eds.) *Image Processing and Analysing With Graphs: Theory and Practice*. Boca Raton (FL): CRC Press, pp. 1-24.

07 How do you get point cloud data from natural images?






Lézoray, O. & Grady, L. 2012. Graph theory concepts and definitions used in image processing and analysis. In: Lézoray, O. & Grady, L. (eds.) *Image Processing and Analysing With Graphs: Theory and Practice*. Boca Raton (FL): CRC Press, pp. 1-24.

Algorithm 4.2 Watershed transform w.r.t. topographical distance based on image integration via the Dijkstra-Moore shortest paths algorithm.

```

1: procedure ShortestPathWatershed;
2: INPUT: lower complete digital gray scale image  $G = (V, E, im)$  with cost function  $cost$ .
3: OUTPUT: labelled image  $lab$  on  $V$ .
4: #define WSHED 0 (* label of the watershed pixels *)
5: (* Use distance image  $dist$ . On output,  $dist[v] = im[v]$ , for all  $v \in V$ . *)
6:
7: for all  $v \in V$  do (* Initialize *)
8:    $lab[v] \leftarrow 0$ ;  $dist[v] \leftarrow \infty$ 
9: end for
10: for all local minima  $m_i$  do
11:   for all  $v \in m_i$  do
12:      $lab[v] \leftarrow i$ ;  $dist[v] \leftarrow im[v]$  (* initialize distances with values of minima *)
13:   end for
14: end for
15: while  $V \neq \emptyset$  do
16:    $u \leftarrow CostMinDist(V)$  (* find  $u \in V$  with smallest distance value  $dist[u]$  *)
17:    $V \leftarrow V \setminus \{u\}$ 
18:   for all  $v \in V$  with  $(u, v) \in E$  do
19:     if  $dist[u] + cost[u, v] < dist[v]$  then
20:        $dist[v] \leftarrow dist[u] + cost[u, v]$ 
21:        $lab[v] \leftarrow lab[u]$ 
22:     else if  $lab[v] \neq WSHED$  and  $dist[u] + cost[u, v] = dist[v]$  and  $lab[v] \neq lab[u]$  then
23:        $lab[v] = WSHED$ 
24:     end if
25:   end for
26: end while
  
```

Meijster, A. & Roerdink, J. B. A proposal for the implementation of a parallel watershed algorithm. *Computer Analysis of Images and Patterns*, 1995. Springer, 790-795.



Holzinger, A., Malle, B. & Giuliani, N. 2014. On Graph Extraction from Image Data. In: Slezak, D., Peters, J. F., Tan, A.-H. & Schwabe, L. (eds.) *Brain Informatics and Health, BIH 2014, Lecture Notes in Artificial Intelligence, LNAI 8609*. Heidelberg, Berlin: Springer, pp. 552-563.

For Voronoi please refer to: Aurenhammer, F. 1991. Voronoi Diagrams - A Survey of a fundamental geometric data structure. *Computing Surveys*, 23, (3), 345-405.


For Delaunay please refer to: Lee, D.-T. & Schachter, B. J. 1980. Two algorithms for constructing a Delaunay triangulation. *Intl. Journal of Computer & Information Sciences*, 9, (3), 219-242.

- More expressive data structures
- Find novel connections between data objects
- Fit for applying graph based machine learning techniques
- New approaches (Belief Propagation, global understanding from local properties)

Bunke, H.: Graph-based tools for data mining and machine learning. In Perner, P., Rosenfeld, A., eds.: *Machine Learning and Data Mining in Pattern Recognition*, Proceedings. Volume 2734 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag Berlin, (Berlin) 7-19

Holzinger, A., Blanchard, D., Bloice, M., Holzinger, K., Palade, V., Rabadan, R.: Darwin, lamarck, or baldwin: Applying evolutionary algorithms to machine learning techniques. In: *The 2014 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2014)*, IEEE (2014) in print


- Topographic maps => landscapes with height structures
- Segmentation into regions of pixels
- Assuming drops of water raining on the map
- Following paths of descent
- Lakes called catchment basins
- Also possible: Flooding based
- Needs Topographical distance measures (MST)



Vincent, L. & Soille, P. 1991. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE transactions on pattern analysis and machine intelligence*, 13, (6), 583-598.

- 1) Transformation into a topographic map
 - Convert gray values into height information
- 2) Finding local minima
 - Inspecting small regions in sequence
- 3) Finding catchment basins
 - Algorithm simulating flooding
 - Graph algorithms such as Minimum Spanning Trees
- 4) Erecting watersheds
 - Artificial divide between catchment basins
 - Final segmentation lines

- We want to find “interesting” novel patterns (rules, anomalies, outliers, similarities, ...)
- Problem #1: How to get a graph?
- Problem #2: How do graphs evolve?
- Problem #3: What tools to apply?
- Problem #4: Scalability to TB, PB, EB ...
- Success is in repeatability and scalability**



Connects each pixel to the lowest neighbor pixel, all pixel connected to same lowest neighbor pixel form a segment

- Study of complex networks started in the 1990s with the insight that real networks contain properties not present in random (Erdős-Renyi) networks.
- Meanwhile networks and network-based approaches form an integral part of many studies throughout the sciences.
- Graph-Theory provides powerful tools to organize data structurally and in combination with statistical and machine learning methods allows a meaningful analysis of underlying processes.
- For instance, a mapping of causal disease genes and disorders as made available by the OMIM database provided novel insights into disease patterns, as recently demonstrated by investigating the diseasesome (<http://diseasome.eu>).

- Region Merging
 - Based on Kruskals MST algorithm
 - Takes input image as natural graph with vertices := pixels and edges := pixel neighborhoods
 - Visits edges in ascending order of weight and merges regions if they satisfy a certain criterion
 - Flexible as merging criterion can be adapted as desired (for amount, size, or shape of resulting regions)

Felzenszwalb, P.F., Huttenlocher, D.P.: Efficient graph-based image segmentation. International Journal of Computer Vision 59 (2004) 167–181

