Capturing and Reusing Experience

2018-04-17

Andrew Ng – The State of AI (December 15, 2017)

- "99% of the economic value created by AI today is through one type of AI: which is learning a mapping A → B, or input to output maps"
  - Falls under category of supervised learning
- Other types (ordered falloff)
  - Transfer learning
  - Unsupervised learning
  - Reinforcement learning

Input	Output
Picture	Is it you? (0/1)
Loan application	Will the applicant repay the loan? (0/1)
Online: (Ad, User)	Will you click? (0/1)
Voice input	Text transcript
English	French
Car: image, radar/lidar	Positions of other cars

https://www.youtube.com/watch?v=NKpuX_yzdYs

ML & Function Approximation

- Arthur Samuel (1959), checkers:
  - Machine Learning: Field of study that gives computer the ability to learn without being explicitly programmed (recognize board patterns that lead to wins/losses)
- Dr. Mark Riedl, what is a deep neural network:
  - Function approximation via stacks of complete acyclic weighted bipartite graphs
- Target function *f* may be known or unknown
  - **If known**, we seek functions that trade accuracy for desirable properties (inexpensive computation, continuity, integral and limit values, etc.)
  - **If unknown**, only a set of points of the form (x, *f*(x)) is provided.

https://www.youtube.com/watch?v=UzxYlbK2c7E

- Several techniques for approximating *f*
  - If *f* is an operation on the real numbers, techniques of **interpolation, extrapolation, regression analysis, and curve fitting** can be used
  - If the codomain (range or target set) of g is a finite set, one is dealing with a **classification** problem instead.
- To some extent, the different problems (regression, classification, fitness approximation) have received a **unified treatment in statistical learning theory, where they are viewed as supervised learning problems**.
- The multilayer perceptron (a class of feedforward ANN) is **a universal function approximator**, as proven by the universal approximation theorem

https://en.wikipedia.org/wiki/Function_approximation

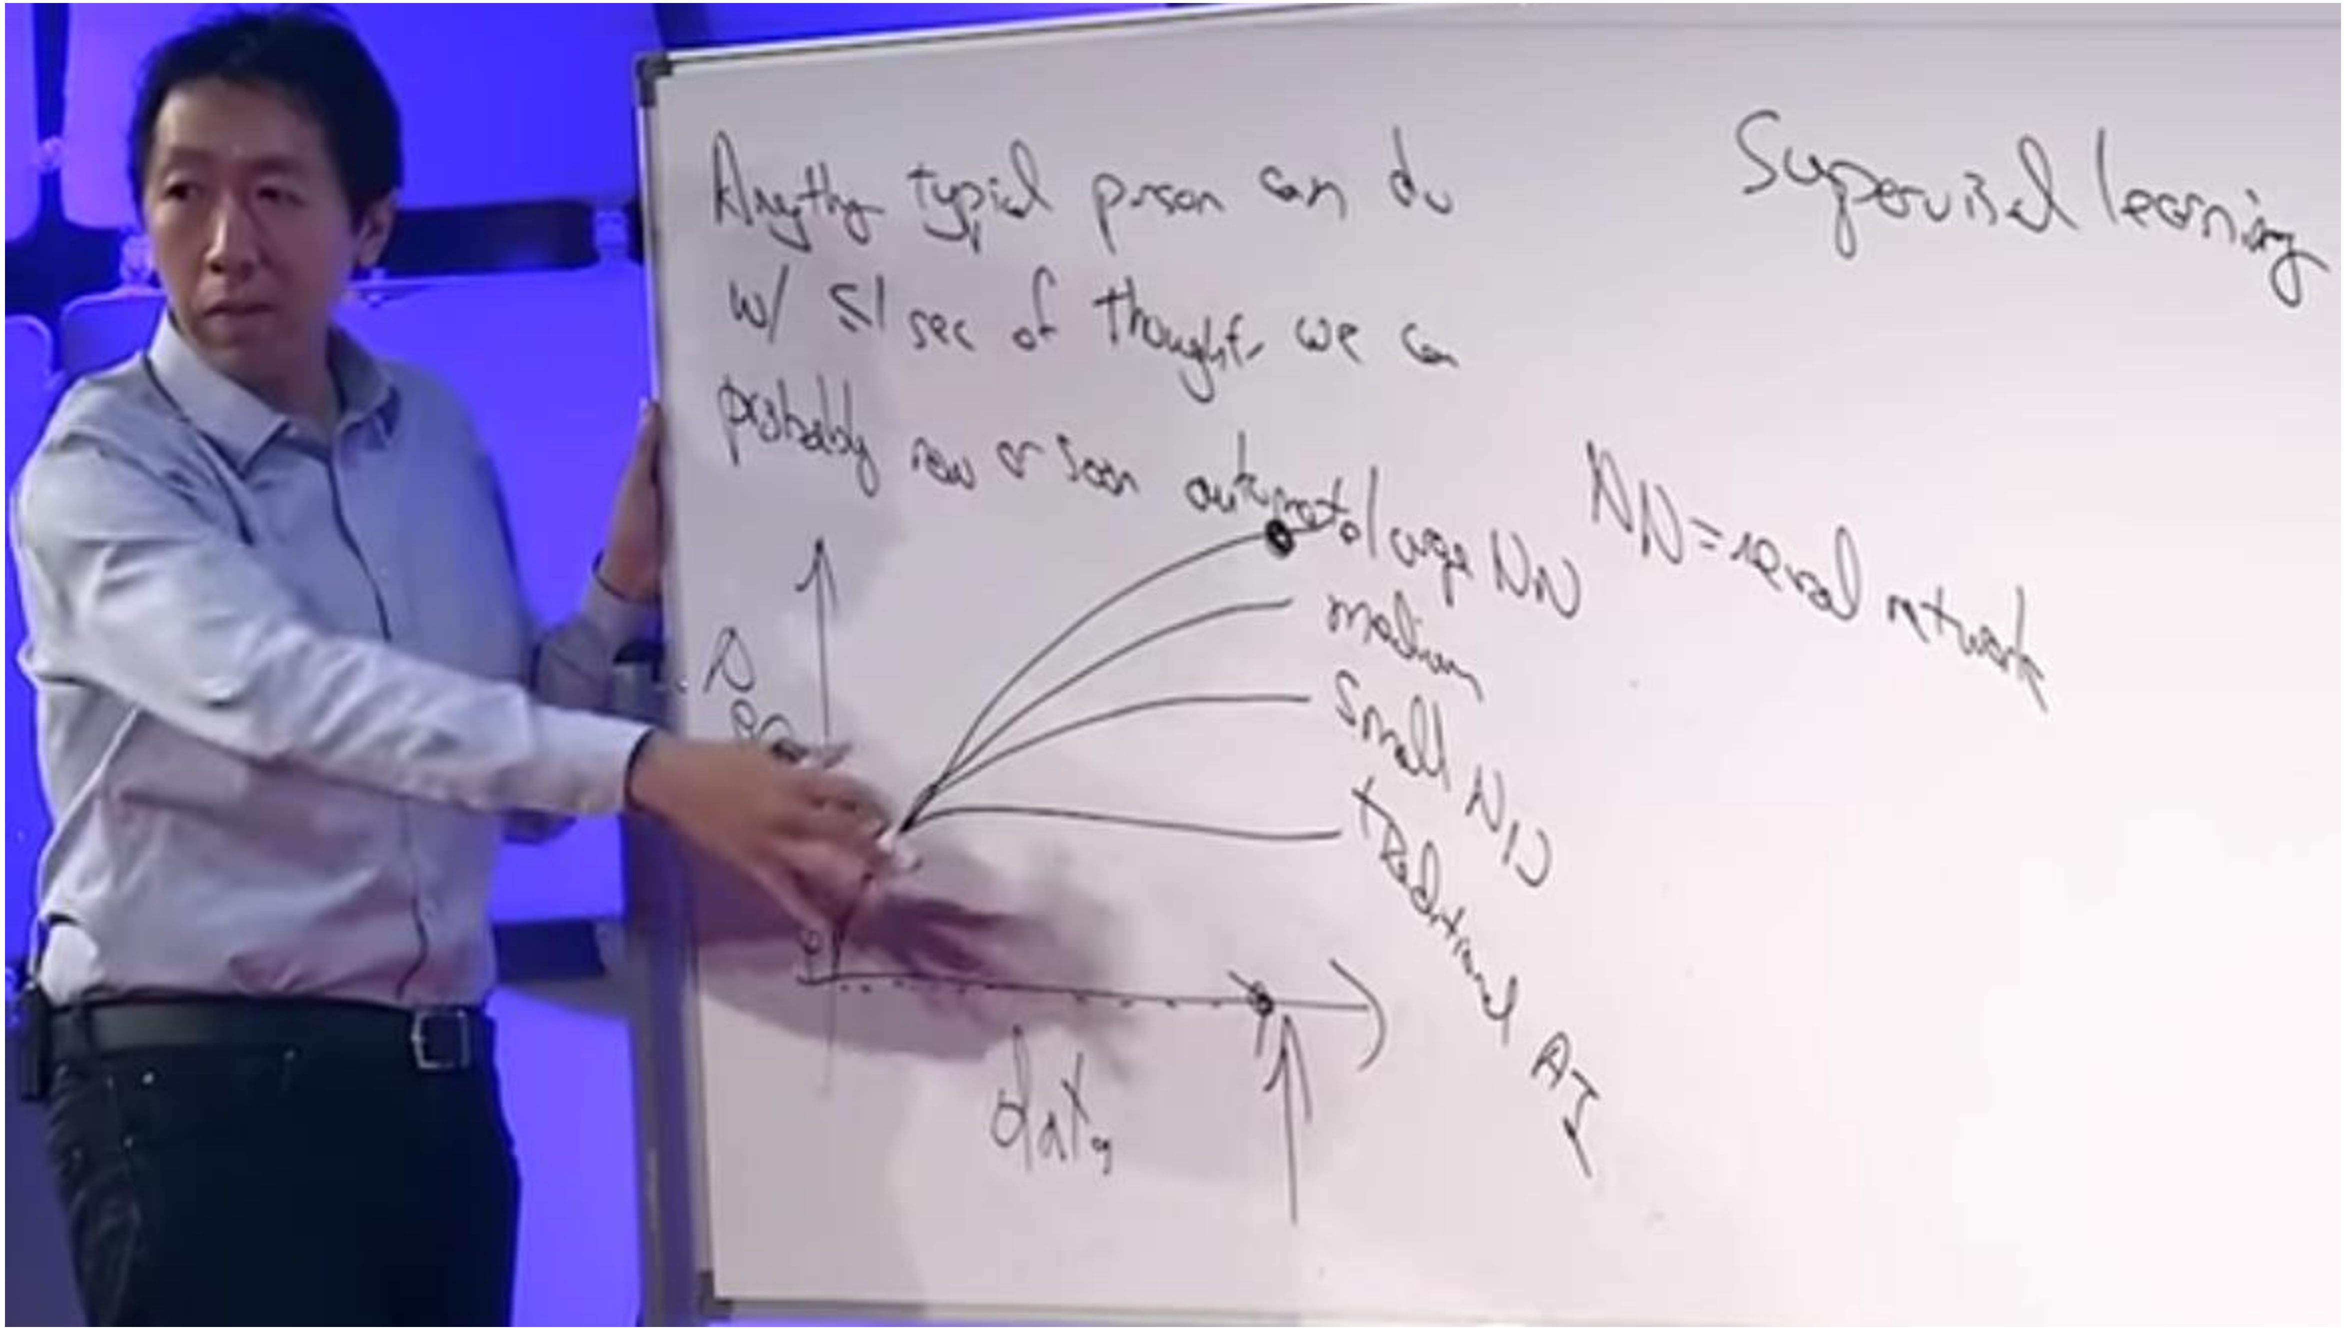
Anything typical person can do
w/ ≤1 sec of thought, we can
probably now or soon automate
Supervised learning
NN = neural network
large NN
medium
small NN
traditional AI
data

