Announcement

- Homework 1 due
- Verify your in-class score on Compass
- Start reading Chapter 10,
  Classical Planning

SEMANTICS

$$\Theta: \forall x \, [\text{Student}(x) \Rightarrow \text{Happy}(x)]$$

Intuitive meaning
in our world

Student means “is a student”
Happy means “is happy”
and all students are joyful in this world

Student means “is a giraffe”
Happy means “has a short neck”
and there are no giraffes in this world

Student means “can drive”
Happy means “can swim”
in our world

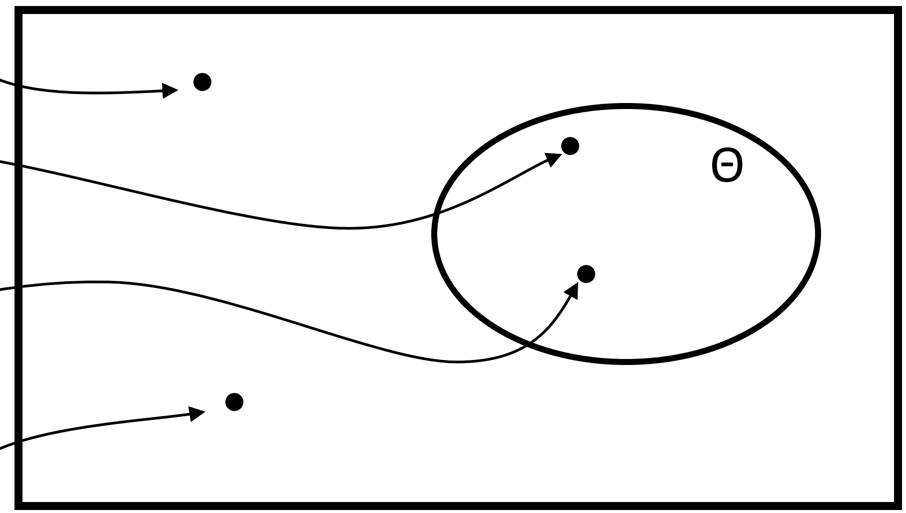

Universe of Possible Worlds

Componential Semantics

- WFFs express constraints
- Meaning of a WFF is the set of possible worlds that satisfy

- $\wedge$ Intersect sets
- $\vee$ Union sets
- $\neg$ Complement set

Entailment $\models$

- An axiom set Δ contains n axioms: Γ_i $_{i=1,n}$
- These are implicitly conjoined:

$$\Delta \equiv \bigwedge_{i=1,n} \Gamma_i$$

- $\Delta \models \theta$ iff θ holds in all possible worlds of Δ

$PW(\Delta) \subseteq PW(\theta)$

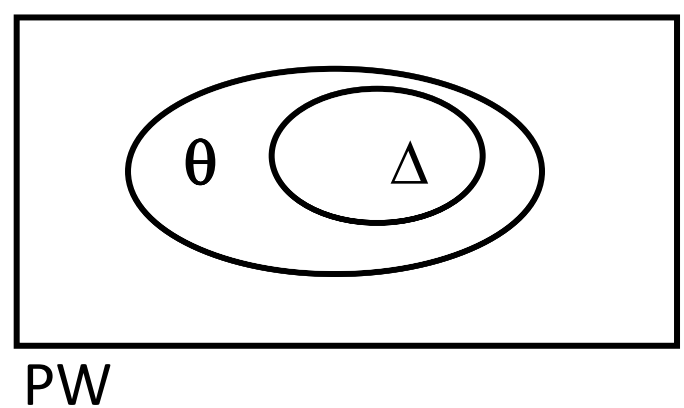

Overloaded “model”

- Logic specific usage
- A model for a sentence is any possible world in which a WFF holds

  (recall that for us a possible world is a world and a denotational correspondence)

- Let $M(\alpha)$ be the models of WFF α
  then $\Delta \models \theta$ is equivalent to
  $M(\Delta) \subseteq M(\theta)$
- Previous slide uses the less formal PW
- To avoid confusion we will minimize this use of “model”

