circumscription

Preface- FOL

- We have KB consist the sentence:
  Violine(stradivari1)
- We want to conclude it have 4 strings, there are exactly 2 possibilities.
- One is to add the explicit sentence:
  4string(stradivari1).
- Two is to add a generic rule:
  $\forall$**x.violine(x)**$\rightarrow$**4string(x)**

Preface- FOL (problem)

- What if we have a 5string violin or 10string?
- Then we should say: *all violins have 4strings except those who do not*.
- Well we just said…
- nothing!

Preface- Default reasoning

- Assuming that birds usually fly, and tweety is a bird, when can we conclude that tweety flies?

  Given that a P is usually a Q, and given P(a) is true, it is reasonable to conclude that Q(a) is true <u>unless there is good reason not to</u>

- Finding that “good reason” is the whole purpose of the all the default reasoning different methods

Preface- Default reasoning

- If all we know that Yoggi is a polar bear then we can reasonably conclude that he is white
- If you have a CS degree then you have studied logic.
- But a bear might be playing with mud, and a CS graduate might have some “Ptor” from logic.
- So this is only a *reasonable defaults*, and generally speaking, this is *default reasoning*.

Preface- Monotonic vs. Non-monotonic reasoning

- Monotonic: if $\mathbf{KB_1} \models \alpha$, then $\mathbf{KB_2} \models \alpha$ for any $\mathbf{KB_1} \subseteq \mathbf{KB_2}$.
- Meaning new facts can only add to early conclusions not contradict them.
- Non-monotonic: New facts can change our conclusions.
- If we know tweety is a bird –we conclude it flies.
- If we find out that tweety is an ostrich we conclude it don’t flies.

Introduction to circumscription

- Circumscription is a powerful non monotonic formalism created by John McCarthy(1977,1980), generalized (in 1984)
- Independently explored by many researchers
- It is the most fascinating and the most controversial of all the formal approaches to non monotonic reasoning.

Extension

- A predicate denoted by an expression U will be called extension of U.
- For example if **U=Bird** (unary constant predicate) and **D=All individuals**. Then the extension of U is a subset of D (intuitively D=All birds).

What is circumscription

- Like we saw in the evolution of the CWA the extension of the predicates became as small as possible.
- It leads us to the natural generalization: **Consider forms of entailment where the extension of certain predicates is as small as possible.**

What is circumscription

- A simple example:
- Let T consist of one axiom **$Red(a) \wedge On(a,b)$**
- A circumscription of Red in T will conclude
  $\forall x.Red(x) \rightarrow x=a$ (because a is the only Red in T)

- A circumscription of On in T will conclude
  $\forall x,y.On(x,y) \rightarrow x=a \wedge y=b$
- We could also jointly circumscribe several predicates, say Red and On for instance.

There is more then one

- There are about 10 different version of circumscription on the AI market today.
- All of then share the following common characteristic:

1.Circumscription allows us to formalize non-monotonic reasoning directly in the language of classic logic.

2.it is always the task of the user to specify which predicates to be minimized. Circumscription provides a general method for it.

3.Circumscription is based on syntactic manipulations.

Few circumscription types

1. Predicate circumscription
2. Formula circumscription
3. Second-order circumscription
4. Non-recursive circumscription
5. Domain circumscription
6. Pointwise circumscription

The predicate Ab

One way to handle default knowledge, suggested by McCarthy(1984), is to assume we have a predicate **Ab** that taking care of all the exceptional or abnormal cases where the default should not apply.

The predicate Ab

For instance, instead of saying that all birds fly we might say:

$$\forall x[\mathbf{Bird(x)} \wedge \neg \mathbf{Ab(x)} \rightarrow \mathbf{Flies(x)}]$$

meaning that all birds that are not in some way abnormal fly, or more succinctly, that all normal birds fly.

The predicate Ab

Now imagine we have this fact in KB along with these facts:

Bird(chilly),Bird(tweety),

(tweety$\neq$ chily), ¬Flies(chilly)

now we would like to conclude that Tweety flies whereas Chilly doesn’t.

But that isn’t the case. WHY?

The predicate Ab

Because **KB$|\neq$ Flies(tweety)**. WHY?