Decision Learning in M&F

- See Millington & Funge
  - 7.4 through 7.8
- Naive Bayes classification (7.5)
  - Try first, simple to implement
  - Good baseline
- Decision tree learning (7.6)
  - Output is interpretable
- Reinforcement learning (7.7)
  - “hot topic in game AI”
  - “a good starting point is the Q-learning algorithm. Q-learning **is simple to implement**, has been widely tested on non-game applications, and can be **tuned without a deep understanding of its theoretical properties**” (p631)
- Neural Networks (7.8)
  - “Very little neural network theory is applicable to games, however” (p646)

Decision Making Questions

1. How can we describe decision making?
2. What do the algorithms we’ve seen share?
3. What are the dimensions we tend to assess?

PCG Questions

1. PCG can be used to p____ or a____ game aspects
2. F(player model, designer constraints, instance) → fitness

RL Questions

Q1: What is the state transition function? Do we need it as input for Q-learning?

Q2: For the multi-armed bandit problem, which statement below best summarizes why we can’t just use exploitation or exploration?

(A.) Exploration means it will take a long time to find the optimal solution

(B.) Both will minimize our total regret

(C.) Pure exploration doesn’t learn, pure exploitation learns too fast

(D.) Pure exploration means we can never do better than chance, pure exploitation can get an agent caught in a local maximum strategy

No free lunch

“One of the greatest challenges in applying reinforcement learning to real-world problems is deciding **how to represent and store value functions and/or policies**. Unless the state set is finite and small enough to allow exhaustive representation by a lookup table […] **one must use a parameterized function approximation scheme**. […]

Most successful applications of reinforcement learning **owe much to sets of features carefully handcrafted based on human knowledge** and intuition about the specific problem to be tackled. […]

in all the examples of which we are aware, the most impressive demonstrations required the network's input to be **represented in terms of specialized features handcrafted for the given problem**”

Sutton & Barto, 2^{nd} ed, section 16.5

Millington 7.3

CAPTURING AND REUSING EXP: ACTION PREDICTION

Action Prediction

- Guess what player will do next
  - E.g. waypoint, weapon, tactic, cover point, melee
  - Make more realistic, challenging (helpful) NPC
  - Can do with little observation
  - Can transfer from other players
- Humans bad at random (Psychology). Furthermore…
  - “We have shared characteristics that run so deep that learning to anticipate one player’s actions can often lead to better play against a completely different player.”

Naïve Algorithm

- Predict using raw probability
  - Keep a tally, use to predict
  - Pro
    - Easy, fast
    - Gives a lot of feedback to player
    - Can learn from many different players
  - Con
    - Player can “game” the system
    - Eventually can reduce to equal probabilities
- “Left or right” game
  - Object in either L or R hand
  - Another persons guesses hand
- Incremental update of average
  - Keep mean (m_{n-1}), and count (n)
  - $m_n = m_{n-1} + (1/n)(v_n – m_{n-1})$

String Matching

- Choice made several times
  - Encode as string “LRRLRLLLRRLRLRR”
  - Predict → find substring, return subsequent choice
  - Example: “RR”. What next?
  - Window size: 2
- Rarely implemented by matching against a string
  - Use a set of probabilities similar to the naïve algorithm

Prediction: *N*-Grams

- String matching + probabilities
  - N is window size + 1 (e.g. 3-gram from before)
  - Record Prob of each move for all windows
  - Must sum to 1
  - E.g. “LRRLRLLLRRLRLRR”

	..R	..L
LL	1/2	1/2
LR	3/5	2/5
RL	3/4	1/4
RR	0/2	2/2

Prediction: *N*-Grams

- String matching + frequencies
  - N is window size + 1 (e.g. 3-gram from before)
  - Record count of each move for all windows
  - Must sum to count
  - E.g. “LRRLRLLLRRLRLRR”

	..R	..L
LL (2)	1	1
LR (5)	3	2
RL (4)	3	1
RR (2)	0	2

Question

How do we choose the window size?

Window Size

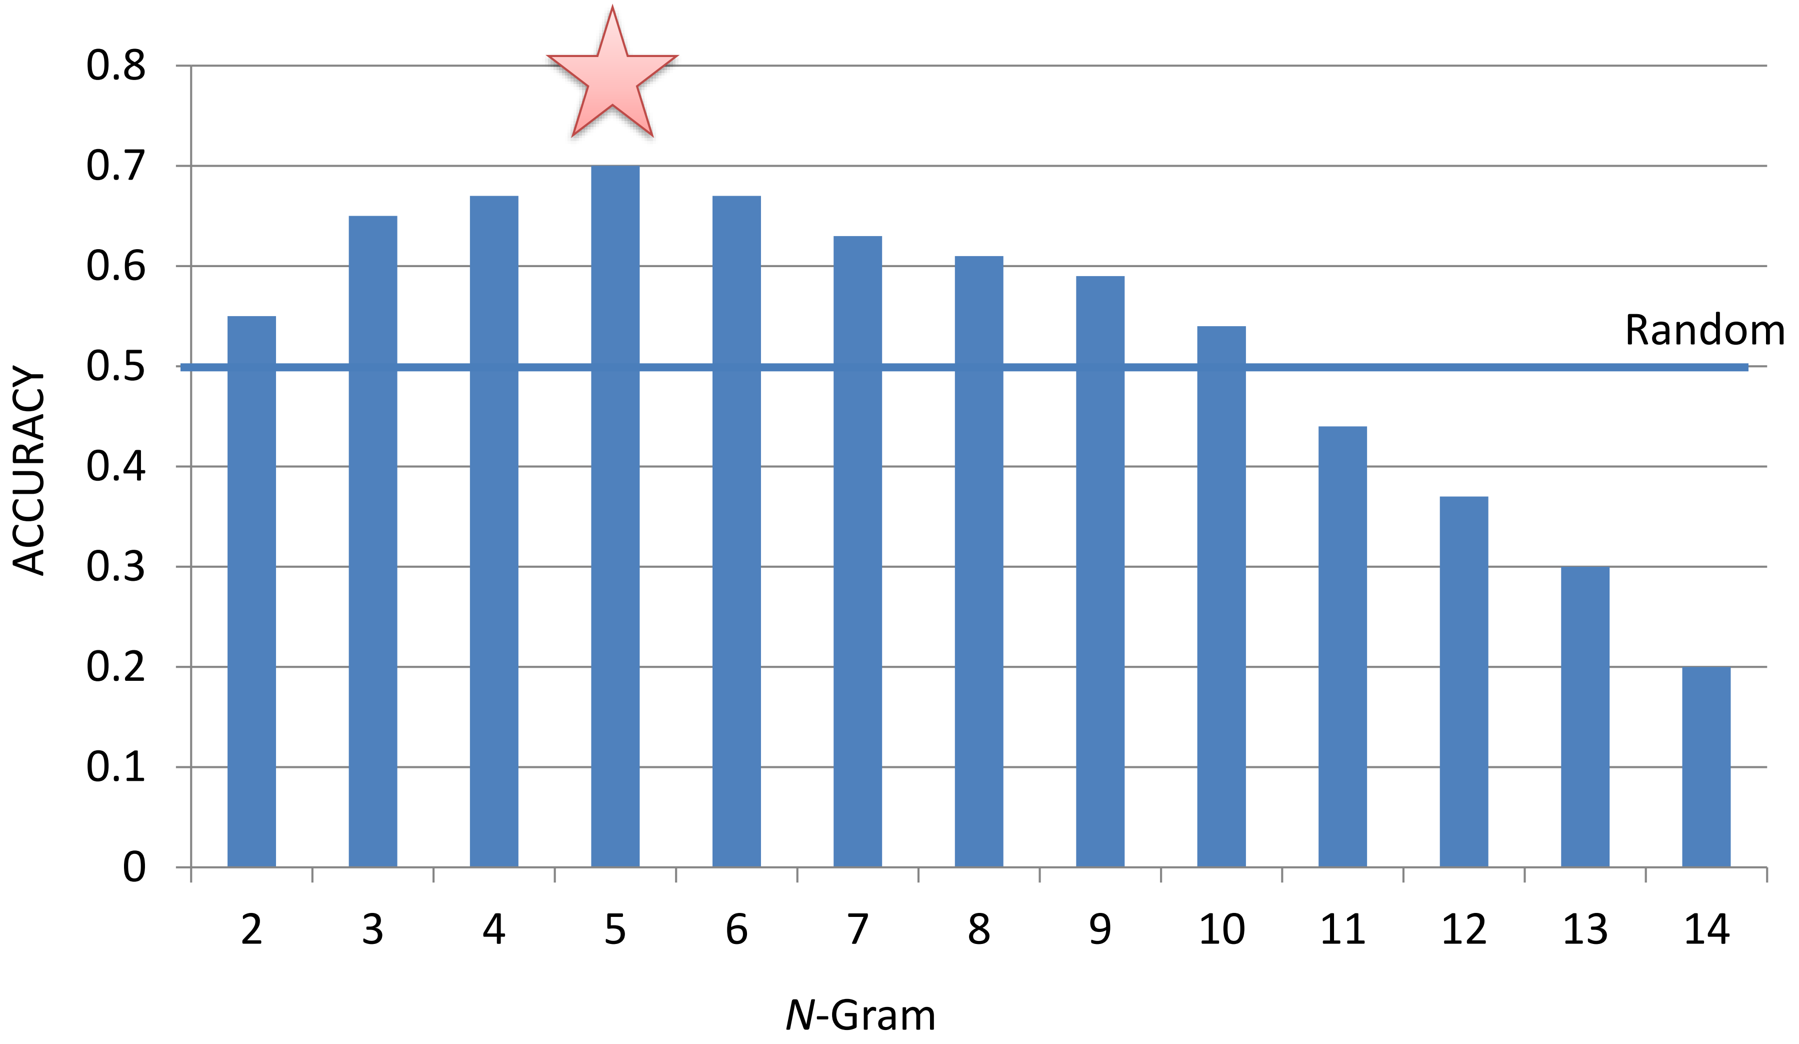

0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
ACCURACY
Random
2
3
4
5
6
7
8
9
10
11
12
13
14
N-Gram

Window Size

- Increase size helps initially, hurts later. Why?
  - Future actions predicted by *short* causal process
  - Similar to Markov assumption?
  - Psychology?
  - Degree of randomness in actions
    - ($\uparrow$ random $\downarrow$window)
- How to tune?

