Explanations in recommender systems

Explanations in recommender systems

Motivation

- “The digital camera *Profishot* is a must-buy for you because . . . .”
- Why should recommender systems deal with explanations at all?
- The answer is related to the two parties providing and receiving recommendations:
  - A selling agent may be interested in promoting particular products
  - A buying agent is concerned about making the right buying decision

Types of explanations

- **(Brewer et al. 1998) distinguishes among**
  - **Functional,**
    - "The car type Jumbo-Family-Van of brand Rising-Sun would be well suited to your family because you have four children and the car has seven seats"
  - **Causal,**
    - "The light bulb shines because you turned it on"
  - **Intentional,**
    - "I washed the dishes because my brother did it last time"
    - "You have to do your homework because your dad said so"
  - **Scientific explanations**
    - Express relations between the concepts formulated in various scientific fields and are typically based on refutable theories

What is an explanation in recommender systems?

Additional information to explain the system's output following some objectives

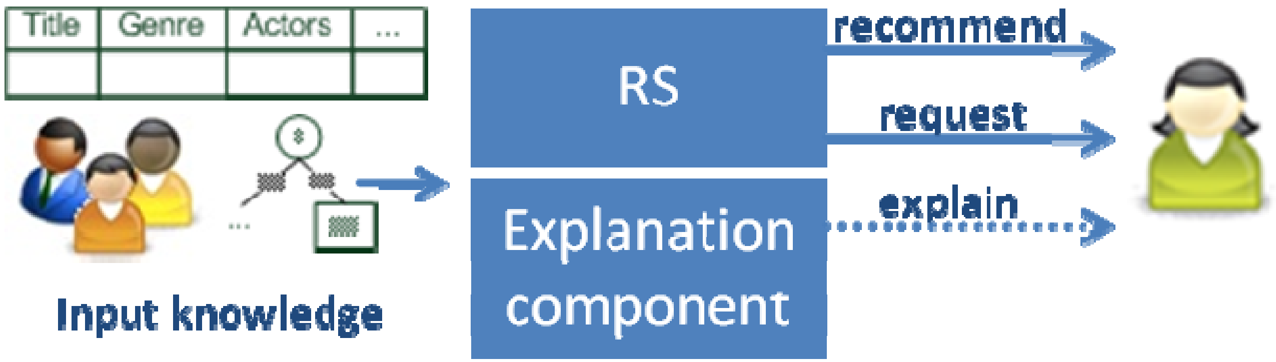

The goals for providing explanations (1) (Tintarev and Masthoff 2007)

- **Transparency**
  - Provide information so the user can comprehend the reasoning used to generate a specific recommendation
  - Provide information as to why one item was preferred over another
- **Validity**
  - Allow a user to check the validity of a recommendation
  - Not necessarily related to transparency
    - E.g., a neural network (NN) decides that product matches to requirements. Transparent disclosure of NN's computations, will not help, but a comparison of required and offered product features allows customer to judge the recommendation's quality.

The goals for providing explanations (2)

- **Trustworthiness**
  - Trust building can be viewed as a mechanism for reducing the complexity of human decision making in uncertain situations
  - Reduce the uncertainty about the quality of a recommendation
- **Persuasiveness**
  - Persuasive explanations for recommendations aim to change the user's buying behavior
  - E.g., a recommender may intentionally dwell on a product's positive aspects and keep quiet about various negative aspects
- **Effectiveness**
  - The support a user receives for making high-quality decisions
  - Help the customer discover his or her preferences
  - Help users make better decisions

The goals for providing explanations (3)

- **Efficiency**
  - Reduce the decision-making effort
  - Reduce the time needed for decision making
  - Another measure might also be the perceived cognitive effort
- **Satisfaction**
  - Improve the overall satisfaction stemming from the use of a recommender system
- **Relevance**
  - Additional information may be required in conversational recommenders
  - Explanations can be provided to justify why additional information is needed from the user

The goals for providing explanations (4)

- **Comprehensibility**
  - Recommenders can never be sure about the knowledge of their users
  - Support the user by relating the user's known concepts to the concepts employed by the recommender
- **Education**
  - Educate users to help them better understand the product domain
  - Deep knowledge about the domain helps customers rethink their preferences and evaluate the pros and cons of different solutions
  - Eventually, as customers become more informed, they are able to make wiser purchasing decisions
- The aforementioned aims for generating explanations can be **interrelated**
  - Persuasiveness**+** → Trust**-**
  - Effectiveness**+** → Trust**+**
  - …

Explanations in general

- ***How?*** **and** ***Why?*** **explanations in expert systems**
- **Form of abductive reasoning**
  - Given: $KB \models_{RS} i$ (item i is recommended by method RS)
  - Find $KB' \subseteq KB$ s.t. $KB' \models_{RS} i$
- **Principle of succinctness**
  - Find smallest subset of $KB' \subseteq KB$ s.t. $KB' \models_{RS} i$

    i.e. for all $KB'' \subset KB'$ holds $KB'' \not\models_{RS} i$
- **But additional filtering**
  - Some parts relevant for deduction, might be obvious for humans

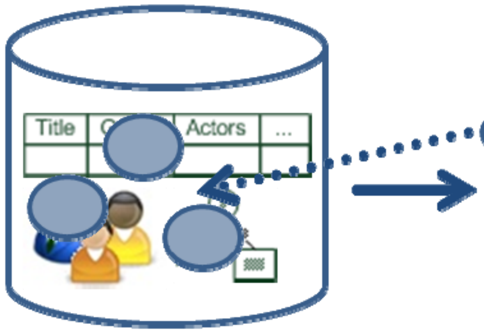

Knowledge base

item	score
i2	0.7
i3	0.3
...	...

