

Lifted Inference Tutorial Notes

UC Irvine

November 4-5, 2014

Rodrigo de Salvo Braz

Outline

- First-order representation
- Splitting Parfactors
- Multiplying Parfactors
- Counting Elimination
- Partial Inversion

Useful materials

- From a Variable Elimination perspective:
thesis by Nima Taghipour (U Leuven)
Very nicely presented
- From a And/Or search perspective:
Probabilistic Theorem Proving,
by Vibhav Gogate

Representation

- A First-Order Probabilistic Model (FOPM) is a set of parfactors
- A parfactor is a “universally quantified factor”:
for all Lvars : Constraint . phi(Atoms)
- An Atom is a relational random variable:
happy(X), friends(X,Y), sunny(), etc.
- Lvars are “logical variables” (indices), discretely valued from a (large) set (the X and Y above)
- Constraint is a boolean formula on equalities on Lvars (for example: “X != bob and Y != X”)

Representation Examples

- A parfactor is a “universally quantified factor”:
for all LVars : Constraint . phi(Atoms)

Examples:

for all X : X != bob . if happy(X) then 0.7 else 0.3

for all X,Y : true . if friends(X,Y) and smoker(X)
then if smoker(Y) then 0.8 else 0.2
else if smoker(Y) then 0.4 else 0.6

Semantics

- Given a set of parfactors and domains for the logical variables, a FOPM defines a joint distribution identical to the one obtained by grounding, or instantiating, or yet propositionalizing, all parfactors into regular factors

Semantics Examples

- for all $X: X \neq \text{bob}$. if happy(X) then 0.7 else 0.3
which X in {ann, bob, carl, dave, ... }
stands for

if happy(ann) then 0.7 else 0.3

(no instantiation for $X = \text{bob}$)

if happy(carl) then 0.7 else 0.3

if happy(dave) then 0.7 else 0.3

and so on.

Joint probability

- A model defines a joint probability defined by the normalization of:
- $\text{prod_parfactor } g$
 $\text{prod_Lvars in } g : \text{Constraint_}g$
 $\text{phi_}g(\text{Atoms_}g)$

Example:

$\text{prod_}\{X: X \neq \text{bob}\} \text{phi1}(\text{happy}(X))$

*

$\text{prod_}\{X, Y: X \neq Y\}$

$\text{phi2}(\text{happy}(X), \text{married}(X, Y), \text{happy}(X))$

Marginalization

- In marginalization, we want to compute:

$\sum_{\{RV \setminus \text{query}\}} \text{parfactors product}$

but the representation of random variables is often intensional, that is, by constraint:

$\sum_{\{\text{happy}(X)\}_{X: X \neq \text{ann}} \}$

$\sum_{\{\text{married}(X,Y)\}_{X,Y : X \neq Y} \}$

... product of parfactors

Note about notation

- Especially in early papers, there was confusion in notation, with
- $\text{sum}_{\{ p(X) \}}$

being used when the meaning was

$\text{sum}_{\{ \{ p(X) \}_{X : \text{true}} \}}$

mostly because the latter is not standard

Marginalization

- $\sum_{RV \setminus \text{query}}$ parfactors product is of the form:

$\sum_{\text{Atom1}} \text{LVars1} : \text{Constraint1}$

$\sum_{\text{Atom2}} \text{LVars2} : \text{Constraint2}$

...

$\text{prod}_{\{\text{Lvars}_{g1} : \text{Constraint}_{g1}\}} \text{phi}_{g1}(\text{Atoms}_{g1})$

$\text{prod}_{\{\text{Lvars}_{g2} : \text{Constraint}_{g2}\}} \text{phi}_{g2}(\text{Atoms}_{g2})$

$\text{prod}_{\{\text{Lvars}_{g3} : \text{Constraint}_{g3}\}} \text{phi}_{g3}(\text{Atoms}_{g3})$

...

Splitting a Parfactor

- We want to split a parfactor
for all $Lvars_g : Constraint_g . \phi(Atom1, \dots, Atom_m)$

into two sets of factors (two parfactors), one parfactor sharing RVs with $Atom:Constraint$, and the other not sharing them.

Example:

for $X, Y : X \neq Y . \phi(\text{happy}(X), \text{married}(X, Y), \text{happy}(Y))$
sharing RVs with $\{ \text{happy}(Z) \}_{Z : Z \neq \text{ann}}$

Splitting a Parfactor

- for $X, Y : X \neq Y . \text{phi}(\text{happy}(X), \text{married}(X, Y), \text{happy}(Y))$
sharing RVs with $\text{happy}(Z) : Z \neq \text{ann}$

What are the instantiations such that $\text{happy}(X)$ or $\text{happy}(Y)$ is in $\{ \{ \text{happy}(Z) \} _ \{ Z : Z \neq \text{ann} \} \}$?

