

IBM Research

QuerioDALI: QA over Knowledge Graphs

Person-Centered Care

Vanessa Lopez, Pierpaolo Tommasi, Spyros
Kotoulas, Jiewen Wu

ISWC'16

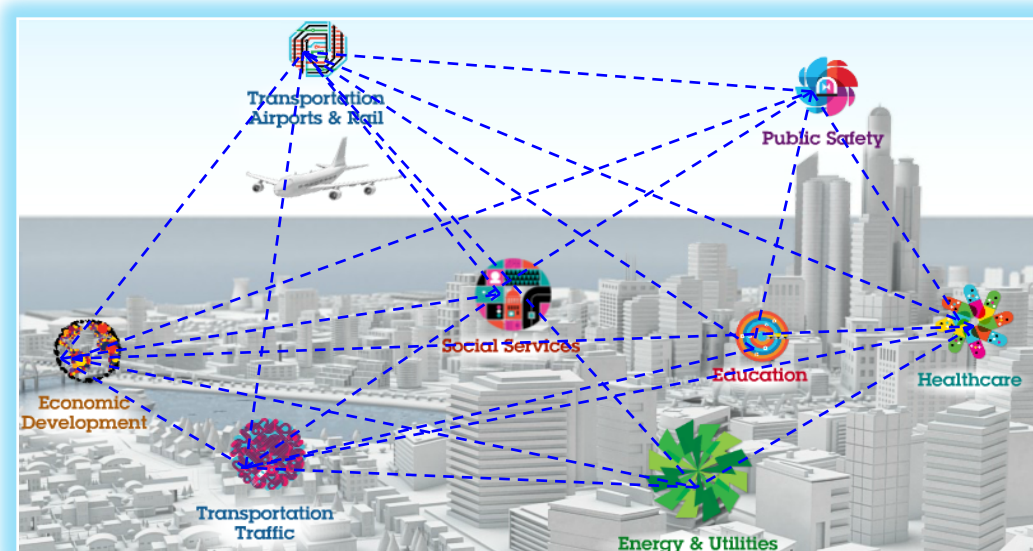
Kobe, Japan

IBM




Goals

- Explore natural ways to **answer complex information needs** over distributed **knowledge graphs**, obtained from Linked Data sources, open and enterprise semi-structured data.



Use Case: Integrated Care




"It cost us one million dollars not to do something about Murray,"



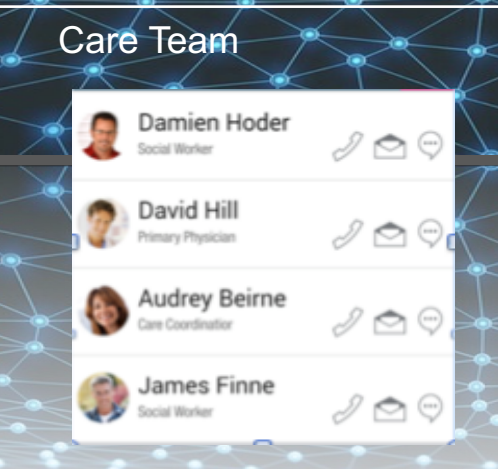
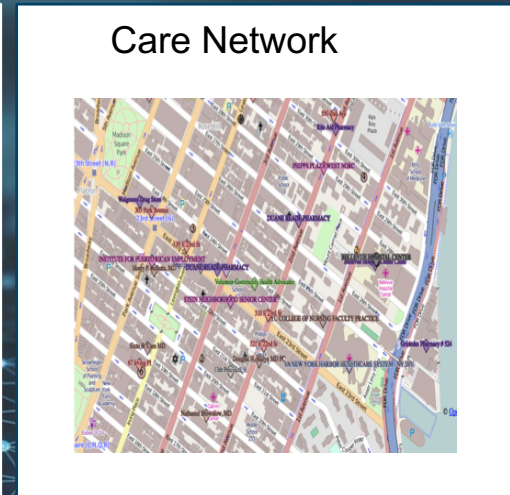
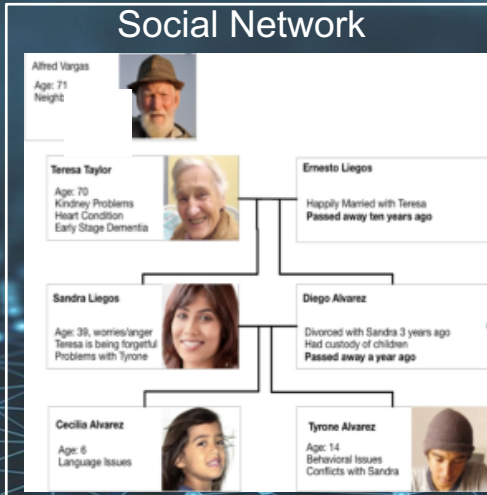
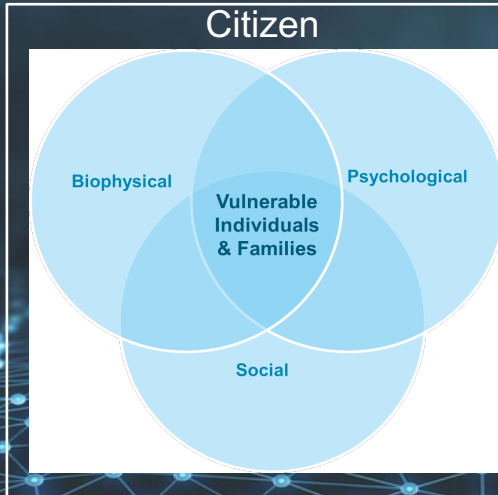
Citizen-centric.