Entailment

An axiom set Δ contains 3 axioms:

$$\Delta \equiv \bigwedge_{i=1,3} \Gamma_i$$

$$(\Delta \equiv \Gamma_1 \wedge \Gamma_2 \wedge \Gamma_3)$$

Any WFF that *completely includes* this intersection is logically entailed

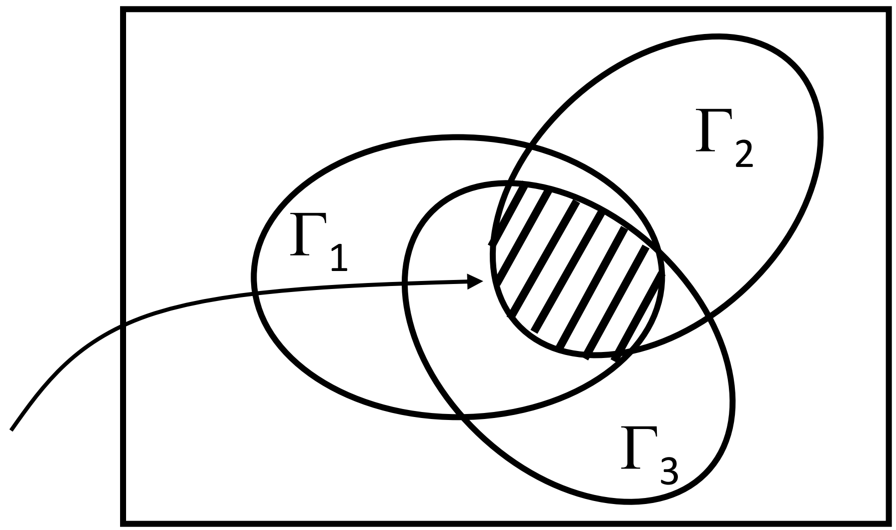

Possible Worlds

Entailment

An axiom set Δ contains 3 axioms:

$$\Delta \equiv \bigwedge_{i=1,3} \Gamma_i$$

$$(\Delta \equiv \Gamma_1 \wedge \Gamma_2 \wedge \Gamma_3)$$

Any WFF that *completely includes* this intersection is logically entailed

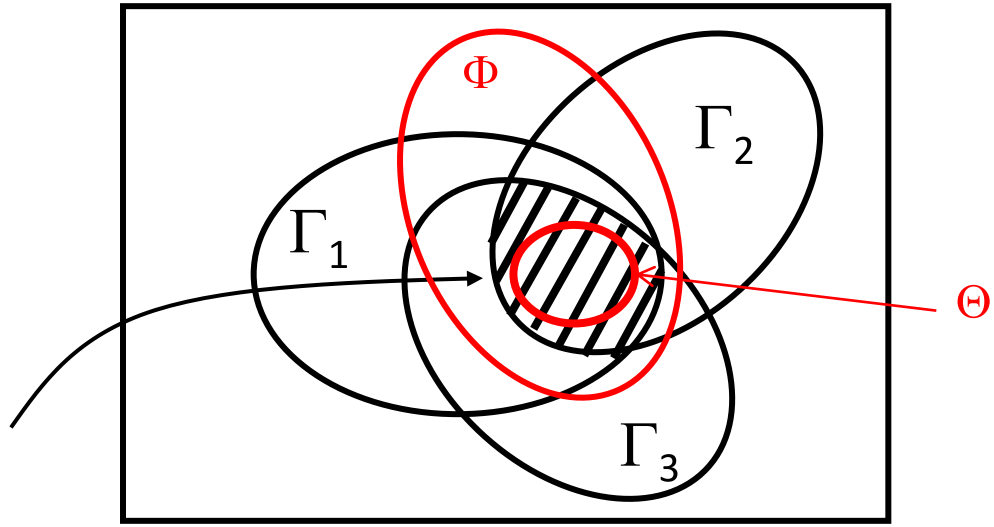

Possible Worlds

Does $\Delta \models \Phi$? Yes

Does $\Delta \models \Theta$? No

Propositional Calculus

- Also "zeroth-order predicate calculus"
- No variables (thus, no quantifiers)
- Unambiguous (FOPC also)
- Not canonical (FOPC also)

Probably encountered as Boolean Logic but be careful!
Boolean "variables" are really Atoms (Atomic WFFs)!
 (do not call them variables in this class!)
Boolean "functions" are really WFFs!
 (do not call them functions in this class!)
WFFs are truth-valuable; Predicates denote relationships in the world
Variables / Functions denote items in the world

Example

B - Fred has blond hair

R - Fred has red hair

"Fred does not have both red and blond hair"

$\neg(B \wedge R)$ $\quad$ $R \Rightarrow \neg B$

$\neg B \vee \neg R$ $\quad$ $(\neg B \vee \neg R) \vee (B \Rightarrow \neg R)$

$B \Rightarrow \neg R$ $\quad$ $(\neg B \vee \neg R) \wedge (B \Rightarrow \neg R)$ ...