[Friedrich & Zanker, AI Magazine, 2011]

Taxonomy for generating explanations in RS

Major design dimensions of current explanation components:

- **Category of reasoning model for generating explanations**
  - White box
  - Black box
- **RS paradigm for generating explanations**
  - Determines the exploitable semantic relations
- **Information categories**

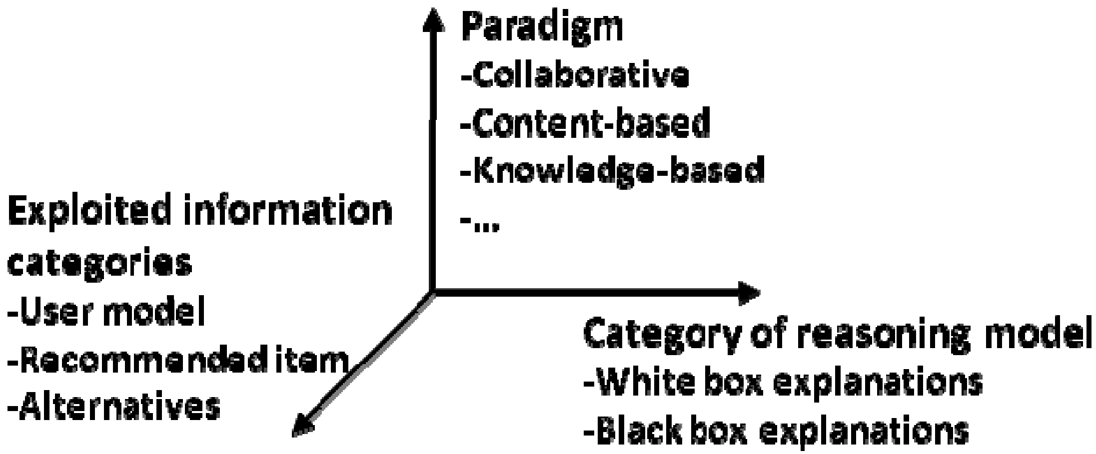

Archetypes of KB

- **Classes of objects**
  - Users
  - Items
  - Properties
- **N-ary relations between them**

- **Collaborative filtering**
  - Neighborhood based CF (a)
  - Matrix factorization (b)
    - Introduces additional factors as proxies for determining similarities

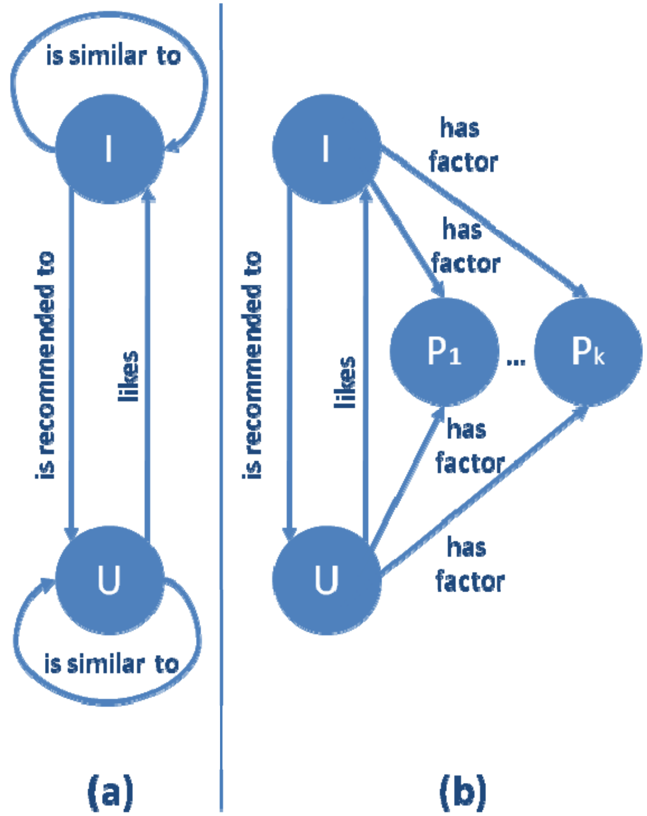

Examples

- **Similarity between items**

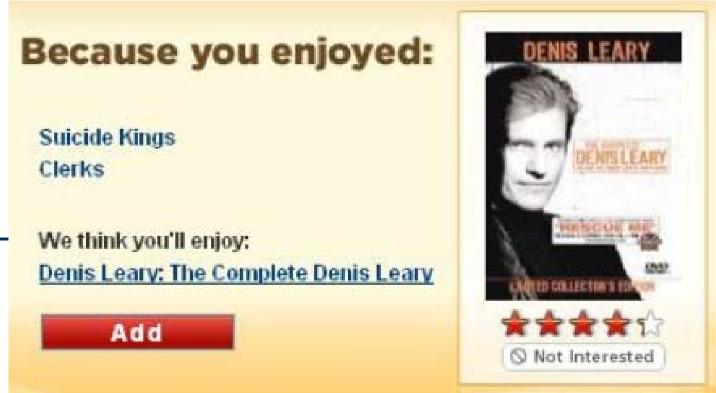

- **Similarity between users**

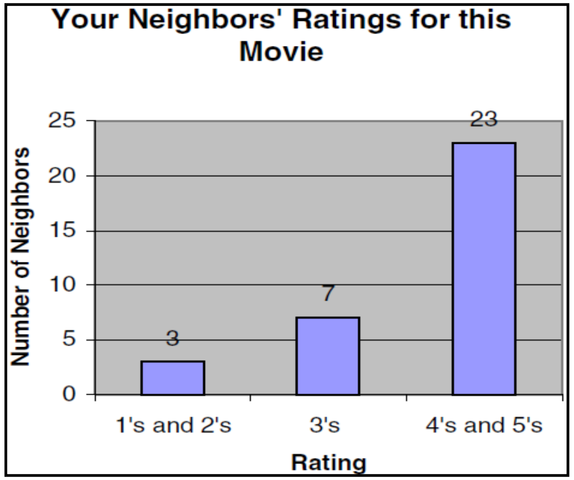

- **Tags**
  - Tag relevance (for item)
  - Tag preference (of user)

Explanations in collaborative filtering recommenders

- Explicit recommendation knowledge is not available
- Recommendations based on CF cannot provide arguments as to why a product is appropriate for a customer or why a product does not meet a customer's requirements
- The basic idea of CF is to mimic the human word-of-mouth recommendation process
- Therefore, give a comprehensible account of how this word-of-mouth approach works:
  - Customers rate products
  - The CF locates customers with similar ratings (i.e., tastes), called neighbors
  - Products that are not rated by a customer are rated by combining the ratings of the customer’s neighbors

Evaluating explanation interfaces (Herlocker et al. 2000)

- Herlocker et al. (2000) examined various implementations of explanation interfaces in the domain of the "MovieLens" system
- Twenty-one variants were evaluated
- Customers were asked, on a scale of 1 to 7, how likely they would be to go to see a recommended movie after a recommendation for this movie was presented and explained by one of the twenty-one different explanation approaches
- They also included the base case in which no additional explanation data were presented
- In addition to the base case, an explanation interface was designed that just output the past performance of the recommendation system – for instance,
  - "MovieLens has provided accurate predictions for you 80% of the time in the past"

The results of the study by Herlocker et al. (2000)

- The best-performing explanation interfaces are based on the ratings of neighbors

