Capturing and Reusing Experience

2016-07-19

OOB

- Capstone
  - Wed 7/6: Team formed
  - Thurs 7/14: Plan and pitch done, implementing
  - Yesterday 7/18: Game runs (bugs & naïve AI ok – stubs ☹)
  - **Thurs 7/21: Rough, runnable demo ready**
  - **Mon 7/25: Impl’n finished. Polishing & small debugging**
  - **Delivery Thursday night 7/28**
  - Capstone demos 7/29 (Fri) 2:50pm-5:40pm, here
- Trajectory
  - CBR & Action prediction
  - Game demos. AI Based games
  - Recap & Evals

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?
4. FSMs/Btrees: _____ :: Planning : ______
5. For the 2nd blank, we need m______s.
6. When is reactive appropriate? Deliberative?
7. What is the ‘hot-potato’ passed around (KE)?
8. H_______ have helped in most approaches.
9. Which approach should you use?

Questions

1. What are the 2 most “complex” decision making techniques we’ve seen?
2. What are their strengths? Weaknesses?
3. What is the key (insight) to their success?
4. What is typically necessary to support this insight (hint: used in Planning + RBS)?
5. What does Planning have that (forward chaining) RBS do not?
6. When do we need a communication mechanism?

PCG Questions

1. What is PCG?
2. PCG can be used to p____ or a____ game aspects
3. Why does industry care about PCG?
4. What are some risks of PCG?
5. Major concerns involving PCG include…
6. What is a player model? What does it allow?
7. What are ways to get a player model?
8. Bartle’s 4-part feature vector: <k,a,e,s>

PCG Desiderata

- Speed (real-time/design time)
- Reliability (catastrophic failures/crashes)
- Controllability (wrt constraints and goals)
- Diversity (variations on a theme)
- Creativity (looks "computer-generated")

See IGDA Webinar, 10 December 2014: PCG in games

PCG as Local Search Questions

- What "search" is happening? Do we seek a path to goal?
- What is the state space? How many states do we save?
- How memory efficient is this search?
- Hill climbing:
  - L____ search
  - What is the "landscape"?
  - Need a function that maps p____ to f_____
- GAs:
  - Good in ______ domains, where _D.K.__ is scarce or hard to encode
  - Can also be used for ____ search
  - Also needs a f_____ function (maps c____ to f_____)
- Other local search techniques
  - Gradient Descent
  - Simulated annealing
  - Local beam
  - Tabu
  - Ant Colony Optimization

GA Steps

1. Create a random set of *n* chromosomes (**population**)
2. Assign a fitness score to each chromosome (**fitness function**)
3. Remove the *m%* (*m* < 100) worst chromosomes
4. Cycle through remaining pairs of chromosomes and *cross-over* (with some probability)
5. Randomly mutate (during?) cross-over (with some probability)
6. Reduce new population to size *n*
7. Repeat steps 2-6 until [stepwise improvement diminishes || one individual is fit enough || # generations reached]

GA Tuning Parameters

- Population size
- Number of generations
- Fitness function
- Representation
- Mutation rate
- Crossover operations
- Selection procedure
- Number of solutions to keep

PCG See also

- Papers linked above & T-square
- IGDA Webinar, 10 December 2014: PCG in games: perspectives from the ivory tower
  - https://www.youtube.com/watch?v=UVRqCK6m7m4
- PCG Book http://pcgbook.com/
  - Grammars: Chapter 5 http://pcgbook.com/wp-content/uploads/chapter05.pdf
- 9.1: Genetic Algorithms and Evolutionary Computing - The Nature of Code
  - https://www.youtube.com/watch?v=6l6b78Y4V7Y

Millington 7.3

CAPTURING AND REUSING EXP: ACTION PREDICTION

Action Prediction

- Guess what player will do next
  - E.g. waypoint, weapon, cover point, melee
  - Make more realistic, challenging (helpful) NPC
  - Can do with little observation
  - Can transfer from other players
- Humans bad at random (Psychology)

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
- Incremental update of average
  - Keep mean, and count
  - $m_n = m_{n-1} + (1/n)(v_n – m_{n-1})$

String Matching

- “Left or Right” coin game
- Choice made several times
  - Encode as string “LRRLRLLLRRLRLRR”
  - Predict → find substring, return subsequent choice
  - Example: “RR”
  - Window size

