Information Interaction in Context: Proceed. of 1st. IIiX Symposion, Copenhagen, 2006. ACM Press, p. 148-162.

The Polyrepresentation Continuum in IR

Birger Larsen[1], Peter Ingwersen[1] and Jaana Kekäläinen[2]

[1]Department of Information Studies, Royal School of Library and Information Science
Birketinget 6, DK- 2300 Copenhagen S, Denmark
[2]Department of Information Studies, University of Tampere,
FIN-33014 University of Tampere, Finland
{blar, pi}@db.dk; jaana.kekalainen@uta.fi

Abstract. The polyrepresentation principle suggests that cognitively and functionally different representations of information objects may be used in information retrieval to enhance quality of results. In the paper, several empirical studies that intentionally or unintentionally have tested the principle are introduced and discussed. The continuum proposed by Larsen (2004; Ingwersen & Larsen, 2005) showing the structural dimension of the retrieval techniques involved in polyrepresentation is further elaborated by adding a novel second dimension consisting of query structure and modus. The new two-dimensional continuum can be seen as a constructive framework for further investigations of polyrepresentative principles in IR.

Symposium themes. Document structure in contextual IIR; Research design

1 Introduction

In a cognitive approach to information interaction, the different actors in the interaction processes contribute interpretations of their situations and pre-suppositions of the world as well as of the information structures involved. Such interpretations take the form of different representations, for instance, author texts, pictures, music tunes, or as database designers' indexing schemes and retrieval algorithms as well as searchers' request formulations and work task descriptions representing their information requirements and problem state (Ingwersen & Järvelin, 2005). In this perspective the representations are manifestations of human cognition, reflection and ideas and thus contextual to one another and interplay over time.

According to Ingwersen (1996) and Ingwersen & Järvelin (2005, p. 206) a *principle of polyrepresentation* can be developed as one of several consequences of a cognitive perspective for Interactive Information Retrieval (IIR). The principle makes deliberate use of the variety of interpretations by means of the evidence that their representations provide. Polyrepresentation encompasses two kinds of representations: *cognitively different* ones deriving from the interpretations by different actors; and *functionally different* representations that derive from the same actor, such as, author generated text structures, image features, diagram captions, and references or out-links (anchors) – Fig. 1. Selectors are special actors responsible for the existence and availability of the information objects, such as editors or publishers.

In addition, documents are made in different presentation *styles* according to the conventions of discourse in domains, genres and media. Further, in a cognitive sense the same group of actors may demonstrate inconsistency or interpretative variation among its members when facing identical information objects or situations.

With reference to Ingwersen & Järvelin (2005, p. 208) the principle of polyrepresentation is based on the following *hypothesis*: "…the more interpretations of different cognitive and functional nature, based on an IS&R situation, that point to a set of objects in so-called cognitive overlaps, and the more intensely they do so, the higher the probability that such objects are *relevant* (pertinent, useful) to a perceived work task/interest to be solved, the information (need) situation at hand, the topic required, or/and the influencing context of that situation." The overlaps of sets of objects created by the divergent cognitive (and functional) representations we name 'cognitive overlaps'.

Essentially, the principle of polyrepresentation attempts to make *simultaneous* combination of evidences (representative features) that are cognitively contextual to one another in a structured way. For instance, the intersection of full text terms and phrases (from authors) by identical or similar index terms, maybe in addition according to different weighting schemes (from indexing algorithms), and by title words from documents citing the full text documents retrieved (from other authors over time), may constitute such a combination, see Fig. 1.

