ngoodman@ stanford.edu





Church and The probabilistic language of thought

> Noah D. Goodman Stanford University

IPAM graduate summer school July 9, 2011

Motivation

The mind is an information processing system.

Motivation

The mind is an information processing system.

But what kind of information processor?

Thought is useful in an uncertain world

Thought is useful in an uncertain world



Thought is useful in an uncertain world

Why did he yell at me?



Thought is useful in an uncertain world

Why did he yell at me?



He wanted to hurt me. He thought I was a telemarketer.



He wanted to hurt me. He thought I was a telemarketer.

Thought is useful in an uncertain world world Why did he yell at me?

He wanted to hurt me. He thought I was a telemarketer.

> Probabilistic inference

Thought is useful in an uncertain world



Probabilistic inference

Thought is useful in an uncertain world



Thought is productive: "the infinite use of finite means"

Probabilistic inference

Thought is useful in an uncertain world



Probabilistic inference



Thought is useful in an uncertain world



Probabilistic inference





Thought is useful in an uncertain world



Probabilistic inference



Thought is useful in an uncertain world



Probabilistic inference



Thought is useful in an uncertain world



Probabilistic inference

Thought is productive: "the infinite use of finite means"



Thought is useful in an uncertain world



Probabilistic inference

Thought is productive: "the infinite use of finite means"



The Language of Thought JERRY A. FODOR

The Language and Thought Series Armid & Kata D. Toward Languation Gauge 8, Millio Million 2017/045

 $\forall x \ King(x) \implies Man(x)$ $\forall y \ Man(y) \iff \neg Woman(y)$

Thought is useful in an uncertain world



Thought is productive: "the infinite use of finite means"



The Language of Thought JERRYA, FODOR

The Language and Thought Series Anniel 3 Acto D. Toward Languadors Courge 5 Millio metarge 2577048

 $\forall x \ King(x) \implies Man(x)$ $\forall y \ Man(y) \iff \neg Woman(y)$

Probabilistic inference

Generative models

Probabilistic language of thought hypothesis

Thought is useful in an uncertain world







Probabilistic inference

Generative models

Probabilistic LoT

- The **probabilistic** language of thought hypothesis:
 - Mental representations are compositional,
 - Their meaning is probabilistic,
 - They encode generative knowledge,
- Hence, they support thinking and learning by probabilistic inference.

Probabilistic LoT

- The **probabilistic** language of thought hypothesis:
 - Mental representations are compositional,
 - Their meaning is probabilistic,
 - They encode generative knowledge,
- Hence, they support thinking and learning by probabilistic inference.

Can this hypothesis be formalized?

 Mental model of the causal process that gives rise to observations.

- Mental model of the causal process that gives rise to observations.
- E.g. Bayes nets.



- Mental model of the causal process that gives rise to observations.
- E.g. Bayes nets.
 - But what's "hidden inside" the nodes and arrows?

Flu TB cough				
cough	ТВ	no TB		
flu	0.8	0.7		
no flu	0.7	0.1		

- Mental model of the causal process that gives rise to observations.
- E.g. Bayes nets.
 - But what's "hidden inside" the nodes and arrows?





- Mental model of the causal process that gives rise to observations.
- E.g. Bayes nets.
 - But what's "hidden inside" the nodes and arrows?





- Mental model of the causal process that gives rise to observations.
- E.g. Bayes nets.
 - But what's "hidden inside" the nodes and arrows?





write talk	want fame	want sleep	
going to IPAM	0.7	0.05	
it's raining	0.002	0.004	
:			

λ calculus

• Notation:

- Function have parentheses on the wrong side:
- Operators always go at the beginning:
- λ makes functions, define binds values to symbols:

$$(define double(\lambda (x) (+ x x))) \qquad (double 3) => 6$$

(define repeat (λ (f) (λ (x) (f (f x)))))

((repeat double) 3) => 12

(define 2nd-derivative (repeat derivative))



Church



Conditioning (inference):





Formalizing the PLoT: Mental representations (concepts) are functions in a **stochastic** lambda calculus (e.g. Church).

- Formalizing the PLoT: Mental representations (concepts) are functions in a **stochastic** lambda calculus (e.g. Church).
 - ...or pieces of probabilistic programs, anyhow.

- Formalizing the PLoT: Mental representations (concepts) are functions in a stochastic lambda calculus (e.g. Church).
 - ...or pieces of probabilistic programs, anyhow.
 - Separate the process of inference from representations and the inferences they license (Cf. Marr's levels).

Outline (next 3 lectures)

- Church: language design considerations.
- Church: inference techniques.
- Social cognition.
- Natural language pragmatics (etc).
- Concept learning as program induction.