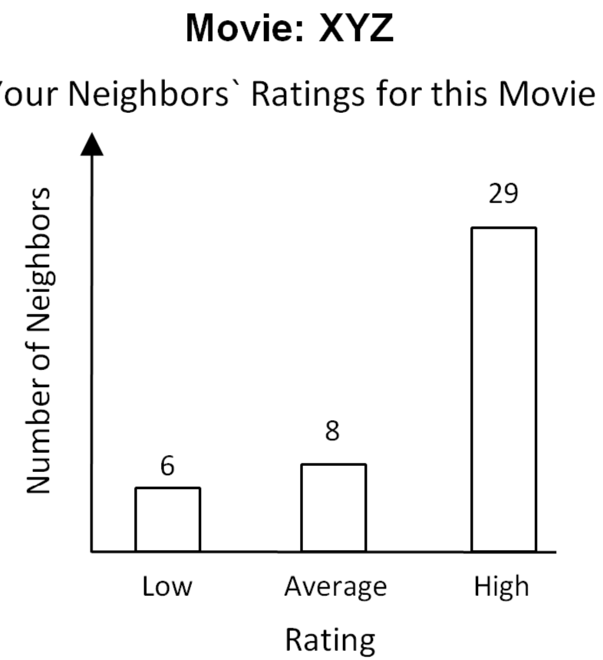

Movie: XYZ

Personalized Prediction: ****

Your Neighbors` Ratings for this Movie

Rating	Number of Neighbors
★	2
★★	4
★★★	8
★★★★	20
★★★★★	9

- In these cases similar neighbors liked the recommended film, and this was comprehensibly presented. The histogram performed better than the table

The results of the study by Herlocker et al. (2000)

- Recommenders using the simple statement about the past performance of MovieLens were the second best performer
- Content-related arguments mentioning the similarity to other highly rated films or a favorite actor or actress were among the best performers
- Poorly designed explanation interfaces decreased the willingness of customers to follow the recommendation, even compared with the base case
- Too much information has negative effects; poor performance was achieved by enriching the data presented in histograms with information about the proximity of neighbors
- Interestingly, supporting recommendations with ratings from domain authorities, such as movie critics, did not increase acceptance

Tagsplanations (Vig et al. 2010)

- **Tag relevance (for item)**
  - X .. Set of (mean-adjusted) ratings for I among users of tag t (U_{ti})
  - Y .. Inferred tag preferences
  - Correlation between user's preference for the tag and their preference for the item

$$tag\_rel(t,i) = \begin{cases} pearson(X,Y), & if \quad t \quad tagged \quad to \quad i \\ 0, & otherwise \end{cases}$$

$$X = \{r_{u,i} - \bar{r}_u : u \in U_{t,i}\}$$

$$Y = \{tag\_pref(u,t) - \bar{r}_u : u \in U_{t,i}\}$$

Tagsplanations (Vig et al. 2010)

- **Tag preference (of user)**
  - I_u .. Items rated by u
  - k .. Smoothing constant, the smoothing constant k accounts for users who have rated few items with a given tag. This smoothing serves to bring the computed tag preference closer to the user’s average rating, because ratings of a small number of items may not properly reflect a user’s tag preference
  - tag_share(t,i) ..The tag_share of a tag t applied to an item i is the number of times t has been applied to i, divided by the number of times any tag has been applied to i
  - Relative importance of tag t for items the user knows and likes compared to the relative importance of t for the known items

$$tag\_pref(u,t) = \frac{(\sum_{i \in I_u} r_{u,i} \times tag\_share(t,i)) + \bar{r}_u \times k}{(\sum_{i \in I_u} tag\_share(t,i)) + k}$$

Archetypes of KB

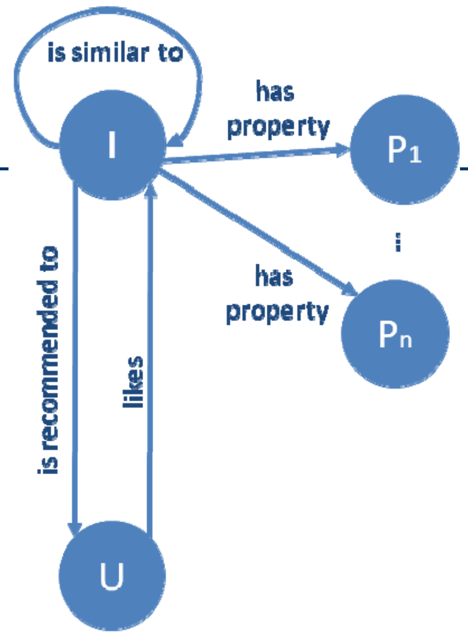

- **Content-based**
  - Properties characterizing items
  - TF*IDF model

- **Knowledge based**
  - Properties of items
  - Properties of user model
  - Additional mediating domain concepts

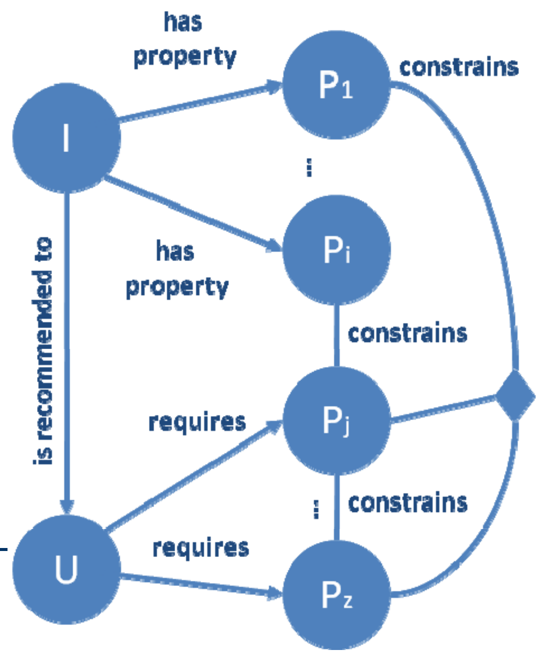

Explanations in case-based recommenders

- The generation of solutions in case-based recommenders is realized by identifying the products that best fit a customer's query
- Each item of a product database corresponds to a case
- Customer query puts constraints on the attributes of products
  - For example, a customer is interested only in digital cameras that cost less than a certain amount of money

Explanations in case-based recommenders

- In particular, given a query Q about a subset A_Q of attributes A of a case (product) description, the similarity of a case C to Q is typically defined (see McSherry 2005) as

$$sim(C,Q) = \sum_{a \in A_Q} w_a sim_a(C,Q)$$

- The function $sim_a(C,Q)$ describes the similarity of the attribute values of the query Q and the case C for the attribute a
- This similarity is weighted by w_a, expressing the importance of the attribute to the customer
- A recommendation set is composed of all cases C that have a maximal similarity to the query Q

Explaining solutions (1)