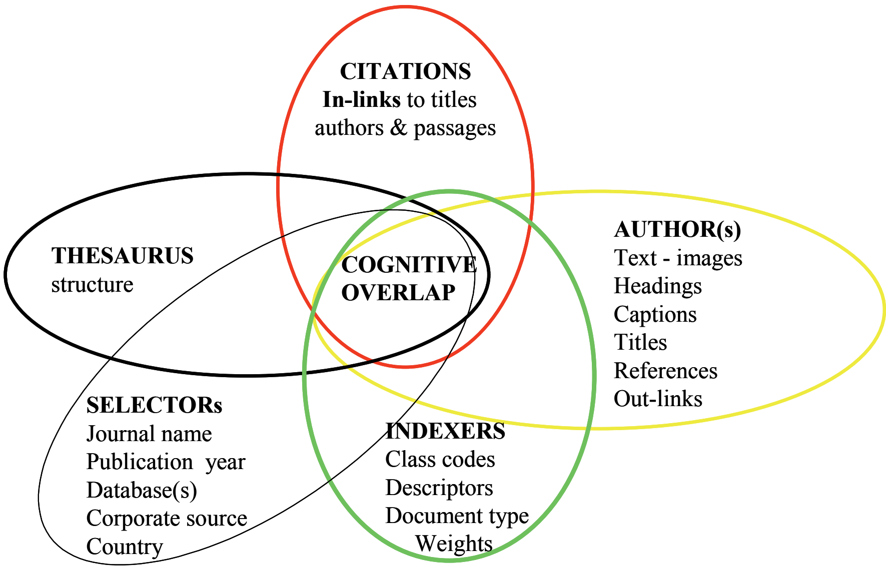

Fig. 1. The principle of polyrepresentation in academic documents. Overlaps of information objects retrieved by representations of cognitively and functionally different information structures, by means of one search engine via search keys associated with one searcher statement (e.g., a work task description). Elaborated from (Ingwersen 1996, p. 28; 2002, p. 294; Ingwersen & Järvelin, 2005, p. 207).

In other words, polyrepresentation is a particularly structured way of carrying out a kind of classic triangulation in the information space *and* in the cognitive space of searchers. The latter aspect constitutes the real novelty of the principle of polyrepresentation. Conceivably, this multidimensionality of the cognitive space can be further exploited by building request models and algorithms (not user models) that extract such evidence of searcher perceptions to be combined with polyrepresentative structures from information space and search engine logics. A more detailed analysis of the scientific background underlying the principle is provided by Ingwersen & Järvelin (2005, p. 206-214 & p. 342-346).

Exactly the structural aspects of the principle made Larsen investigate the 'continuum of polyrepresentation' (Larsen 2004; Larsen & Ingwersen, 2005) – Fig. 2. The continuum attempts to outline and explain to which *extent* the involved kinds of information are structured by retrieval logics and principles.

The present paper starts with a discussion of the polyrepresentation continuum (Larsen, 2004; Larsen & Ingwersen, 2005), followed by an analysis and discussion of a series of empirical studies that are *explicitly* based on the principle of polyrepresentation. Their results lead to an extension of the continuum, by adding a novel second dimension consisting of query/request structure and modus. The studies mentioned are briefly discussed and mapped in relation to the continuum framework. The paper ends by outlining some trends for future research on polyrepresentation.

2 The Polyrepresentation Continuum

Only a few empirical studies have so far looked into *which kinds* of cognitively different representations that best lead to good retrieval (and seeking) results. For instance, it is known that retrieval of documents by search keys found in titles and abstracts *and* by involving the citations to such documents (made by some other cognitive agents at a later time) produces much higher odds for finding relevant documents in the constructed overlap than in each of the retrieved sets independently (Pao, 1994), see also Skov and colleagues (2004) below. A few other matching combinations have been tried so far, generally with improved IR performance as result. According to Larsen (Larsen 2004; Larsen & Ingwersen, 2005) there seems to exist a continuum of polyrepresentative solutions from one extreme of hard structured exact match-like combinations to unstructured bag-of-words modes of the principle (see Fig. 2 below). Ontology (or thesaurus) support seems to improve the outcome along the continuum (Skov et al., 2004). The cloud in the middle signifies hitherto unknown territory concerning the degree of structure empirically tested.

One should note that the model, Fig. 2, assumes the application of at least two cognitively and/or functionally different kinds of representations of information space. According to Larsen and Ingwersen (2005, p. 57), at the *structured* pole of the continuum the implementations are based on exact match principles, leading to sets of retrieved documents for each representation from which overlaps can be formed and a pseudo-ranking be constructed. At the *unstructured* pole of the continuum the implementations are based on best match principles leading to a rank of the retrieved documents per representation as input for polyrepresentation. Rather than generating overlaps between sets, the latter implementations will fuse the ranks to produce a final ranked output. Evidently, several different IR engines (or logics) can be used to provide polyrepresentative outputs per representation to be fused, located somewhere within the mid-section of the continuum (the cloud). This cloud illustrates the main challenge for future research on the principle of polyrepresentation: to identify flexible and effective matching methods that can generate high quality cognitive overlaps from a variety of the most promising representations.