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

Window Size

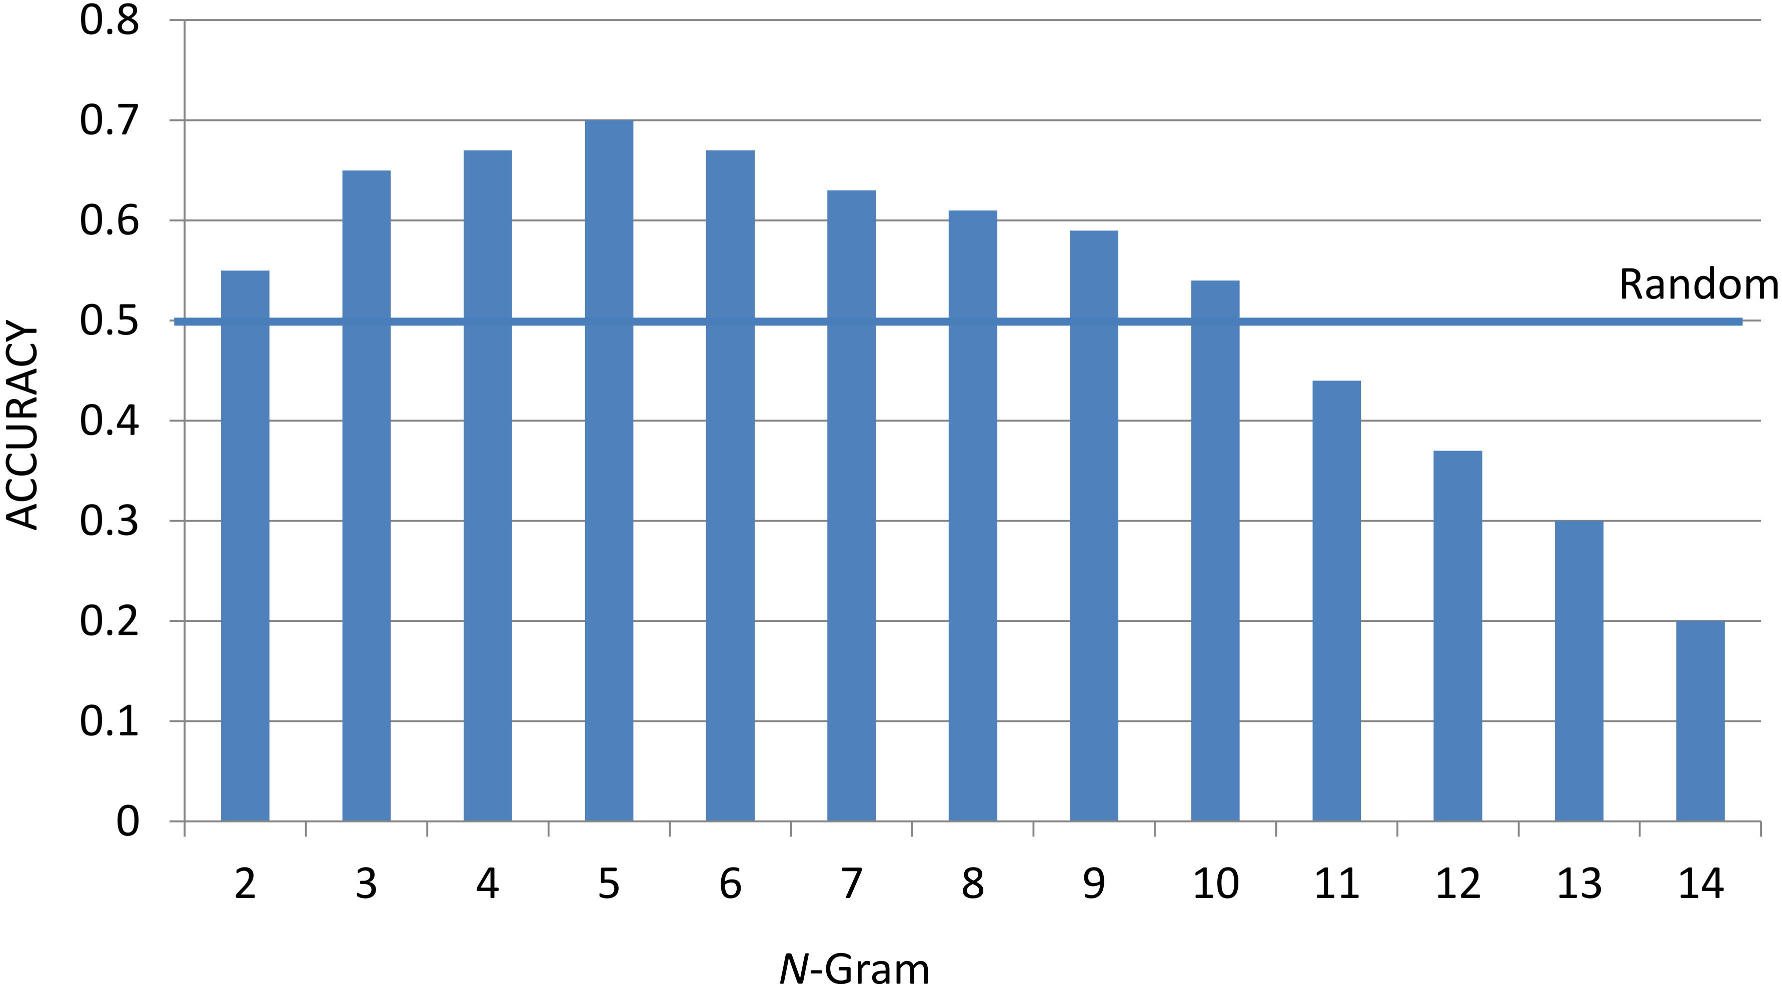

0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
ACCURACY
Random
2
3
4
5
6
7
8
9
10
11
12
13
14
N-Gram

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
    - (↑ random ↓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
- Synon. with combo-based melee games
  - Can make unbeatable (no fun) AI
  - Often is “gimped”
- Many other uses
  - statistical analysis techniques (e.g. language)
  - [Weapon, location, unit] selection…

CAPTURING AND REUSING EXP: CASE-BASED REASONING

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

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)

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)
  - http://www.youtube.com/watch?v=sITkmOefamc