- The typical approach used to answer a why-question is to compare the presented case with the customer requirements and to highlight which constraints are fulfilled and which are not (McSherry 2003b)
- Example:

id	price	mpix	Opt-zoom	LCD-size	movies	sound	waterproof
p1	148	8.0	4x	2.5	no	no	yes
p2	182	8.0	5x	2.7	yes	yes	no
p3	189	8.0	10x	2.5	yes	yes	no
p4	196	10.0	12x	2.7	yes	no	yes
p5	151	7.1	3x	3.0	yes	yes	no
p6	199	9.0	3x	3.0	yes	yes	no
p7	259	10.0	3x	3.0	yes	yes	no
p8	278	9.1	10x	3.0	yes	yes	yes

Explaining solutions (2)

- If a customer is interested in digital cameras with a price less than 150, then p1 is recommended.

id	price	mpix	Opt-zoom	LCD-size	movies	sound	waterproof
p1	148	8.0	4x	2.5	no	no	yes
p2	[illegible]	8.0	5x	2.7	yes	yes	no
p3	[illegible]	8.0	10x	2.5	yes	yes	no
p4	196	10.0	12x	2.7	yes	no	yes
p5	151	7.1	3x	3.0	yes	yes	no
p6	199	9.0	3x	3.0	yes	yes	no
p7	259	10.0	3x	3.0	yes	yes	no
p8	278	9.1	10x	3.0	yes	yes	yes

Why?

Explaining solutions (3)

- The weights of the attributes can be incorporated into the answers
- If the customer requires a price less than 160 and LCD size of more than 2.4 inches, where LCD size is weighted much more than price, then p5 is recommended

id	price	mpix	Opt-zoom	LCD-size	movies	sound	waterproof
p1	148	8.0	4x	2.5	no	no	yes
p2	182	8.0	5x	2.7	yes	yes	no
p3	189	8.0	10x	2.5	yes	yes	no
p4	196	10.0	12x	2.7	yes	no	yes
p5	151	7.1	3x	3.0	yes	yes	no
p6	199	9.0	3x	3.0	yes	yes	no
p7	[illegible]	[illegible]	[illegible]	[illegible]	yes	yes	no
p8	278	9.1	10x	3.0	yes	yes	yes

Why?

Explaining solutions (4)

- The requirements of a customer might be too specific
- Why-explanations provide information about the violated constraints
- For example, if the customer requires a price less than 150 and a movie function, then no product fulfills these requirements.

id	price	mpix	Opt-zoom	LCD-size	movies	sound	waterproof
p1	148	8.0	4x	2.5	no	no	yes
p2	182	8.0	5x	2.7	yes	yes	no
p3	189	8.0	10x	2.5	yes	yes	no
p4	196	10.0	12x	2.7	yes	no	yes
p5	151	7.1	3x	3.0	yes	yes	no
p6	199	9.0	3x	3.0	yes	yes	no
p7	259	10.0	3x	3.0	yes	yes	no
p8	278	9.1	10x	3.0	yes	yes	yes

Most similar products

Explaining solutions (5)

- p1 and p5 can be considered as most similar products for a given similarity function, although one of the user requirements is not satisfied
- A why-explanation for p1 would be,
  - "p1 is within your price range but does not include your movie requirement."
- Techniques can be used to generate minimal sets of customer requirements that explain why no products fit,
  - Employed to propose minimal changes to the set of requirements such that matching products exist

Explanations in constraint-based recommenders

- **Answering two typical questions a customer is likely to pose (Buchanan and Shortliffe 1984)**
- **Why-explanations**
  - In the case that the recommendation process requires input from the customer, the customer could ask why this information is needed
- **How-explanations**
  - When a recommender proposes a set of solutions (e.g., products) then a customer might ask for an explanation why a proposed solution would be advantageous for him or her
- **Next we examine methods for answering these two types of questions by exploiting the reasoning methods of constraint-based systems**

An example from the car domain (1)

- **Two different packages are available for a car**
  - Business package
  - Recreation package
- **The customer expresses requirements based on functions:**
  easy parking, towing and hands-free mobile communication
- **The customer can decide if he or she wants one, both, or neither of these packages**
- **Recreation package**
  - Includes a coupling device for *towing* trailers and a video camera on the rear of the car, which allows the driver to ascertain the distance to an obstacle behind the car
  - This camera supports the customer oriented product function *easy parking*

An example from the car domain (2)

- **Business package**
  - Includes a radio with a GSM telephone (GSM radio), which supports *hands-free mobile communication*
  - Includes a sensor system in the back bumper, which also supports *easy parking*
  - However, the sensor system is incompatible with the recreation package for technical reasons
  - From the customer's point of view, the video camera and the sensor system provide the same functionality. Therefore, if the customer orders the business package and the recreation package, the car includes the video camera, which implements the *easy parking* function
  - In this configuration, the sensors are not only forbidden (because they are incompatible with the coupling device), but also dispensable

Model domain as a constraint satisfaction problem (1)

- Set of variables *V* = {*biz-pack, rec-pack, GSM-radio, sensor, video, coupling-device, free-com, easy-parking, towing*}
- It is assumed that for each variable the domain is {*y, n*}
- Further constraints are specified by the following tables:
- $c_{r,v}$: If *rec-pack* is chosen, then *video* must also be included (and vice versa).

rec-pack	video
y	y
n	n

Model domain as a constraint satisfaction problem (2)

- $c_{b,r,s}$: If *biz-pack* is chosen and *rec-pack* is not chosen, then *sensor* is included. *rec-pack* and *sensor* are incompatible:

biz-pack	rec-pack	sensor
y	y	n
y	n	y
n	y	n
n	n	n
n	n	y

Model domain as a constraint satisfaction problem (3)

- $c_{v,s,e}$: If *video* or *sensor* is included, then *easy-parking* is supported (and vice versa)

