Case-Based Recommendation

Case-Based Recommendation

Barry Smyth

P. Brusilovsky, A. Kobsa, and W. Nejdl (Eds.): The Adaptive Web, LNCS 4321, pp. 342376, 2007

Recommendation

Smyth describes two approaches to recommendation:

- Collaborative filtering: Each user rates items, the system recommends items users with similiar rating patterns have liked in the past. No data about items themselves.
  - Explicit
  - Implicit
- Content-based: Recommend based on similiar items (eg. to what the user has liked in the past).

Case-based Recommendation

Case-based recommenders implement a particular style of content-based recommendation, distinguished by:

- Product representation: Structured instead of unstructured (eg. recommending news articles based on keywords and textual search).

Case-based Recommendation

- Similarity: Case-based can use more sophisticated similarity because of structured data, compare keyword-based search for the query "$1000 6 mega-pixel DSLR"
- Specialised feature level similarity knowledge

Very well suited to product recommendation domains (esp. e-commerce) where detailed feature-based product "cases" are readily available.

Aside: A similarity metric

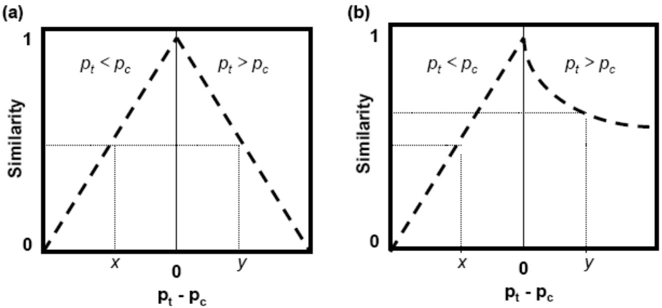

$$s_{\text{price}}(p_t, p_c) = 1 - \frac{p_t - p_c}{\max(p_t, p_c)}$$

Feature weight learning

Obvious place to apply machine learning of weights: The user can play the role of trainer: the product the user selects in the list of recommendations should have been placed at the top of the list.

Collaborative filtering applications

Normally a collaborative filtering recommender system can only evaluate the similarity between two profiles if they share ratings.

Consider a TV program recommender. If one user rates ER and another Frasier they can't be directly compared.

O'Sullivan et al. point out that global ratings patterns can be analysed to estimate the similarity between programmes like ER and Frasier. Using data-mining techniques shows, eg., 60% of the people who have liked ER also liked Fraiser, and they use this as a proxy for the similarity between these two programmes.

Similarity vs. Diversity

$$Diversity(c_1,...c_n) = \frac{\sum_{i=1..n}\sum_{j=i..n}(1 - Similarity(c_i,c_j))}{\frac{n}{2} * (n-1)}$$

The most similar cases will usually lack diversity. Eg. top vacation recommendations might all be at the same hotel in different weeks.

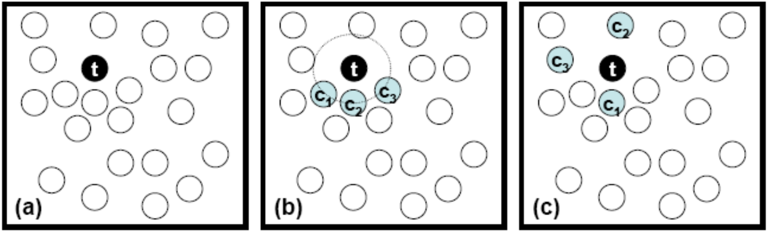

Diverse selection strategies

- Bounded random selection: select at random from *kb* nearest. Performs poorly.
- Similarity layers: adds little diversity.
- Bounded greedy selection
- Replace nearest neighbour case retrieval (more later)
- Others . . .

Similarity layers

When some of the nearest cases have near-equal similarity to query, we can increase diversity without trading off similarity.
Better suited to returning large number of recommendations.

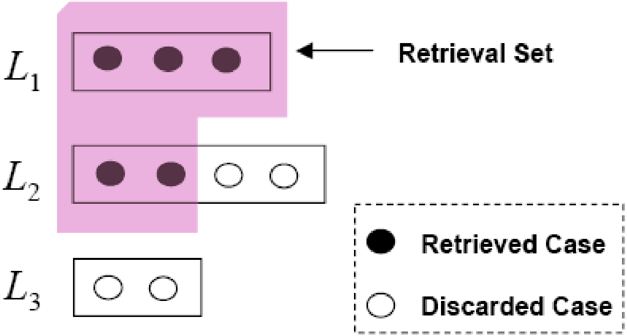

Bounded greedy selection 1

$$Quality(t,c,R) = Similarity(t,c) * RelDiversity(c,R)$$

$$\begin{aligned} RelDiversity(c,R) &= 1 \; if \, R = \{\}; \\ &= \frac{\sum_{i=1..m}(1 - Similarity(c,r_i))}{m}, otherwise \end{aligned}$$

Bounded greedy selection 2

For k diverse recommendations from the bk nearest (where b is the *bound*)

$C' \leftarrow nk$-nearest neighbours to query
$R \leftarrow \{\}$
for $i = 1$ to k **do**
  move $c \in C'$ with highest $\mathrm{Quality}(t, c, R)$ to R
end for

Compromise-driven retrieval

Definition : "A given case is more *acceptable* than another if it is more similar to the user's query and it involves a subset of the compromises that the other case involves"
Build a list of recommendations so that no recommendation is more acceptable than any other.
Provides "full coverage": when something is left out a better recommendation is included.
Unknown number of cases need to be retrieved.