Hierarchical *N*-Grams

- Online learning approach
- Balances max predictive power and alg. perf.
  - Large window, better potential, slower coverage
- Essentially several parallel *N*-grams
  - E.g. Hierarchical 3-gram: 1, 2, and 3 gram
  - When prediction requested, look up window with
    - sufficient examples
    - highest predictive accuracy
  - What is sufficient number of examples?

N-gram summary

- Simple, effective prediction mechanism
- Synonymous with combo-based melee games
  - Can make unbeatable (no fun) AI
  - Often is intentionally weakened
- Many other uses
  - statistical analysis techniques (e.g. language)
  - [Weapon, location, unit] selection…

According to (Smith 2016), CBR has led **to more practical applications than any other AI family of techniques with the exception of expert systems and machine learning**. IBM's Watson system (Ferruci et al. 2010) is a famous example of the power of memory-based reasoning.
-- Ashok K. Goel and Belen Diaz-Agudo, AAAI-17. "What's Hot in Case-Based Reasoning"
https://www.aaai.org/ocs/index.php/AAAI/AAAI17/paper/viewFile/15041/14020

CAPTURING AND REUSING EXP: CASE-BASED REASONING

Videos: CBR in games

- Many Games (Tetris, Soccer, RTS, Poker, …)
  - http://youtu.be/-EPb-zxbEhw
- Xdomain (AAAI 2010 best video)
  - http://www.youtube.com/watch?v=fwYfkCu4mFI
- Imitation in soccer (AAAI 2008 winner)
  - http://www.youtube.com/watch?v=zNjyXLWVSWI
- Football (Casey’s Quest)
- The Killer Groove | The Shadow AI of Killer Instinct (@~9:00)
  - https://www.youtube.com/watch?v=Etj5ykJugwU

Sources

- **Many(!) slides from Dr. Hector Munoz-Avila**
- cbrwiki.fdi.ucm.es/
- www.iiia.csic.es/People/enric/AICom.html
- www.cse.lehigh.edu/~munoz/CSE335/
- www.aic.nrl.navy.mil/~aha/slides/
- www.csi.ucd.ie/users/barry-smyth
- www.csi.ucd.ie/users/lorraine-mcginty

How would we solve the traffic problem given known techniques?

- Planning approach
  - Replan. As expensive as whatever planning approach you use
- MDP/RL
  - Already have a perfect optimal strategy to take in every state
  - But then we would have had to anticipate this traffic, and that’s not always possible

Alternative

Case-based reasoning

1. Analyze: Identify the current situation
2. Retrieve: The closest case(s) in our case-base
3. Adapt: Change the case(s) to match the current situation
4. Test: Ensure that these cases match our
5. Repair: Make any changes from test (test again if so)
6. Store: Store this new solution in case base

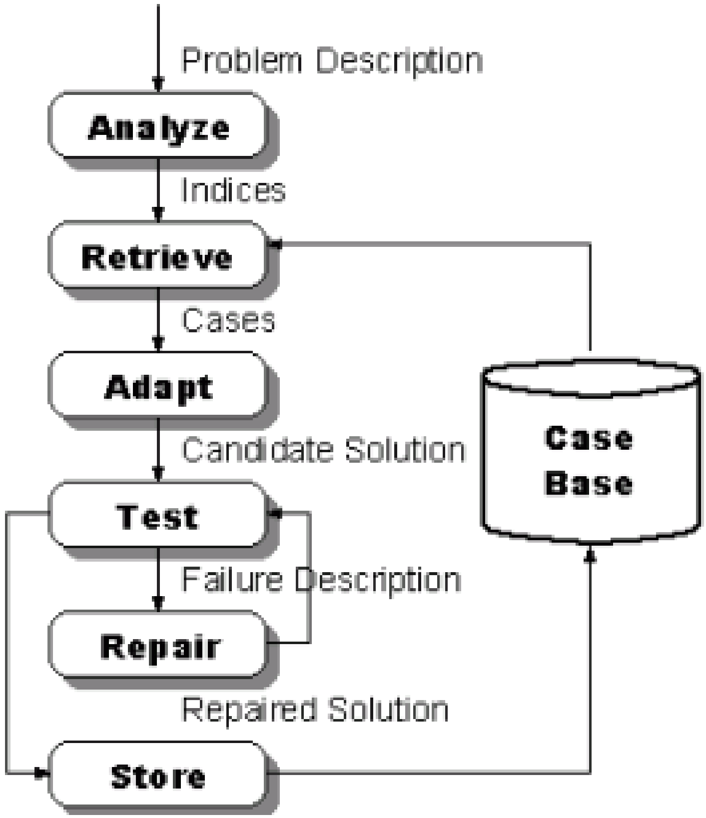

Comparison to other learning techniques

- Machine learning techniques we have seen so far develop a model based on training data before being tested.
- In CBR, the case base constitutes one part of the model
- CBR uses information in the problem to adapt its knowledge base in real time

The Hot Potato

Requires a good chunk of domain knowledge authoring

- State representation
- Distance function
- **Adaptation**
- Revision

Requires a lot of training data or to be be tested iteratively in ways to encourage general learning

Learning can be slow if it gets stuck in a test/repair loop

Overview of Case-Based Reasoning

CBR is […] reasoning by remembering. (Leake 1996)

A case-based reasoner solves new problems by adapting solutions that were used to solve old problems. (Riesbeck & Schank 1989)

CBR is both […] the ways people uses cases to solve problems, and the ways we can make machines use them. (Kolodner 1993)

Bergmann, Survey of CBR, 2000

CBR in one slide

- CBR is a **methodology**
  - to model human reasoning and thinking
  - for building intelligent computer systems
- Basic idea
  - Store known past experiences (cases) in memory (case-base)
  - Given a new problem…
    - Retrieve most similar experience (similarity assessment)
    - Reuse it for the new problem (adaptation)
    - Revise it based on efficacy (feedback)
    - Retain for future use (learning)

CBR: Definition

A problem-solving methodology where solutions to similar, previous problems are reused to solve new problems.

Notes:

- Intuitive
- AI focus (e.g., search, knowledge representation, inference)
- Case = <problem, solution>
- Lazy, incremental, sustained approach to learning

