Introduction to Logic
Introduction

Michael Genesereth
Computer Science Department
Stanford University

Lecture will begin at ~13:35 PDT.

Greek Trivium

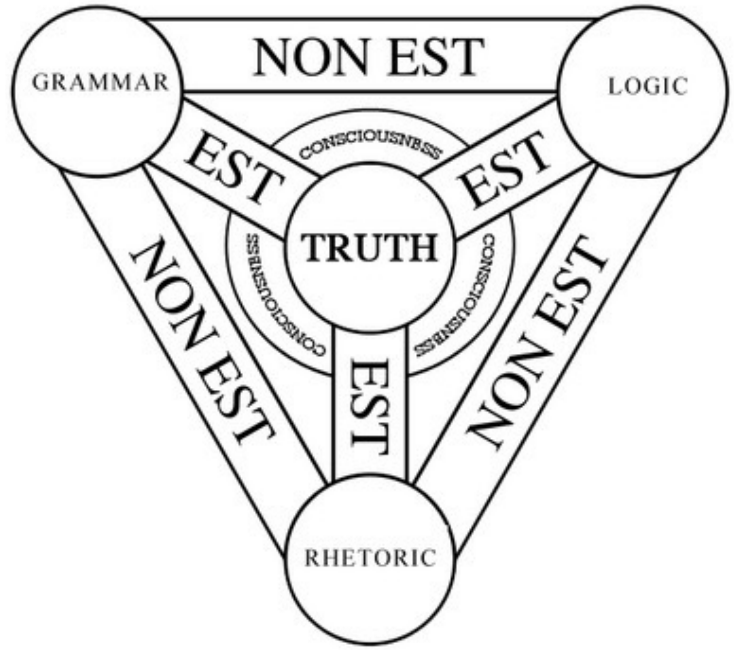

History

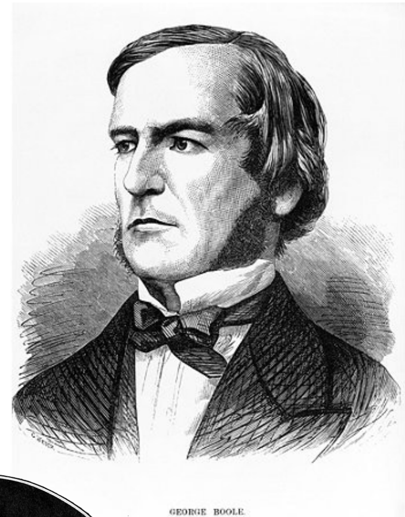

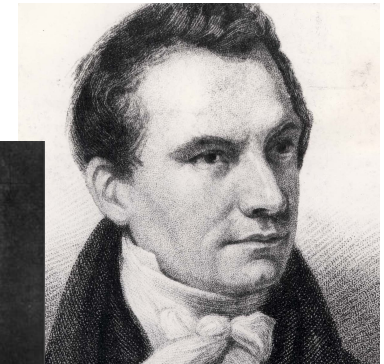

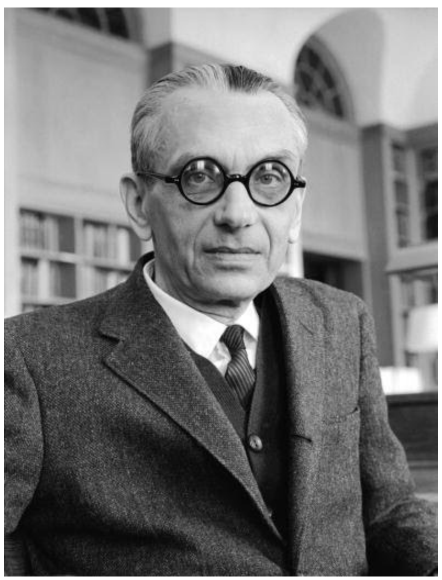

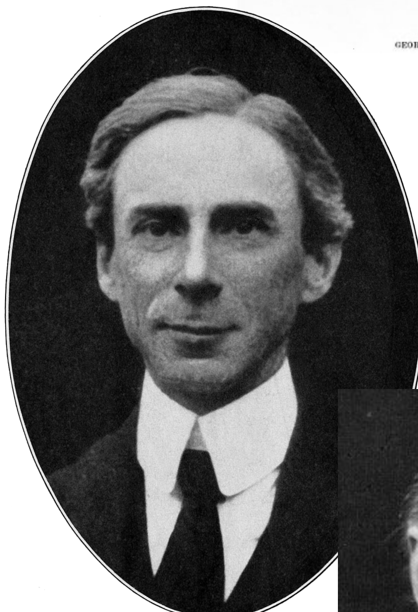

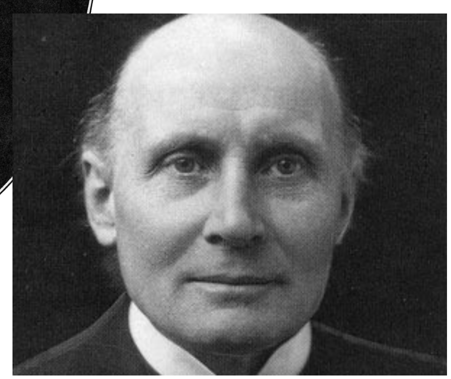

History of Logic

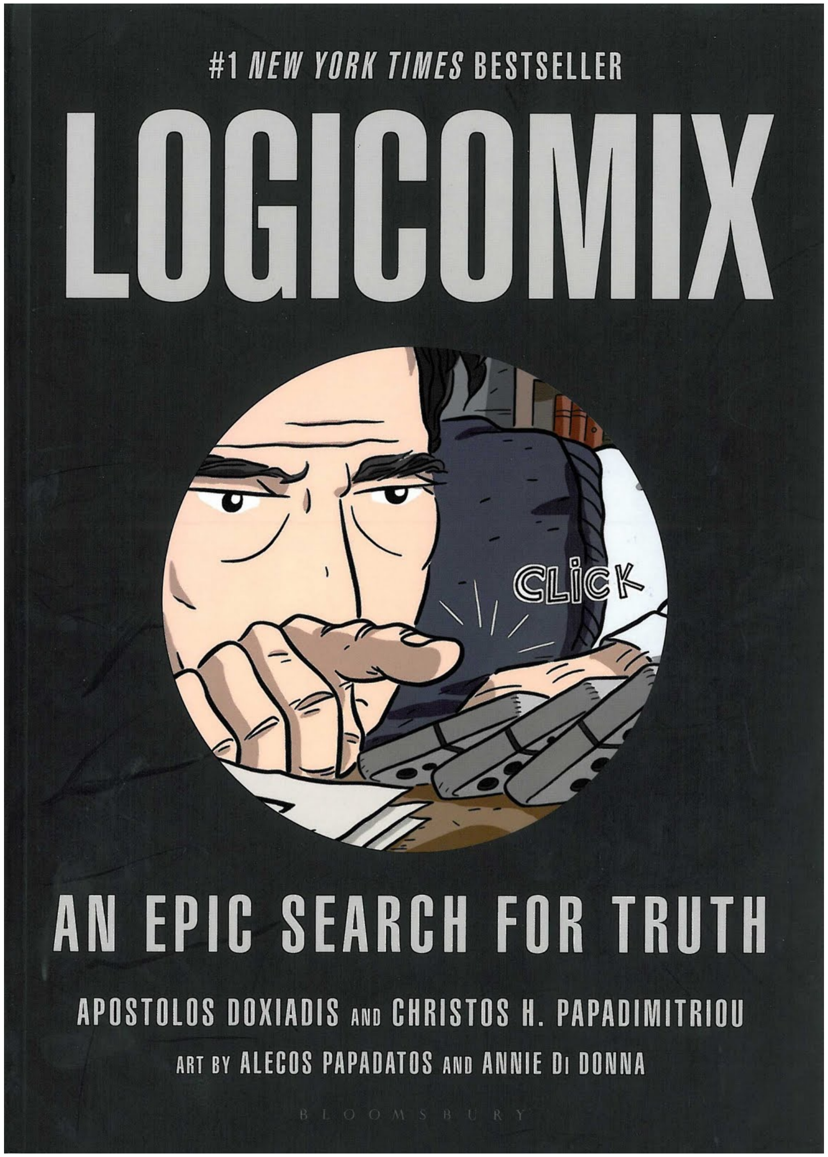

History of Logic

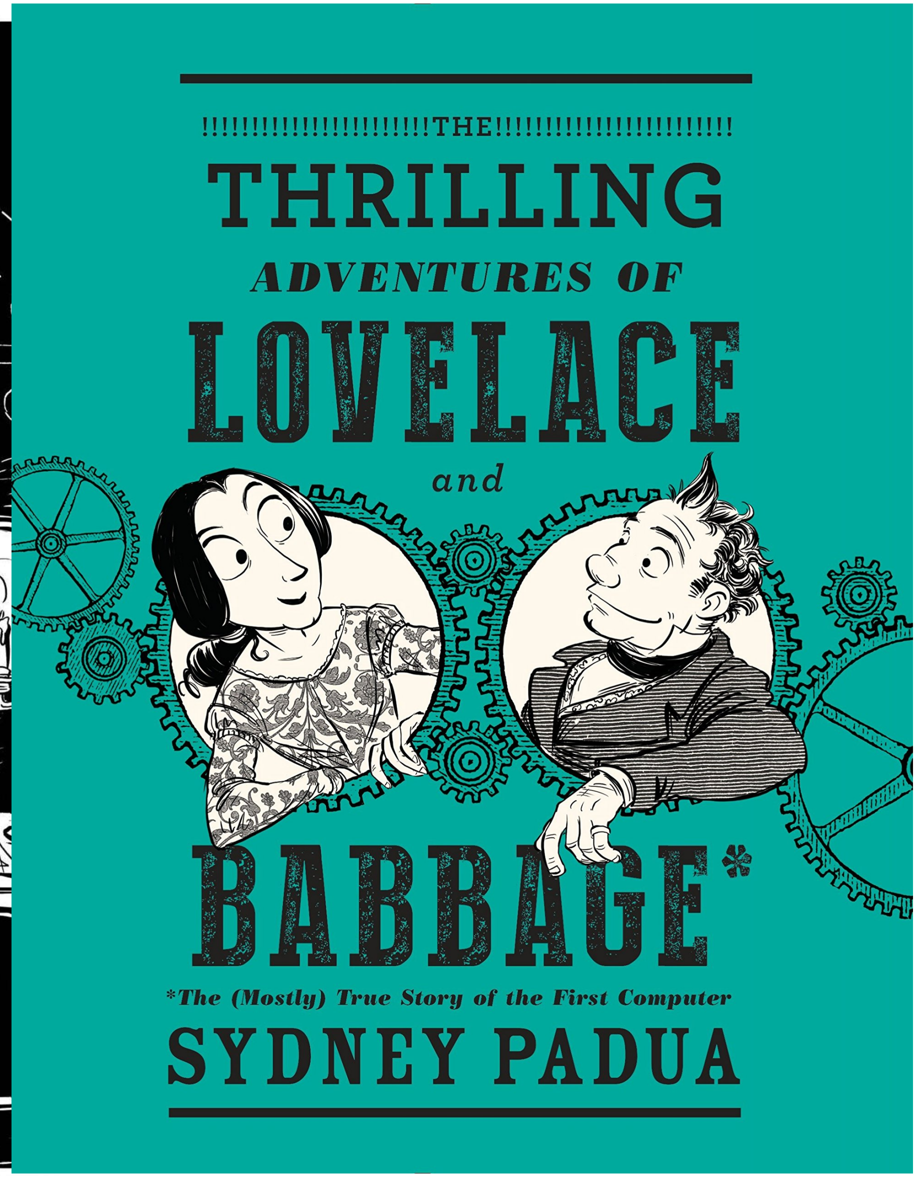

History of Logic

Uses of Logic

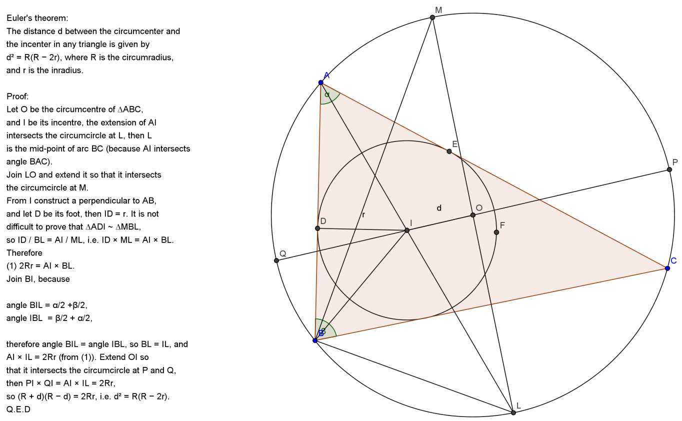
Euler's theorem:
The distance d between the circumcenter and the incenter in any triangle is given by $d^2 = R(R - 2r)$, where R is the circumradius, and r is the inradius.

Proof:
Let O be the circumcentre of ΔABC, and I be its incentre, the extension of AI intersects the circumcircle at L, then L is the mid-point of arc BC (because AI intersects angle BAC).
Join LO and extend it so that it intersects the circumcircle at M.
From I construct a perpendicular to AB, and let D be its foot, then ID = r. It is not difficult to prove that $\Delta ADI \sim \Delta MBL$, so $ID / BL = AI / ML$, i.e. $ID \times ML = AI \times BL$.
Therefore
(1) $2Rr = AI \times BL$.
Join BI, because

angle $BIL = \alpha/2 + \beta/2$,
angle $IBL = \beta/2 + \alpha/2$,

therefore angle BIL = angle IBL, so BL = IL, and $AI \times IL = 2Rr$ (from (1)). Extend OI so that it intersects the circumcircle at P and Q, then $PI \times QI = AI \times IL = 2Rr$, so $(R + d)(R - d) = 2Rr$, i.e. $d^2 = R(R - 2r)$.