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

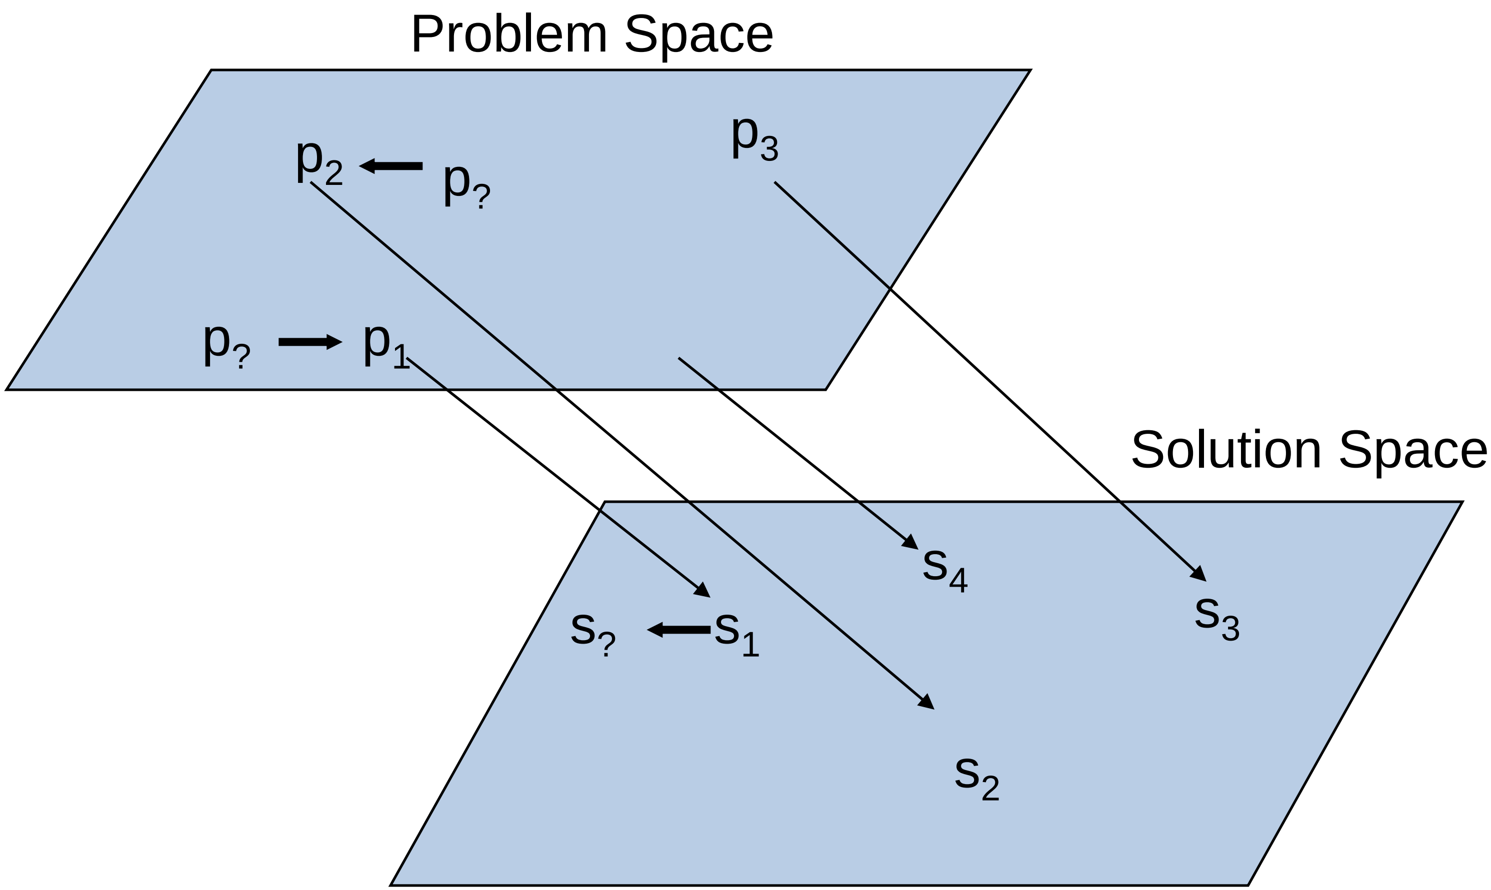

Courtesy of David W. Aha

Example: Slide Creation

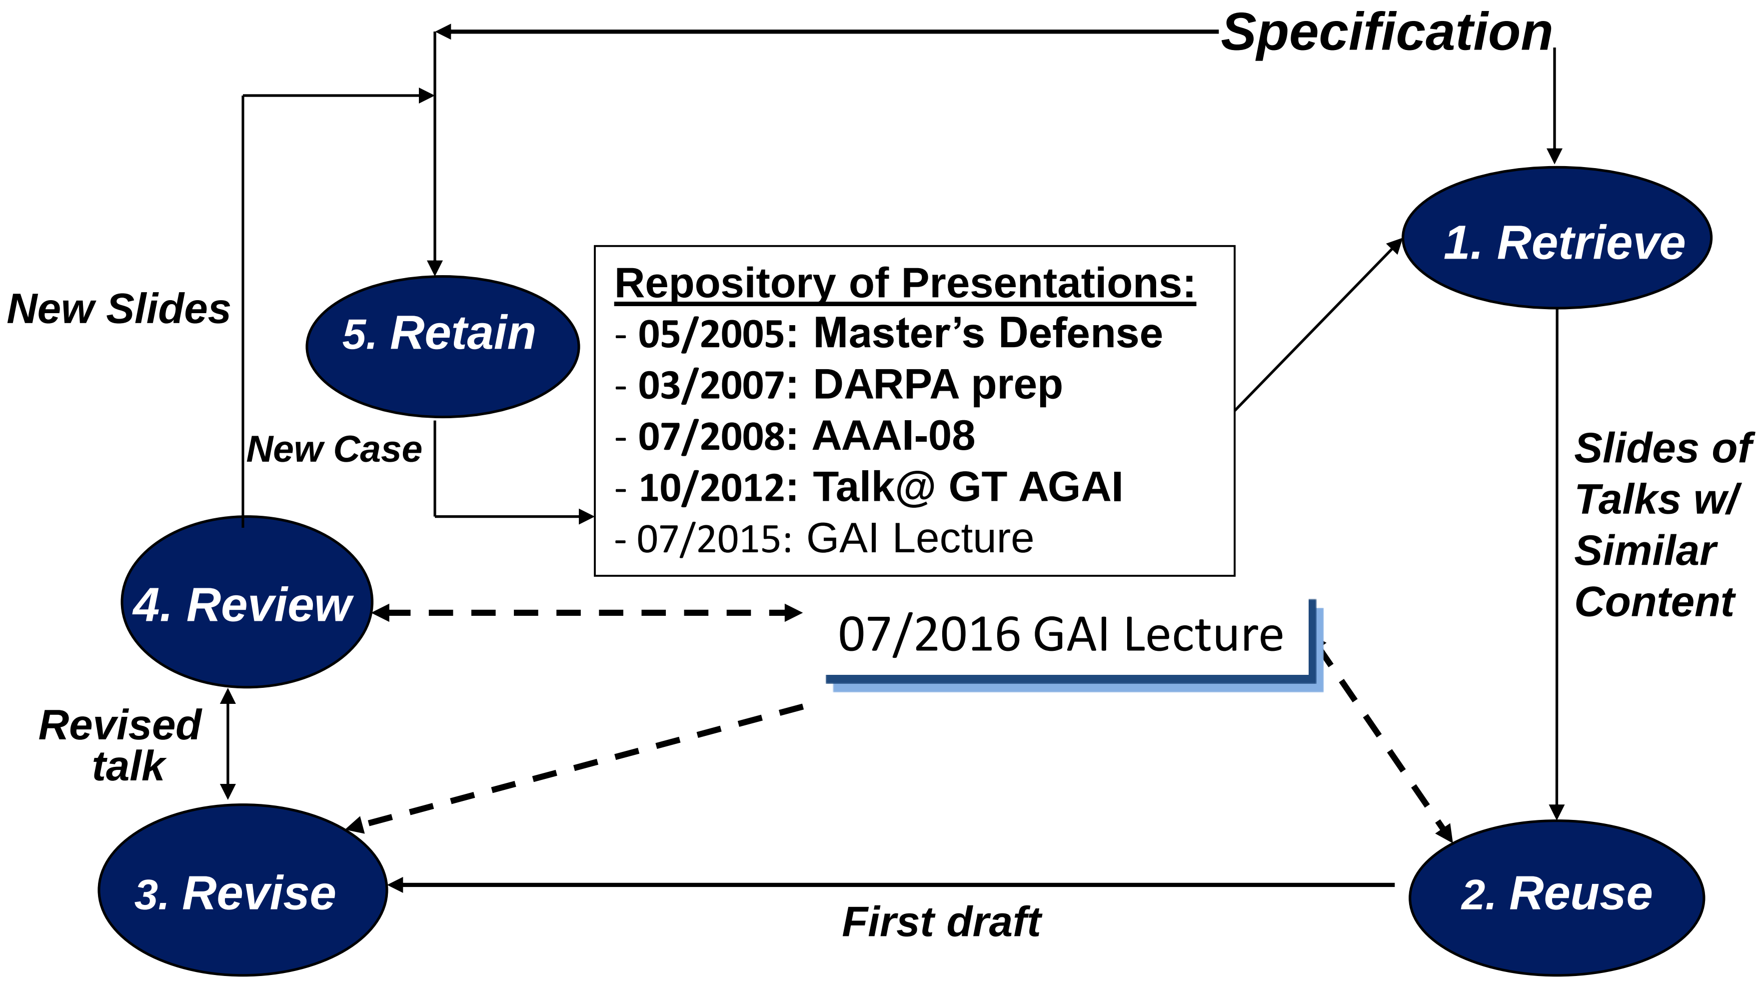

Courtesy of David W. Aha

Problem Solving Cycle of CBR

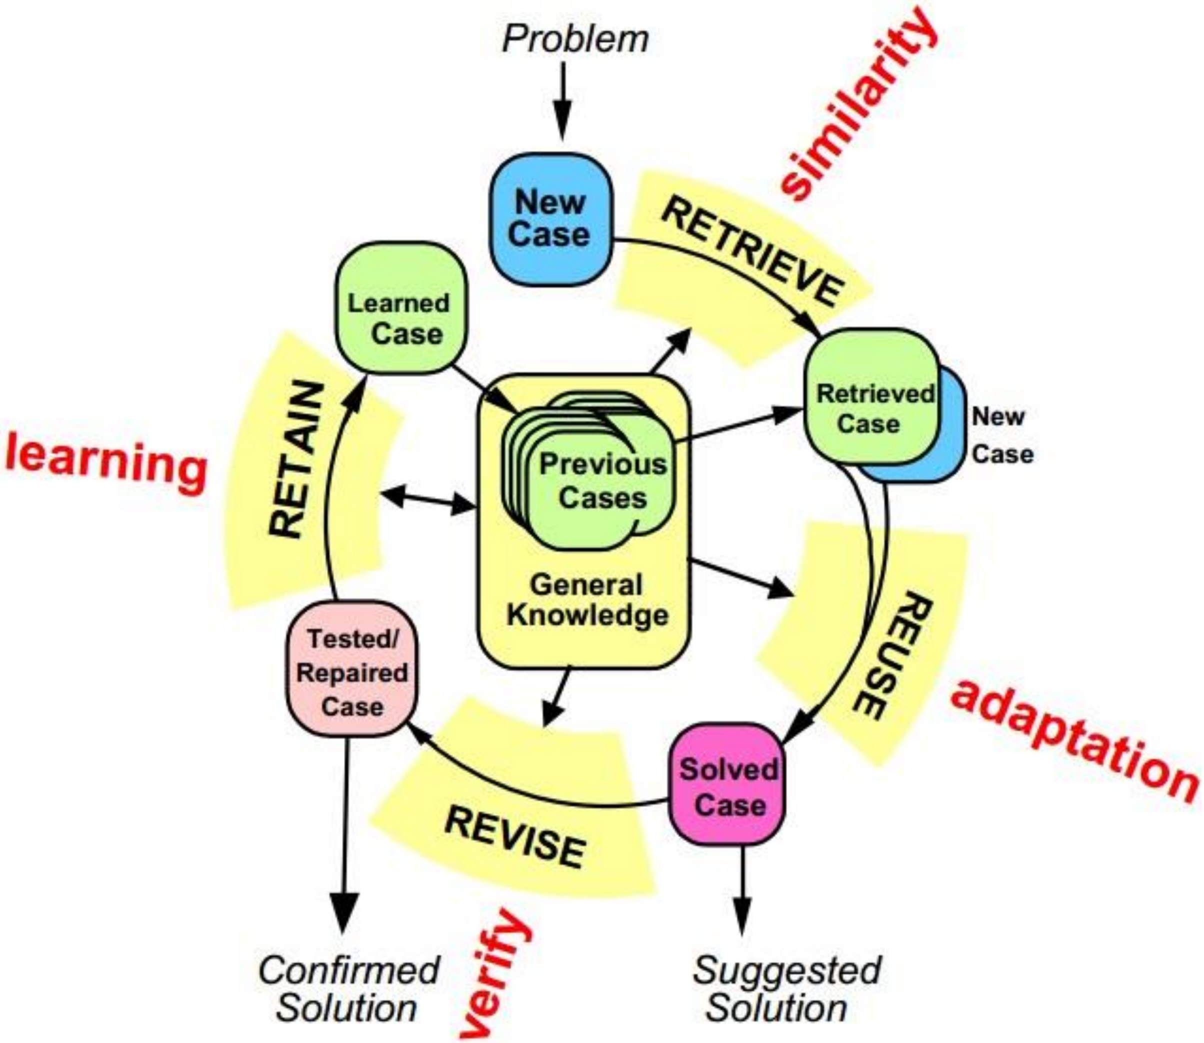

Aamodt & Plaza 1994

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 to be held at GT late October

You Have Seen this Before!
(A consumer's Customer Service Experience)

Have you called a customer service support line lately?

It goes something like this (automatic machine):

1. If you want to speak to a sales representative, please press one
2. ….

…

9. If you are experiencing technical difficulties with our wonderful product **Neutronious-L** please press nine

You Have Seen this Before!

(A consumer's Customer Service Experience- part 2)

Welcome to our costumer support menu (automatic machine):

1. If you want to listen to the FAQ please press one
2. ….

…

9. If none of the above help you please press nine.

After 40 minutes of hearing music meant to drive you insane…

You Have Seen this Before!

(A consumer's Customer Service Experience- part 3)

Yes this is Felix may I have the serial number of Neutronious-L, please? (a person reading from an automatic machine):

1. Is Neutronious-L ringing? *You:* no
2. Is a red light Neutronious-L blinking? *You*: no

…

9. How many green lights are on on Neutronious-L? *You*: 3
10. Are you sure? *You*: yes

Well, in that case you should call the company that constructed your building. If you ask me that must be excessive moisture… Now let me ask you a few questions about our service…
sir? Hello? Are you still there?