Courtesy of David W. Aha

Problem-Solving with CBR

CBR(problem) = solution

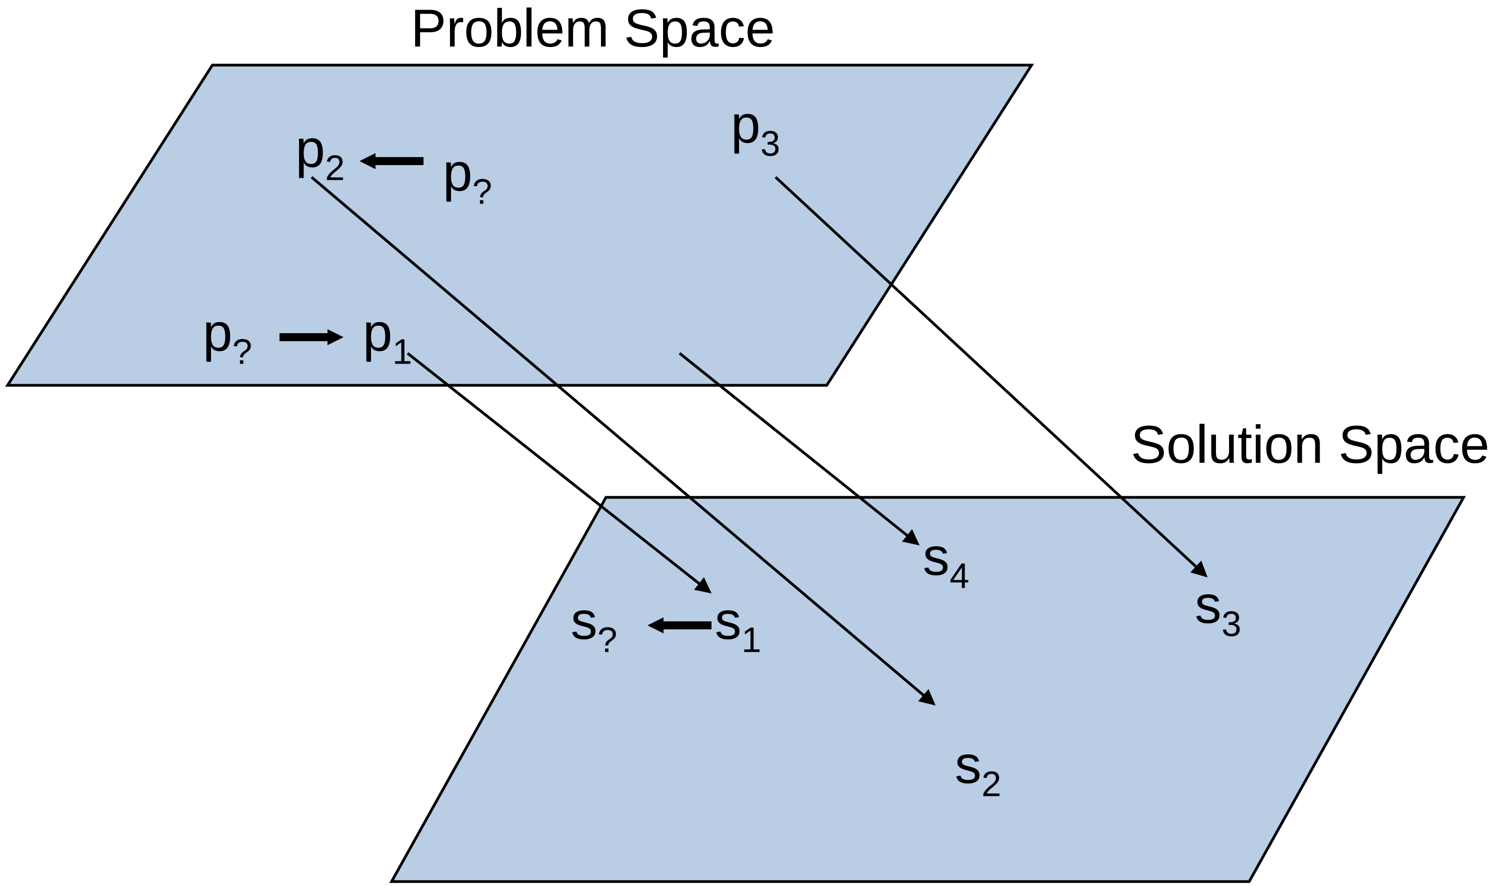

Courtesy of David W. Aha

Example: Slide Creation

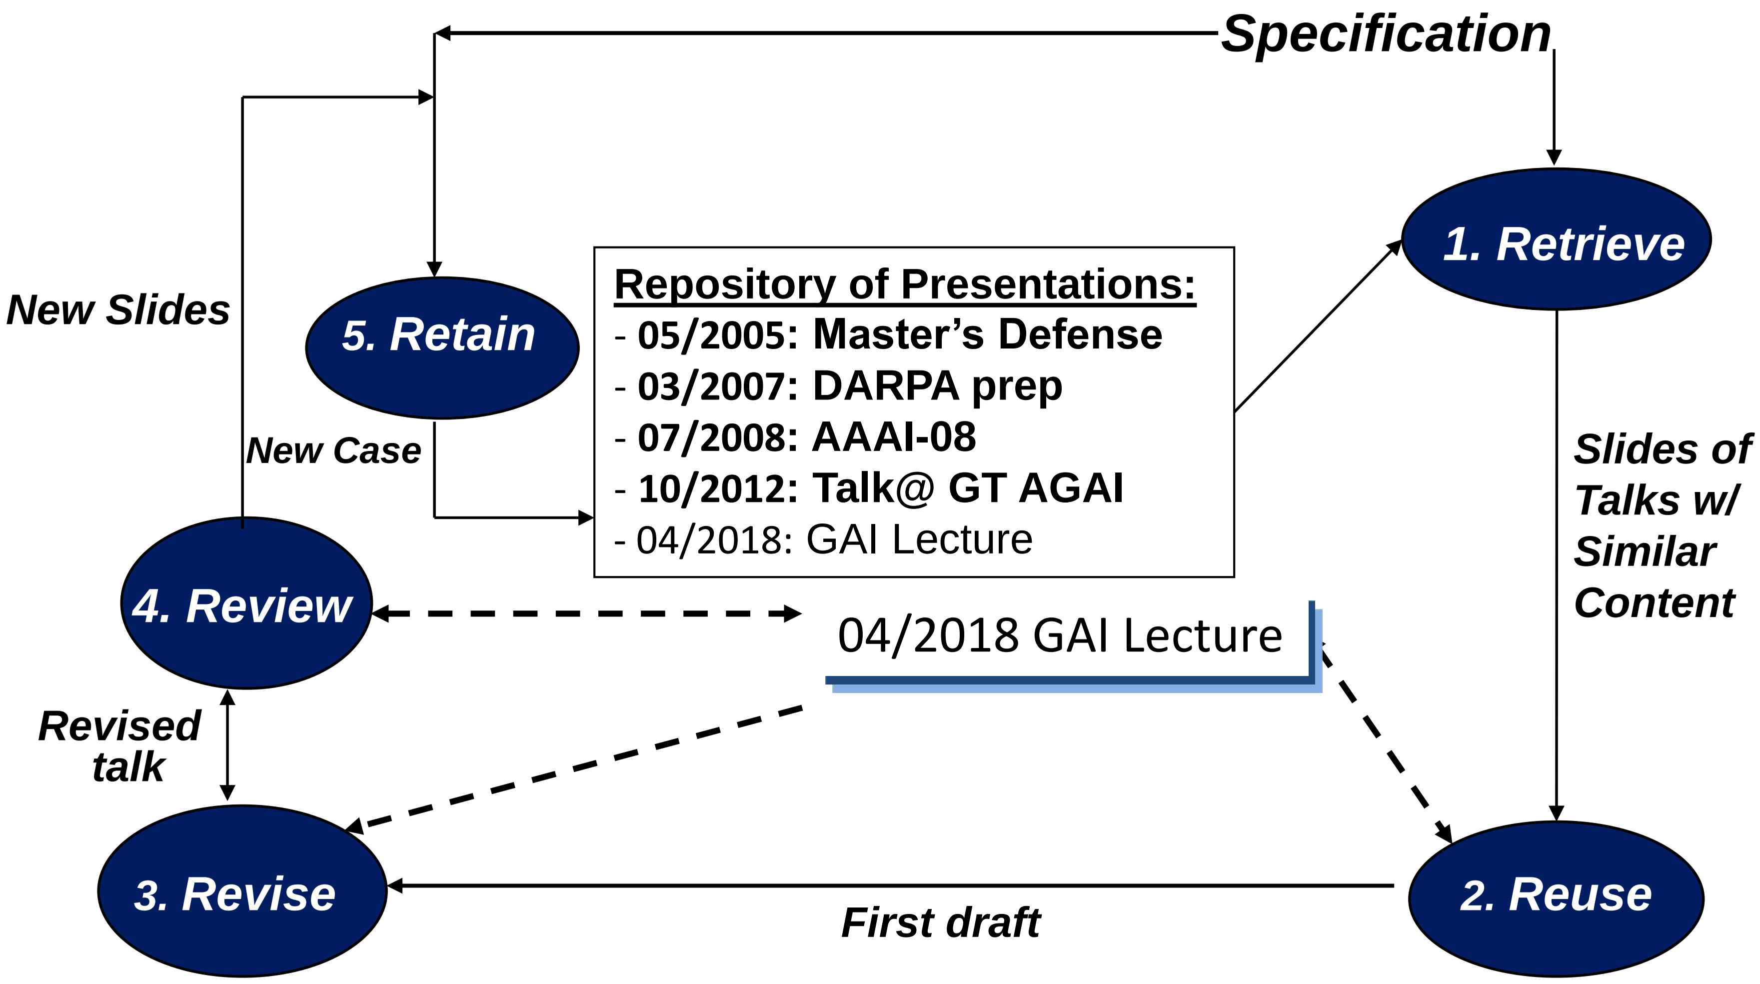

Courtesy of David W. Aha

Problem Solving Cycle of CBR

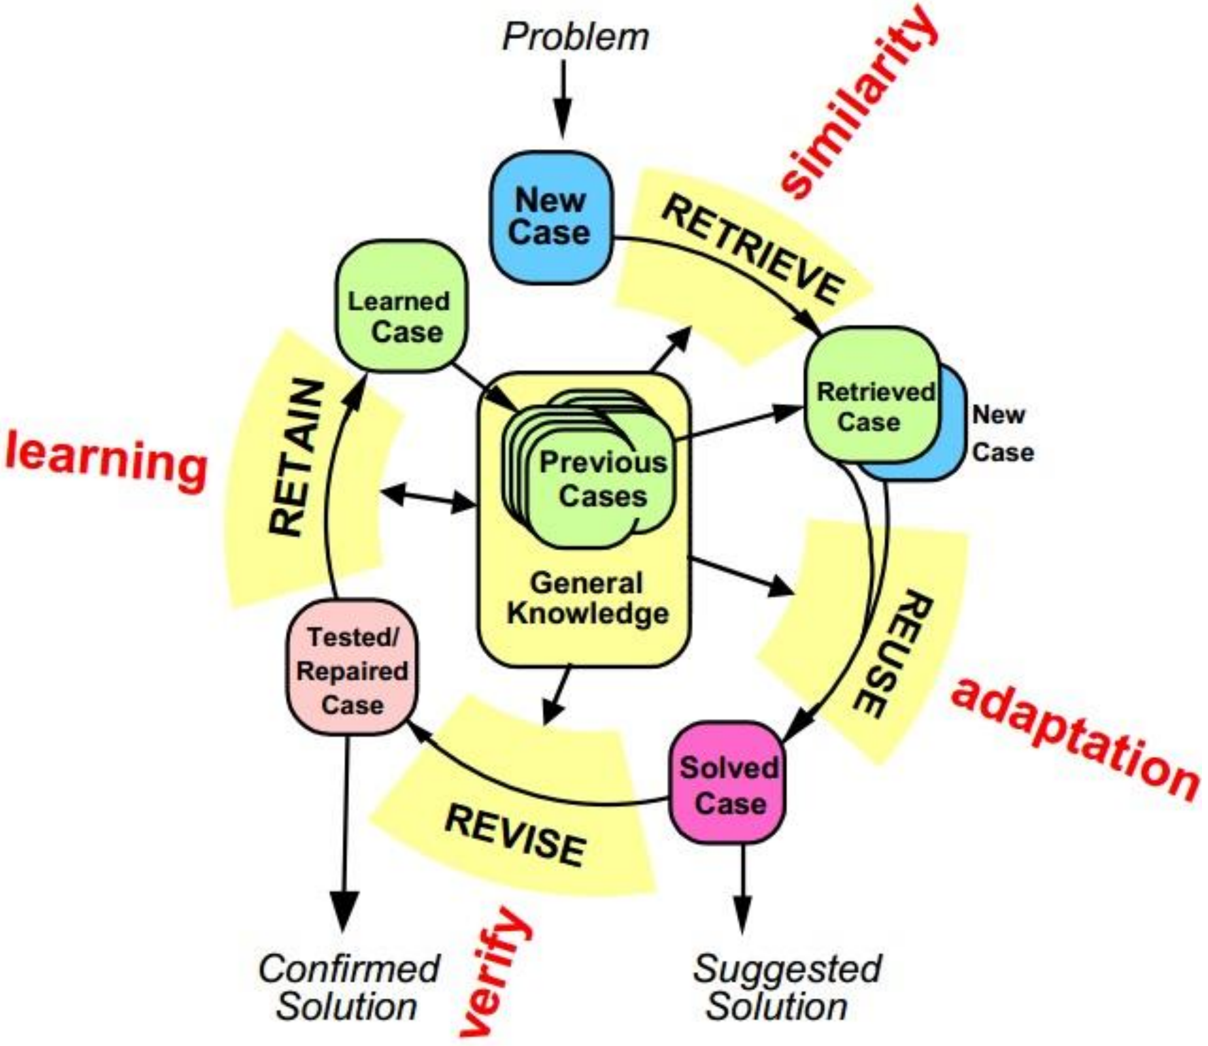

Aamodt & Plaza 1994

Key ideas

- “Similar problems have similar solutions”
- Observations define a new problem
  - Not all feature values must be known
  - A new problem is a case without solution part
- Similarity computation is essential (retrieval)
- Adaptation can be essential (reuse)