Outline (next 3 lectures)

- Church: language design considerations.
- Church: inference techniques.
- Social cognition.
- Natural language pragmatics (etc).
- Concept learning as program induction.
Church language design

- Start with (pure subset of) LISP.
- Add random primitives.
- Add mem.
- Add query.

Why functions?

- Composition of probabilistic functions represents *directed* generative models.
 - Captures the intuition that much human knowledge is about causal process.
 - For example, explaining away: conditioning an undirected model can introduce new independence but not new dependence.
 - Hierarchical models, etc.

- Complex conditions in query represent (any) undirected constraints.
- For example an Ising model:

- Complex conditions in query represent (any) undirected constraints.
- For example an Ising model:

```
(query

(define a (flip))

(define b (flip))

(define c (flip))

(list a b c)

(and (flip (if (equal? a b)) 1.0 0.3) e^{\ln(0.3)\delta_{a=b}}

(flip (if (equal? b c)) 1.0 0.3)))
```

- Complex conditions in query represent (any) undirected constraints.
- For example Markov Logic:

```
(query
 (define world (repeat N flip))
 world
 (and
  (log-flip (if (clause1 world) 0.0 weight1)))
  (log-flip (if (clause2 world) 0.0 weight2)))
 ...))
```

- Complex conditions in query represent (any) undirected constraints.
- So, Church can represent directed, undirected, and hybrid models.
 - Complex conditions can be built out of directed pieces. (This seems important for natural language. Cf. Montegue.)

Why functional?

• First class abstraction.

- The pieces needed for model specification are representable in the modeling language.
- Abstraction permits domain specific languages
 -- changes the effective language of thought.
- Higher-order functions are useful. (E.g. the higher-order distribution DPmem.)
- Recursion.
- Models are parsimonious.

Why LISP?

- Small core language.
 - Simple semantics.
 - Very little syntax.
 - Simple type system.
- Code as data (meta-circular, homoiconic).
- Well-worked out concepts for computation (see SICP).

Random primitives

- Just the fair-coin flip primitive is enough ($\psi\lambda$ -calculus),
- But it is convenient (and well-formed) to add other random primitives.
- An ERP is a primitive (with arguments) that can sample values and score.
 - For each set of arguments, multiple samples from an ERP must be independent identically distributed (iid)
 - Examples: flip, sample-discrete, gaussian.

Exchangeability

- If all the ERPs return iid samples, then is any Church thunk (procedure without arguments) an iid distribution?
- No. Consider:
- But they are exchangeable.

(define weight (beta 1 1))
(define (mycoin) (flip weight))
(repeat 10 mycoin)

=> (t t t t t t t t t t)

 Theorem (Freer & Roy, 2010): Any church thunk is exchangeable and any exchangeable distribution may be represented by a Church thunk.

Random primitives

- Relax the interface for random primitives: an XRP can sample and score. Multiple calls exchangeable.
- Examples:
 - Uniform draw without replacement.
 - Gensym.
 - Beta-binomial.
 - CRP.

Memoization

 $(= (flip) (flip)) \longrightarrow True with probability 0.5$

 $(define mem-flip (mem flip)) \longrightarrow True with probability 1.0$ (= (mem-flip) (mem-flip))

- The mem primitive memoizes a procedure.
 - This changes the semantics (unlike in deterministic languages).
 - Interacts with symbols and gensym, to enable "BLOG style" style world models.

```
(define people (repeat (poisson 1.0) gensym))
(define eye-color
   (mem (\lambda (person) (uniform-draw '(blue brown)))))
(define blue-eyed-people
 (filter (\lambda (person) (equal? 'blue (eye-color person)))
         people))
```

Stochastic memoization

- DPmem: stochastically reuse return values from functions.
 - A higher-order distribution: it takes a stochastic function (of any type) and samples a new function (of same type) that concentrates the original.
 - Adapts a generative process to balance reuse with re-computation (see Johnson, O'Donnell).
 - Implemented in Church via stick-breaking or using an XRP.

(define mem-flip (DPmem 1.0 flip))
(= (mem-flip) (mem-flip))



DP mixture model

```
(define draw-class
  (DPmem 1.0 gensym))
(define class
  (mem (lambda (obj) (draw-class))))
(define class-weight
  (mem (lambda (obj-class feature)
                          (beta 1.0 1.0))))
(define observe-feature
  (lambda (obj feature)
                          (flip (class-weight (class obj) feature)))))
```

DP mixture model

(define	draw-class	A pool of cluster symbols
(DPme	em 1.0 gensym))	with DP prior.
(define	class	
(mem	(lambda (obj) (draw-cl	ass))))
(define	class-weight	object and
(mem	(lambda (obj-class fea	ature) cluster parameters
	(beta 1.0 1.0)	())) are persistent.
(define observe-feature		
(lamk	oda (obj feature)	
(f]	Lip (class-weight (clas	ss obj) feature)))))

Infinite relational model

The IRM model from: Kemp, et al, 2006

Infinite relational model

The IRM model from: Kemp, et al, 2006



Nested CRP

(category-hierarchy 3) x4...