Q.E.D

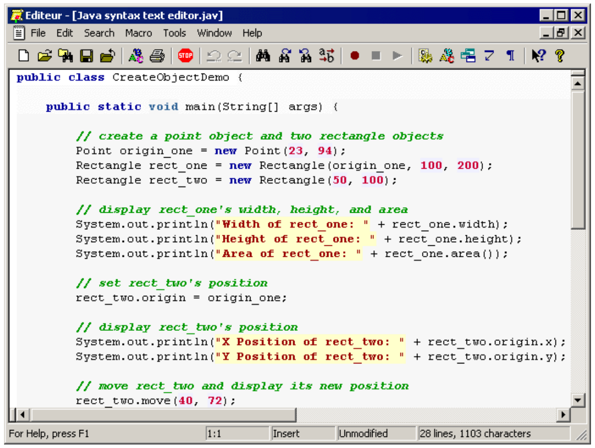

Uses of Logic

*"Whether I am on a **soccer field** or at a **robotics competition**, I face a lot of situations where logic is necessary to make decisions."*

*"I have always loved **puzzles** and like to solve challenging problems."*

*"**Math** classes aren't the only classes that require logic; in AP **History** I am often called upon to recognize patterns and cycles spanning centuries, while in **English** classes I need to write persuasive essays."*

Space-indexed Dynamic Programming: Learning to Follow Trajectories

J. Zico Kolter KOLTER@CS.STANFORD.EDU
Adam Coates ACOATES@CS.STANFORD.EDU
Andrew Y. Ng ANG@CS.STANFORD.EDU
Yi Gu GUYINET@STANFORD.EDU
Charles DuHadway DUHADWAY@STANFORD.EDU
Computer Science Department, Stanford University, CA 94305

Theorem 3.1 [following (Bagnell et al., 2004)] *Suppose $\pi = (\pi_1, \ldots, \pi_D)$ is a policy returned by an ϵ-approximate version of state-indexed PSDP where on each step the algorithm obtains π_d such that*

$$E_{s\sim\mu_d}[V_{\pi_d,\pi_{d+1},\ldots,\pi_D}(s)] \geq \arg\max_{\pi\in\Pi} E_{s\sim\mu_d}[V_{\pi,\pi_{d+1}\ldots\pi_D}(s)] - \epsilon$$

Then for all $\pi_{\text{ref}} \in \Pi^D$,

$$V_\pi(s_0) \geq V_{\pi_{\text{ref}}}(s_0) - D\epsilon - Dd_{\text{var}}(\mu, \mu_{\pi_{\text{ref}}})$$

where μ is the baseline distribution over space-index states (without the time component) provided to SI-PSDP, d_{var} denotes the average variational distance, and $\mu_{\pi_{\text{ref}}}$ is the state distribution induced by π_{ref}.