There are interpretations (Models) satisfying the KB where Flies(tweety) is false. For example

M=$\langle$D,I$\rangle$, D={KB, $\neg$Flies(tweety)}

however note that in these interpretations, the denotation of tweety is contained in the extension of Ab.

Minimizing Abnormality

As mentioned before and in CWA as well our goal is to consider only the smallest interpretations, meaning those that has a minimum Ab predicate appearances. This strategy called *minimizing abnormality*.

in the previous example we saw that chilly is an abnormal bird but we have no such information about tweety.

Circumscribing the predicate Ab

The default assumption we want to make is that the **extension of Ab is only as large as it has to be given what we know**.

Therefore in our last example we include chily but exclude tweety (in our extension), this is called **circumscribing the predicate Ab**. And the technique called **circumscription**.

The minimal extension

- The minimal extension of Ab in our example is:
- **$\forall x[x \neq chilly \rightarrow \neg Ab(x)]$,**
- **$\forall x[Bird(x) \wedge x \neq chilly \rightarrow Flies(x)]$.**

- Lets add two more facts to our KB:
- **$Ostrich(Joe) \wedge Joe \neq chilly \wedge Joe \neq Tweety$**
- **$\forall x[Ostrich(x) \rightarrow Bird(x)]$,**
- Now we conclude that Joe is normal and that he flies

The minimal extension

- Ostriches don't fly (as far as we know)!
- Solution 1: $\forall$**x[Ostrich(x)$\rightarrow\neg$Flies(x)],**

- But what if we are not sure they don't fly?
- Solution2: $\forall$**x[Ostrich(x)$\rightarrow$Ab(x)],**
- now we can point out some special ostrich that flies, and yet we can not conclude it by default.

There is more then Ab

Note that while chily is abnormal in its flying abilities it might be perfectly normal in other “birdy” aspects like having two legs, lying eggs and so on.

Ab is not enough so we need a multiple Ab predicate variations:

$\mathbf{Ab_i}$, so chily might be in extension $\mathbf{Ab_1}$ but not in extension $\mathbf{Ab_2}$.

Minimal entailment

- Circumscription is intended to be a more accurate and convenient tool then the CWA
- And we want to apply it in much broader settings
- So we don’t add negative literals to the KB like CWA. Instead we characterize a new form of entailment.

≤ definition

Let **P** be a fixed set of unary predicates, which we will understand as Ab predicates.

Let $\mathbf{M_1} = \langle D, \mathbf{I_1} \rangle$ and $\mathbf{M_2} = \langle D, \mathbf{I_2} \rangle$. We define the relation, ≤**:**

$$\boldsymbol{M_1 \leq M_2 \Leftrightarrow \text{for every } p \in P,\ I_1[p] \subseteq I_2[p]}$$

Also $\boldsymbol{M_1 < M_2}$ iff $\boldsymbol{M_1 \leq M_2}$ and not ($\boldsymbol{M_2 \leq M_1}$ **).**

More Normal

Intuitively, given two interpretation over the same domain, one is less then another if it makes the extension of all the abnormal predicates smaller.

We can think of interpretation that is less then another as **more normal**.

Minimal entailment $\models_{<}$

Now we can define a new form of minimal entailment $\models_{<}$:

KB $\models_{<} \alpha \Leftrightarrow$ ***for every interpretation M ,such that M*** $\models$***KB, either M*** $\models\alpha$ ***or there is M' such that M'<M and M'***$\models$***KB***

Minimal entailment

This is similar to the definition of entailment itself: we require each interpretation that satisfies KB to satisfy α **except that** it may be excused if there is a smaller (more normal) interpretation that satisfies it.

Generally speaking we don't ask that every interpretation that satisfies M will satisfy α only the **most normal** one(s).

Back to tweety and chilly example

- As noted (16th slide) **KB$|\neq$Flies(tweety)**.
- But **KB $|=_{<}$ Flies(tweety).**
- Because: if **M $|=$KB** but **M $|\neq$Flies(tweety),** then **M$|=$Ab(tweety).**
- So let M' be exactly M except that we remove the denotation of tweety from the extension of Ab, M'<M and M'|=KB
- Thus in the most normal form KB$|=_{<}\neg$Ab(tweety)
- So tweety is normal but chilly is not, because chilly is abnormal in any model of KB, minimal or not.