Nested CRP

(category-hierarchy 3) X4...

```
=> ((g0 g1 g2 top)
    (g3 g1 g2 top)
    (g4 g1 g2 top)
    (g5 g6 g7 top))
```

HDP-HMM

first-class query

- How should we specify observations in probabilistic programs?
 - Separate model from data / data variables?
 - 'Observe' statements within model?
 - Better: make inference an ordinary function.
 - Can define inference *within* the language (e.g. by rejection).
 - Gives proper scoping, complex conditions, nested query ('inference about inference').

Nested query

Alice and Bob arrange to meet at the bar, each later realizes they didn't fix which bar and must guess where to meet. (A coordination game, see Schelling, 1960)

Recursive social inference:

```
(define (sample-location)
  (if (flip .55) 'popular-bar 'unpopular-bar))
(define (bob depth)
  (query
    (define bob-location (sample-location))
   bob-location
    (equal? bob-location (alice (- depth 1))))
(define (alice depth)
  (query
    (define alice-location (sample-location))
   alice-location
    (or (= depth 0))
         (equal? alice-location (bob depth))))
```

Nested query

Alice and Bob arrange to meet at the bar, each later realizes they didn't fix which bar and must guess where to meet. (A coordination game, see Schelling, 1960)

Recursive social inference: (define (sample-location) (if (flip .55) 'popular-bar 'unpopular-bar) Probability (define (bob depth) (query (define bob-location (sample-location)) bob-location (equal? bob-location (alice (- depth 1))) (define (alice depth) (query (define alice-location (sample-location)) alice-location (or (= depth 0))(equal? alice-location (bob depth))))



Nested query

Policy-free MDP planning as a recursively optimal action selection.

```
(define (choose-action state)
  (query
   (define action (action-prior))
   action
   (flip (normalize-reward
           (sample-reward action state)))))
(define (sample-reward action state)
  (let ((next-state (state-transition state action)))
      (+ (reward next-state)
         (if (terminal? next-state)
             (sample-reward
              (choose-action next-state)
              next-state)))))
```

Church vs. ...

- Bugs: Less expressive language for directed models. More stable engine.
- BLOG: Nice constructs for unknown objects. No abstraction.
- IBAL (etc): Similar. No continuous variables, mem, xrps. Obs not query.
- MLN: Undirected, based on FoL. Good implementations.
- Csoft, Factorie, Figaro, Hansei, ProbLog....

Church inference

- Universal representation (PLoT) needs universal inference.
 - The mind can't have a separate algorithm for every task.
- Rejection sampling is sound, but slow.
- Efficient universal inference?

Metropolis-Hastings



- To sample from P(x), repeatedly:
 - propose x' from $Q(x \rightarrow x')$,
 - accept with probability: $\min\left(1, \frac{P(x)Q(x \to x')}{P(x')Q(x' \to x)}\right)$
- This converges in distribution to P(x) (even if P(x) isn't normalized).

- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...

- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...

```
(and
 (if (flip)
  (flip)
  true)
 (flip))
```

- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...



- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...



- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...



- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...



- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...



- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...
 - Named via dynamic calling address:



- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...
 - Named via dynamic calling address:

```
 \begin{split} &\mathcal{A}^{top}[\![E]\!] = ((\texttt{lambda}(\texttt{addr}) \mathcal{A}[\![E]\!]) \ '(\texttt{top})) \\ &\mathcal{A}[\![(\texttt{lambda}(I_{i=1}^n) \ E_{body})]\!] = (\texttt{lambda}(\texttt{addr} \ I_{i=1}^n) \ \mathcal{A}[\![E_{body}]\!]) \\ & \text{where } \mathcal{S} \text{ is a globally unique symbol.} \\ &\mathcal{A}[\![(\texttt{mem} \ E)]\!] = ((\texttt{lambda}(\texttt{maddr} \ f)(\texttt{lambda}(\texttt{addr} \ \texttt{args})(\texttt{apply} \ f(\texttt{cons} \ \texttt{args} \ \texttt{maddr}) \ \texttt{args}))) \ \texttt{addr} \ \mathcal{A}[\![E]\!]) \\ &\mathcal{A}[\![(\texttt{begin} \ E_{i=1}^n)]\!] = (\texttt{begin} \ \mathcal{A}[\![E_i]\!]_{i=1}^n) \\ &\mathcal{A}[\![(\texttt{letrec}((I_i \ E_i)_{i=1}^n) \ E_{body})]\!] = (\texttt{letrec}((I_i \ \mathcal{A}[\![E_i]\!])_{i=1}^n) \ \mathcal{A}[\![E_{body}]\!]) \\ &\mathcal{A}[\![(\texttt{letrec}((I_i \ E_i)_{i=1}^n) \ E_{body})]\!] = (\texttt{letrec}((I_i \ \mathcal{A}[\![E_i]\!])_{i=1}^n) \ \mathcal{A}[\![E_{body}]\!]) \\ &\mathcal{A}[\![(\texttt{define} \ I \ E_i)]\!] = (\texttt{if} \ \mathcal{A}[\![E_i] \ \mathcal{A}[\![E_i]\!] \ \mathcal{A}[\![E_i]\!] \\ &\mathcal{A}[\![(\texttt{define} \ I \ E)]\!] = (\texttt{define} \ I \ \mathcal{A}[\![E]\!]) \\ &\mathcal{A}[\![(\texttt{cop} \ E_{i=1}^n)]\!] = (\mathcal{A}[\![E_{op}]\!] \ (\texttt{cons} \ 'S \ \texttt{addr}) \ \mathcal{A}[\![E_i]]_{i=1}^n) \\ &\mathcal{A}[\![E]\!] = E, \text{ otherwise.} \end{split}
```
MH for Church

- What is a state?
 - The set of random choices made in executing (forward) the query expression.
- How can choices be individuated?
 - Not by evaluation order (can change)...
 - Named via dynamic calling address:



MH for Church

- MH over 'execution traces':
 - propose: change a single random choice,
 - trace update: re-execute to update the state, reusing existing choices (new choices from conditional prior),
 - collect, the (now) unused choices,
 - check condition and compute score,
 - accept/reject.
- Initial state must satisfy condition (rejection, annealing, constraint propagation).

MH for Church

- Very similar to BLOG algorithm.
 - Evaluation automatically instantiates minimal "partial world" (but cf. lazy evaluation).
- "Lightweight" technique, addressing by program transformation, works for almost any programming language -enter Stochastic X.
- Current implementation doesn't scale well (due to software engineering issues).

Goodman, et al. (2008) Wingate, Stuhlmueller, Goodman (2011)

Hamiltonian Monte Carlo

- Add an auxiliary "momentum" variable for each continuous variable.
- Hamiltonian evolution is $\phi(x) = -\ln(P(x))$ reversible and conserves $H(x,p) = \phi(x) + mp^2$ probability (hence satisfies $\frac{dx}{dt} = \frac{\partial H}{\partial p}, \ \frac{dp}{dt} = -\frac{\partial H}{\partial q}$
 - Numeric integration introduces conservation errors, so add an MH accept step.
 - Need the gradient of the score....

Automatic differentiation

- Since programs are compositions of functions, can compute derivatives automatically (remember the chain rule?)
 - A "non-standard interpretation" -implemented by operator overloading.
 - Exact (to machine precision) gradient of score at each point. (No silly errors!)
 - Low overhead (roughly 2x in Bher).
 - Generalizes back-prop.

HMC via AD

- We use AD to provide gradients needed in HMC.
 - HMC can be much more efficient.
 - Sampling an eccentric 3-dim gaussian:



Wingate, Stuhlmueller, Siskind, Goodman (under review)

HMC via AD

- We use AD to provide gradients needed in HMC.
 - Conditioning Perlin noise on symmetry:



Wingate, Stuhlmueller, Siskind, Goodman (under review)

Other techniques

- Sequential Monte Carlo
 - Particle filter with rejuvenation,
 - Any model, and any sequentialization (specified via free variable in query).
- Dynamic programming (UDP/cosh)
 - Inference as marginalization,
 - Builds compact system of polynomial equations...

Other techniques



Many directions open..



Bonus: counterfactuals

- Structured models do more than specify distributions. See Pearl, next.
 - Can build a counterfactual operator in Church from the same "update" operator needed for MH.

```
(counterfactual
((dad-eyes 'brown))
(define dad-eyes (uniform-draw '(blue brown)))
(define mom-eyes (uniform-draw '(blue brown)))
(if (or (eq? dad-eyes 'brown)
       (eq? mom-eyes 'brown)
       (flip 0.1))
       'brown
       'blue))
```

With Tobi Gerstenberg and Andreas Stuhlmueller.

Conclusion

- The PLoT: mental representations as stochastic functions.
- Church: an expressive probabilistic programming language suited for cognitive modeling.
- Next lectures: applications to social cognition, natural language, concept learning.