Logic and Computer Science

Logic is the mathematics of Computer Science

as

Calculus is the mathematics of Physics.

This Course

Target Audience

University students

Talented High School Students

Interested professionals

Prerequisites

Sets and Set Operations (union, intersection, complement)

Symbolic Manipulation (e.g. high school algebra)

Elements of Logic

Elements of Logic:

- Logical Sentences
- Logical Entailment
- Logical Proofs

Symbolic Logic

- Problems with Natural Language
- Benefits of Formal Language

Course Logistics

- Logics covered in the course
- Advanced Logics *not* covered in the course

Logical Sentences

Friends

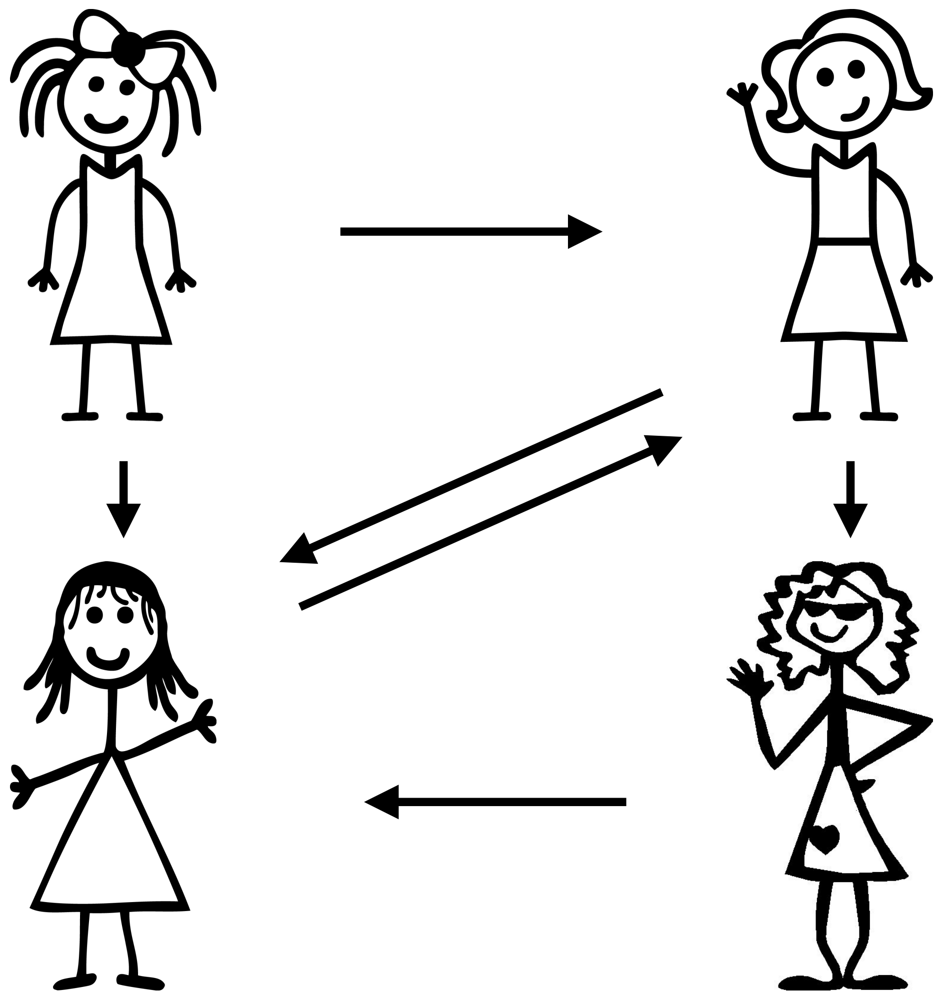

Friends

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

Friends

	Abby	Bess	Cody	Dana
Abby	✓		✓	
Bess		✓		✓
Cody	✓		✓	
Dana		✓		✓

Possible Worlds

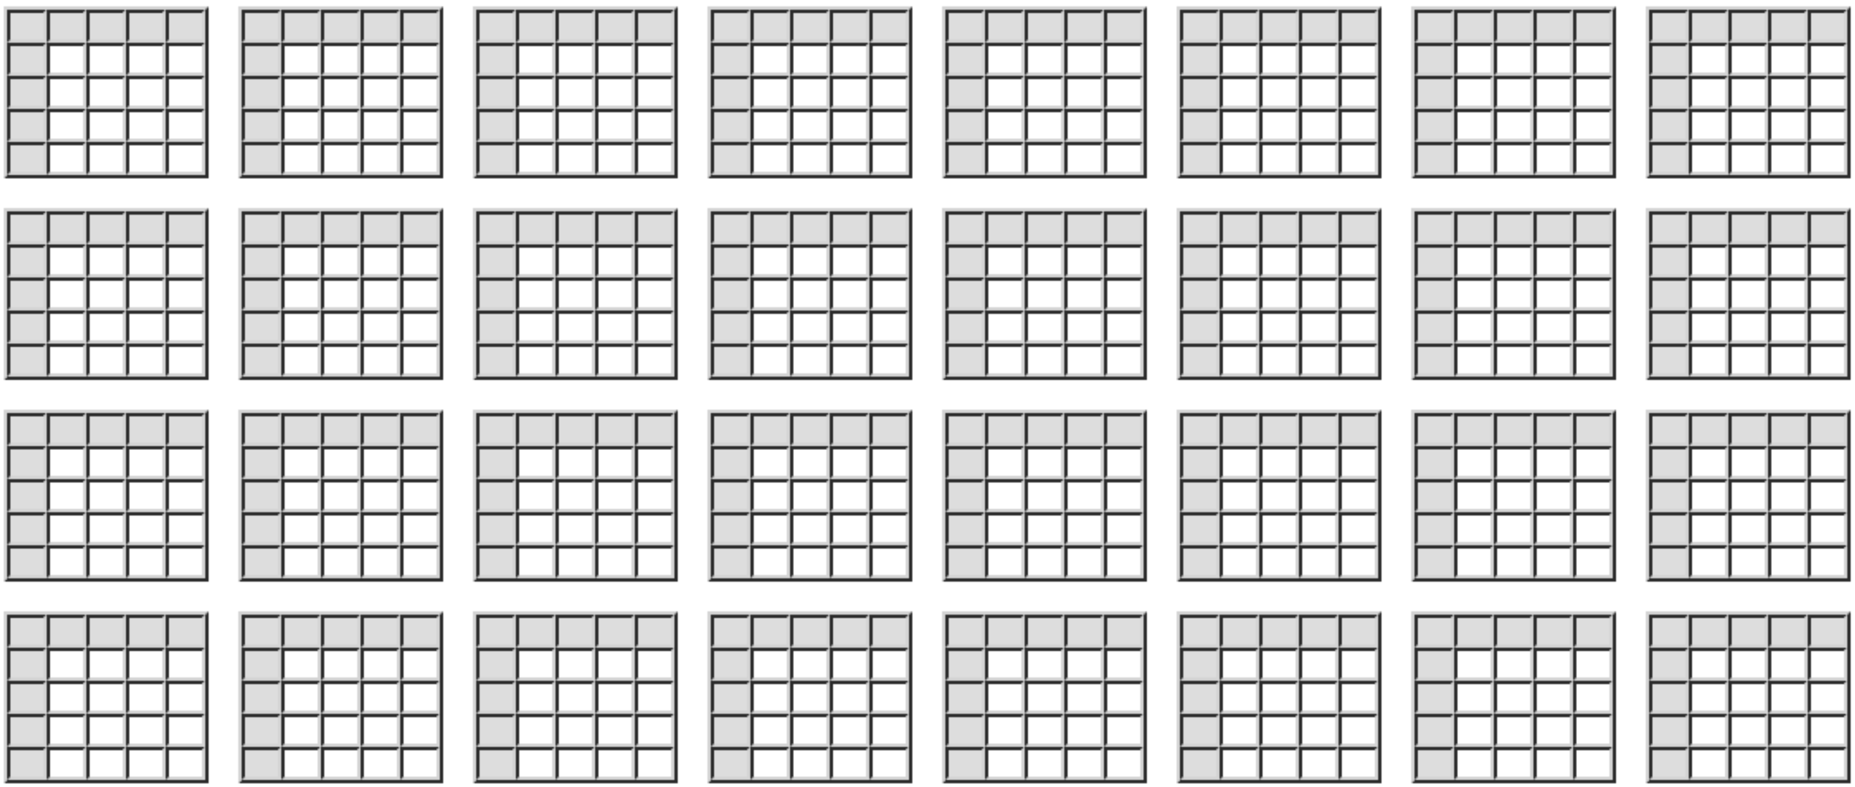

One Specific World

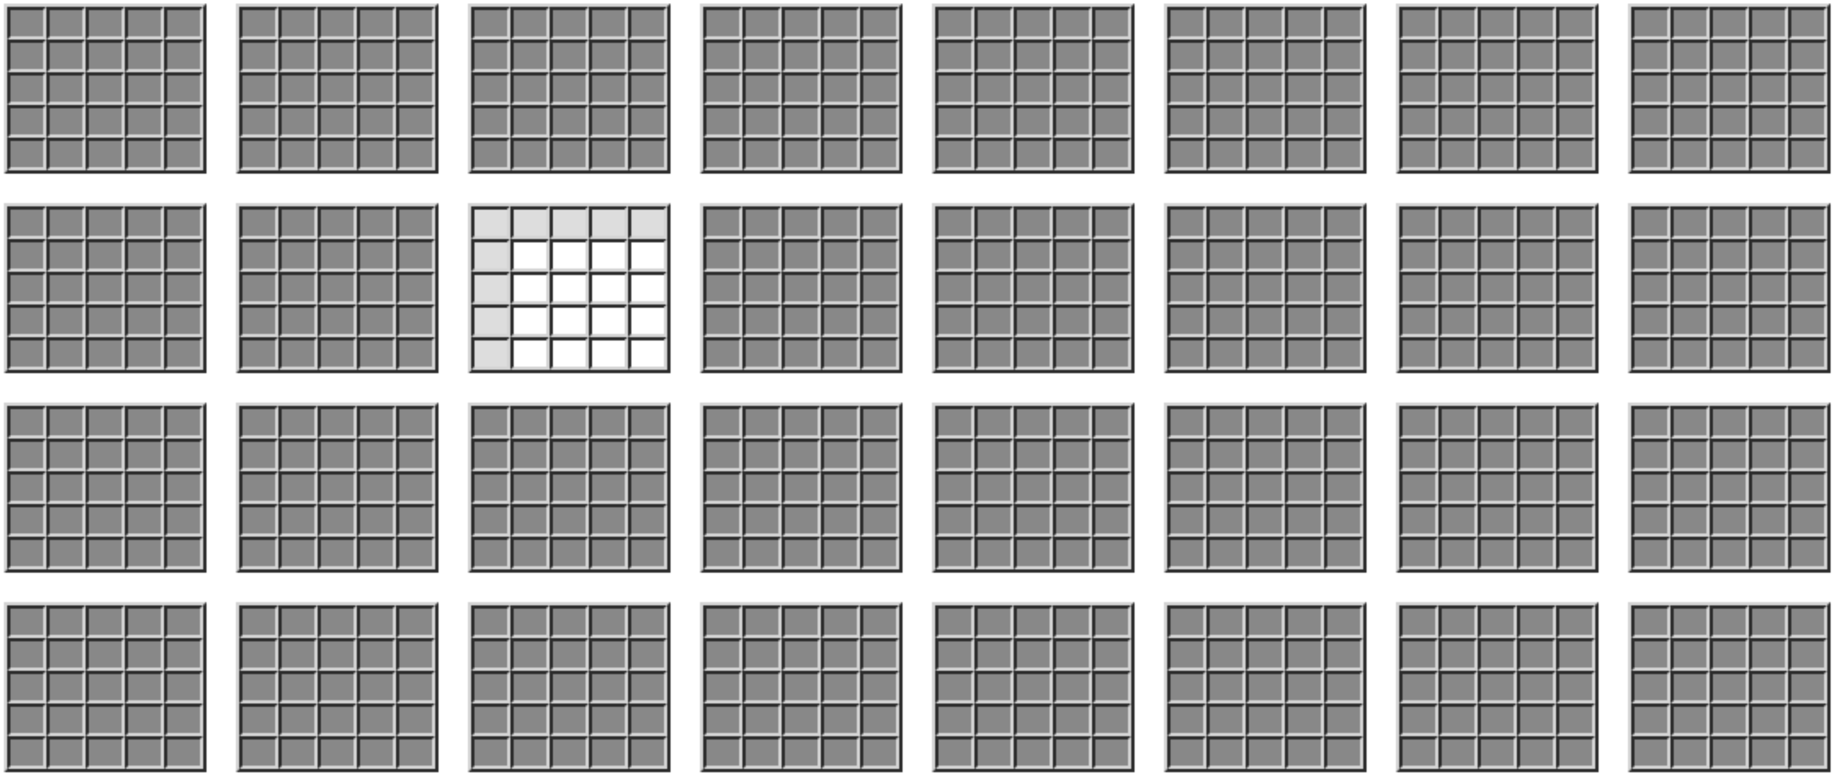

Logical Sentences

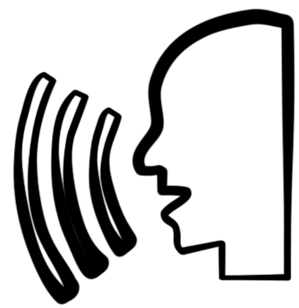

Logical Language

Dana likes Cody.