These are the ones with X, Y such that

$\exists Z : Z \neq \text{ann}$ and ($\text{happy}(X)$ unifies with $\text{happy}(Z)$ or
 $\text{happy}(Y)$ unifies with $\text{happy}(Z)$)

which is

$\exists Z : Z \neq \text{ann}$ and ($X = Z$ or $Y = Z$)

Let's call that IC (intersection constraint)

A fundamental tool: Equality Logic

- Often shoved aside as a detail in papers, but important and not very simple
- Model Counting:
How many assignments to Lvars in equality boolean formula?

$\#X, Y : X \neq \text{bob and } Y \neq X$

=

$(|\text{Dom}(X)| - 1) * (|\text{Dom}(Y)| - 1)$

A fundamental tool: Equality Logic

- We also need *quantifier elimination*:

$$\exists X . Y = a \text{ and } X \neq Y \text{ and } X \neq a$$

\Leftrightarrow

$$Y = a$$

- This can be seen as a projection (from relational algebra or databases) of tuples on X,Y on tuples on Y alone

Splitting a Parfactor

- for $X, Y : X \neq Y . \text{phi}(\text{happy}(X), \text{married}(X, Y), \text{happy}(Y))$
sharing RVs with $\text{happy}(Z) : Z \neq \text{ann}$

IC (intersection constraint) on X, Y is:

$\exists Z : Z \neq \text{ann}$ and ($\text{happy}(X)$ unifies with $\text{happy}(Z)$ or
 $\text{happy}(Y)$ unifies with $\text{happy}(Z)$)

=

$\exists Z : Z \neq \text{ann}$ and ($X = Z$ or $Y = Z$)

=

$X \neq \text{ann}$ or $Y \neq \text{ann}$

Then we obtain two parfactors:

for $X, Y : X \neq Y$ and IC . $\text{phi}(\dots)$

for $X, Y : X \neq Y$ and not IC . $\text{phi}(\dots)$

Splitting a Parfactor

- In general, for all $Lvars_g : Constraint_g . \phi(Atom1, \dots, Atom_m)$

sharing RVs with $\{ Atom \}_{Lvars : Constraint}$

is the pair of parfactors:

for all $Lvars_g : (Constraint_g \text{ and } IC) . \phi(Atom1, \dots, Atom_m)$

for all $LVars_g : (Constraint_g \text{ and not } IC) . \phi(Atom1, \dots, Atom_m)$

(the second one being called *residual*)

where IC is

$\exists Lvars .$

$Constraint$

and $(Atom \text{ unifies with } Atom_1 \text{ or } \dots \text{ or } Atom \text{ unifies with } Atom_m)$

$(A \text{ unifies with } B) \text{ iff } A \text{ and } B \text{ share predicate, and arguments are equal}$

Up to now

- So, up to now, we have learned how to represent the problem and how to factor parfactors out of a summation.
- Now we need to learn how to actually perform the summation, which is over an intensionally defined (that is, by constraint) set of random variables:

```
sum_{ Atom }_Lvars:Constraint  
  prod_Lvars_g1:Constraint_g1 phi_g1(Atoms_g1)  
  ...  
  prod_Lvars_gn:Constraint_gn phi_gn(Atoms_gn)
```

Algorithm Overview

1. Pick a set of random variables
2. Split parfactors involving those
3. Multiply them into one parfactor
4. Apply Partial Inversion and Counting variable conversion as much as possible
5. If you get summations without products inside, perform them, otherwise ground one of the variables in the product and go to 4.

Multiplying Parfactors

- Example:

How can I replace the product of two parfactors potentials:

$(\text{prod_X } \text{phi_1} (p(X))) . (\text{prod_Y } \text{phi_2}(q(Y)))$

by the potential of a single parfactor (a necessary condition for later operations)?