CBR: History

1982-1993: Roger Schank’s group, initially at Yale
- Modeling cognitive problem solving (Janet Kolodner, 1993)
- New topics: Case adaptation, argument analysis, …

1990: First substantive deployed application (Lockheed)

1991-: Help-desk market niche (Inference/eGain)

1992: Derivational analogy (Veloso, Carbonell); CBP

1993: European emergence (EWCBR’93)

1993-1998: INRECA ESPRIT projects

1995: First international conference (ICCBR’95)
- Knowledge containers (M. Richter)
- First IJCAI Best Paper Award (Smyth & Keane: Competence models)

1997-: Knowledge management / CB maintenance

1999-: e-Commerce

2001-: Recommender Systems

2003-: *Readings in CBR*

2016: International Conference on CBR held at GT

Representing Cases

- Cases contain knowledge about a previous problem solving experiences

- Typically a case contains the following information:
  - **Problem/Situation**
  - **Solution**
  - Adequacy (**utility**)

- Scope of the information:
  - Complete/partial solution
  - Detailed/abstracted solution

- Representation formalism (depends upon domain/task):
  - Attribute-value vector: Case = $(V_1, \ldots, V_k, V_{k+1}, \ldots, V_n)$
  - Structured representation: Objects, graphs
  - High-order: predicate logic formula, plans

Similarity and Utility in CBR

•The goal of the similarity is to select cases that can be easily adapted to solve a new problem

Similarity = *Prediction* of the utility of the case

•**Utility**: measure of the improvement in efficiency as a result of a body of knowledge

•However:

- The similarity is an a priori criterion
- The utility is an a posteriori criterion

• Sample similarity metric: aggregating local similarity metrics, SIM():

- $SIM(V_{1..n}, Y_{1..n}) = \alpha_1 sim_1(V_1, Y_1) + \ldots + \alpha_n sim_n(V_n, Y_n)$
- $sim_i()$ is a local similarity metric, values in [0,1]

Case Retrieval

Problem description:

•**Input**: a collection of cases CB = $\{C_1, \ldots C_n\}$ and a new problem P

•**Output**:

- The most similar case: A case C_i in CB such that sim(C_i, P) is minimal, *or*
- A collection of *m* most similar cases in CB $\{C_1,\ldots, C_m\}$, *or*
- A *sufficiently* similar case: case C_i in CB such that
  sim(C_i, P) > *th*

Solutions:

•Sequential retrieval: $O(|CB| \times \log_2 (k))$
•Two-step retrieval: (1) select subset S of cases. (2) Sequential retrieval on S.
•Retrieval with indexed cases

Case Adaptation

CaseA (new)

Car type: sport
Color: red
Seating: 2
Valves: 48
Type: 5.7L

Model name: name1
Price: 200,000
Year: 2003

Feedback: not successful
Cause: price is too high

Retrieve

CaseB (old)

Car type: sport
Color: red
Seating: 2
Valves: 48
Type: 5.7L

Model name: name1
Price: 200,000
Year: 2003

Feedback: successful

Copy

Adapt

CaseC (adapted)

Car type: sport
Color: red
Seating: 2
Valves: 40
Type: 3.6L

Model name: name 2
Price: 150,000
Year: 2000

Feedback: successful

Problem description:

- **Input**: A retrieved case C and a new problem P
- **Output**: A solution for P obtained from C

Considered an open problem

Trade-off between Retrieval and Adaptation Effort

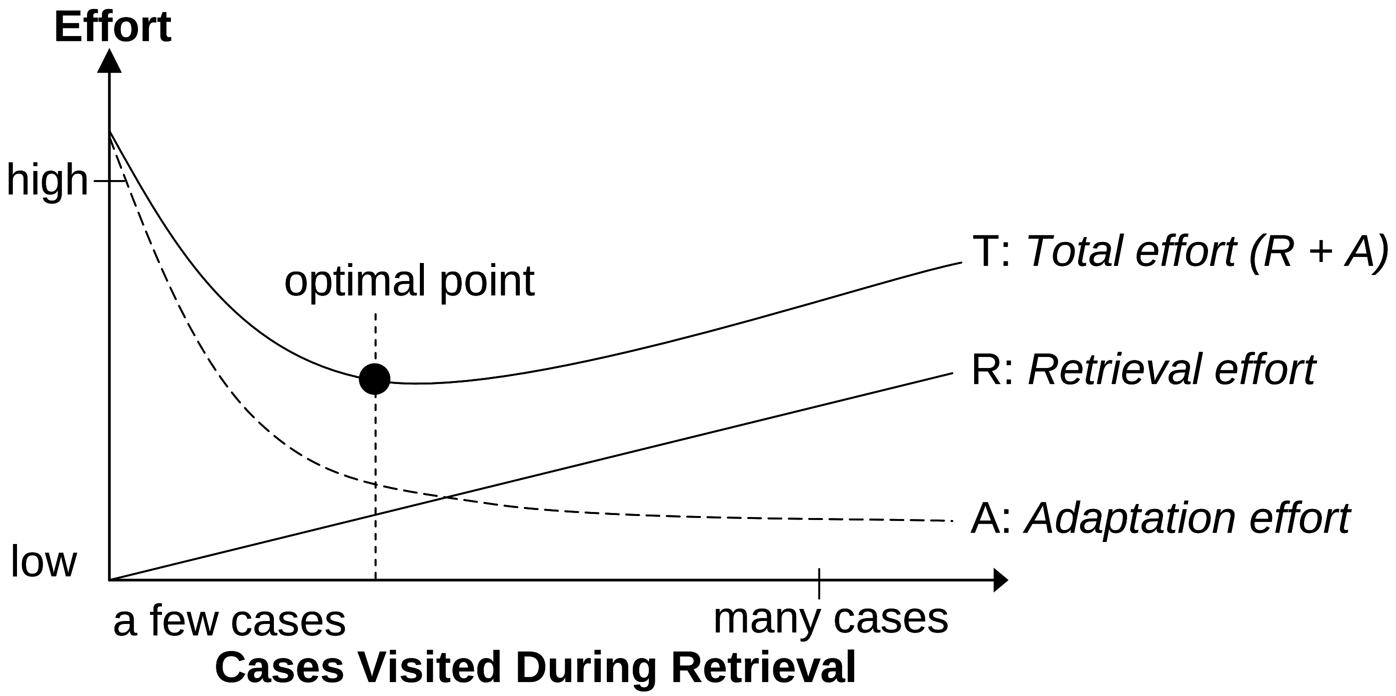

- If little time is spent on retrieval, then the adaptation effort is high
- If too much time is spent on retrieval, then the adaptation effort is low
- There is an optimal intermediate point between these two extremes

Taxonomy of Problem Solving and CBR

For which of these CBR have been shown to be effective?

- **Synthesis**:
  - constructing a solution
  - Methods: planning, configuration
- **Analysis**:
  - interpreting a solution
  - Methods: classification, diagnosis

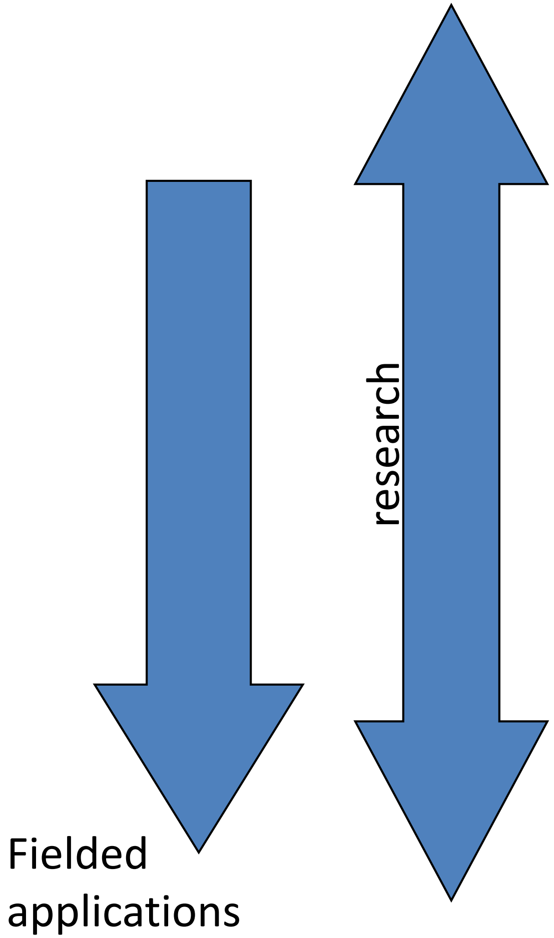

Main Topics of CBR Research ~ 10yr

- Study by Derek Greene, Jill Freyne, Barry Smyth, Pádraig Cunningham
- Social network analysis based on co-citations links
- Sources:
  - Bibliographic data from Springer about ICCBR, ECCBR
  - Citation data from Google scholar
- **An Analysis of Research Themes in the CBR Conference Literature**. ECCBR'08
- Next two slides from http://mlg.ucd.ie/cbr