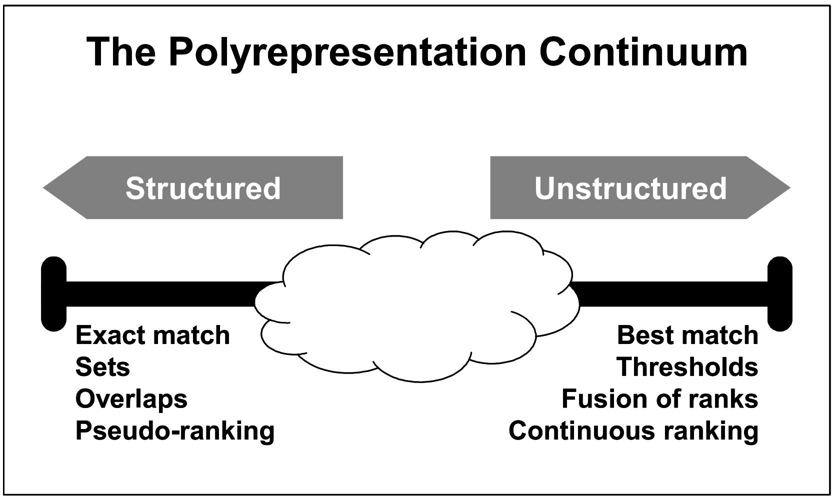

Fig. 2. The polyrepresentation continuum
(Larsen 2004, p. 38; Larsen & Ingwersen, 2005, p. 57).

However, the entire principle of polyrepresentation does not rely on one request formulation from a searcher. It assumes that functionally different representations of the searcher's cognitive structures come into play too, such as current knowledge state of domain, information gap or work task perception (see Ingwersen, 1996). Again, one may regard such representations as highly structured, e.g., by their intersections, or more loosely combined (Kelly et al., 2005) – as discussed in detail in the next section. As a consequence, the continuum, Fig. 2, lacks a dimension concerned with query structure. This second dimension is necessary in order fully to understand the characteristics of the polyrepresentation principle, and is developed below.

3 Polyrepresentation of the Cognitive Space

Several scenarios exist by which the cognitive space of information searchers may be utilized for IR. An extensive list of possibilities is displayed in Ingwersen & Järvelin (2005, p. 335). For instance, recommender systems make use of longitudinal assessments by searchers and/or their search profile contents. A more direct way is to apply different formulations extracted simultaneously from a current searcher. Such searcher statements could target current lack of domain knowledge (as in a request), describe the current problem state or/and the perceived work task situation.

Kelly and colleagues (2005) investigated this type of polyrepresentation. They had 13 searchers supplying 45 topics to the 2004 TREC HARD track. The searchers assessed relevance and used an online clarification form as a front-end to the retrieval system. The form consisted of four questions (Q1-Q4) that allowed a searcher to (Q1) state the times in the past he/she had searched that topic; (Q2) describe what he/she *already knows* about the topic (knowledge state); (Q3) state *why* he/she wants to know about the topic; and (Q4) add *any keywords* that further describe the topic. The TREC collection and search engine principle (Lemur, BM25) were controlled variables and the baseline run (BL) used terms from the TREC topic title and description elements combined as queries. No distinction was made between different representations of the TREC documents. Mean Average Precision (MAP) and significance tests were used. Kelly's and colleagues' (2005) experimental (single) runs consisted of BL + different combinations of relevance feedback modes, and Q2-Q4. Whereas BL yielded on average 9.33 words, Q2 yielded 16.18, and Q3 10.67 different words. Q4 did only produce 3.3 words. This result stresses the usefulness of attempting extractions of searcher statements, in particular if initial requests (the TREC topic title) are very short statements. When combining the Q2-Q4 searcher statements, Kelly and colleagues (2005) did not retrieve documents by each statement separately and then intersect the retrieval results in order to define overlaps of documents (equal to somewhat structured query representations on the continuum, Fig. 5). Instead they applied the statements in union with weights for term repetition (term overlap), Fig. 3. This signifies a much less structured approach.