Other retrieval/ranking techniques

- Shimazu's: Find a set of similar cases, then pick 3:
  c_1 : most similar
  c_2 : most dissimilar to c_1
  c_3 : most dissimilar to c_1 and c_2
- Order-based retrieval

Single-shot recommendation

Recommendations based on a single query. If users don't find what they want, they have to start again.

Conversational Recommenders

- Navigation by asking. (Typical Conversational CBR) "How much optical zoom do you need?" Narrow down case base until finally reaching recommendations.
  Symth says not always appropriate: users may not tolerate long exchanges of direct questions, or may not know answers.
- Navigation by proposing. Iteratively present possible recommendations.

Navigation by proposing

3 feedback alternatives:

- Ratings-based feedback (untypical)
- Preference-based feedback: Just pick a recommendation.
- Critique-based feedback: Specify how to modify recommendation

Critique-based feedback

Recommender provides recommendations one-at-a-time, the user provides feature constraints, eg. "cheaper", which act as filters on the set of cases most similar to the *current recommendation*. Critiques need not map directly onto single features. Suggest compound critiques based on remaining cases to speed progress.

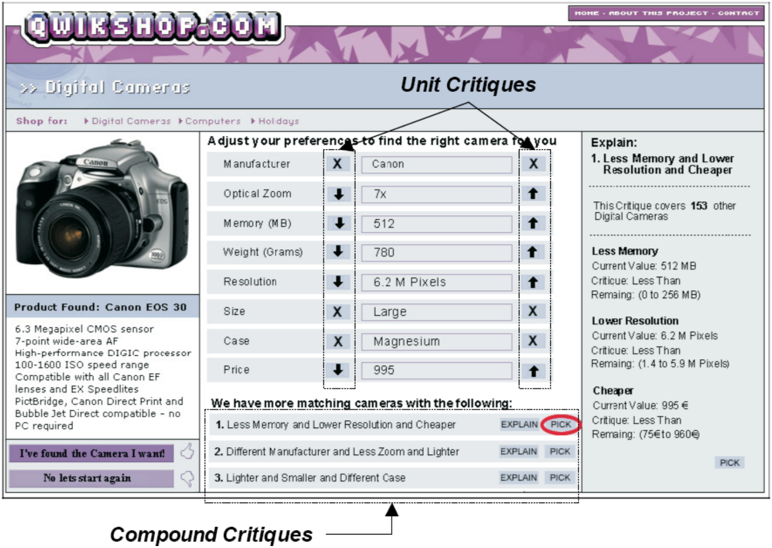
QWIKSHOP.COM
HOME · ABOUT THIS PROJECT · CONTACT
>> Digital Cameras
Unit Critiques
Shop for: Digital Cameras Computers Holidays
Adjust your preferences to find the right camera for you
Manufacturer X Canon X
Optical Zoom 7x
Memory (MB) 512
Weight (Grams) 780
Resolution 6.2 M Pixels
Size X Large X
Case X Magnesium X
Price 995
Product Found: Canon EOS 30
6.3 Megapixel CMOS sensor
7-point wide-area AF
High-performance DIGIC processor
100-1600 ISO speed range
Compatible with all Canon EF lenses and EX Speedlites
PictBridge, Canon Direct Print and Bubble Jet Direct compatible - no PC required
I've found the Camera I want!
No lets start again
We have more matching cameras with the following:
1. Less Memory and Lower Resolution and Cheaper EXPLAIN PICK
2. Different Manufacturer and Less Zoom and Lighter EXPLAIN PICK
3. Lighter and Smaller and Different Case EXPLAIN PICK
Explain:
1. Less Memory and Lower Resolution and Cheaper
This Critique covers 153 other Digital Cameras
Less Memory
Current Value: 512 MB
Critique: Less Than
Remaing: (0 to 256 MB)
Lower Resolution
Current Value: 6.2 M Pixels
Critique: Less Than
Remaing: (1.4 to 5.9 M Pixels)
Cheaper
Current Value: 995 €
Critique: Less Than
Remaing: (75€ to 960€)
PICK
Compound Critiques

Preference-based feedback

It's hard to infer reasons for the user's preference of one choice over another. Also want to be efficient: not waste user time.

- Most straighforward: take the user's choice as the new query and find similar cases. Don't gain much.
- Transfer features from the user's selection to the query if they distinguish the selection.
- Modify feature weights by guessing which features are reponsible for the selection, eg. statistical inferencing.

Technique combination

“Natural” to combine critique-based and preference-based feedback.

Can add increased diversity to recommendations if the user doesn’t think they are improving.

Are recommendations failing to improve?

Include the last product which was selected/recommended in the new batch. If the user reselects, supposedly there was no improvement. Shown to potentially reduce recommendation sessions.

Personalisation

If a recommender can learn a repeat user's long-term preferences, should be able to determine some constraints or query features automatically. Also, can personalise the ranking of retrieved cases eg. CASPER uses the similarity of each recommendation to recommendations the users has previously rated or accepted, based on their rating.