Proof by Truth Table

$$\neg(p \vee q) \overset{?}{\equiv} \neg p \wedge \neg q$$

p	q	$p \vee q$	$\neg(p \vee q)$	$\neg p$	$\neg q$	$\neg p \wedge \neg q$
F	F	F	T	T	T	T
F	T	T	F	T	F	F
T	F	T	F	F	T	F
T	T	T	F	F	F	F

$$\neg(p \vee q) \overset{!}{\equiv} \neg p \wedge \neg q$$

Possible World
Equivalence Classes

p, q, r are WFFs

$p \wedge q \equiv q \wedge p$

$p \vee q \equiv q \vee p$

$p \wedge (q \wedge r) \equiv (p \wedge q) \wedge r$

$p \vee (q \vee r) \equiv (p \vee q) \vee r$

$p \Rightarrow q \equiv \neg q \Rightarrow \neg p$

$\neg(\neg p) \equiv p$

$p \Rightarrow q \equiv \neg p \vee q$

$p \Leftrightarrow q \equiv (p \Rightarrow q) \wedge (q \Rightarrow p)$

$p \vee p \equiv p$

$p \wedge p \equiv p$

$\neg(p \vee q) \equiv \neg p \wedge \neg q$

$\neg(p \wedge q) \equiv \neg p \vee \neg q$

$p \wedge (q \vee r) \equiv (p \wedge q) \vee (p \wedge r)$

$p \vee (q \wedge r) \equiv (p \vee q) \wedge (p \vee r)$

Inference

Making explicit what you already know

All men are mortal

Socrates is a man

Socrates is mortal

Modus Ponens

(well-known inference rule)

$$\frac{\begin{array}{c}\Theta \Rightarrow \Psi \\ \Theta\end{array}}{\Psi}$$

Δ is a set of sentences (axioms)

ϕ is a sentence (goal)

A *derivation* of ϕ from Δ is a sequence of sentences culminating with ϕ in which each sentence is either a member of Δ or concluded by a rule of inference whose conditions match sentences earlier in the sequence.

$$\Delta \vdash_{m.p.} \phi$$

Axiom set Δ:

1. **$P \Rightarrow Q$**
2. **$L \Rightarrow R$**
3. **$Q \Rightarrow R$**
4. **$\neg L \Rightarrow Z$**
5. **$S \Rightarrow L$**
6. **$P \Rightarrow G$**
7. **$\neg L$**
8. **A**
9. **P**
10. **G**

Is **R** true?

Can we prove it using MP?

Derivation of R

1.	$P \Rightarrow Q$	Δ
2.	$L \Rightarrow R$	Δ
3.	$Q \Rightarrow R$	Δ
4.	$\neg L \Rightarrow Z$	Δ
5.	$S \Rightarrow L$	Δ
6.	$P \Rightarrow G$	Δ
7.	$\neg L$	Δ
8.	A	Δ
9.	P	Δ
10.	G	Δ
11.	Q	MP: 1,9
12.	R	MP: 3,11

Database Δ:

1. $P \Rightarrow Q$
2. $L \Rightarrow R$
3. $Q \Rightarrow R$
4. $\neg L \Rightarrow Z$
5. $S \Rightarrow L$
6. $P \Rightarrow G$
7. $\neg L$
8. A
9. P
10. G

What about $\neg S$?

Can we prove it using MP?

Interesting Observation #1

It may not be possible for a set of inference rules to infer a sentence even though the sentence is entailed by the database

An inference procedure is *complete* iff any sentence entailed by a database can be derived from the database using the inference procedure .

Recall: A database Δ entails a sentence ϕ ($\Delta \models \phi$) iff every possible world that satisfies Δ also satisfies ϕ.

Interesting Observation #2

We don’t know what P, Q or any other proposition means, and yet we deduced R.