What is Going on the Other Side

Space of known problems for Neutronious-L

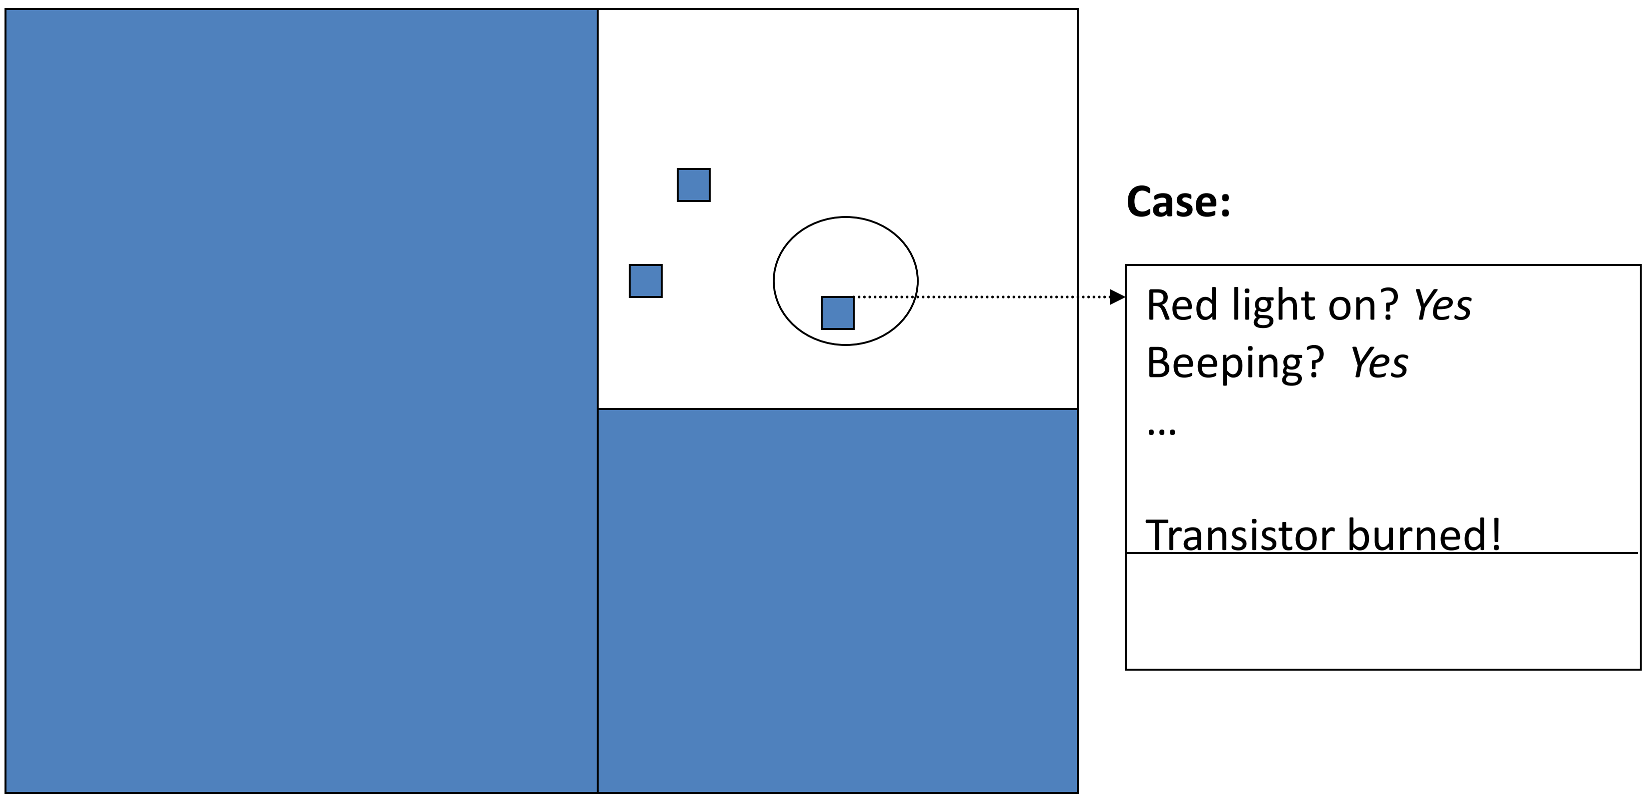

This is an example of **a Conversational Case-Based Reasoning Process**

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

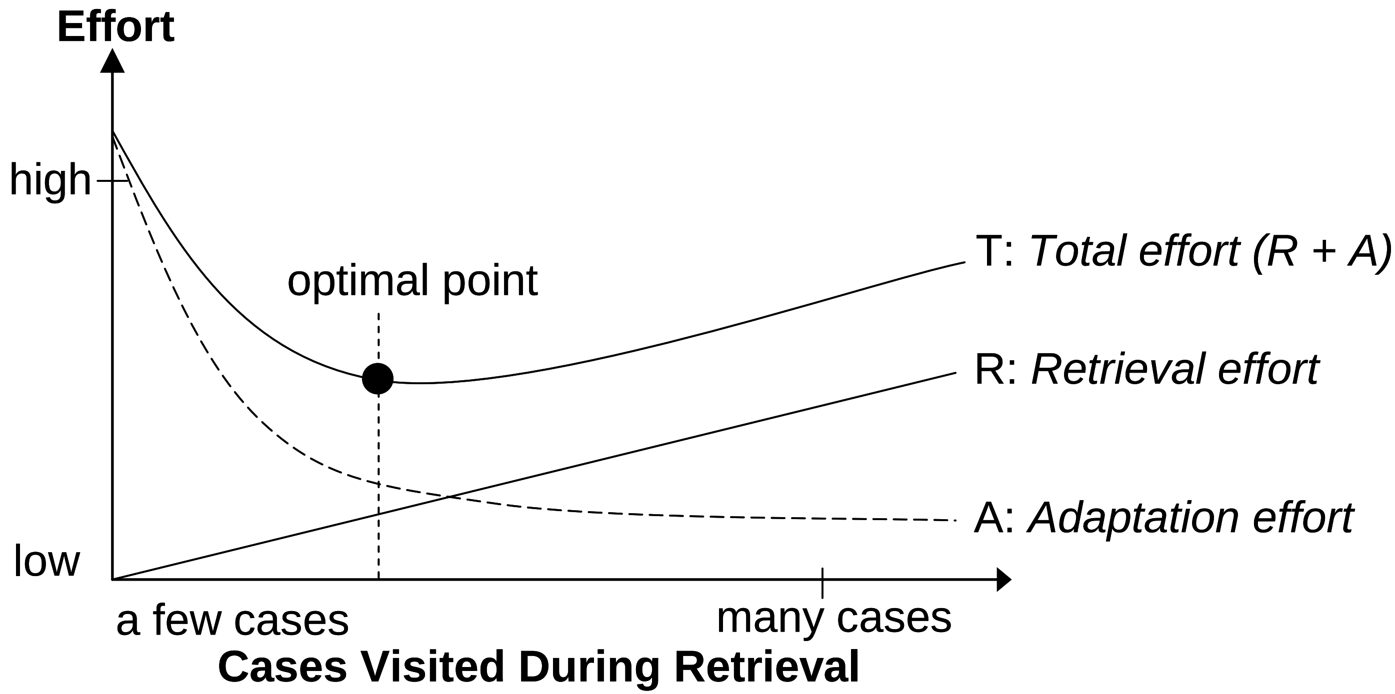

- If little time is spent on retrieval, then the adaptation effort is high
- If too much time is spent on retrieval, then the adaptation effort is low
- There is an optimal intermediate point between these two extremes