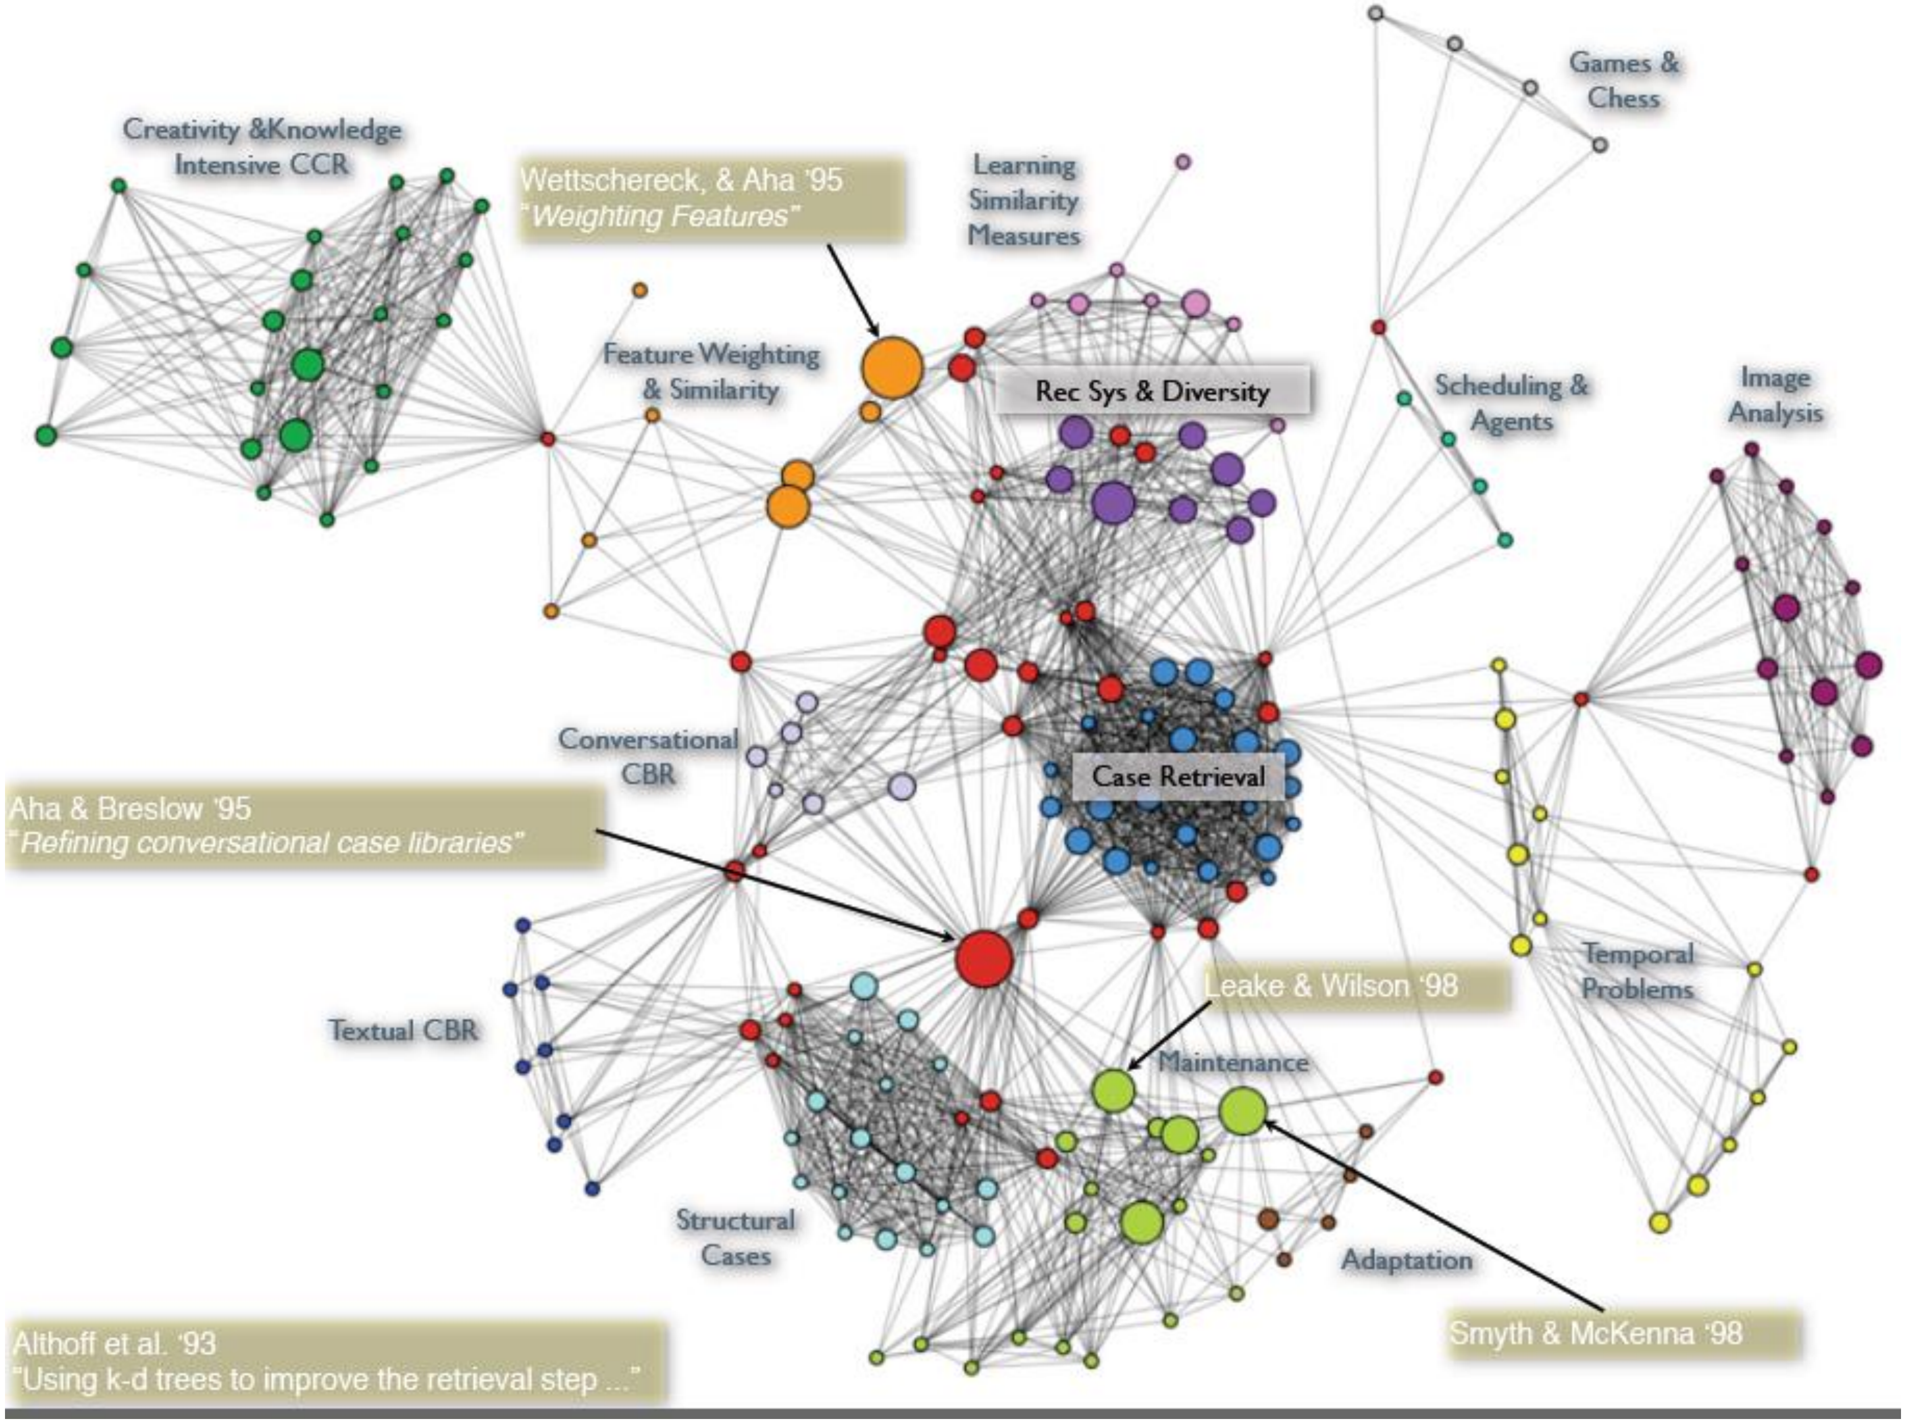

Games & Chess
Creativity &Knowledge Intensive CCR
Wettschereck, & Aha '95
"Weighting Features"
Learning Similarity Measures
Feature Weighting & Similarity
Rec Sys & Diversity
Scheduling & Agents
Image Analysis
Conversational CBR
Case Retrieval
Aha & Breslow '95
"Refining conversational case libraries"
Textual CBR
Leake & Wilson '98
Maintenance
Temporal Problems
Structural Cases
Adaptation
Smyth & McKenna '98
Althoff et al. '93
"Using k-d trees to improve the retrieval step ..."