Back to tweety and chilly example

The default step was to conclude that tweety is normal (or not abnormal) all the rest was a simple deductive reasoning using what we know a bout birds.

That is circumscription proposal for default reasoning

There is more then one champion

- Not all the minimal models satisfies the exact same sentences
- For instance

  KB={Bird(c),Bird(d),¬flies(c)∨¬flies(d)}
- The minimal model will have only a **Flies(c)** or **Flies(d).**
- So we get **KB** $|\neq_{<}$**Flies(c)** and **KB**$|\neq_{<}$**Flies(d)**
- Meaning we cannot conclude by default that c is a normal bird nor that d is.
- …but we can conclude one of them is: **KB**$|=_{<}$**Flies(c)** ∨ **Flies(d)**

CWA vs. Circumscription- round 1

- Unlike the CWA, we are **consistent with what we know**
- CWA in the previous case would add all the negation of what we don’t know, meaning **¬Ab(c), ¬Ab(d),** what would lead to inconsistency.
- thus circumscription is more cautious the CWA in its assumption
- But less cautious then GCWA which wouldn’t conclude anything about normality of d or c, while circumscription concludes that at least one of them flies.

CWA vs. Circumscription- round 2

- By using interpretation directly instead of adding literals circumscription works equally well with unnamed individuals.

- For example:

  $\exists$**x[Bird(x)$\wedge$(x$\neq$chilly)$\wedge$(X$\neq$tweety)]**

- Circumscription concludes:
  $\exists$**x[Bird(x)$\wedge$(x$\neq$chilly)$\wedge$(X$\neq$tweety)$\wedge$Flies(x)]**

CWA vs. Circumscription- round 2

- This also carries over to unnamed abnormal individuals. Suppose we have: $\exists$**x[Bird(x)**$\wedge$**(x**$\neq$**chilly)**$\wedge$**(X**$\neq$**tweety)**$\wedge\neg$**Flies(x)]**
- In this case the minimal model will have exactly 2 abnormal individuals. Thus we conclude: $\exists$**x**$\forall$**y[(Bird(y)**$\wedge\neg$**Flies(y))**$\leftrightarrow$**(y=chilly**$\vee$**y=x)]**
- unlike CWA and GCWA we don't have to name all exceptions explicitly to avoid inconsistency.

Fixed and variable predicates

- Although the assumption made by circumscription are usually weaker than those of the CWA, sometimes they appear too strong!
- Suppose for example we have KB:

  $\forall$**x[Bird(x)** $\wedge$ $\neg$**Ab(x)** $\rightarrow$ **Flies(x)],**

  Bird(tweety),

  $\forall$**x[Penguin(x)** $\rightarrow$**(Bird(x)**$\wedge$$\neg$**Flies(x))].**
- We can conclude: $\forall$**x[Penguin(x)**$\rightarrow$**Ab(x)]**

Problem

- When we make default assumption with circumscription we minimize the set of abnormalities and by doing so in this KB, we conclude that:

  $$\mathbf{KB} \models_{<} \neg\exists \mathbf{x.Ab(x)}$$

- and by that we conclude:

  $$\mathbf{KB} \models_{<} \neg\exists \mathbf{x.Penguins(x)}$$ (there is a model that satisfies it)

- **WHAT JUST HAPPENED???**

Problem

All we wanted is to conclude that tweety is not a penguin (perhaps), and that it flies, but by doing so **We have just eliminated the existence of penguins.** Which is definitely not what we wanted.

Solution

- To improve that undesirability we would want to be able to conclude by default that penguins are the only abnormal birds (in this example).
- One way of doing so, was proposed by Vladimir Lifschits, is to redefine $|=_{<}$ so we will look for the smallest model **with exactly the same amount of penguins** (or any other exceptional classes in general)

Fixed extension vs. variable extension

- In the previous example we say that the extension of penguins remains **fixed**
- But in the same example we can make an extension in which we conclude Flies(tweety). So Flies is **variable** in the minimization

Redefinition $\leq$

With respect to a set of unary predicates P (variables- to be minimized)

Q (fixed predicates)