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

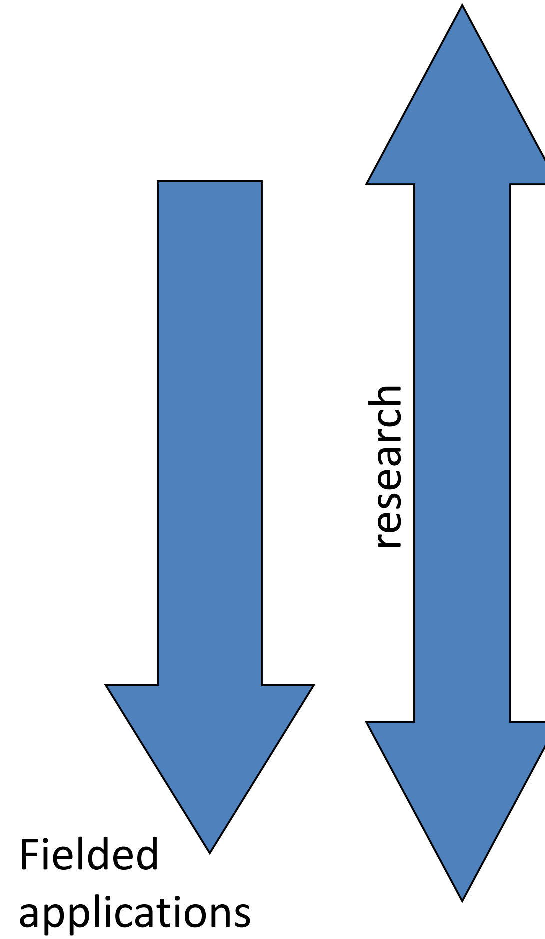

Main Topics of CBR Research ~ 10yr

- 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

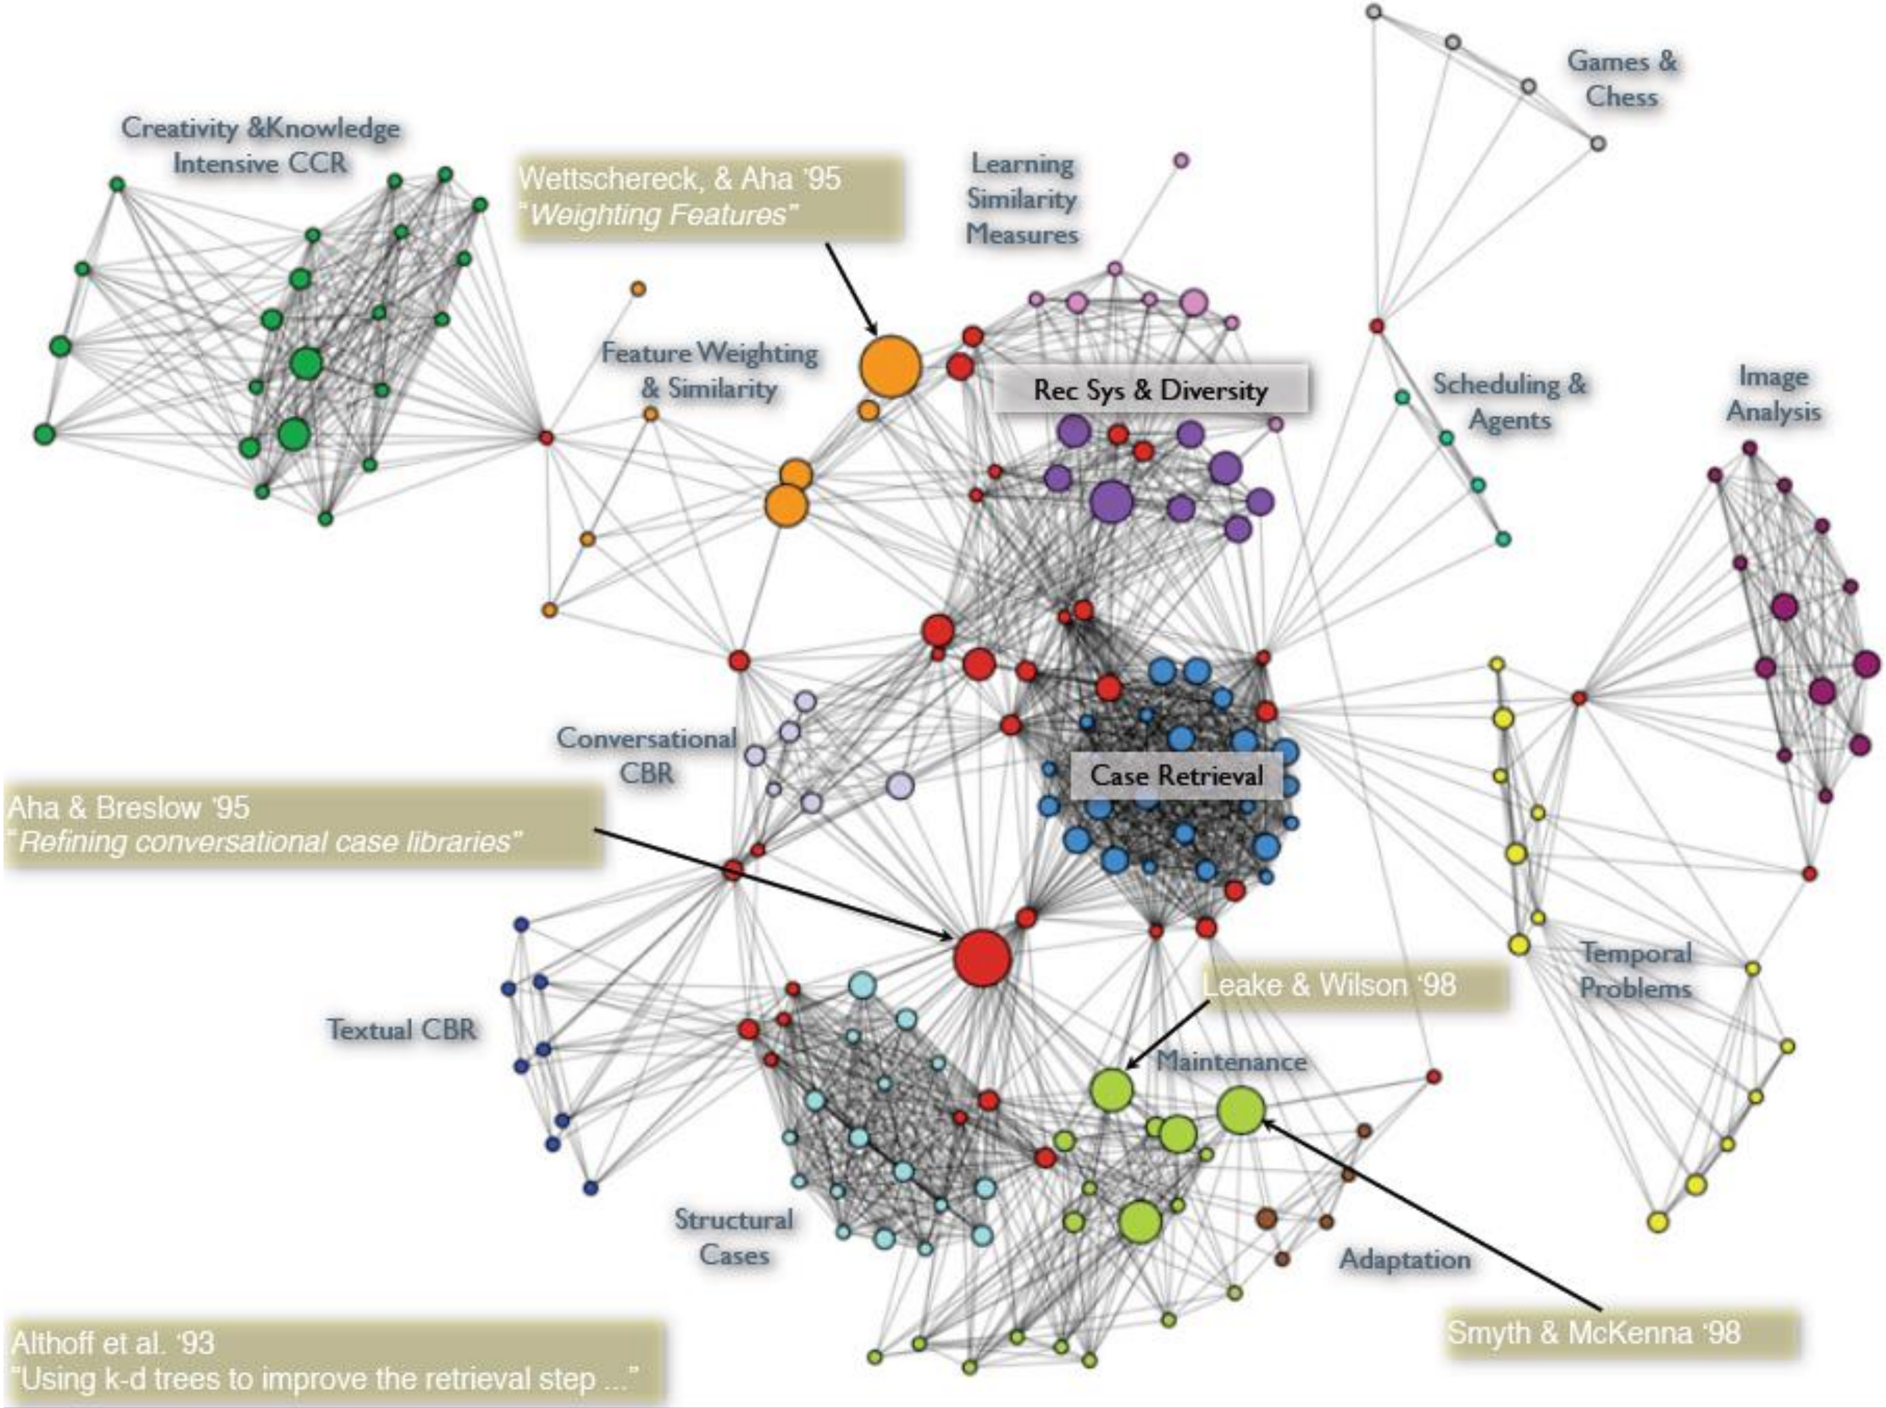
Games & Chess
Creativity &Knowledge Intensive CCR
Wettschereck, & Aha '95
"Weighting Features"
Learning Similarity Measures
Feature Weighting & Similarity
Rec Sys & Diversity
Scheduling & Agents
Image Analysis
Conversational CBR
Case Retrieval
Aha & Breslow '95
"Refining conversational case libraries"
Temporal Problems
Leake & Wilson '98
Textual CBR
Maintenance
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
management
classification
support
software
integrating
analysis
application
representation
problem
model
diagnosis
selection
collaborative
measures
prediction
structured
techniques
information
framework
conversational
supporting
textual
reuse
decision
explanation
feature
image
domain
strategies
architecture
memory
development
study
search
quality
tasks
solution
combining
experience
requirements
incremental
fuzzy
problems
process
hybrid
interactive
electronic
indexing
building
towards
time
data
bases
multiple
solutions
improving
applications
modelling
filtering
solving
adaptive
rule
applying
maintaining
web
re
domains
metric
environment
different
metrics
networks
autonomous
engineering
evaluation
integrated
general
cooperative
machine
method
new
competence
context
methods
assessment
diagnostic
theory
failure
semantic
dynamic
critiquing
interpretation
reinforcement
personalised
knowledge-intensive
bayesian
integration
evaluating
formal
medical
description
agents
generation
plan
task
object
improvement
algorithm
managing
function
improve
role
analogical
hierarchical
library
reasoner
complex
processing
comparison
neighbor
confidence
use
tool
product