video	sensor	easy-parking
n	n	n
y	n	y
n	y	y
y	y	y

Constraint network of car example

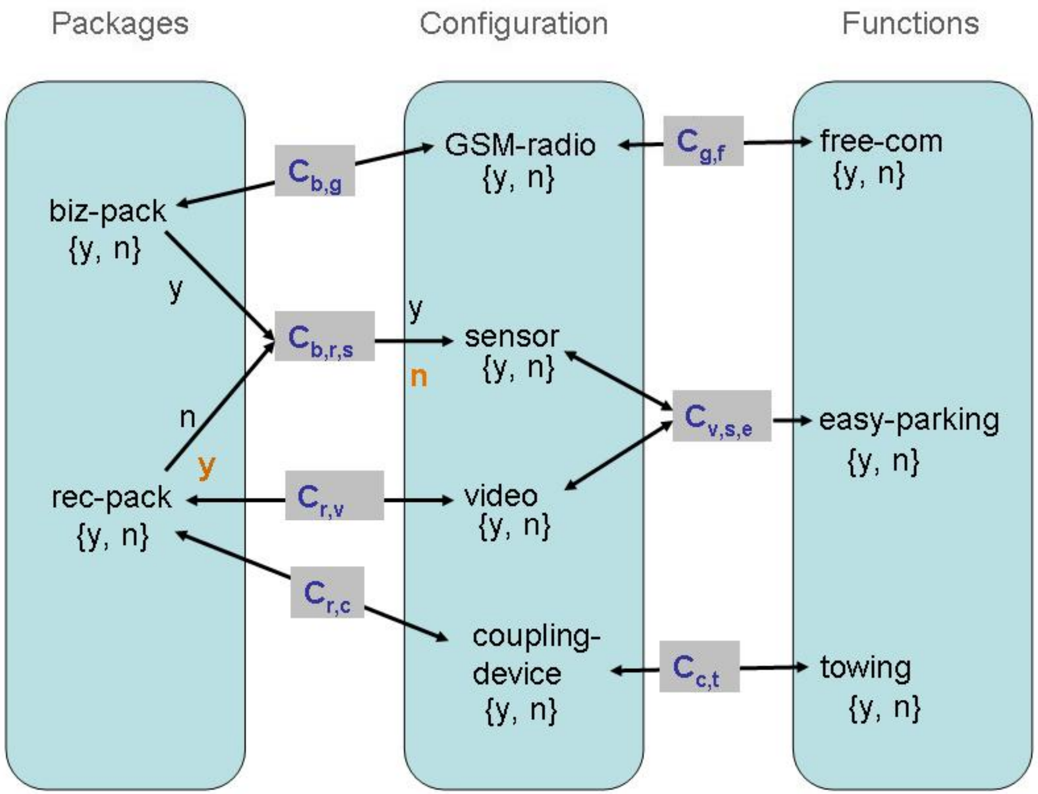

Explaining solutions

- If the customer sets {*free-com* = *y*, *towing* = *y*}, then the solution to the constraints $C = \{c_{r,v}, c_{b,r,s}, c_{v,s,e}, c_{b,g}, c_{g,f}, c_{r,c}, c_{c,t}\}$ representing the configured car would be to assign
  – {*video* = *y*, *sensor* = *n*, *GSM-radio* = *y*, *coupling-device* = *y*, *easy-parking* = *y*, *free-com* = *y*, *towing* = *y*, *biz-pack* = *y*, *rec-pack* = *y*}
  – This solution is presented to the customer
- If the customer asks which package led to the parking capabilities of the specific configured car…
  – ***Easy parking*** is supported because the car comes with a ***video camera***
  – This ***video camera*** is included because it is included in the ***recreation package***