*Abby does **not** like Dana.*

*Dana does **not** like Abby.*

*Bess likes Cody **or** Dana.*

*Abby likes **everyone** that Bess likes.*

*Cody likes **everyone** who likes her.*

***No one** likes herself.*

After One Sentence

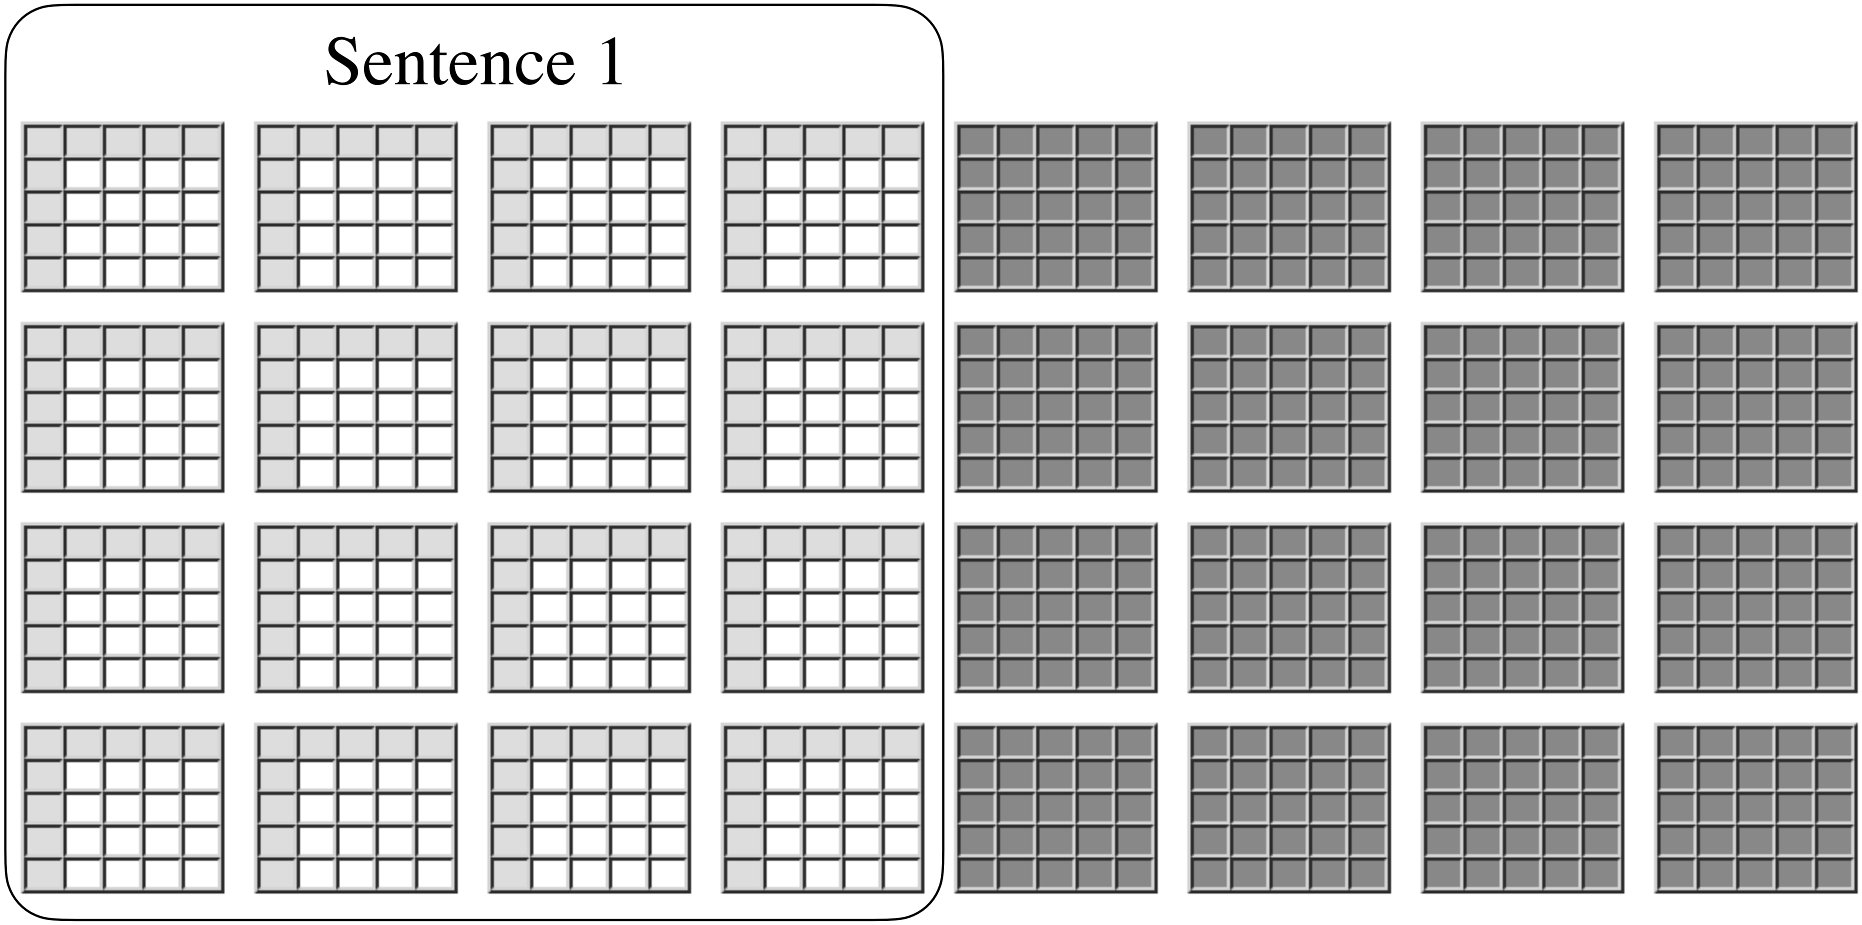

After Two Sentences

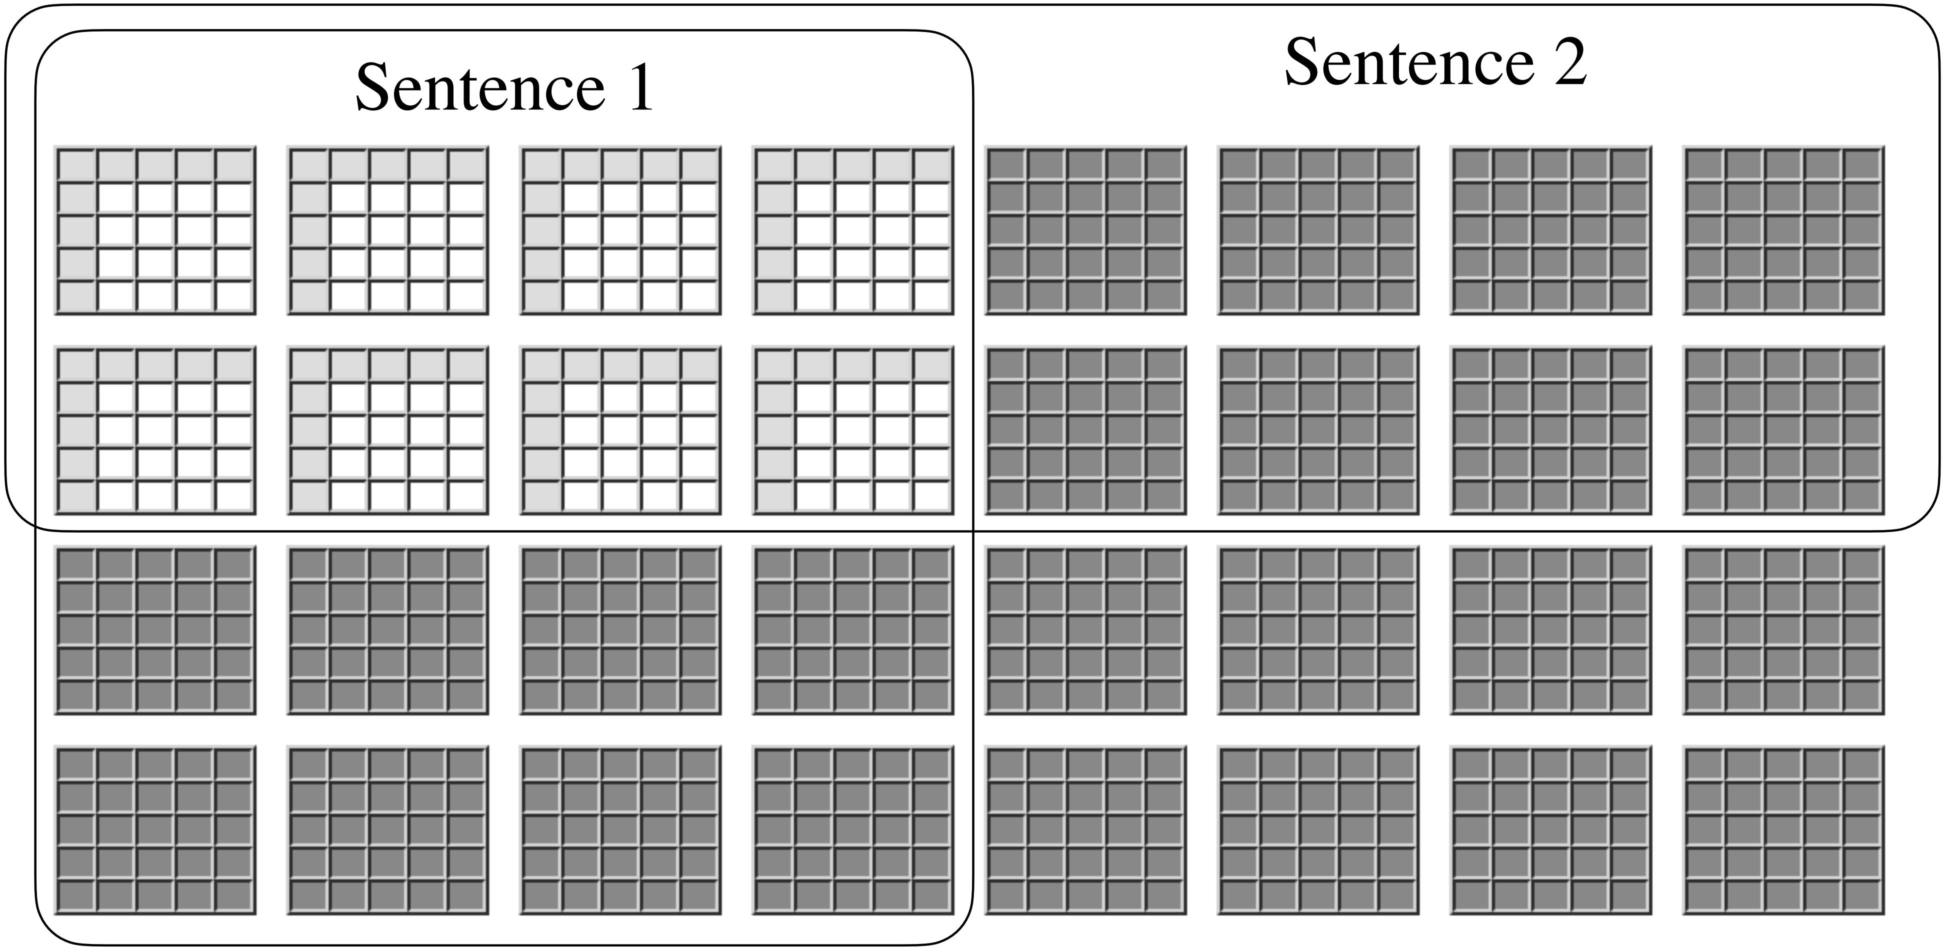

One Specific World

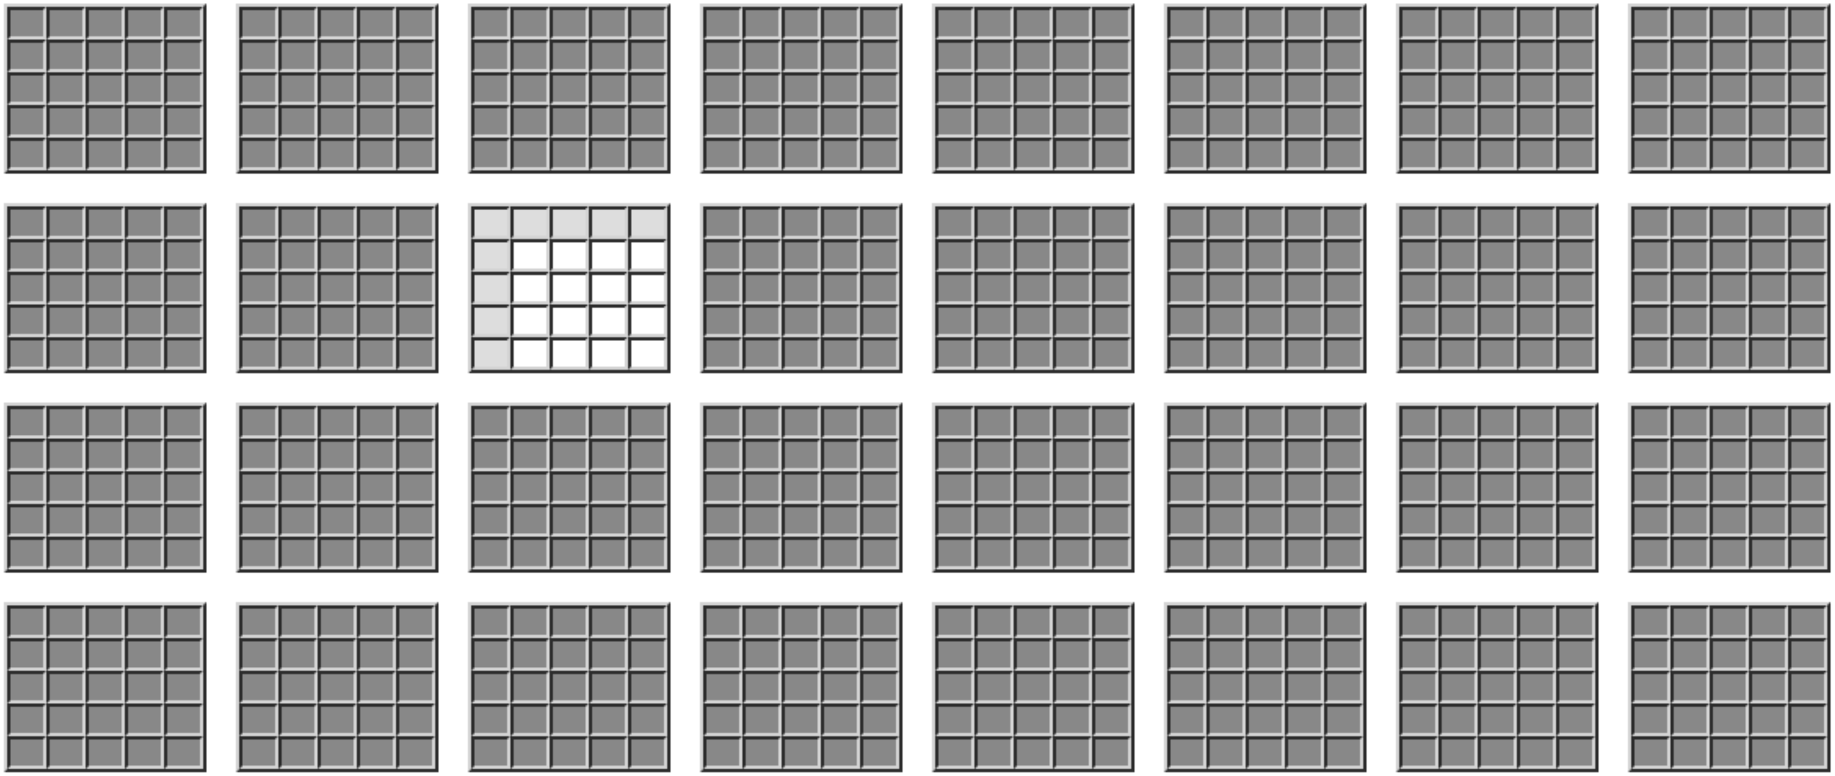

Complete Information

Premises:

Dana likes Cody.
Abby does not like Dana.
Dana does not like Abby.
Abby and Dana do not like Bess.
Bess likes Cody or Dana.
Abby likes everyone that Bess likes.
Cody likes everyone who likes her.
No one likes herself.

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

Conclusions:

Does Bess likes Cody?
Does Abby like Bess?
Does Dana like Bess?

Logical Entailment

Complete Information

Premises:

Dana likes Cody.
Abby does not like Dana.
Dana does not like Abby.
Abby and Dana do not like Bess.
Bess likes Cody or Dana.
Abby likes everyone that Bess likes.
Cody likes everyone who likes her.
No one likes herself.

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

Conclusions:

Does Bess likes Cody?
Does Abby like Bess?
Does Dana like Bess?

Logical Conclusions

Premises:

Dana likes Cody.
Abby does not like Dana.
Dana does not like Abby.