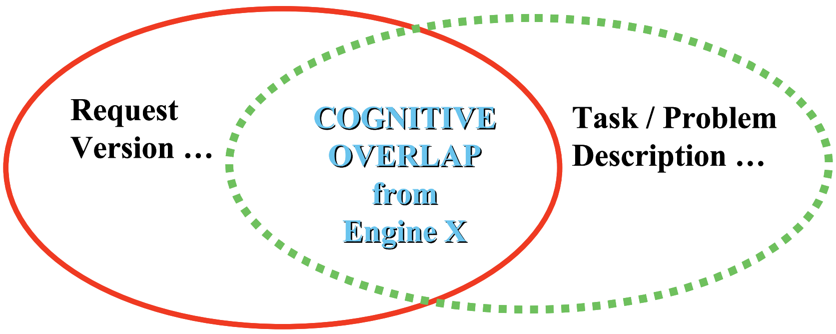

Fig. 3. Example of polyrepresentation of the cognitive space with overlap (of terms or documents depending on degree of structural approach).
From Ingwersen & Järvelin (2005, p. 348).

Kelly's and colleagues' (2005) results are very promising in relative terms. All single Q-versions and Q-combinations outperform BL. Also, there is a significant correlation between query length and performance.

4 Contexts for Polyrepresentation of Search Engines

Lund (2005) investigated the 12 most effective TREC 5 search engines' combined retrieval power in different combinations based on 30 TREC-5 topics (threshold: > 45 relevant documents per topic). The most detailed tests were performed on the top-4 search engines. They consisted of two different versions of the SMART system (based on the vector space model) from Cornell University (Cor5M2rf) and Swiss Federal Institute of Technology (ETH), the former using human relevance feedback for query expansion, a third one mixing natural language processing and vector space with query expansion from an US laboratory group (genrl3), and a fourth engine running on very different principles applying GCL (structured) query language from University of Waterloo (uwgcx1). The former two systems (Buckley, Singhal & Mitra, 1997; Ballerini et al., 1997) are hence regarded as rather functionally different whilst the third (Strzalkowski et al., 1997) and fourth systems (Clarke & Cormack, 1997) are cognitively different from the former and one another.

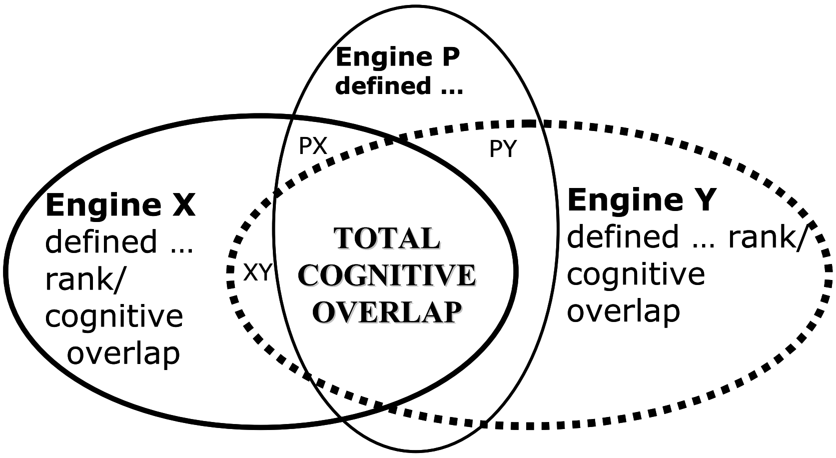

Fig. 4. Polyrepresentation of three different search engine's retrieval results in the form of overlapping documents. Variation of Ingwersen & Järvelin (2005, p. 347).

One may see the tests as instances of data fusion structured according to polyrepresentative principles. For each of the four engines the top-100 ranked documents per TREC topic were re-assigned rank weights different from their original scores in order to avoid scaling problems between engines: from value 100 to 1 at rank positions 1 to 100. By boosting further the sum of the re-assigned weights for documents located in the engine overlaps (the duplicates) the polyrepresentative principle is achieved. Multipliers were: 1, 25, 50 and 100, with 100 in the case of a document found in all 4 systems within top-100 (named Fuse-4). Since a Fuse-4 ranking often retrieve much less than 100 documents (the document cut-off value, DCV), the best performing documents from the three Fuse-3 overlaps, i.e., from overlap *px, py or xy*, were added to the rankings. This constitutes the so-called *Super System* – see Fig. 4. In addition, a fusion of all 12 search engines, named Fuse-12, was also tested (Lund, 2005).