Multiplying Parfactors

$$(\text{prod}_X \text{ phi}_1(p(X))) \cdot (\text{prod}_Y \text{ phi}_2(q(Y)))$$

=

$$(\text{prod}_X \text{ phi}_1^{\{1/|Y|\}}(p(X)))^{\{|Y|\}} (\text{prod}_Y \text{ phi}_2(q(Y)))$$

= (move each of $|Y|$ instances of prod_X inside prod_Y , one per Y)

$$\text{prod}_Y (\text{prod}_X \text{ phi}_1^{\{1/|Y|\}}(p(X)) \text{ phi}_2(q(Y)))$$

= (now do the analogous operations to get $\text{phi}_2(q(Y))$ inside prod_X)

$$\text{prod}_Y \text{ prod}_X \text{ phi}_1^{\{1/|Y|\}}(p(X)) \text{ phi}_2^{\{1/|X|\}}(q(Y))$$

=

$$\text{prod}_{X,Y} \text{ phi}_3(p(X), q(Y))$$

which is the potential of a new parfactor with logical variables X, Y , potential function

$$\text{phi}_3(v1, v2) = \text{phi}_1^{\{1/|Y|\}}(v1) \text{ phi}_2^{\{1/|X|\}}(v2)$$

Multiplying Parfactors

- In general:

$$\text{sum}_{\{ \text{Atom} \}}_{\text{Lvars:Constraint}} \left(\text{prod}_{\text{Lvars}_1 : C_1} \text{phi}_1(A_1) \right) \dots \left(\text{prod}_{\text{Lvars}_n : C_n} \text{phi}_n(A_n) \right)$$

these parfactors are multiplied into one:

$$\text{sum}_{\{ \text{Atom} \}}_{\text{Lvars:Constraint}} \text{prod}_{\text{Lvars}_1 \dots \text{Lvars}_n : C_1 \text{ and } \dots \text{ and } C_n} \text{phi}_1^{c_1}(A_1) \dots \text{phi}_n^{c_n}(A_n)$$

where $c_i = 1 / \# \text{Lvars}_j : j \neq i$. C_1 and ... and C_n
to compensate for the replication of phi_i due to its being inside
a product of the Lvars that belong to the *other* parfactors

Alignment Multiplication

- Multiplication produces a parfactor with all logical variables from all multiplied factors.
- For complexity reasons, we want to keep the number of logical variables down.
- If two logical variables in distinct parfactors being multiplied have the same assignments, we can unify them. If X have the same domain:
- $$\begin{aligned} & (\text{prod}_X \text{phi1}(p(X))) . (\text{prod}_Y \text{phi2}(q(Y))) \\ & = \\ & \text{phi1}(p(a)) \dots \text{phi1}(p(z)) . \text{phi2}(q(a)) \dots \text{phi2}(q(z)) \\ & = \\ & \text{phi1}(p(a)) . \text{phi2}(q(a)) \dots \text{phi1}(p(z)) . \text{phi2}(q(z)) \\ & = \\ & \text{prod}_X \text{phi1}(p(X)) . \text{phi2}(q(X)) \end{aligned}$$

Counting Variables (Milch et al)

- Example:

$\text{prod}_{X,Y,Z} : Z \neq b . \text{phi}(p(X), p(Y), q(Z)) =$

$\text{prod}_{Z} : Z \neq b . \text{prod}_v \text{phi}(v1, v2, q(Z))^{ \{ (\#X : p(X) = v1)(\#Y:p(Y)=v2) \} }$

$\text{prod}_{Z} : Z \neq b . \text{phi}' (\#X : p(X), \#Y : p(Y), q(Z))$

where $\#X : p(X)$ and $\#Y:p(Y)$ are histogram-value random counting variable (two-bar histograms for boolean $p(\cdot)$, more for multi-value $p(\cdot)$)

- In general, if constraints $C_{x_1} \dots C_{x_k}$ and C' do not share logical variables and X_i does not appear in A_j , for $j \neq i$, and not in atoms B :

$\text{prod}_{X_1, \dots, X_k, Y} : (C_x \text{ and } C') . \text{phi}(A_1(\dots X_1 \dots), \dots, A_k(\dots X_k \dots), B(Y))$

=

$\text{prod } Y : C' . \text{phi}' (\#X_1 : C_{x_1} . A_1(\dots X_1 \dots), \#X_k : C_{x_k} . A_k(\dots X_k \dots), B(Y))$

where each

$\#X_i : C_{x_i} . A_i(\dots X_i \dots)$

is a histogram-valued *counting variable* with dimension equal to the number of values in the range of $\text{Atom}(\dots X \dots)$

Why Counting Variables are Useful

- $$\sum_{\{p(U)\}_U} \sum_{\{q(V)\}_{V:V \neq b}} \prod_{X,Y,Z : Z \neq b} \text{phi}(p(X), p(Y), q(Z))$$

has complexity dependent on domain size because of the summation on number of random variables dependent on domain size, and products also iterating over domain.
- But that's the same thing as

$$\sum_{\{p(U)\}_U} \sum_{\{q(V)\}_{V:V \neq b}} \prod_{Z : Z \neq b} \text{phi}'(\#X : p(X), \#Y : p(Y), q(Z))$$