Bess likes Cody or Dana.
Abby likes everyone that Bess likes.
Cody likes everyone who likes her.
No one likes herself.

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

	Abby	Bess	Cody	Dana
Abby		✓	✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana		✓	✓	

	Abby	Bess	Cody	Dana
Abby		✓	✓	
Bess			✓	
Cody	✓	✓		✓
Dana		✓	✓	

In Logic, we typically need to deal with ***incomplete information****.*

Incomplete Information

Premises:

Dana likes Cody.
Abby does not like Dana.
Dana does not like Abby.
Bess likes Cody or Dana.
Abby likes everyone that Bess likes.
Cody likes everyone who likes her.
No one likes herself.

Questions:

Does Bess likes Cody?
Does Abby like Bess?
Does Dana like Bess?

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

	Abby	Bess	Cody	Dana
Abby		✓	✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana		✓	✓	

	Abby	Bess	Cody	Dana
Abby		✓	✓	
Bess			✓	
Cody	✓	✓		✓
Dana		✓	✓	

Logical Entailment

A set of premises *logically entails* a conclusion if and only if *every* world that satisfies the premises also satisfies the conclusion.

Logical Conclusions

Premises:

Dana likes Cody.
Abby does not like Dana.
Dana does not like Abby.
Bess likes Cody or Dana.
Abby likes everyone that Bess likes.
Cody likes everyone who likes her.
No one likes herself.