User input and solution

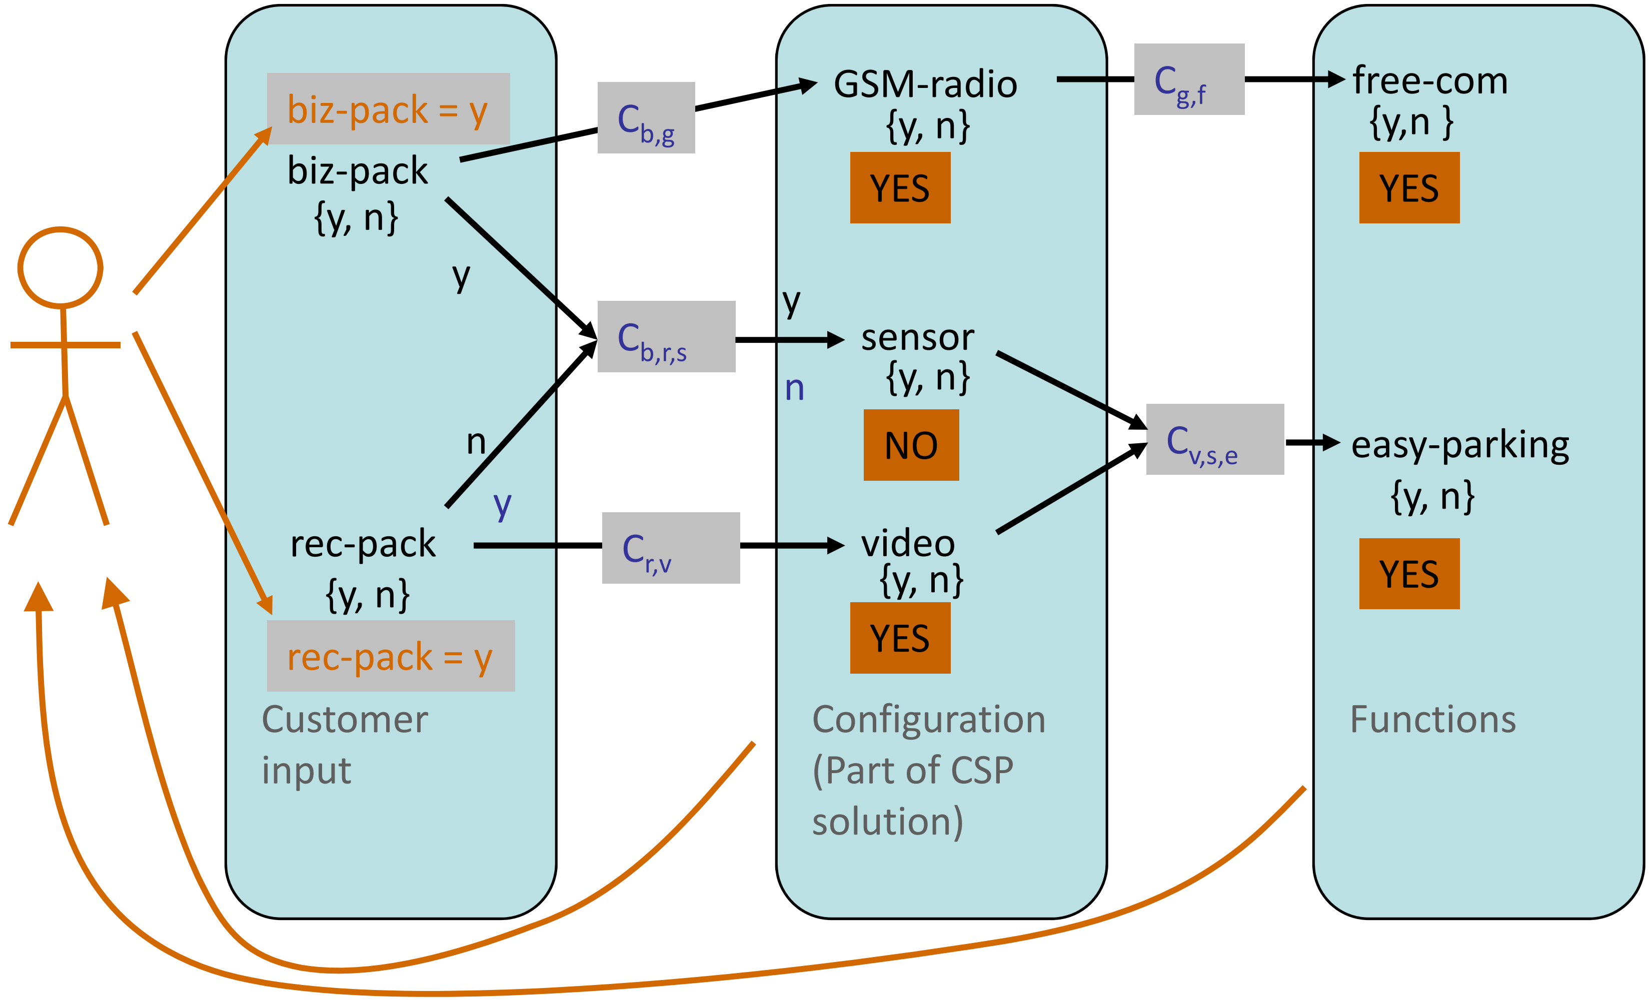

Why easy-parking? Some basics:

Constraint satisfaction Problem (CSP) defined by:
$\mathcal{C}$ (Constraints), $\mathcal{V}$ (Variables), $\mathcal{D}$ (Domain)

Let $(\mathcal{C},\mathcal{V},\mathcal{D})$ be a consistent CSP, **Φ** a constraint:

- A subset **C** of $\mathcal{C}$ is called an explanation for **Φ** iff **C** $\models$ **Φ**.
  (**Φ** is true in all models of **C**.)
- **C** is a minimal explanation for **Φ** iff
  no proper subset of **C** is an explanation for **Φ**.

Why easy-parking?

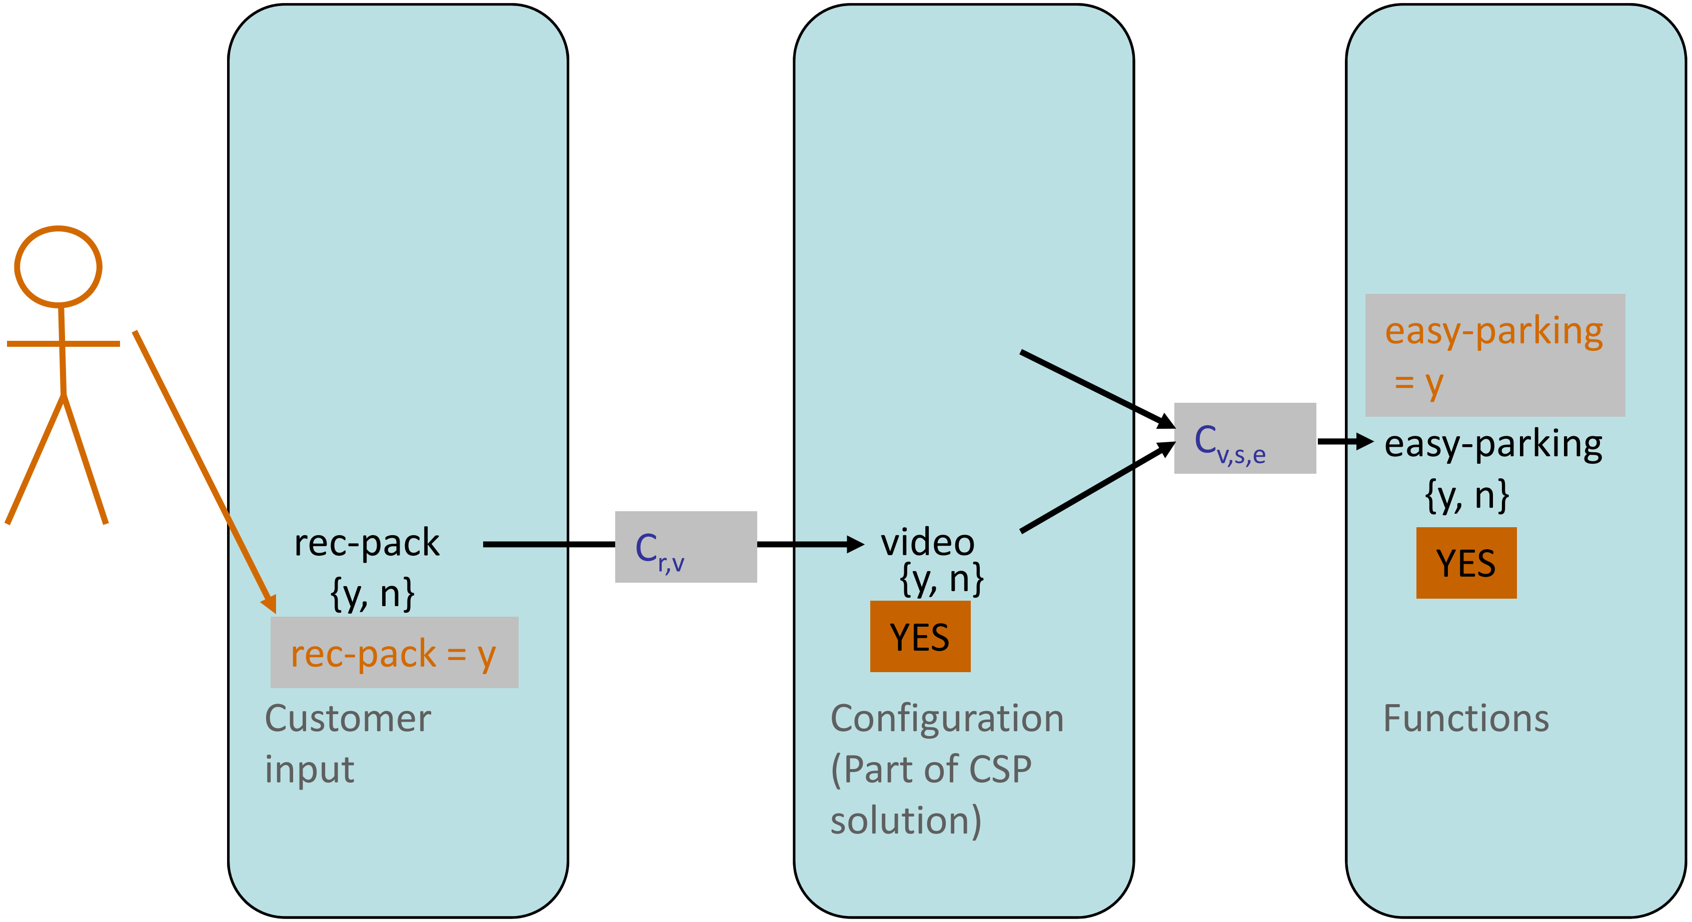

Because rec-pack = y

Why easy-parking?