Social context.





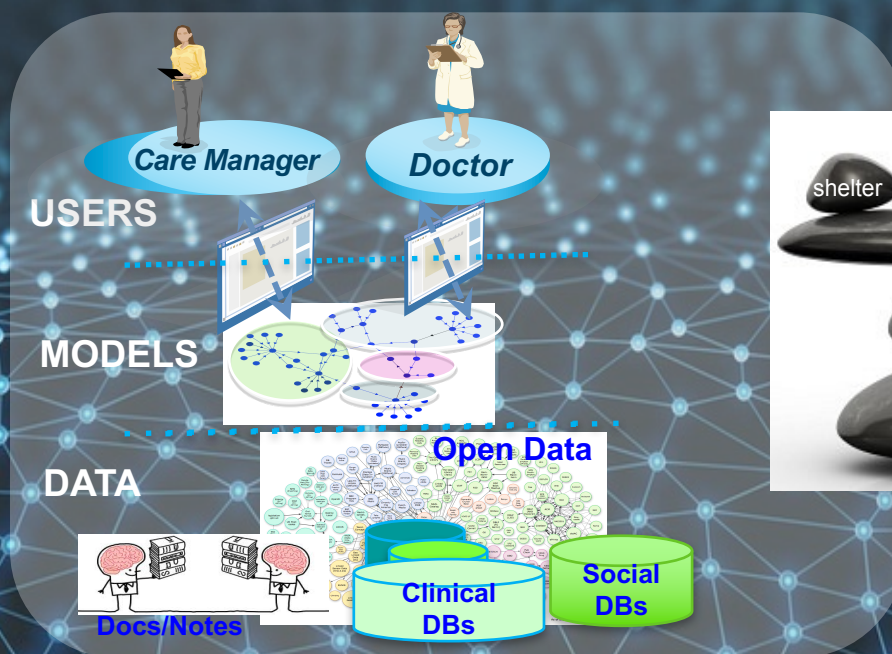
A 360° person view





A 360° person view

- **Semantics** for understanding and interlinking complex data
- **Contextual** what is relevant depends on the user task
- **NL** for users to **interact** with the system





- How can we make this information available to a care team in a natural way?
Right info at right time

The screenshot displays the BlueLENS web application interface for a patient named Teresa Taylor. The interface is organized into several panels:

- Header:** Patient name (Teresa Taylor), gender (Female, 80 years old), and address (5F-3555 Kings College Pl, Bronx, NY 10467, United States). Focus on: Domains (Bio, Functional, Psycho, Social).
- Condition Summary:** Lists medical conditions with dates: heart failure (11 August 2015), mobility issues (18 February 2015), osteoarthritis (8 June 2014), and prediabetes (7 November 2013).
- Area Deprivation Index:** A composite measure of neighborhood socioeconomic disadvantage. Includes a bar chart for 9 digit zip (10467 - 1550) with a score of 122, 5 digit zip (10467) with a score of 109, and State (New York) with a score of 95.
- Services:** Lists various services and their durations: Dietician Consultation (1 January 2014 - 1 January 2016), Medication Pickup (1 July 2014 - 1 July 2015), Daily Living Assistance (1 March 2015 - 1 March 2016), GP Home Visit (1 March 2015 - 1 December 2015), and Medication Pickup (1 August 2015 - 31 July 2016).
- Social Network:** A network diagram showing relationships between Teresa Taylor and other individuals: Maria Ward (Social Worker), Karen Lawson (Social Worker), Cecelia Alvarez (Grandchild), Tyrone Alvarez (Grandchild), Sandra Liegos (Child), and Carol M. Erlington (Doctor).
- Rehospitalisation Risk:** A line graph showing the association between neighborhood socioeconomic disadvantage and 30-day rehospitalization. Includes a note: "The image depicts the association between neighborhood socioeconomic disadvantage at the census block group level, as measured by the Singh validated area deprivation index (X-axis), and 30-day rehospitalization (Y-axis)."
- Medication Summary:** Lists medications: Paracetamol 500mg tablet, Daily; and Eplerenone 25 mg tablet, Daily.
- Cost Summary (Teresa Taylor):**