Let $\mathbf{M_1} = \langle D, \mathbf{I_1} \rangle$ and $\mathbf{M_2} = \langle D, \mathbf{I_2} \rangle$. We redefine the relation, $\leq$**:**

$M_1 \leq M_2 \Leftrightarrow$ for every $p \in P$, $I_1[p] \subseteq I_2[p]$

and for every $q \in Q$ $I_1[q]=I_2[q]$

Redefinition $|=_<$

- The rest of the $|=_<$ definition stays the same
- In our example we should take P={Ab}, Q={Penguin} and we get our goal by preserving the only abnormal bird, the penguin.
- In previous $|=_<$ definition Q={}.
- With our new definition we have now minimal models (not all but exist!!) in which there are penguins, so: KB$|\neq \neg\exists$**Penguin(x)**

Problem1

- As we can see in this method of circumscription we must **point out** which predicates to minimize and also which to leave fixed.
- The solution should be able to declare **automatically** that flying is a variable predicate, for example. And it is far from clear how.

Problem 2

- **KB** $|\neq_{<}$ **Flies(tweety).** WHY?
- There is a minimal model in which tweety happens to be a penguin. We can no longer find a lesser model where tweety flies because we will have to change the set of penguins, which must remain fixed.
- What we do get is **KB**$|=\neg\exists$**Penguin(tweety)**$\rightarrow$**Flies(tweety)**

Problem 2

- So if we know that tweety is not a penguin as in:
  Canary(tweety),$\forall$x[Canary(x)$\rightarrow\neg$Penguin(x)]
- Now we can conclude, Flies(tweety)
- BUT this is not derivable by default, even if we add something saying that birds are **normally** not penguins:

  $$\forall x[\mathbf{Bird(x)} \wedge \neg \mathbf{Ab}_2\mathbf{(x)} \rightarrow \neg \mathbf{Penguin(x)}]$$
- Tweety still doesn’t fly because we cant reduce the set of penguins!

Names are abstract

- If we remove from the KB (slide 15) the sentence **(tweety$\neq$ chily), we will get the same problem we have just introduced.**
- KB $|\neq_{<}$Flies(tweety)
- But only KB$|\neq_{<}$(tweety$\neq$chilly)$\rightarrow$Flies(tweety)
- We don’t know yet how to conclude that names are unique by default.

…One bad Apple…

- It doesn't have to be chilly.
- Enough to say: $\exists x[Bird(x) \wedge \neg Flies(x)]$
- So we wont be able to conclude that tweety is a flying bird by default.

The beginning

- The original problem considered by McCarthy was “missionaries and cannibals”
- His goal wasn’t getting the solution but rather excluding the solutions that are not explicitly stated
- For instance, taking a bigger boat or going over a bridge.
- Our assumption should be that the problem contains all the information needed for its solving.
- Stating that there is no bridge is obviously not a solution because there are infinitely many other solutions theoretically possible, like taking a helicopter.

Formula circumscription

- Its circumscription generalization in which extension of a Formula minimize rather then a predicate.
- For example predicate expression ($\forall$x.Bird(x)) instead of predicate constant (Bird(tweety)).

Domain circumscription

- This circumscription type is based on minimizing the domain of first-order logic, rather then the extension of predicates.
- To be exact model considered less then another if it has a smaller domain.

Ray Reiter's coin

- Circumscription does not always handle well disjunctive sentences.
- An example provided by Ray Reiter is:
- A coin tossed over a check board
- As a result it should fall on ether black or white.
- By circumscription we can conclude all the places it cant fall on
- Like: the floor, the moon and so on…
- So we exclude predicates like on(coin,moon)

problem

- What if it falls on both?
- We can not conclude by circumscription a a model with extension of on(coin,black) and on(coin,white).
- Because this extension won’t be minimal
- The minimal extension of on is a model with only ONE of them true.

Theory curbing

- Theory curbing is a solution proposed by Thomas Eiter, Georg Gottlob, and Yuri Gurevich.
- The idea is that the model with on(coin,black)$\wedge$on(coin,white) is a model of a greater formula (w.r.t to predicate on) then both possible models
- Meaning that this model is the least upper bound of the 2 selected models
- So beside the circumscription we select (by curbing) the least upper bound models of all sets of models the circumscription contains.