Questions:

Does Bess likes Cody? *Yes*
Does Abby like Bess? *No*
Does Dana like Bess? *Maybe*

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

	Abby	Bess	Cody	Dana
Abby		✓	✓	
Bess			✓	
Cody	✓	✓		✓
Dana			✓	

	Abby	Bess	Cody	Dana
Abby			✓	
Bess			✓	
Cody	✓	✓		✓
Dana		✓	✓	

	Abby	Bess	Cody	Dana
Abby		✓	✓	
Bess			✓	
Cody	✓	✓		✓
Dana		✓	✓	

Naive Model Checking Algorithm

Iterate through *all* possible worlds. For every world that satisfies the premises, check if it satisfies the conclusion.

Problem with Naive Model Checking

Lots of Worlds (sometimes infinitely many)

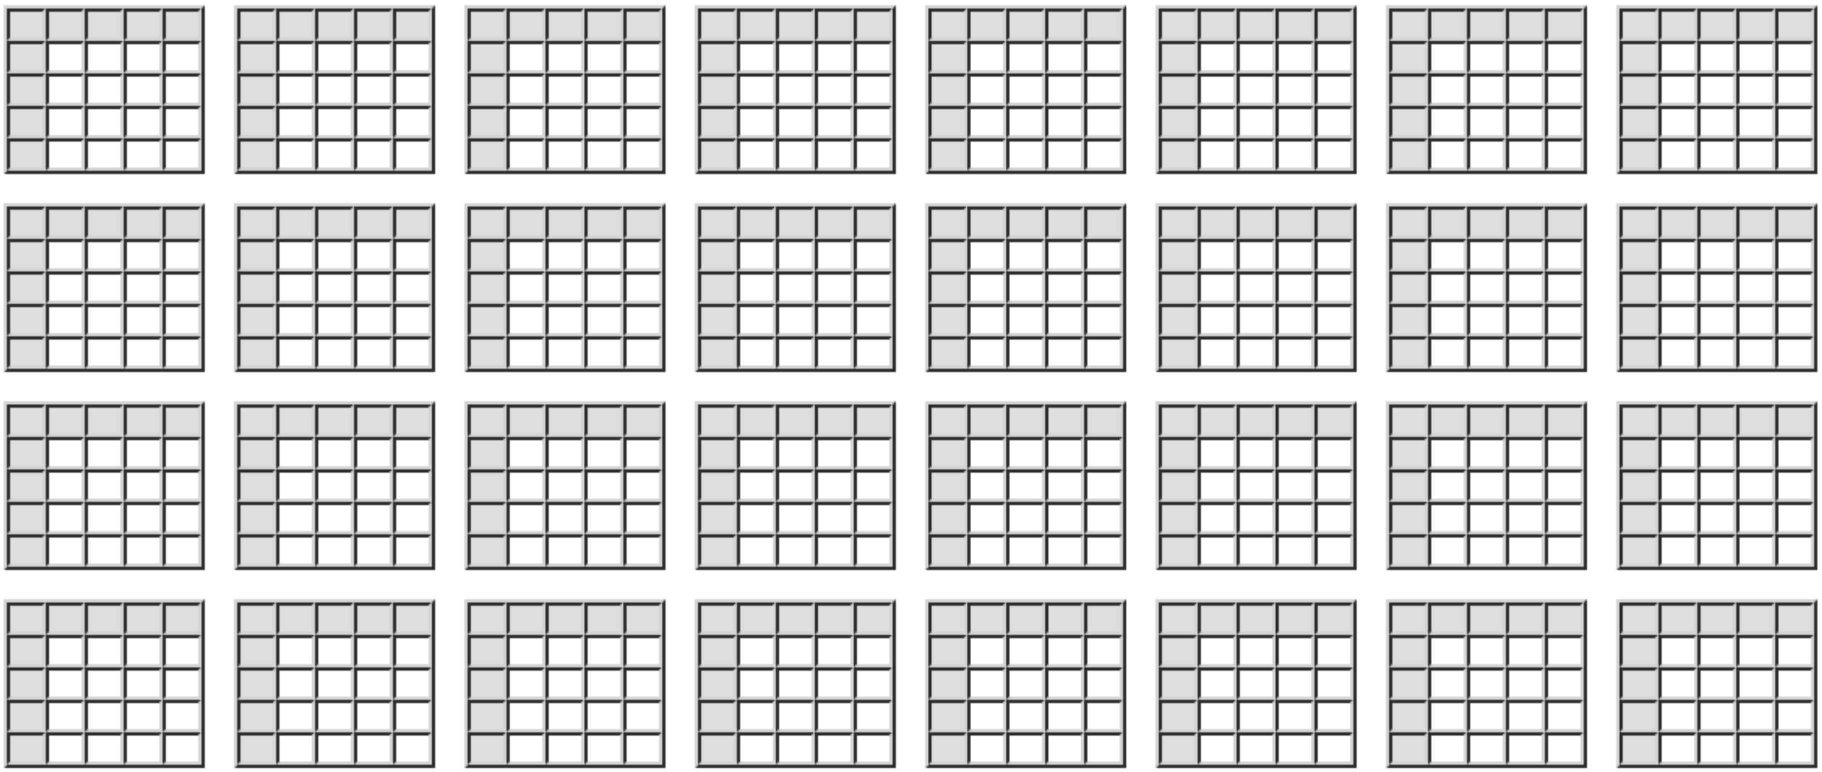

Model Checking is like solving polynomial equations by enumerating all possible values for the variables.

Logical Proofs

Symbolic Manipulation in Algebra

$$x - 3y = 0$$
$$x + y = 12$$

$$\downarrow$$

$$-4y = -12$$

$$\downarrow$$

$$y = 3$$
$$x = 9$$

Symbolic Manipulation in Logic

Dana likes Cody.
Abby does not like Dana.
Dana does not like Abby.
Abby and Dana do not like Bess.
Bess likes Cody or Dana.
Abby likes everyone that Bess likes.
Cody likes everyone who likes her.
No one likes herself.

↓

Bess likes Cody.
Abby does not like Bess.

Rules of Inference

A *rule of inference* is a reasoning pattern consisting of some premises and some conclusions.

A *proof* is a sequence of sentences in which every sentence is either a premise or the result of applying a *rule of inference* to earlier elements of the sequence.

Sample Rule of Inference

All *of Abby's friends are Bess's friends.*
All *of Bess's friends are Cody's friends.*
Therefore, ***all*** *of Abby's friends are Cody's friends.*

Sample Rule of Inference

All *Accords are Hondas.*
All *Hondas are Japanese.*
Therefore, ***all*** *Accords are Japanese.*

Sample Rule of Inference

All *borogoves are slithy toves.*
All *slithy toves are mimsy.*
Therefore, ***all*** *borogoves are mimsy.*

General Rule of Inference

All *x are y.*
All *y are z.*
Therefore, ***all*** *x are z.*

Bertrand Russell

Logic "may be defined as the subject in which we never know what we are talking about nor whether what we are saying is true."

- Bertrand Russell

Unsound Rule of Inference

All *x are y.*
Some *y are z.*
Therefore, ***some*** *x are z.*

No! No!! No!!!

Using Unsound Rule of Inference

All *Toyotas are Japanese cars.*
Some *Japanese cars are made in America.*
Therefore, ***some*** *Toyotas are made in America.*

Sometimes produces a result that *happens* to be true.

Using Unsound Rule of Inference

All *Toyotas are cars.*
Some *cars are Porsches.*
Therefore, ***some*** *Toyotas are Porsches.*

Sometimes produces a result that *happens* to be false.

Deduction

A rule of inference is *sound* if and only if the conclusion is true whenever the premises are true.

The application of sound rules of inference is called *deduction*.

Induction

Induction is reasoning from the specific to the general.

> *I have seen 1000 black ravens.*
> *I have never seen a raven that is not black.*
> *Therefore, every raven is black.*

If induction is incomplete, it is not necessarily sound (but it can be useful).

Induction versus Deduction

Induction is the basis for **Science** (and machine learning)
Deduction is the subject matter of **Logic**.

Science aspires to discover / propose **new** knowledge.
Logic aspires to apply and/or analyze **existing** knowledge.