Home Care	\$960/month
Dietician Visit	\$60/month
Medication Pickup Service	\$55/month
Eplerenone (25 mg tablet) cost	\$58.52/month
Medication Pickup Service	\$55/month
Paracetamol (500mg tablet) cost	\$9/month
GP Home Visit	\$120/month
Total	\$1317.52/month
- Cost Summary (Sandra Liegos):**

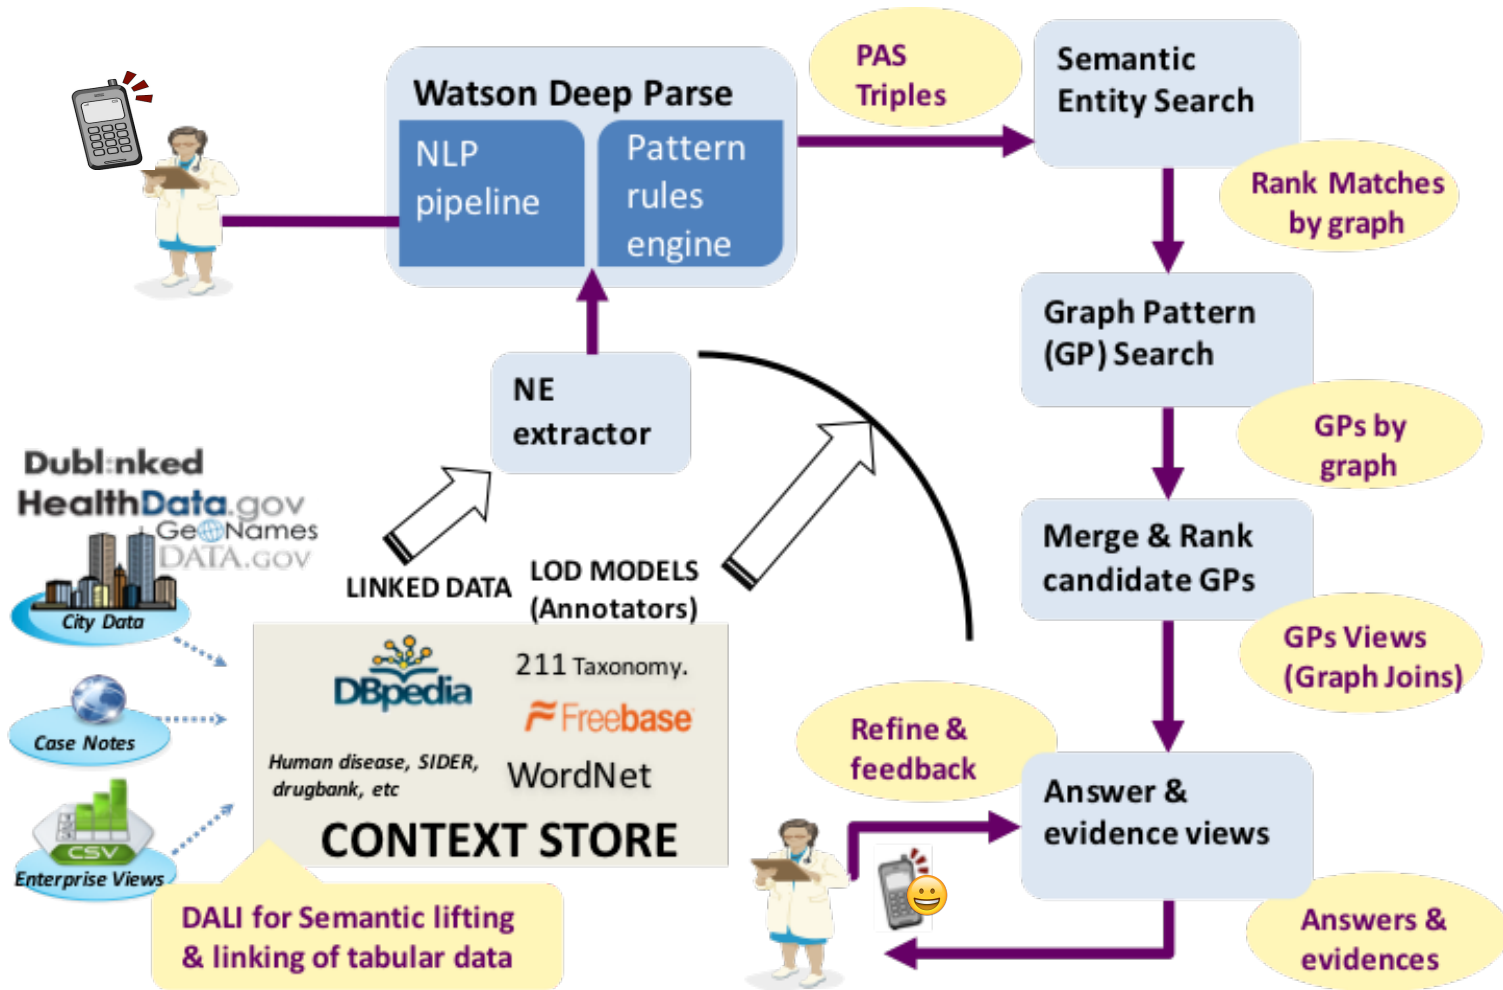
Psychologist Visit	\$95/month
	\$95/month
- Condition Summary (Sandra Liegos):** depression and anxiety (7 January 2015).
- Note:** Teresa required moderate assistance with activities of daily living and with her self-care skills. The patient lives by herself. She lives in a house, wheelchair accessible.