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.

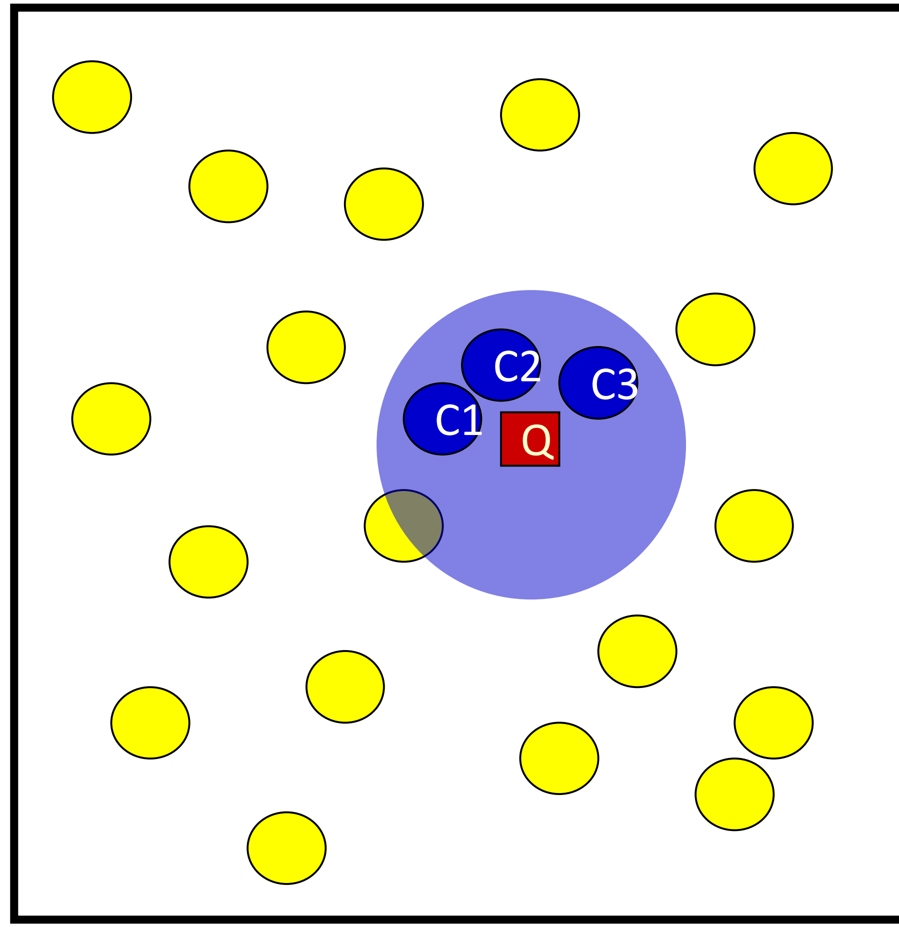

Lorraine McGinty and Barry Smyth
Department of Computer Science, University College Dublin

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.