=

$$\sum_{\{p(U)\}_U} \sum_{\{q(V)\}_{V:V \neq b}} (\#Y : p(Y) \text{ same as } \#X : p(X)) \prod_{Z : Z \neq b} \text{phi}''(\#X : p(X), q(Z))$$

=

$$\sum_{\#U : p(U)} \sum_{\#V : V \neq b} \text{phi}'''(\#U : p(U), \#V : V \neq b \cdot q(V))$$

which is now linear in domains of U and V

Partial Inversion

- $$\sum_{U,W} \{p(U,W)\} \sum_{V:V \neq b} \{q(V)\} \prod_{R,X,Y,Z:Z \neq b} \text{phi}(p(R,X), p(R,Y), q(Z))$$
- Cannot convert to counting because R is shared
- Note that R “slices” the set $\{p(U,W)\}_{U,W}$ in disjoint sets
- That would not be the case if we had $p(R,X), p(Z,R)$ (mixing “columns and rows”)
- $$\sum_{V:V \neq b} \{q(V)\} \prod_R \sum_W \{p(R,W)\} \prod_{X,Y,Z:Z \neq b} \text{phi}(p(R,X), p(R,Y), q(Z))$$

(X is now bound)
- $$\sum_{V:V \neq b} \{q(V)\} \prod_R \sum_W \{p_R(W)\} \prod_{X,Y,Z:Z \neq b} \text{phi}(p_R(X), p_R(Y), q(Z))$$

($p(R,.)$ can be seen as new predicate)
- Simpler sub-problem, can be solved once for all R

Partial Inversion

- $$\begin{aligned} & \text{sum}_{\{ \{q(V)\}_{V:V \neq b} \}} \\ & \quad \text{prod}_R \text{sum}_{\{ \{p_X(W)\}_W \}} \\ & \quad \quad \text{prod}_{X,Y,Z} \text{phi}(p_R(X), p_R(Y), q(Z)) \\ & = \\ & \text{sum}_{\{ \{q(V)\}_{V:V \neq b} \}} \\ & \quad \text{prod}_R \text{sum}_{\{ \#W : p_X(W) \}} \\ & \quad \quad \text{prod}_Z \text{phi}'((\#W : p_R(W)), q(Z)) \\ & = \\ & \text{sum}_{\{ \#V : V \neq b . q(V) \}} \\ & \quad \text{prod}_R \text{sum}_{\{ \#W : p_R(W) \}} \\ & \quad \quad \text{phi}'((\#W : p_R(W)), \#V: q(V)) \\ & = (\text{defining } \text{phi}''(v2) = \text{sum}_{v1} \text{phi}'(v1, v2)) \\ & \text{sum}_{\{ \#V: V \neq b . q(V) \}} \text{prod}_R \text{phi}''(\#V: q(V)) \\ & = (\text{defining } \text{phi}'''(v) = \text{phi}''(v)^{| \text{Dom}(R) |}) \\ & \text{sum}_{\{ \#V: V \neq b . q(V) \}} \text{phi}'''(\#V: q(V)), \text{ linear in } |D(V)| \end{aligned}$$

Partial Inversion

- In general:
- We can invert on a set of logical variables appearing in the same argument position in a set of literals

- $\text{sum}_{\{ A_1(\dots) \}} \dots \text{sum}_{\{ A_m(\dots) \}}$
 $\text{prod}_{X_1, \dots, X_k, Y_1, \dots, Y_n} \text{phi}(\dots)$

and each X_j appear in A_i inside phi in same argument positions (this will cause them to slice the sets)

=

$$\text{prod}_{X_1, \dots, X_k}$$

$$\text{sum}_{\{ A_1(\dots) \text{ with } X_1, \dots, X_k \text{ bound in their positions} \}}$$

...

$$\text{sum}_{\{ A_m(\dots) \text{ with } X_1, \dots, X_k \text{ bound in their positions} \}}$$

$$\text{prod}_{Y_1, \dots, Y_n} \text{phi}(\dots)$$

Generalizations and optimizations

- Nima Taghipour et al relaxed several conditions for both counting variables and inversion.
- Multiplying parfactors not always necessary, for example if your parfactors are clauses (Gogate)

Algorithm Overview

1. Pick a set of random variables
2. Split parfactors involving those
3. Multiply them into one parfactor
4. Apply Partial Inversion and Counting variable conversion as much as possible
5. If you get summations without products inside, perform them, otherwise ground one of the variables in the product and go to 4.