Requirements for QA

- Open Domain
- No fixed schema
- Dynamic and evolving
- No training data
- Combining facts across sources



Approach: QA Pipeline

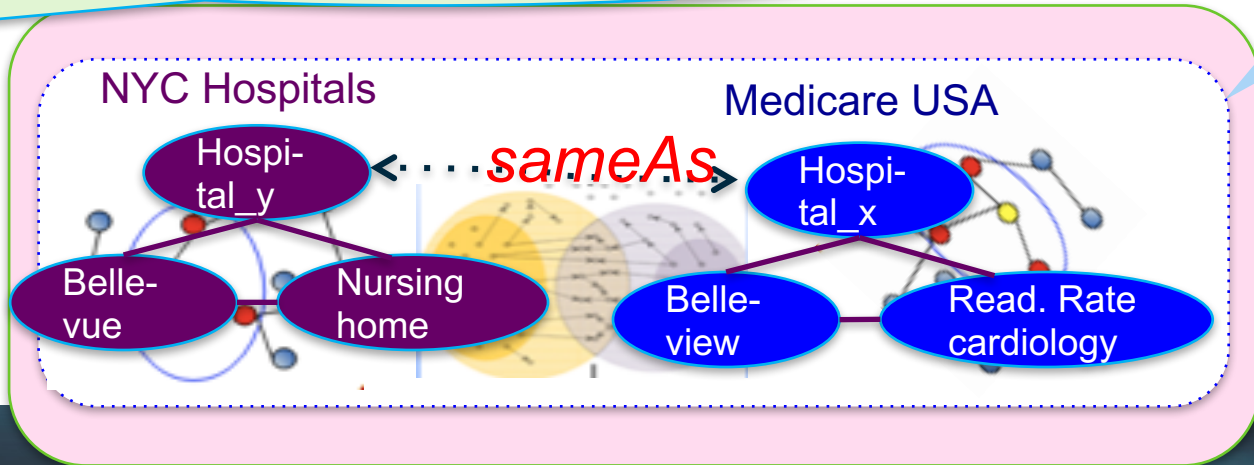


Linked Repositories & Annotators

Which hospitals with elderly care have the lowest readmission rates for cardiology in NYC?



Views




- Context
- Enterp. Data
- Open Data



1.9 billion

RDF triples



SIDER

DRUGBANK

12 million


Concept names in UMLS

580 million


RDF triples in English DBPedia

Diseasome

WordNet



NYC Safety Net
3.5M Data points,
39 Data Sets

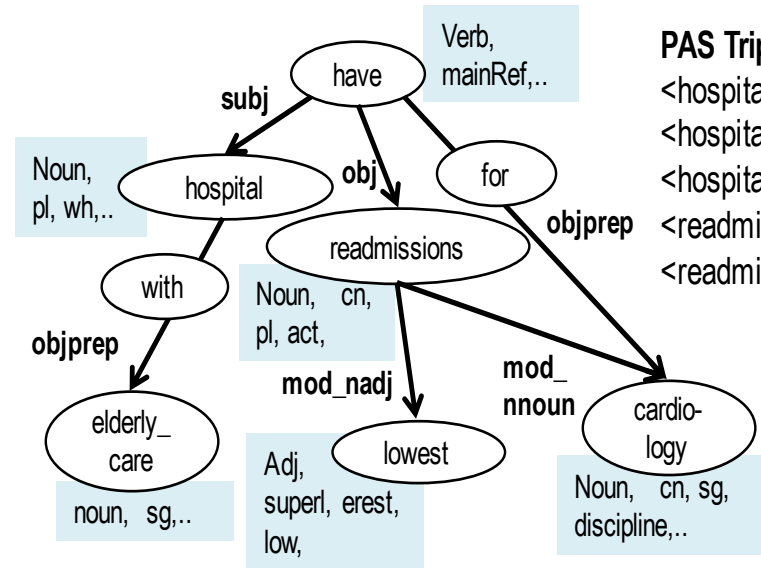


Step 1: NL Deep Parsing

Which hospitals with elderly care have the lowest readmission rates for cardiology in NYC?