An inference procedure is *sound* iff any sentence derivable from a database using the inference procedure is entailed by the database.

Venn Diagrams

set of axioms

inference procedure

E: entailed $\models$

D: derivable $\vdash$

E: entailed WFFs D: derivable WFFs

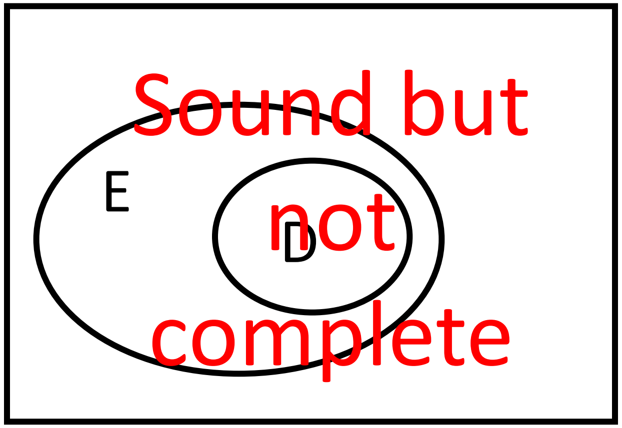

All WFFs

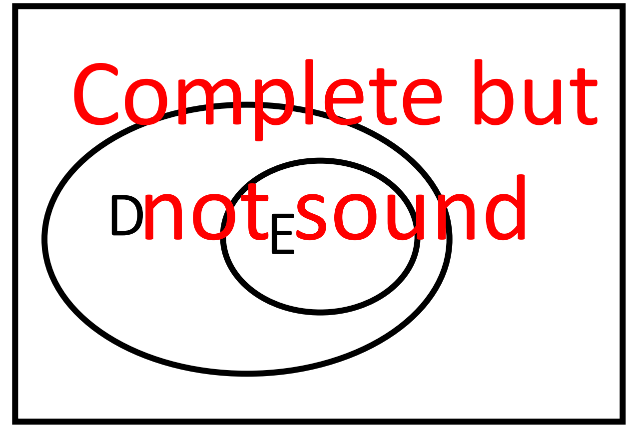

All WFFs

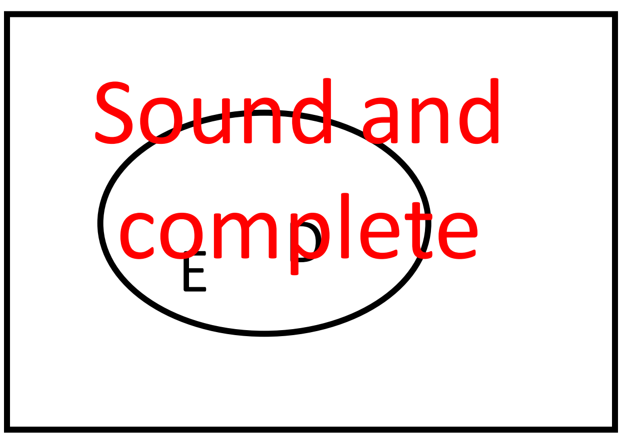

All WFFs

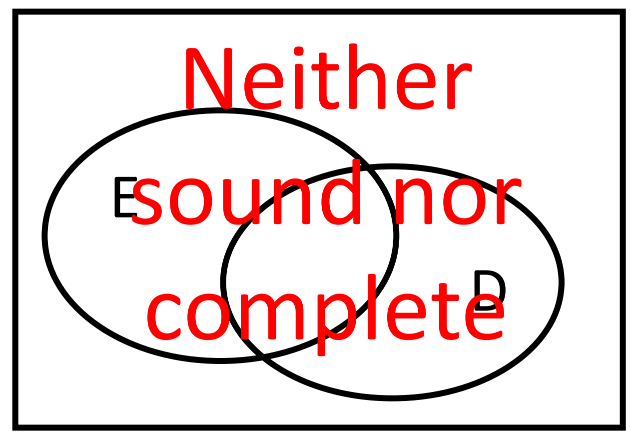

All WFFs

Unsound Inference Rule

$$\frac{\Theta \Rightarrow \Psi \qquad \Psi}{\Theta}$$

Known in AI as “abduction”

Resolution

$$\alpha \vee \beta \longrightarrow \neg\alpha \Rightarrow \beta$$

$$\neg\beta \vee \gamma \longrightarrow \beta \Rightarrow \gamma$$

$$\overline{\alpha \vee \gamma} \longleftarrow \overline{\neg\alpha \Rightarrow \gamma}$$

Sound? Yes

Complete? No (why not?)