Lund's results are very promising (2005). First she found that when combining the four search engines according to polyrepresentation principles, "the performance in terms of Recall and Precision depends on how many relevant documents potentially exist in the search task". The more relevant documents in the topics over the four engines in all combinations the higher the precision (Lund, Schneider & Ingwersen, 2006). This result implies that in comparisons of performance with other retrieval principles the test collections involved, for fairness reasons, should contain topics with 'substantial numbers' of relevant documents. Secondly, in all performance indicators the Super System outperforms the other configurations, except in terms of the Ranked Half-Life indicator (Borlund & Ingwersen, 1998) – see Table 1. We observe that the *uwgcx1* system from University of Waterloo, like the *ETHme1*, perform best as single systems and in line with or better than *Fuse-12*. Lund's results indicate strongly that 3-5 systems may be better performers in combination than, for instance, 12 different IR engines (2005).

Table 1. Top-100 performance measures over four single TREC-5 systems and combinations for 30 topics, each with more than 45 relevant documents (Lund, 2005).

	Precision	Recall	F-measure	nRHL	MAP
S-system	0.56	0.46	0.47	90.3	0.40
Fuse4	0.48	0.29	0.36	60.7	-
Fuse12	0.41	0.31	0.33	144.3	0.27
Cor5M2rf	0.34	0.27	0.28	241.1	0.23
ETHme1	0.40	0.32	0.33	102.4	0.31
genrl3	0.35	0.28	0.29	144.1	0.24
uwgcx1	0.42	0.32	0.33	138.5	0.29

Among the *Fuse-3* combinations over the 30 topics at DCV=100 (not shown in table) Lund's results demonstrate that when *cognitively different systems* are combined (genrl; uwg; and ETH (precision = 0.40); or Cornell (precision = 0.39)) the precision is higher than when only functionally different systems are combined: Cornell; ETH; uwg (precision = 0.38); Cornell; ETH; genrl (precision = 0.27). The latter result is statistically significant (Friedman test). This positive result for the polyrepresentative principle also corresponds to the fact that when the best single system participates in combinations (the 'uwg' system), that combination performs well. Further, when the two 'familiar' systems (the vector space systems Cornell & ETH) are fused into one retrieval system, the precision drops significantly to 0.20 over all 30 topics at DCV = 100 (Lund 2005, p. 55).

5 Polyrepresentation of Information Space

There are many possibilities of representing information space in cognitively different ways. They range from combining two or more complementing databases, over applying a variety of different indexing methods to document contents and structure in the same database, including the use of an ontology, to applying different external features that are contextual to document contents, such as academic citations and inlinks – see Fig. 1.

In a study similar to that by Pao (1994) combining *several collections*, Christoffersen (2004) applied Medline, Embase and SCI in order to test the relevance proportions in any of the overlaps created online between indexer of MeSH (Medical Subject Headings in Medline), author text (Embase) and citing authors (SCI). Expert relevance assessments were used. He found that "[t]he degree of overlap strongly correlates with the percentage of relevant items in a set" (p. 391). The results were statistically significant (Ingwersen & Järvelin, 2005).

5.1 Combining different document representations

Skov and colleagues (2004) investigated the use of different functional representations, like document titles, reference title words (RF) and abstracts (author-derived) combined with cognitively different ones, like major (MJ) and minor (MN) MeSH subject headings (indexers). (Un)structured queries served as another independent variable in their tests. The Cystic Fibrosis test collection (Shaw et al., 1991) of 1,239 documents from MedLine indexed under InQuery retrieval system, 100 requests and tripartite relevance assessments (highly, partial & non-relevance) constituted the experimental setting. Skov and colleagues used 29 requests and made use of 15 different overlap combinations. The same 29 requests were also formed into highly structured queries, modified in a number of ways in association with the Kekäläinen & Järvelin approach to query structuring (1998; 2000).

The results demonstrate that, in general, overlaps generated from three or four different representations have higher precision than overlaps generated from two or one representation. (Skov & al., 2004). These findings support the principle of polyrepresentation.

Further, the highly structured queries result in higher precision than queries entered as bag-of-words – although both follow the principle of polyrepresentation. Further results and discussion of findings are provided in (Skov, Ingwersen & Larsen, 2006).