- From dependency tree to PAS triples
 - Watson deep parsing (UIMA)
- NE extraction (multi-words: elderly care)



PAS Triples:
 <hospital, have, cardiology>
 <hospital, elderly care>
 <hospital, have, readmission>
 <readmission, cardiology>
 <readmission, low>

```

pattern=nounAdjSup ->
nsubjVar[hasPartOfSpeech("noun")]
{ mod_nadj}mod_nadv ->
ndetVar[hasParseFeature("superl")]
}
    
```

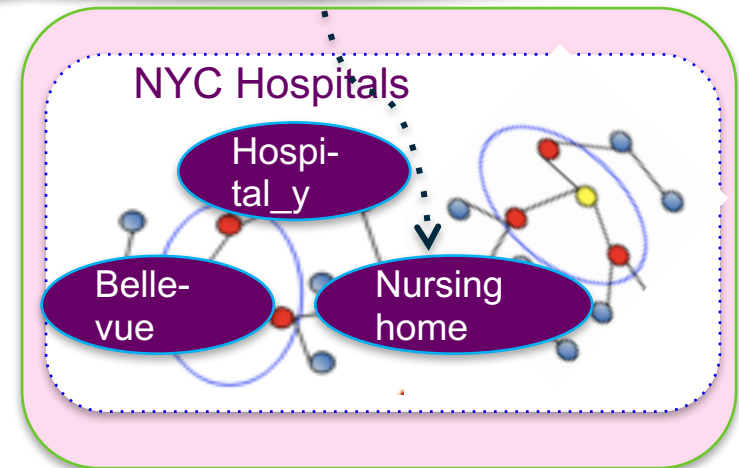



Step 2: Semantic Entity Search

Which **hospitals** with **elderly care** have the lowest readmission rates [↑] for cardiology in NYC?



- Candidate URIs for each query term
 - Mix of schema / non-schema: dateOfBirth, birthDate, birthdate, ..
 - Duplicated terms / missing types: spouse, husband, married to, ..
- Balancing P/R:
 - Semantic expansion
 - Confidence score





Step 3: Graph Pattern Search

- Candidate entities in the PAS are map (partially or completely) into *Graph Patterns*, which are combined into *views* that answer the query
- Large space of mapping combinations:
 - Find most precise (coverage, less expensive) interpretations first

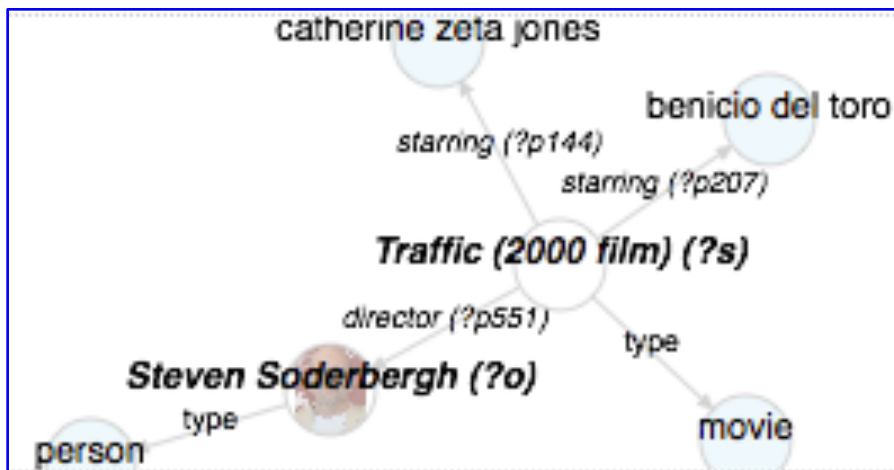
Question: Who is the director of movies starring Benicio del Toro and Catherine Zeta-Jones?

PAS Triples: <person/org., direct, movies>
 <movies, starring, Benicio del Toro>
 <movies, starring, Catherine Zeta-J>

Focus: person/org

Join Term: movies

Answer: Steven Soderberg (movie: traffic)



A GP can be translated it into SPARQL and consists of:

- BGP that belong to the same graph
- JOINS AND UNIONS between the BGPs
- Solution modifiers such as as : ORDER BY DESC/ASC , COUNT, OFFSET and LIMIT
- A confidence score
- The variables that are the focus of the GP

GP Templates and Semantic Relatedness

IBM Research Querio Answering



Which mountains are higher than the Nanga Parbat?