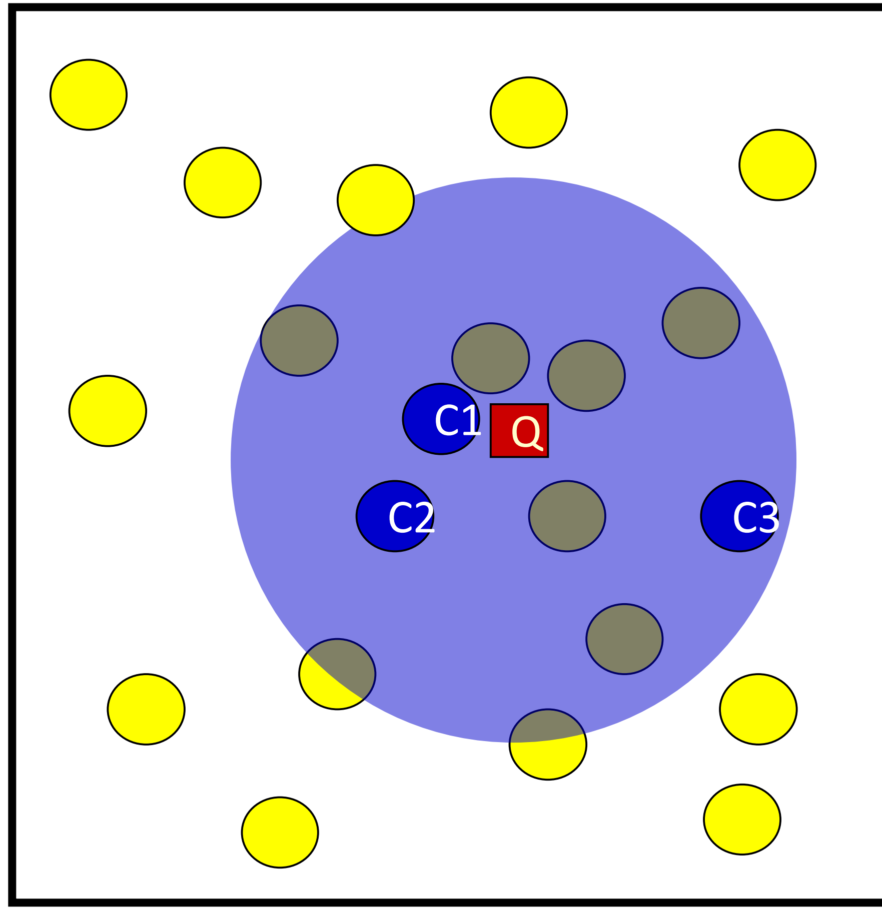

Lorraine McGinty and Barry Smyth
Department of Computer Science, University College Dublin

Dangers of Diversity Enhancement

- **Leap-Frogging the Target**
  - Problems occur when the target product is rejected as a retrieval candidate on diversity grounds.
  - $\Rightarrow$ Protracted dialogs.
- **Diversity is problematic in the region of the target product.**
  - Use similarity for fine-grained search.
- **Similarity is problematic when far from the target product.**
  - Use diversity to speed-up the search.

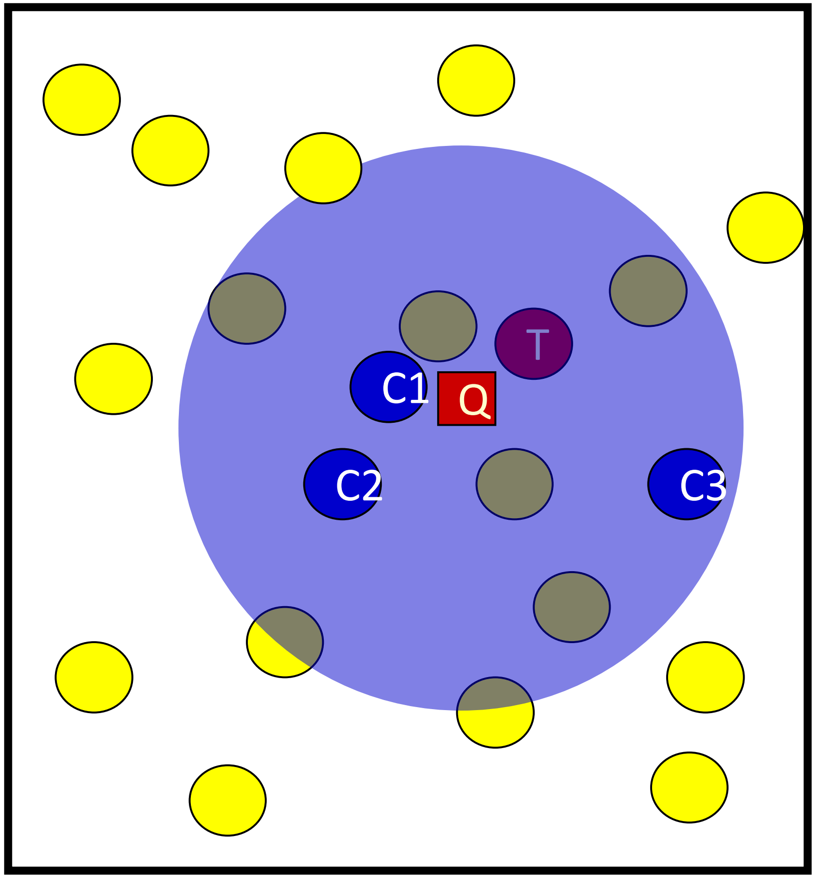

Lorraine McGinty and Barry Smyth
Department of Computer Science, University College Dublin

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

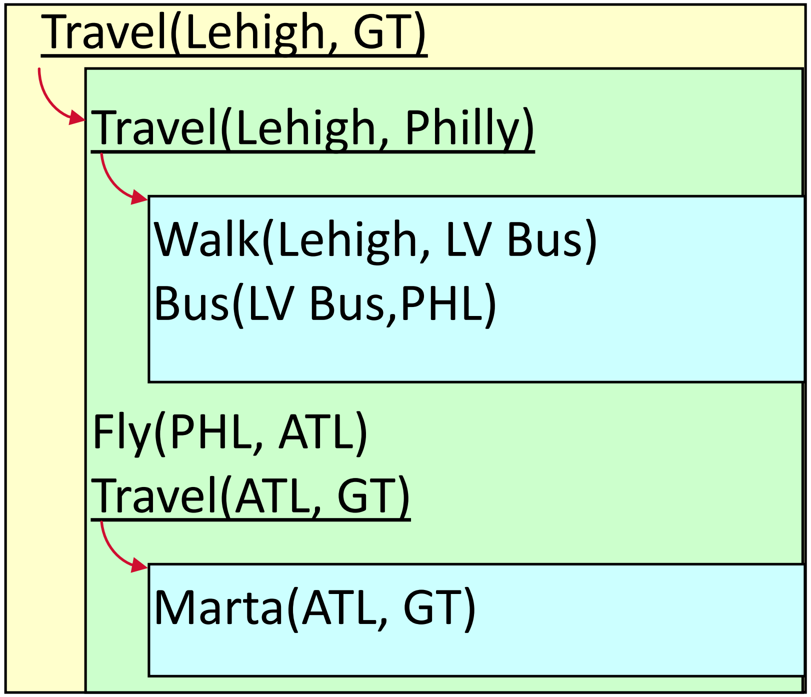

Knowledge Sources

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.

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

Recent uses at GT

- 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.
- 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.
- 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.
- 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

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 _________ __________".