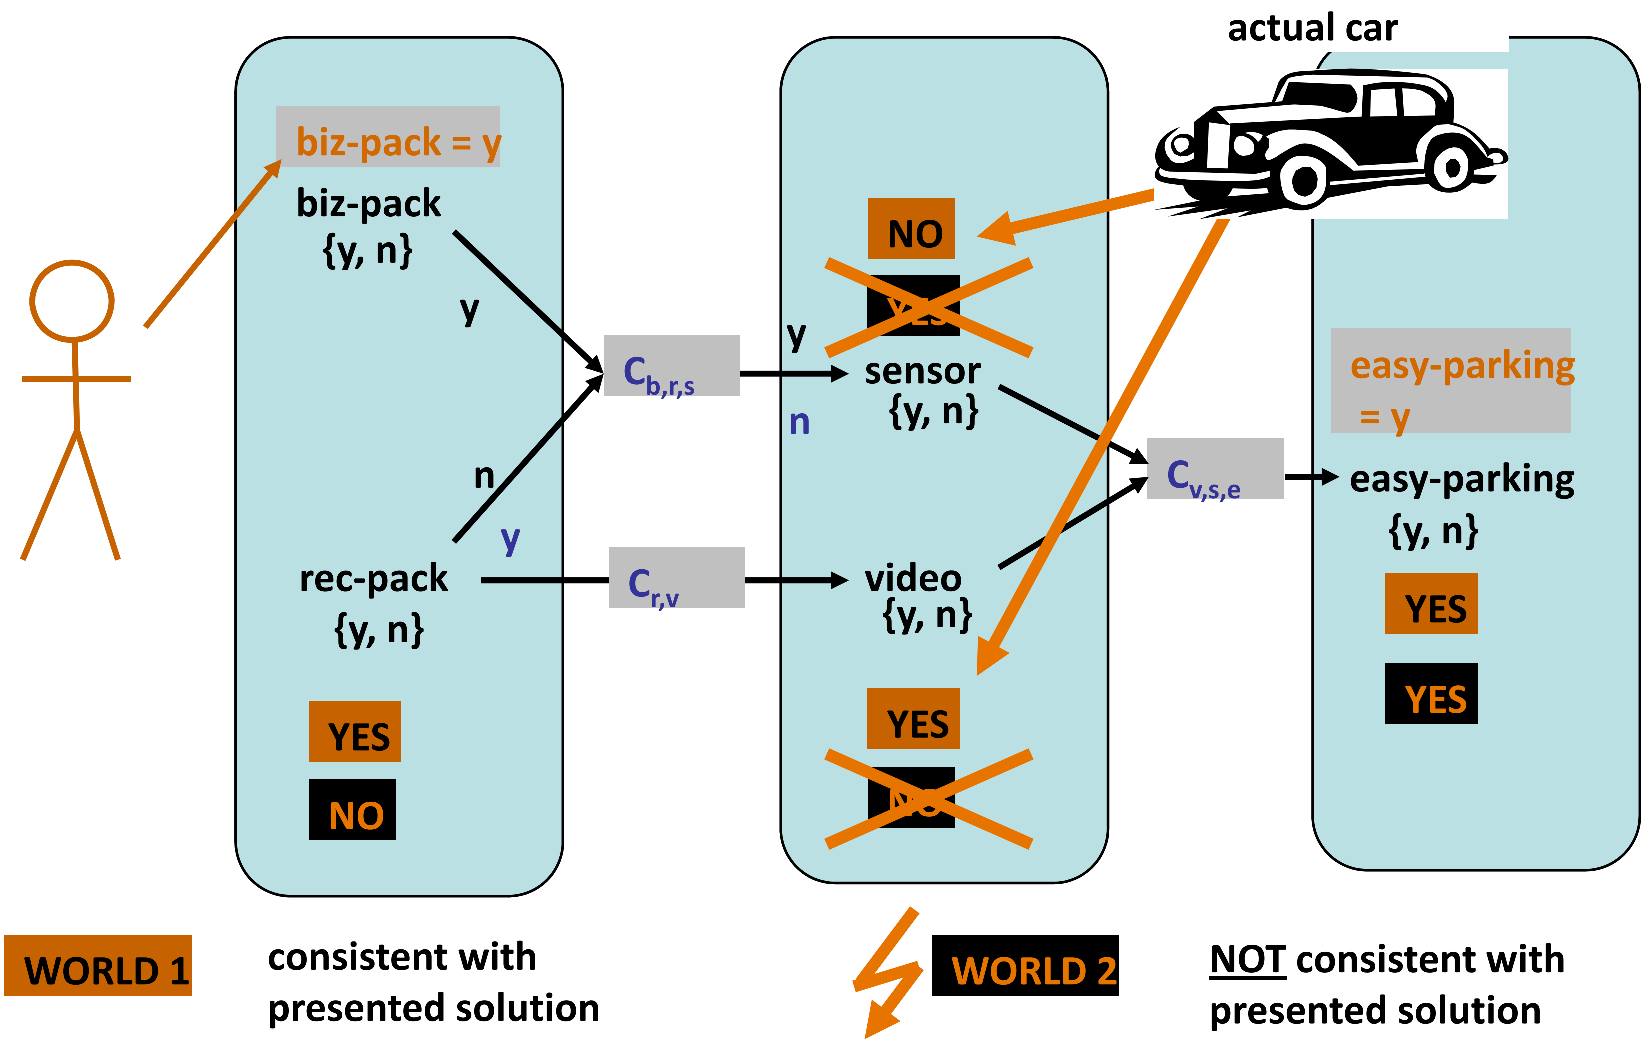

Because biz-pack = y, in world 1 and 2 easy-parking = y,
but world 2 does not correspond to original solution

Principal idea

Let $(\mathcal{C},\mathcal{V},\mathcal{D})$ be a consistent CSP, Φ a constraint:

$\mathcal{S}$: a specific solution (presented to the user).

A subset C of $\mathcal{C}$ is called an "explanation" for Φ iff $C \models \Phi$

and

the solutions s of C are consistent with $\mathcal{S}$ (i.e. agree on the value assignment) with respect to the **relevant** variables.