Answer (score: 0.69, percent: 100)

- Kangchenjunga, 8586
- K2, 8611
- Mount_Everest, 8848
- Cho_Oyu, 8201
- Lhotse, 8516
- Makalu, 8481
- Dhaulagiri, 8167
- Manaslu, 8156
- Banua_Wuhu, -5

<http://dbpedia.org>

Boolean Comparative Pattern

```
Select distinct ?s ?o1 where { {
?s rdf:type <dbo:Mountain> . ?s ?property ?o1.
<dbr:Nanga_Parbat> ?property ?o2.
FILTER ( ?o1 > ?o2). FILTER (?property = <dbp:elevationM> ) } }
```

Find hidden connections on the fly

dbpedia:Elevation

dcterms:subject

Category:Length

dcterms:subject

dbpedia:Height

dbpedia:disambiguates

dbpedia:High



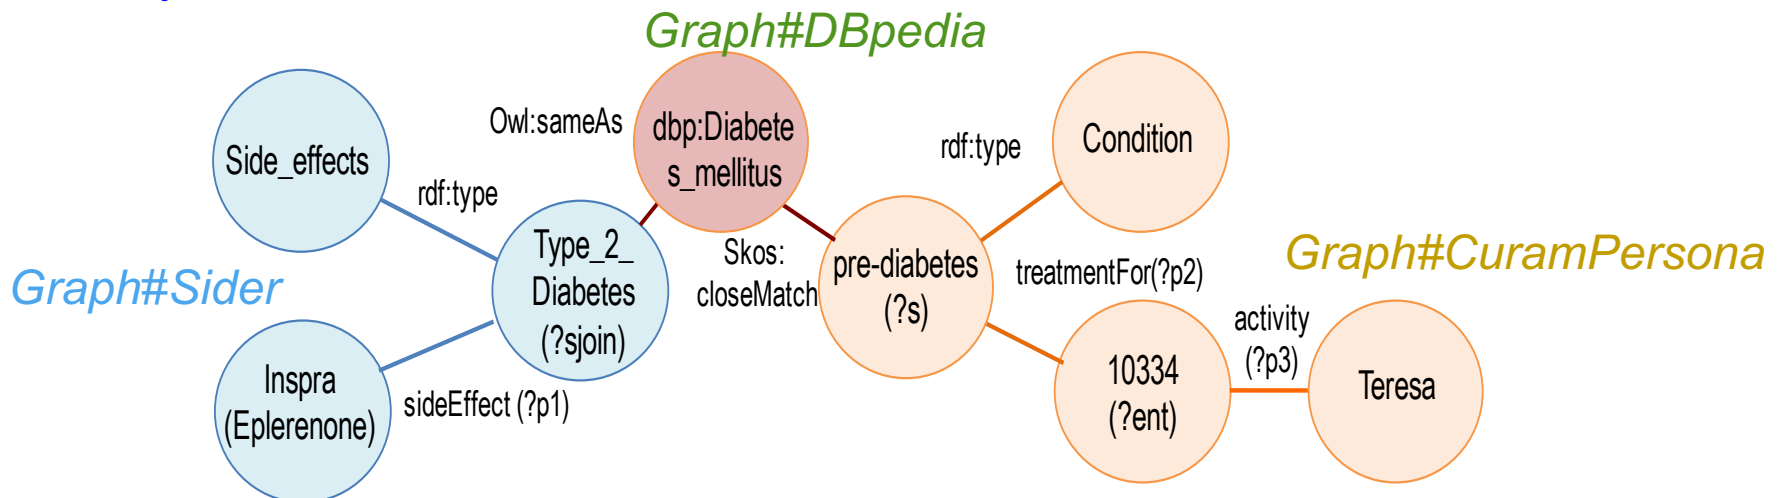
Step 4: Merging & Ranking

- Merging facts and partial answers across sources
- Entity co-reference (join term): different URIs for same *world* entity
 - Syntactic (similar labels) & Semantic Merging (entity linkage)
- Rank alternative translations

Question: Is Eplerenone having side effects for Teresa's conditions?