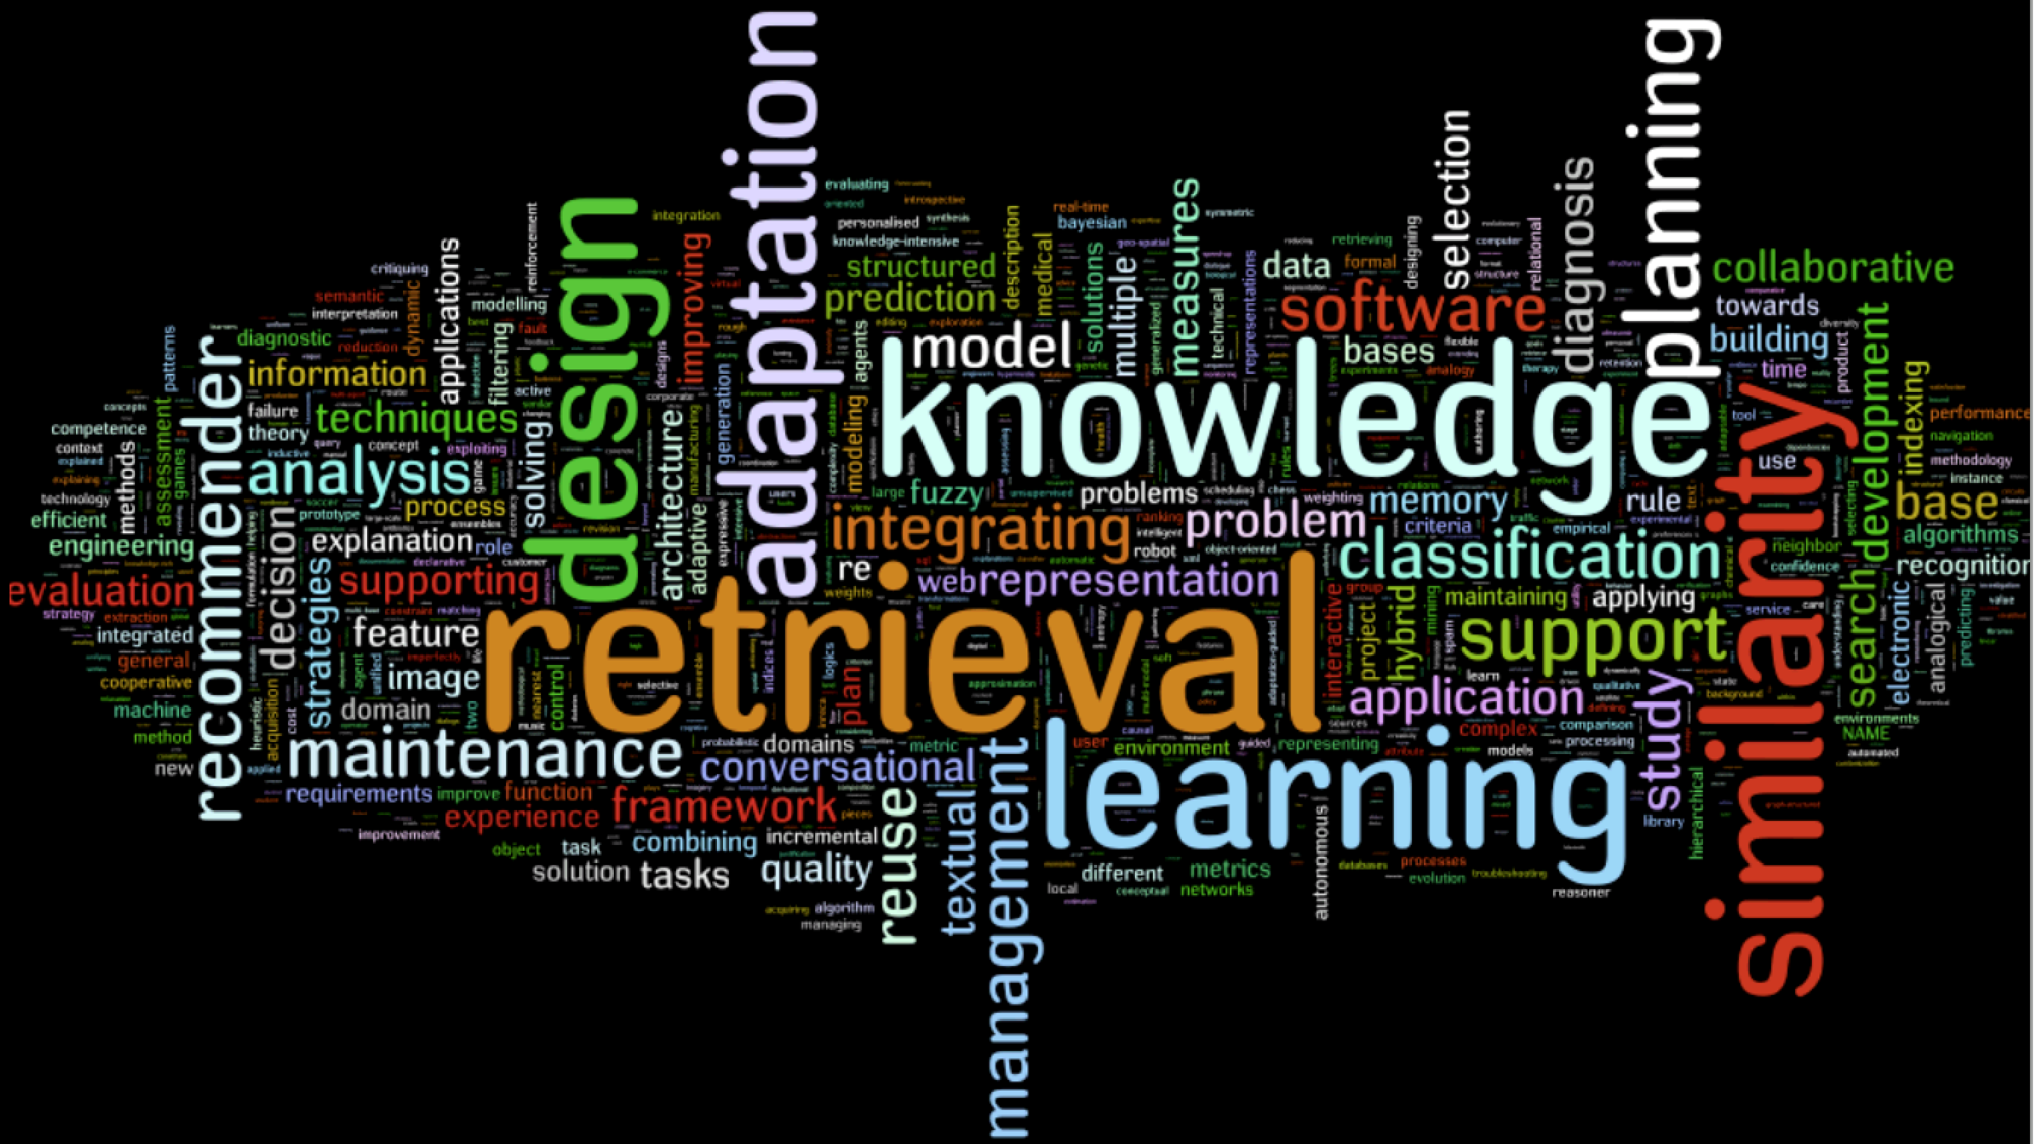

retrieval
knowledge
learning
similarity
adaptation
design
planning
recommender
maintenance
software
integrating
classification
support
representation
analysis
application
management
reuse
textual
conversational
framework
evaluation
techniques
information
model
problem
memory
collaborative
diagnosis
selection
prediction
structured
measures
supporting
explanation
feature
image
domain
decision
strategies
architecture
development
search
base
study
experience
requirements
quality
tasks
solution
combining
incremental
fuzzy
process
engineering
problems
bases
data
building
towards
rule
maintaining
applying
recognition
algorithms
indexing
electronic
analogical
hybrid
interactive
project
environment
metrics
networks
different
domains
web
re
multiple
solutions
medical
description
agents
modeling
improving
adaptive
generation
solving
filtering
applications
dynamic
diagnostic
semantic
critiquing
modelling
failure
theory
competence
assessment
methods
games
efficient
patterns
integrated
general
cooperative
machine
method
new
strategy
extraction
technology
context
explained
concept
exploiting
role
prototype
function
improve
object
task
improvement
algorithm
managing
local
autonomous
reasoner
library
hierarchical
complex
criteria
empirical
time
use
tool
performance
navigation
methodology
instance
neighbor
confidence
service
product
NAME
environments
automated
representing
evolution
processes
troubleshooting
models
comparison
processing
qualitative
learn
state
background
robot
weights
large
users
formal
technical
representations
generalized
retrieving
designing
computer
relational
evaluating
integration
personalised
knowledge-intensive
introspective
bayesian
real-time
oriented
synthesis
symmetric
reinforcement
interpretation
reduction
ensembles
active
game
fault
mining
group
weighting
scheduling
chess
ranking
rules
editing
exploration
geo-spatial
flexible
analogy
therapy
retention
experiments
user
metric
probabilistic
logics
plan
indices
selective
approximation
guided
measure
conceptual
databases
sources
attribute
two
control
nearest
acquisition
heuristic
cost
music
soccer
applied
agent
unified
imperfectly
constraint
matching
customer
declarative
accuracy
expressive
revision
manufacturing
corporate
designs

Major Themes in CBR

- Recommender systems and diversity
- Case-Based Maintenance
- Case Retrieval
- Learning similarity measures
- Adaptation
- Image analysis
- Textual & Conversational CBR
- Feature weighting and similarity

http://www.csi.ucd.ie/files/ucd-csi-2009-03.pdf

Some Interrelations between Topics

- Retrieval
  - Information gain
  - Similarity metrics
  - Indexing
- Reuse
  - Rule-based systems
  - Plan Adaptation
- Revise & Review
  - Constraint-satisfaction systems
- Retain
  - Induction of decision trees

Focus Point: Diversity in CBR

Traditional Retrieval Approach

- Similarity-Based Retrieval
  - Select the *k* most similar items to the current query.

    Query
    Available case
    Similar case

- Problem
  - Vague queries.
  - Limited coverage of search space in every cycle of the dialogue.

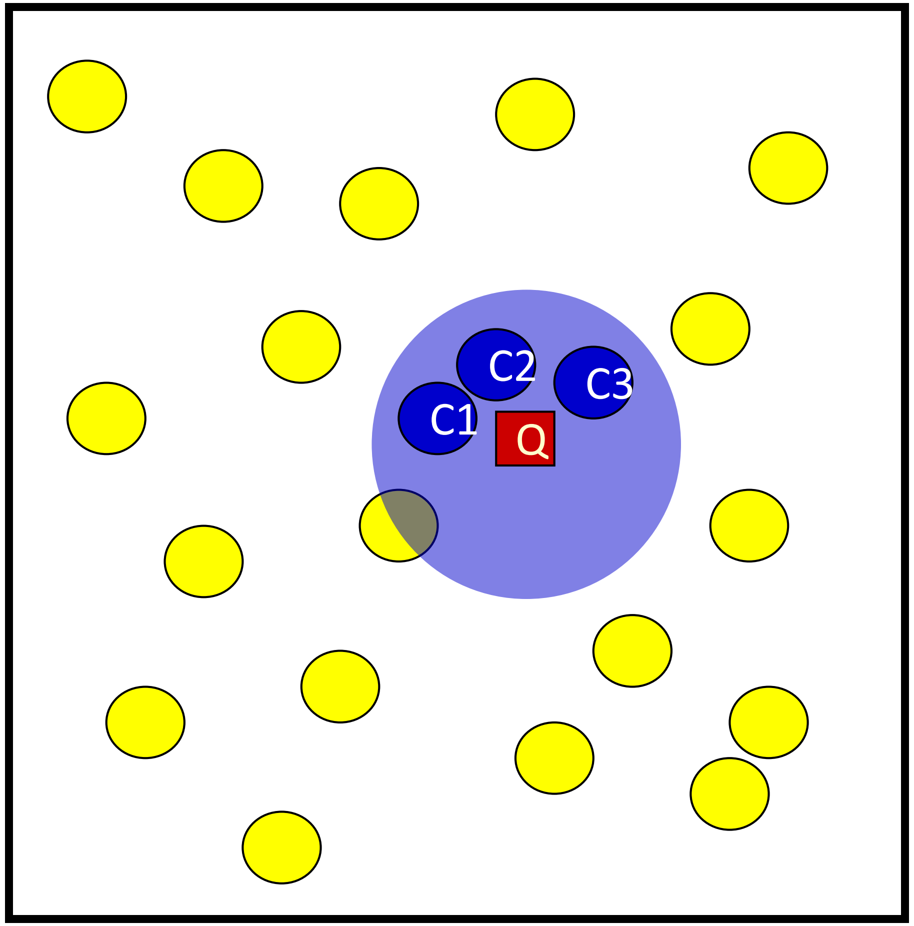

Lorraine McGinty and Barry Smyth
Department of Computer Science, University College Dublin

Diversity Enhancement

- Diversity-Enhanced Retrieval
  - Select *k* items such that they are both **similar to the current query but different from each other**.