Details see (G. Friedrich 2004)

Application case for KB recommendation and explanation

Thermencheck.com (hot spring resorts)

Recommendation

10 EMPFOHLENE THERMEN

schließen

Österreich
Längenfeld

» zur Therme

Diese Therme entspricht zu 83% den von Ihnen gesuchten Kriterien

Warum wurde Ihnen diese Therme empfohlen:

Die Therme AQUA DOME - Tirol Therme Längenfeld ist gut für Familien mit Kindern geeignet. Der Service umfasst unter anderem Kleinkinderbereich, Kinderanimation, eigener Kinder- und Familienbereich mit 90m Rutsche, Unterwasserkamera, kostenloser Kinderanimation und -betreuung.. Spass und Fun kommen bei Wasserrutschen, Strömungskanal, Wasserfall nicht zu kurz. Die Therme kann als mittelgroÃŸ bezeichnet werden mit rund 1000.0 Liegebetten. Sie erfüllt laut Angaben des Betreibers die geforderten Suchkriterien: Familie, Fun, Wellness, Kulinarik, Kuscheln. Besonders hervorzuheben ist die Wirkung des Wassers zur Linderung folgender Beschwerden: Muskelerkrankungen. Hinzuweisen ist, dass über die Linderung von Herzerkrankungen nichts bekannt ist. Kulinarisch bietet sie biologisches Essen, aber leider nicht wie gewünscht koscheres Essen. Im Detail ist das Wellnessangebot wie folgt: Saunen, Dampfbäder, Faltenunterspritzungen, Gymnastikprogramme, aber leider nicht wie gewünscht Hautglättungen.

Explanation

It offers services for families with small children, such as X, Y and Z.

It is a spa resort of medium size offering around 1000 beds.

The water has favorable properties for X, but it is unknown if it also cures Y.

It offers organic food, but no kosher food.

Results from testing the explanation feature

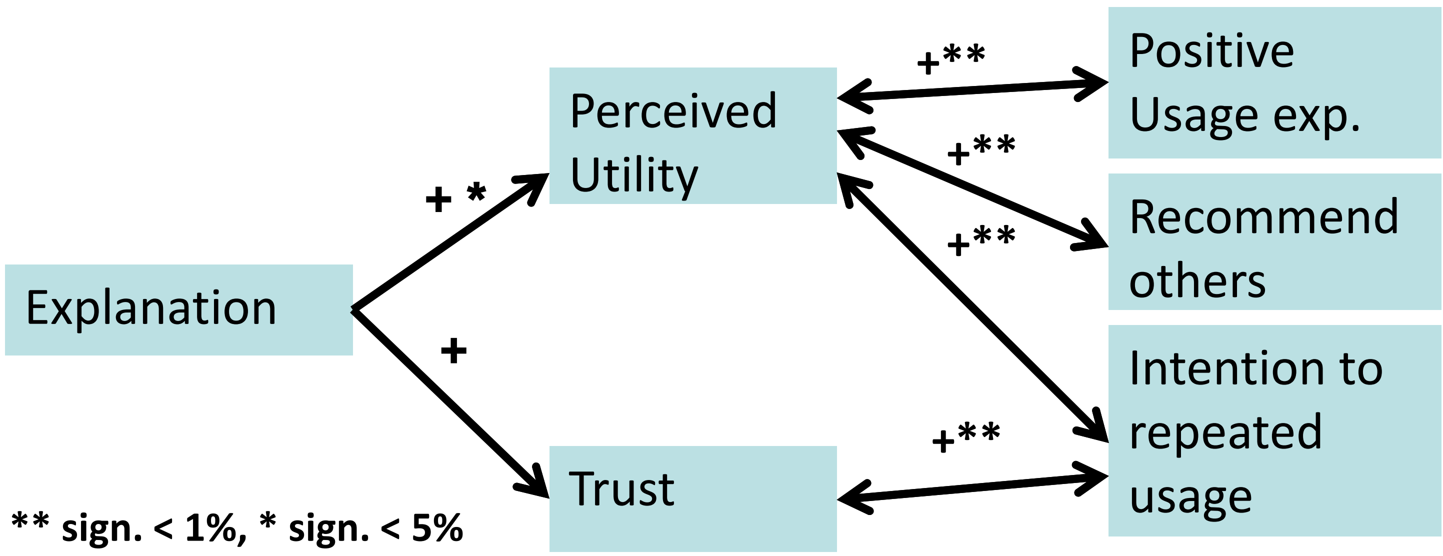

- **Knowledgeable explanations significantly increase the users' perceived utility**
- **Perceived utility strongly correlates with usage intention etc.**

Explanations in recommender systems: Summary

- **There are many types of explanations and various goals that an explanation can achieve**
- **Which type of explanation can be generated depends greatly on the recommender approach applied**
- **Explanations may be used to shape the wishes and desires of customers but are a double-edged sword**
  - On the one hand, explanations can help the customer to make wise buying decisions;
  - On the other hand, explanations can be abused to push a customer in a direction which is advantageous solely for the seller
- **As a result a deep understanding of explanations and their effects on customers is of great interest.**

Literature

[Brewer et al. 1998] Explanation in scientists and children, Minds and Machines **8** (1998), no. 1, 119–136

[Buchanan and Shortliffe 1984] Rule-based expert systems: The Mycin experiments of the Stanford Heuristic Programming Project (the Addison-Wesley series in artificial intelligence), Addison-Wesley Longman, Boston, 1984

[G. Friedrich 2004] Elimination of spurious explanations, Proceedings of the 16th European Conference on Artificial Intelligence, IOS Press, 2004.

[Herlocker et al. 2000] Explaining collaborative filtering recommendations, Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW ’00) (Philadelphia), ACM, 2000, pp. 241–250

[Friedrich and Zanker 2011] A Taxonomy for Generating Explanations in Recommender Systems, AI Magazine, Vol 32(3), 2011.

[McSherry 2003b] Similarity and compromise, Proceedings of the 5th International Conference on Case-Based Reasoning (ICCBR ’03) (Trondheim, Norway) (Kevin D. Ashley and Derek G. Bridge, eds.), LNCS, vol. 2689, Springer, June 2003b, pp. 291–305

[McSherry 2005] Explanation in recommender systems,Artificial Intelligence Review**24** (2005), no. 2, 179–197

[Tintarev 2007] Explanations of recommendations, Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys ’07) (Minneapolis, MN), ACM, 2007, pp. 203–206

[Tintarev and Masthoff 2007] Effective explanations of recommendations: user-centered design, Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys’07) (Minneapolis, MN), ACM, 2007, pp. 153–156

[Vig et al., 2009] Tagsplanations: Explaining Recommendations Using Tags, In IUI '09: Proceedings of the 13th international conference on Intelligent user interfaces (IUI’09) , pp. 47–56