Join terms: "side effects" in SIDER ontology ("the side effects of Eplerenone") and conditions in the patient data ("conditions of Teresa")

Answer: yes



QALD & Free917 Benchmarks

- **Prototype 1: DBpedia** (triples: 204,954,931)
 - QALD-5 50 questions
- **Prototype 2: Freebase** (triples: ????)
 - Free917 101 questions



	P@1	R@1	P@2	R@2	P@3	R@3	F-1	F-2	F-3
QALD5	0.64	0.69	0.53	0.73	0.52	0.73	0.61	0.55	0.55
Free917	0.72	0.72	0.63	0.73	0.62	0.73	0.72	0.65	0.64

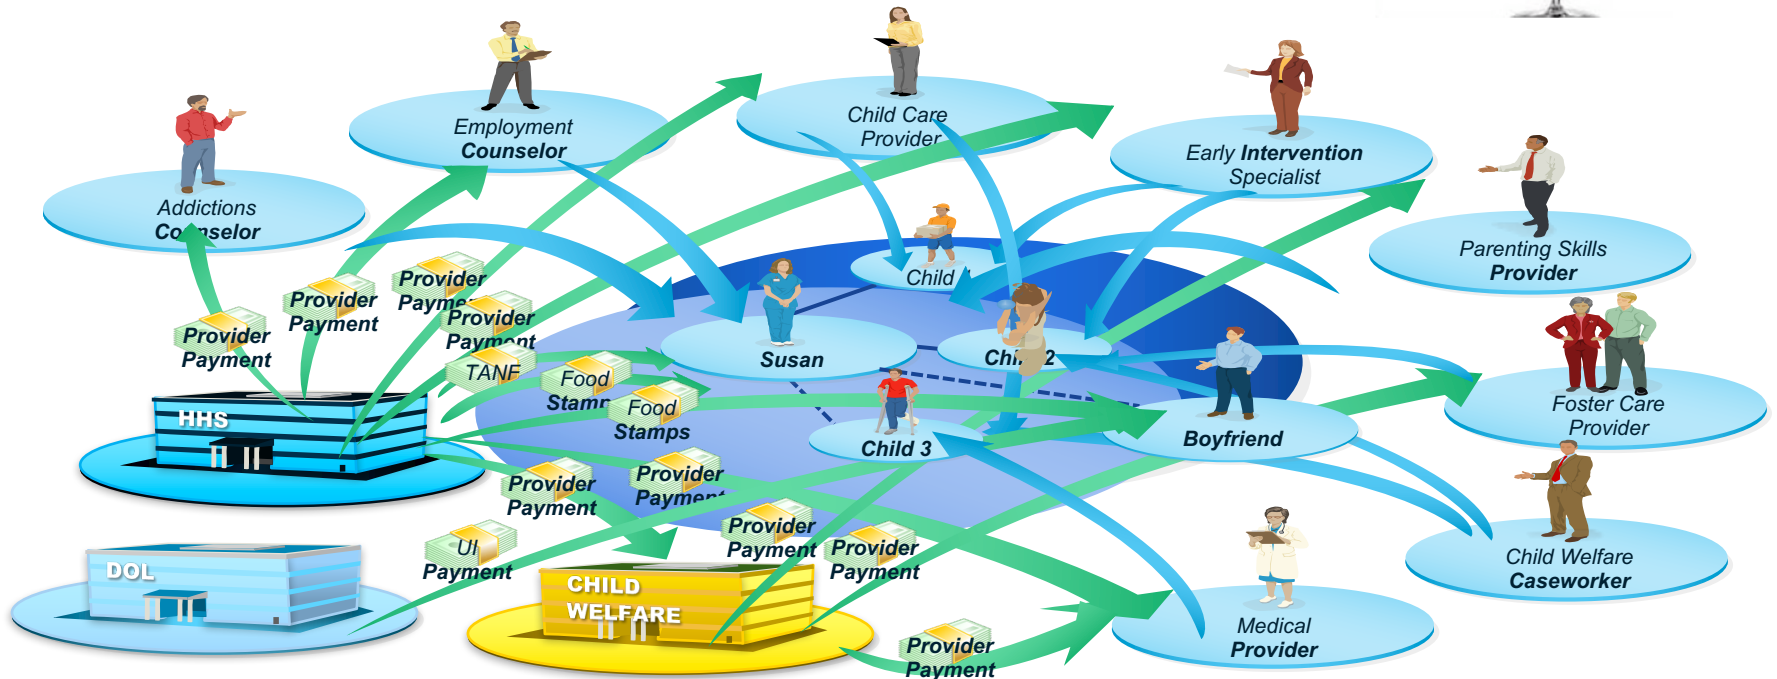
- **Prototype 3: Biomedical Linked Open Data**
 - QALD-4 25 questions
 - Sider (number of triples: 101,558)
 - Drugbank (number of triples: 517,144)
 - Dis easome (number of triples: 92,079)




P@1	R@1	P@2	R@2	P@3	R@3	F-1	F-2	F-3
0.85	0.88	0.78	0.92	0.78	0.92	0.85	0.83	0.83

Many Open Challenges

- Open Domain QA is very hard
- Brittleness and scalability
- Hybrid QA
- More cognitive: Active learning from user interactions
- Personalization: roles





Thank you!