Scientific theories *may* be false (even if premises true).
Logical conclusions *must* be true (if premises true).

Niels Bohr to Albert Einstein

“You are not thinking; you are just being logical.”

Entailment versus Provability

A set of premises *logically entails* a conclusion if and only if every world that satisfies the premises satisfies the conclusion.

A conclusion is *provable* from a set of premises if and only if there is a finite sequence of sentences in which every element is either a premise or the result of applying a *sound* rule of inference to earlier members in the sequence.

Soundness and Completeness

As we shall see, for well-behaved logics, logical entailment and provability are identical - a set of premises **logically entails** a conclusion *if and only if* the conclusion is **provable** from the premises.

This is a very big deal.

Symbolic Logic

Logical Sentences

Dana likes Cody.

*Abby does **not** like Dana.*

*Dana does **not** like Abby.*

*Bess likes Cody **or** Dana.*

*Abby likes **everyone** that Bess likes.*

*Cody likes **everyone** who likes her.*

***Everyone** likes herself.*

Complexity of Natural Language

One grammatically correct sentence:

The cherry blossoms in the spring.

Another grammatically correct sentence:

The cherry blossoms in the spring sank.

Grammatical Ambiguity

There's a girl in the room with a telescope.

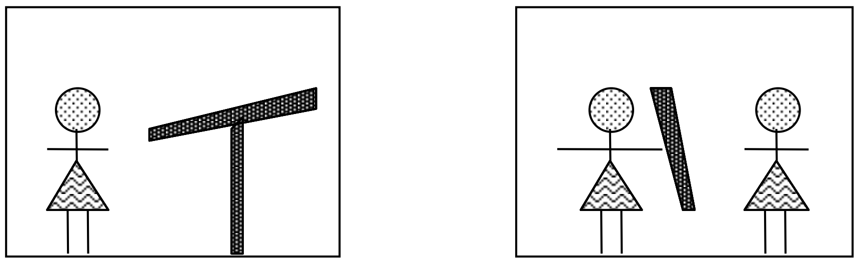

Newseum Headlines

Crowds Rushing to See Pope Trample 6 to Death

Newseum Headlines

Crowds Rushing to See Pope Trample 6 to Death

Scientists Grow Frog Eyes and Ears

Newseum Headlines

Crowds Rushing to See Pope Trample 6 to Death

Scientists Grow Frog Eyes and Ears

Fried Chicken Cooked in Microwave Wins Trip

Newseum Headlines

Crowds Rushing to See Pope Trample 6 to Death

Scientists Grow Frog Eyes and Ears

Fried Chicken Cooked in Microwave Wins Trip

British Left Waffles on Falkland Islands

Newseum Headlines

Crowds Rushing to See Pope Trample 6 to Death

Scientists Grow Frog Eyes and Ears

Fried Chicken Cooked in Microwave Wins Trip

British Left Waffles on Falkland Islands

Indian Ocean Talks

Newseum Headlines

Crowds Rushing to See Pope Trample 6 to Death

Scientists Grow Frog Eyes and Ears

Fried Chicken Cooked in Microwave Wins Trip

British Left Waffles on Falkland Islands

Indian Ocean Talks

Misteaks in Print

Residents report that a hole was cut in the fence surrounding a nudist colony. Police are ***looking into*** *it.*

Doug Lenat's Logic

Champagne is better than beer.
Beer is better than soda.
Therefore, *champagne is better than soda.*

X is better than Y.
Y is better than Z.
Therefore, *X is better than Z.*

Bad sex is better than nothing.
Nothing is better than good sex.
Therefore, *bad sex is better than good sex.*
Really?

Logistics

Elements of Logic

Logical Language

Definitions: *A triangle is a polygon with three sides*.

Constraints: *Parents are older than their children*.

Partial Information: *Abby likes one of Cody or Dana*.

Model Checking

Truth tables

Logic grids

Symbolic Manipulation

Formula transformations

Proofs

"Metalevel" Concepts and Analysis

Properties of Logical Sentences

Validity, Contingency, Unsatisfiability

Relationships Between Sentences

Equivalence, Entailment, Consistency

Computational Analysis

Soundness, Completeness, Decidability

Logical Extensions

Language

Probabilities

Metaknowledge - knowledge about knowledge

Paradoxes, e.g. *This sentence is false*.

Reasoning

Negation as Failure - *knowing not* versus *not knowing*

Induction, Abduction, Analogical Reasoning

Paraconsistent Reasoning - reasoning with inconsistency

Multiple Logics

Propositional Logic (logical operators)

If *it is raining* ***and*** *it is cold,* ***then*** *the ground is wet.*

Relational Logic (variables and quantifiers)

If ***x*** *is younger than* ***y****, then* ***y*** *is older than* ***x****.*

Term Logic (compound terms)

$\{a, b\}$ *is a subset of* $\{a, b, c\}$.

Schedule

Week	Tuesday	Thursday
1	September 24 Introduction	September 26 Propositional Logic
2	October 1 Propositional Analysis	October 3 Direct Proofs
3	October 8 Natural Deduction	October 10 Refutation Proofs
4	October 15 Review	October 17 Quiz 1
5	October 22 Relational Logic	October 24 Relational Analysis
6	October 29 Fitch Proofs	October 31 Review
7	November 5 No Class	November 7 Quiz 2
8	November 12 Term Logic	November 14 Induction
9	November 19 Equality	November 21 Review
	Thanksgiving Week	
10	December 3 No Class	December 5 Quiz 3
	December 10 Optional Final	

Numerical Scores

Quizzes

Propositional Logic (Week 4)
Relational Logic (Week 7)
Term Logic (Week 10)

Final Exam

Optional
Three one-hour exams, one on each topic

Overall Score

40% - Propositional Logic
30% - Relational Logic
30% - Term Logic
We will use the higher of quiz score and exam score

Letter Grades

Letter Grade

Based on Quiz and Exam Scores (see above)

No curve - i.e. independent of number of students

A, B, C distributed uniformly over 70% - 100%

Discretionary Extra Credit

class attendance, Ed forum, puzzles, ...

taken into account at boundaries of letter grades

discretionary

Secret Word

"Fallacy"

Secret Word

spoken, not written

Course Website

http://cs157.stanford.edu

Hints on How to Take the Course

Materials of the Course

- Lectures
- Textbook / Lessons
- Exercises
- Puzzles
- Tools

Read the notes.

Do the exercises.

Do the exercises!

Do the exercises!!

*Learn **actively**.*

Ed Discussion

- Read discussion
- Post questions
- Answer questions

Working in groups is okay / recommended!!

Biggest Mistake

Quiz 1 Mean Score

80.8

Value of Theory

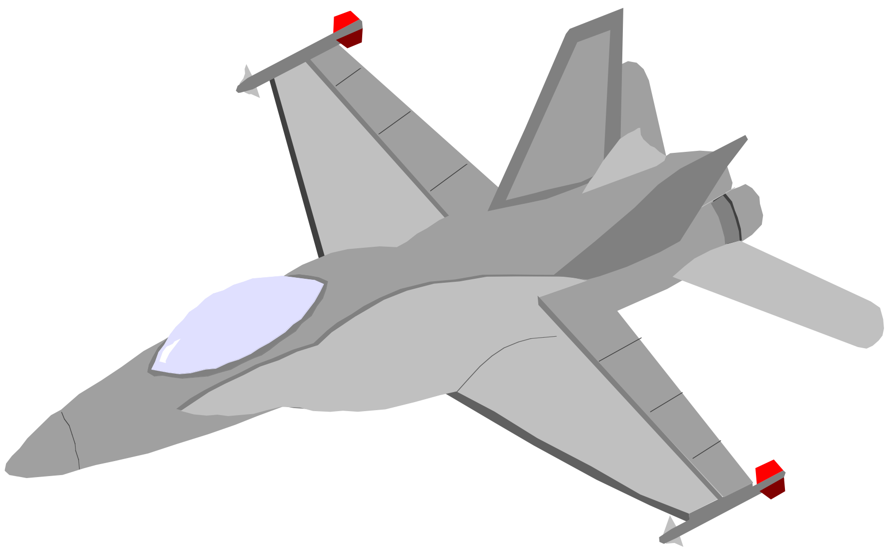

Value of Practice