Resolution

$\alpha \vee \beta$	$\longrightarrow$	$\neg\alpha \Rightarrow \beta$
$\neg\beta \vee \gamma$	$\longrightarrow$	$\beta \Rightarrow \gamma$
$\alpha \vee \gamma$	$\longleftarrow$	$\neg\alpha \Rightarrow \gamma$

Sound? Yes

Complete? No (why not?)

Inference Requires Matching

M.P.

$$\frac{\Theta \Rightarrow \Psi \qquad \Theta}{\Psi}$$

Database Δ:

1. **$P \Rightarrow Q$**

2. **...**

7. **$\neg L$**

8. **A**

9. **P** **...**

11. **Q**

Matching in propositional logic is easy…

Matching can be Hard / Problematic

Depends on matching criteria.

Should these match?

President(US, 2010)	Barack-Obama1
President(US, 1989)	George-Bush2
Best(Restaurant, C-U)	Timpone's47
Best-item(Music)	Longest-Piece(Telemann39)
Favorite-class	CS 440

No, then matching would depend on

- Interpretation
- World
- Our knowledge
- Our opinions...

Unification: Popular and efficient matcher for complex statements; works only on certain forms.

Matching in First-Order: Unification

$\forall$x [Man(x) $\Rightarrow$ Mortal(x)]

Man(Socrates)

$\therefore$ Mortal(Socrates)

Matching Man(x) with Man(Socrates)
succeeds provided x=Socrates

Unifier

A *unifier* (also *substitution*, *binding list**) is a set of pairings of variables with terms:

$\{v_1 = e_1, v_2 = e_2, v_3 = e_3, \ldots v_n = e_n\}$

such that

- each variable is paired at most once
- a variable's pairing term may not contain the variable directly or indirectly

{x = Socrates}

* Do not confuse with bound / free variables!!!

Are These Acceptable Unifiers?

{x = y}	YES
{x = y, z = F(y)}	YES
{x = y, z = F(y), x = A}	NO
{x = y, z = F(y), y = A}	YES
{x = y, y = F(z), z = G(x)}	NO

Applying a unifier to an expression results in a *unification instance.*

A set of expressions *unify* (are *unifiable*) iff there exists a unifier that when applied results in **identical** unification instances.

Do These Unify?

(Single lower case letters are variables)

P(x,y,z)   P(w,w,Fred)

Yes, consider

θ = {x=Fred, y=Fred, z=Fred, w=Fred}

Equivalently: {x=Fred, w=y, z=Fred , y=x}

Both yield

P(Fred,Fred,Fred)   P(Fred,Fred,Fred)

Are there others?

P(x,y,z)   P(w,w,Fred)

Yes, consider
θ = {x=Mary, y=Mary, z=Fred, w=Mary}

Equivalently: {x=Mary, w=y, z=Fred , y=x}

Both yield
P(Mary,Mary,Fred)   P(Mary,Mary,Fred)

Most General Unifier MGU

The MGU imposes the fewest constraints, specifying the *weakest* conditions for matching

MGU is unique

order is not important

variable names are not important (alphabetic variants)

Applying the MGU to an expression yields a *most general unification instance*.

Variable substitutions are always interpreted with the unifier applied

What is the MGU?

P(x,y,z)   P(w,w,Fred)

{x=w, y=w, z=Fred}

Yields P(w,w,Fred)

Equivalently, {x=u, y=u, w=u, z=Fred}

Yields the alphabetic variant
P(u,u,Fred)

What is the MGU?

M(Ann,x,Bob)	M(Ann,x,Bob)
M(Ann,x,Bob)	M(y,x,Chuck)
M(Ann,x,Bob)	M(y,x,Father-of(Chuck))
P(w,w,Fred)	P(x,y,y)
Q(r,r)	Q(x,F(x))
Q(r,r)	Q(x,F(y))
R(G(x,Bob),y,y)	R(z,G(Fred,w),z)