    Query

    Available case

    Retrieved case

  - Providing a wider choice allows for broader coverage of the product space.
  - Allows many less relevant items to be eliminated.

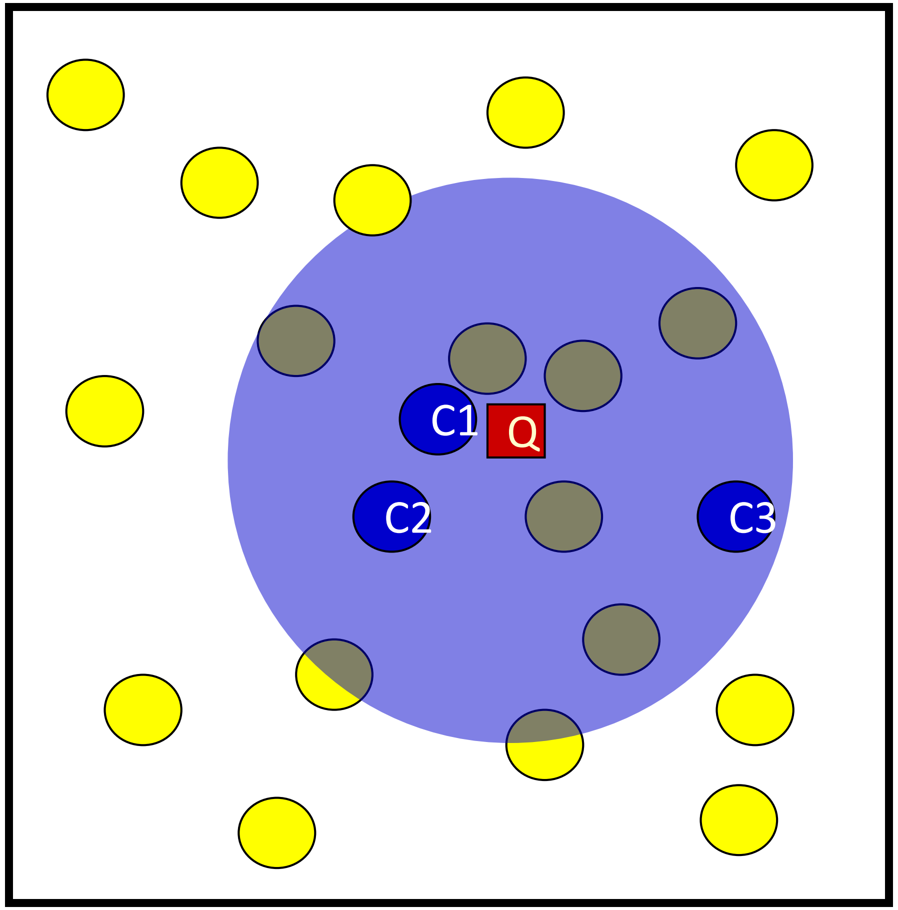

Focus Point: Augmenting General Knowledge with Cases

Why Augment General Knowledge With Cases?

- In many practical applications, **encoding complete domain knowledge is unpractical/infeasible and episodic knowledge is available**

Example: Some kinds of military operations where two kinds of knowledge are available:

- **General guidelines and standard operational procedures** which can be encoded as a (partial) <u>general domain knowledge</u>
- Whole compendium of **actual operations and exercises** which can be captured as <u>cases</u>

Hierarchical Problem Solving

Hierarchical case-based planning techniques combine domain knowledge and episodic knowledge (cases)

Knowledge source

domain

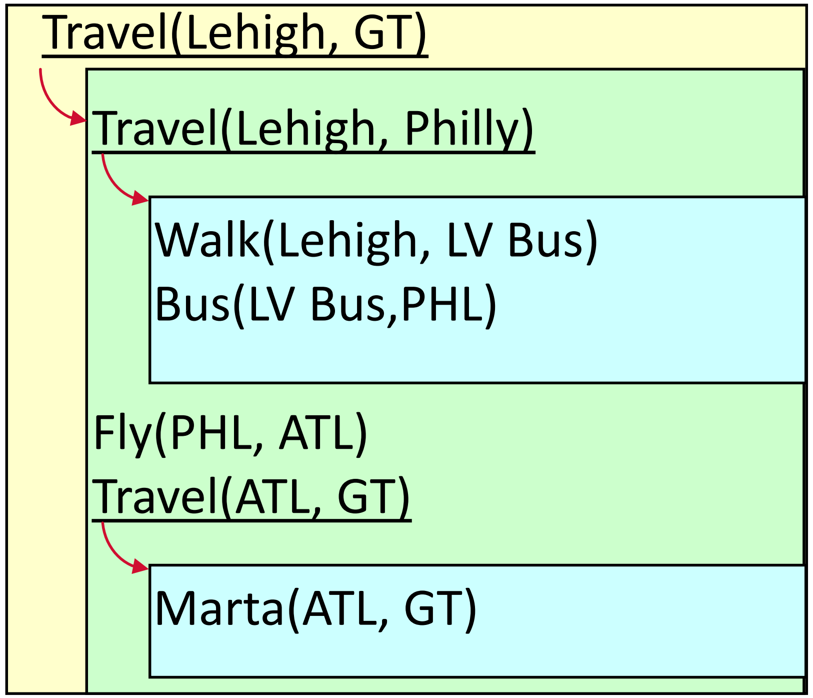

Knowledge Sources

General

Methods denote generic task decompositions and *conditions* for selecting those decompositions:

Task: travel(?A,?B)

Decomposition:

travelC(?A, ?Airp1)
travelIC(?Airp1,?Airp2)
travelC(?Airp2, ?B)

Conditions:

in(?A,?City1)
in(?B,?City2)
airport(?Airp1,?City1)
airport(?Airp2,?City2)

Episodic

Cases denote concrete task decompositions:

Task: travelC(Lehigh, PHL)

Decomposition:

take(bus, Lehigh, PHL)

Conditions:

enoughMoney()

CBR: Final Remarks

Advantages of CBR

- Reduces knowledge acquisition effort
- Requires less maintenance effort
- Reuse of solutions improves problem solving performance
- Can make use of existing data
- Improves over time, adapts to changes
- Has enjoyed high user acceptance

Bergmann, Survey of CBR, 2000

Why not cbr?

In fact, this is the crux of the argument: if you have a good scripting language, or even a visual tree editor to capture sequences, **you'll be orders of magnitude more productive (and more reliable) than an expert trying indirectly to get the system to induce specific sequences from examples**. As such, it's fair to claim that CBR isn't particularly well suited to these kinds of problems in game AI.

(2008) http://aigamedev.com/open/editorial/critique-case-based-reasoning/

Past uses at GT for Digital Entertainment

Dialog

- Sanjeet Hajarnis, Christina Leber, Hua Ai, Mark O. Riedl, and Ashwin Ram (2011). **A Case Base Planning Approach for Dialogue Generation in Digital Movie Design**. Proceedings of the 19th International Conference on Case Based Reasoning, London, UK.

Decision making

- Santiago Ontañón and Ashwin Ram (2011) **Case-Based Reasoning and User-Generated AI for Real-Time Strategy Games**. In Pedro Antonio González-Calero and Marco Antonio Gómez-Martín (Editors), Artificial Intelligence for Computer Games, pp. 103-124. Springer-Verlag.

Drama Management

- Manu Sharma and Santiago Ontañón and Manish Mehta and Ashwin Ram (2010) **Drama Management and Player Modeling for Interactive Fiction Games**, in Computational Intelligence Journal, Volume 26 Issue 2, pp. 183-211.

PCG

- Manish Mehta and Santiago Ontañón and Ashwin Ram (2008) **Adaptive Computer Games: Easing the Authorial Burden**. in Steve Rabin (Editor), AI Game Programming Wisdom 4. pp. 617-632

Other applications in games

- Gillespie, K., Karneeb, J., Lee-Urban, S., and Munoz-Avila, H. (2010) **Imitating Inscrutable Enemies: Learning from Stochastic Policy Observation, Retrieval and Reuse**. Proceedings of the 18th International Conference on Case Based Reasoning (ICCBR 2010). AAAI Press.
- Auslander, B., Lee-Urban, S., Hogg, C., and Munoz-Avila, H. (2008**) Recognizing The Enemy: Combining Reinforcement Learning with Strategy Selection using Case-Based Reasoning**. Proceedings of the 9th European Conference on Case-Based Reasoning (ECCBR-08).
- Hogg, C., Lee-Urban, S., Auslander, B., and Munoz-Avila, H. (2008) **Discovering Feature Weights for Feature-Based Indexing of Q-Tables**. Proceedings of the Uncertainty and Knowledge Discovery in CBR Workshop at the 9th European Conference on Case-Based Reasoning (ECCBR-08).

CBR: Takeaway

1. Sometimes natural (e.g., law, diagnosis)

2. Cases simplify knowledge acquisition
   - Easier to obtain than rules
   - Captures/shares people's experiences

3. Good for some types of tasks
   - When perfect models are not available
     - Faulty equipment diagnosis
     - Online sales
     - Legal reasoning
     - Games

4. Commercial applications
   - Help-desk systems (e.g., Inference corp.: +700 clients)

5. Similar problems have similar solutions.
   - Retrieve, Reuse, Revise, Retain

Questions?

- http://cbrwiki.fdi.ucm.es/
- http://aitopics.net/CaseBasedReasoning
- http://www.cse.lehigh.edu/~munoz/CSE335/
- http://mlg.ucd.ie/cbr

- http://gaia.fdi.ucm.es/research/colibri/jcolibri

CBR: Recap

1) What are the 4 processes, each beginning with an "R", commonly used to describe the CBR methodology?

2) The ______ metric is used to find the problem/solution pair in the casebase most related to the new problem, by comparing the relatedness of the features of the new problem to the features of known problems in the casebase.

3) In case-based reasoning, problem solving cannot commence without the ability to compute this, which is a number indicating how related an existing case is to the new problem.

4) A foundational assumption in CBR is that "Similar problems have _________ __________".