CareDB: A Context and Preference-Aware Database System

Justin Levandoski Microsoft Research Mohamed Khalefa Alexandria University

Mohamed Mokbel University of Minnesota



Why a Context and Preference-Aware DBMS?

Consider a location-based restaurant finder





Why a Context and Preference-Aware DBMS?

- Existing applications will return k
 closest restaurants
- Consider the following five closest restaurants
 - 1) Restaurant 1:
 - Hour and a half wait
 - 2) Restaurant 2:
 - Does not meet my dietary restrictions
 - 3) Restaurant 3:
 - Way too expensive
 - 4) Restaurant 4:
 - Closed for remodeling
 - 5) Restaurant 5:
 - 30-minute drive time due to accident along route



The answers are <u>NOT</u> useful and detached from:

- 1. Personal preferences (dietary restrictions, budget)
- 2. Extra contextual data (time of day, traffic, waiting times)

A Context and Preference Aware Environment



Input data contains user preferences and dynamic contextual information; each type can uniquely affect the answer, be uncertain, expensive to derive,



A database that is aware of *user preferences* and *surrounding contextual information*, and uses this information to give *personalized* query answers to the user.



CareDB Overview



Three Classes of Input to CareDB:

User Context/Preference User location User status User salary User preferences

Environmental Context

Traffic Weather Road Network Transportation

Database Context

Restaurant waiting time Restaurant location Restaurant rating Restaurant price

.



"Towards Context and Preference-Aware Location-Based Database Systems", MobiDE 2009

.

CareDB Project Summary



Outline

- Background
- The FlexPref Framework
- Handling Contextual Data in CareDB
- Handling Uncertain Data in CareDB
- Demonstration
- Conclusion



Preference Methods



Preference Methods

Quick Exercise

- Go to Google Scholar
- Search for papers on preference evaluation methods
- How many results do you get back?





Preference Evaluation in a DBMS





The FlexPref Architecture

- FlexPref is a set of extensible relational operators
- Examples:







Writing Preference Queries in FlexPref

SELECT *	FROM	Restaurants	R	WHERE	[Where	clause]
----------	------	-------------	---	-------	--------	---------

PREFERRING [Attribute List]

USING [Pref Method]

OBJECTIVES [Preference Objectives]

SELECT * FROM Restaurants R WHERE ... PREFERRING Price P, Distance D, Rating R

```
USING Skyline OBJECTIVES MIN P, D, MAX R
```

```
SELECT * FROM Restaurants R
WHERE ...
PREFERRING Price P, Distance D, Rating R
USING TopKDom WITH K=5 OBJECTIVES MIN P, D, MAX R
```

```
SELECT * FROM Restaurants R
WHERE ...
PREFERRING Price P, Distance D, Rating R
USING TopK WITH K=5 OBJECTIVES MIN Func(P,D,R)
```



Adding a New Preference Method to FlexPref

- Adding a preference evaluation method "MyPref" to FlexPref requires the implementation of <u>three functions</u> and <u>two macros</u> in a separate file "MyPref.c" outside the DB engine.
- Once implemented, the preference method is registered using a
 DefinePreference command

DefinePreference MyPref with MyPref.c







FlexPref Generic Functions & Macros

FlexPref Macros

#define DefaultScore
Default score assigned to each
object
#define IsTransitive
Whether preference function
is transitive or not

PairwiseCompare(Object P, Object Q)

INPUT : Two objects P and Q
ACTION : Update the score of P
RETURN : 1 if Q can never be a preferred object
-1 if P can never be a preferred object
0 otherwise

IsPreferredObject(Object P, PreferenceSet S)

INPUT: A data object P and a set of preferred objects S
RETURN: True if P is a preferred object and can be added to S
False otherwise

AddPreferredToSet(Object P, PreferenceSet S)

INPUT: A data object P and a set of preferred objects S **ACTION**: Add P to S and remove or rearrange objects from S



FlexPref Operator Implementation

 FlexPref operators written in terms of the three generic functions and two macros:



```
Input: Single Table T
Output Preference set S
Preference Set S \leftarrow NULL
For each object P in T
   P.score = #DefaultScore
   for each Object Q in T
     cmp \leftarrow PairwiseCompare(P,Q)
     if (cmp ==1)
        if Q is in S then remove Q
        if #isTransitive
           then discard Q from T
     if (cmp == -1)
        if #isTransitive
           then discard P from T
        read next object P
     if (IsPreferredObject(P,S))
        then AddToPreferredSet(P,S)
Return S
```



FlexPref Join Operation





Performance Analysis

Comparison of preference-aware join operator





Outline

- Background
- The FlexPref Framework
- Handling Contextual Data in CareDB
- Handling Uncertain Data in CareDB
- Demonstration
- Conclusion