Evaluation

- Errors led to noisy answers (precision) or missing answers (recall)
- **Linguistic**: inaccurate parse tree or not enough rule coverage (e.g., what is the height difference between Mount Everest and K2?)
- **Anchoring**: not finding or filtering out the right mapping or multi-word term (e.g., democrat to “Democratic _party (United States))
- **Patterns coverage**: to bridge the lexical and structural gap (e.g., queries with temporal reasoning – since 2000 – indirect - *who was born in Halloween?*)
- **Merging**: failed to do the linking across join terms
- **Ranking**: not enough context (e.g., popularity)



Graph Pattern Search: Templates

Pattern1: E.g.: Is anorexia an eating disorder?

Direct relation entity to entity `<entity1> ?p <entity2>`
 (Focus: ?p) UNION `<entity2> ?p <entity1>`

Boolean (is-a)
query (yes/no)

Pattern2: E.g.: rivers in Russia

Direct instance to type `?s ?property <type>. {?s ?px <instance>}`
 (Focus: ?s) UNION `{<instance> ?px s}`

Factoid (answer:
list of entities)

Pattern 4: E.g.: population in countries

Direct type to property `?s ?property <type>. {?s <prop> ?o}`
 (Focus: ?s ?o) UNION `{?o <prop> ?s}`

Factoid (answer:
entities/datatypes)

Pattern 6: E.g.: highest mountain; who produced the most films

Superlative pattern Pattern 4+ order by desc(?o) offset 1 limit 0
 (Focus: ?o) Pattern4/5+ order by desc(count(?o)) offset 0 limit 1

Factoid
(superlative)

Pattern 7: E.g.: Is Lake Baikal bigger than the Great Bear Lake?

Comparative pattern Instance to Instance `<instance1> <prop> ?o1. <instance2> <prop> ?o2.`
 FILTER `(?o1 > ?o2)`

Query : Which artists were **born** on the same **date** as Rachel Stevens? (JOIN: date – property-)

SPARQL: select distinct ?s ?o where { {
 ?s ?property <http://dbpedia.org/ontology/Artist> .
 ?s <http://dbpedia.org/property/birthDate> ?o.
 <http://dbpedia.org/resource/Rachel_Stevens> <http://dbpedia.org/property/birthDate> ?o. } }


QA over KGs: Scope

**Global-view
(proprietary) KG**
Evi, Wolfram Alpha,
Facebook KG, Google KG

Domain-specific
ontology-based QA, NLIDBs

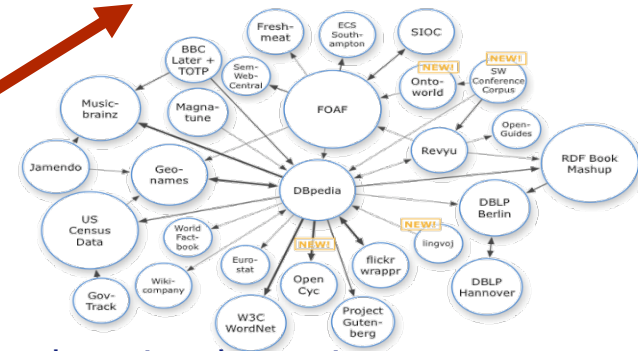
**DESIGN
CRITERIA:**

- Domain-specific
- Portability (new twist on traditional NLDB issues)
- Medium-sized homogeneous Knowledge Base



- Open domain, dynamic
- Build their own large scale KG
- Homogeneous trusted

Open Semantic Web content



- Open domain, dynamic
- Heterogeneous and distributed
- Large scale, noisy data
- Varying levels of quality and trust



Research Agenda

- How can we make this information available to a care team in a natural way? **Right info at right time**
- How to capture knowledge and best practices from care workers to provide **actionable insights**?

The screenshot shows the 'IBM Watson Care Manager' interface. On the left is a patient profile for Laurie Thompson, a 78-year-old female. The main area is titled 'Plan Workspace' and contains several program cards: 'Healthier Living' (7/2/2015 - 7/2/2015), 'Improve Mental Health' (7/2/2015 - 7/2/2015) with sub-items 'Attend Individual Counselling' by Linda Smith and 'Engage in Community Activities' by Jonathan Richman, and 'Improve Social Life' (7/2/2015 - 7/2/2015). Each card has a progress indicator of colored dots.

The screenshot shows the 'Watson Insights Care Plan Recommendation' interface. It features three recommendation cards, each with a 'Yes' and 'No' response option and a red 'X' icon in the top right corner:

- Card 1: "Has Laurie been suffering from tremors?" with 'Yes' and 'No' buttons.
- Card 2: "Do you want to add the HADS assessment to Laurie's program?" with 'Yes' and 'No' buttons.
- Card 3: "What floor of Laurie's apartment building does she live on?" with a '1' button and a dropdown arrow.