Contextual Data

- Data retrieved in "real time" during preference query
- Likely retrieved from third party (e.g., web, "cloud")





Contextual Data: Expensive Attributes

- Contextual data is expensive to retrieve relative to local data
- Experiment implemented in PostgreSQL prototype



Outline

- Background
- The FlexPref Framework
- Handling Contextual Data in CareDB
- Handling Uncertain Data in CareDB
- Demonstration
- Conclusion



Contextual Data: Uncertain Data





Contextual Data: Handling Uncertain Data

• Some attributes represented as range

	Restaur	ants	
ld	Price	Dist	Rating
а	5-10	5	6
b	7-15	6	8
С	20-30	1	7
d	15-25	2	4

• The **UPref** framework:

• Answers are probabilistic

SELECT * FROM Restaurants PREFERRING P, D, R USING Skyline OBJECTIVES MIN Price, MIN Dist, MAX Rating

	Answers
ld	Preference Probability
а	75%
b	50%

System Parameters Input Threshold T: All answers must have probability > T Output Tolerance R: Allowed probability calculation error Output Preference query P UPref Answer to P with preference probability: (1) Greater than or equal to T (2) Calculated within error R



Outline

- Background
- The FlexPref Framework
- Handling Contextual Data in CareDB
- Handling Uncertain Data in CareDB
- Demonstration
- Conclusion



CareDB Prototype Demonstration



Full video demonstrating the CareDB is available online:

http://www.cs.umn.edu/~justin/publications.html





Outline

- Background
- The FlexPref Framework
- Handling Contextual Data in CareDB
- Handling Uncertain Data in CareDB
- Demonstration
- Conclusion



Summary





Thank You



Questions



Extra Slides



"Any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and application themselves." [Key, 2001]



Embedding Preference in a DBMS

- Notion of "preference" is subjective
 - Each method challenges notion of "preferred"
 - No limit to the number of proposed methods!
 - We would like to support all methods in CareDB

How do we embed these methods in a DBMS?

- Must handle arbitrary queries
 - Selection over single table
 - Data may reside in multiple tables
 - Data may be sorted
- How do we do this efficiently and in a scalable manner?



Implementing Preference Methods in a DBMS

The Layered Approach



Implementing Preference Methods in a DBMS

The Built-In Approach





Implementing Preference Methods in a DBMS



- Simplicity: easy to implement
- Limited Efficiency: cannot interact with DBMS internals, thus no query optimization

The Built-In Approach

- Efficient: methods tightly coupled with DBMS
- X Infeasible: cannot provide custom implementation for <u>every</u> preference method



The FlexPref Architecture

FlexPref





FlexPref Sorted List Operator







FlexPref: Query Optimization

• Why optimize?

We study the three fundamental properties of FlexPref necessary for query optimization:

Algebraic properties

Cost

Cardinality estimation



FlexPref Theoretical Framework Examples

Implies

Example 1: Symmetric dominance

B

39

Methods exhibiting symmetric dominance not allowed

Β

• Example 2: Conditional Preferences

"I prefer cheap Mexican restaurants when it is sunny"

"I prefer expensive French restaurants when rainy"



• Example 3: Irreflexive, asymmetric, transitive dominance



Preference Queries Over Expensive Attributes

- Preference queries with a mix of "local" and "expensive" attributes
- Goal: retrieve the least amount of expensive attributes

```
SELECT *
FROM Restaurants R
PREFERRING MIN R.Price, MIN R.Distance, MAX R.Rating
```

IdPriceDistanceRatinga2b3
a 2
h 3
c 5



Contextual Data

- You cannot "download" this data
- Data may change from query to query
- Legal reasons from **yelp** API terms:



← → C (www.velp.com/

Phone API

API v2.0 (beta)

General Categor

41

yelper for DEVELOPERS

Yelp Review Search API Overview

Searching by Map Bounding Bo

 Searching by Geo-Point and Radiu Searching by Neighborhood, Address a Narrowing Result Set By Category Response Values Response Codes Searching by Man Bounding Boy

ontional

Preference Query Processing Over Expensive Attributes





Multi-Objective Query Experiment





Future Research Directions – Long Term

Geo-social systems (Application)

 Build a scalable framework to share and analyze user-generated and geo-tagged multimedia data in social networking environments

Real-time recommender systems (Application)

 Adapt existing, high-quality recommendation techniques to new dynamic environments where "recommendations" change rapidly (e.g., social network data, online news/blog posts).

Data management anywhere (Platform)

 Explore system architectures to integrate and perform query processing over existing and emerging disparate data sources (e.g., web services, Amazon Mechanical Turk, local data).