In earlier investigations Larsen had tested the so-called *Boomerang Effect* on a variety of inter and intra-document features. He applied different document representations from the INEX test collection and INSPEC thesaurus terms also involving *citation cycling* strategies, i.e., backward chaining followed by forward citation chaining (Larsen, 2004).

The best precision result was achieved by tests combining functionally different representations, such as article titles, section headings and the cited titles in the references. Reasonable effectiveness was obtained by combining those representations intersected with indexer descriptors and the Boomerang Effect (Larsen, 2004).

Different *weighting schemes* were tested, e.g., by application of the frequency of documents cited. Unstructured queries were used and no super system-like aggregation (e.g., Lund, 2005) was introduced as a result of the citation cycling. The effect was compared to 1) best match bag-of-words used for each representation separately and the results fused in a simplistic polyrepresentative way; and to 2) common bag-of-words based baselines in the INEX test collection. The results showed that the Boomerang Effect did not decrease performance, but pure polyrepresentation was slightly better. However, bag-of-words obtained the best overall performance (Larsen, 2004).

White and colleagues (2005a+b, 2006) have recently investigated the principle of polyrepresentation applied to *interface functionality* and *implicit relevance feedback* (IRF) algorithms for interactive IR (2006). White proposes to apply "content-rich search interfaces that implement an aspect of polyrepresentation theory, and are capable of displaying multiple representations of the retrieved documents simultaneously in the results interface" (White, 2006, p. 1). The prototype interface implements a *progressive revealment* strategy where searchers can access an increasing amount of retrieved document content by following interactive relevance paths between different representations created from the same document. Such representations are top-ranking sentences from a retrieved document, its title, its query-biased summary (commonly four sentences), single summary sentences, or summary sentences in context. By hovering over specific representations or by clicking on icons the interface guides the searcher further on, and the traversal of these paths is then used by an IRF model to select terms for query modification (White, 2006, p. 3). Three 'simulated search' scenarios were tested for retrieval performance by means of searcher simulations of all possible combinations of representations and paths available. The best performing combination of representations consisted of document title, its query-biased summary and summary sentence in context.

6 The Extended Structural Polyrepresentation Continuum

In common to several of the polyrepresentation studies is the observation of two kinds of 'query structure': search keys logically structured and value added by synonyms from an ontology – the Kekäläinen and Järvelin approach (1998; 2000) – and the structural dimensions adhering from the polyrepresentation principle (Ingwersen & Järvelin, 2005). Compared to traditional bag-of-words query forms structured queries seem to perform better. It is consequently logical to extend the structural Polyrepresentation Continuum (Larsen, 2004; Larsen & Ingwersen, 2005) by adding a second dimension concerned with query structure – see Fig. 6 below. As in the original model, Fig. 2, the new continuum, Fig. 6, assumes the application of at least two cognitively and/or functionally different kinds of representations of information space – like author text and indexer keywords. The horizontal axis represents the degree of structuring by means of the retrieval technique(s) and weighting schemes applied to the information space. If two or more best-match retrieval engines (e.g., Lund, 2005; Lund, Schneider & Ingwersen, 2006) are applied this is a move towards the clouded area from the right pole.

With respect to the vertical axis, a highly structured query for polyrepresentation would have Boolean characteristics, as in the traditional query languages applied to bibliographic databases, such as, Thomson Dialog, or in InQuery. The logic behind is related to Quorum search logic, constantly applying intersections of separate inverted database fields. In Thomson Dialog a typical retrieval command of polyrepresentative nature would look like (1), assuming that one wants to find documents in which both the indexer descriptor field (/DE), the authors' title (/TI) *and* the abstract (/AB) fields contain the search key "food" – see also Fig. 5, right-hand side:

(1) SELECT food/TI AND food/AB AND food/DE

(2) SELECT (food OR nutrition)/TI AND (food OR nutrition)/DE

Command (2) insures that relevant synonyms or other associated terms deriving from a thesaurus (e.g., nutrition) may serve as alternative search keys in the required polyrepresentative fields. The common way of searching topical keywords in operational systems is, however, quite different – see Fig. 5, left-hand side. The search key(s) would be entered and searched in the inverted basic index in a union fashion.

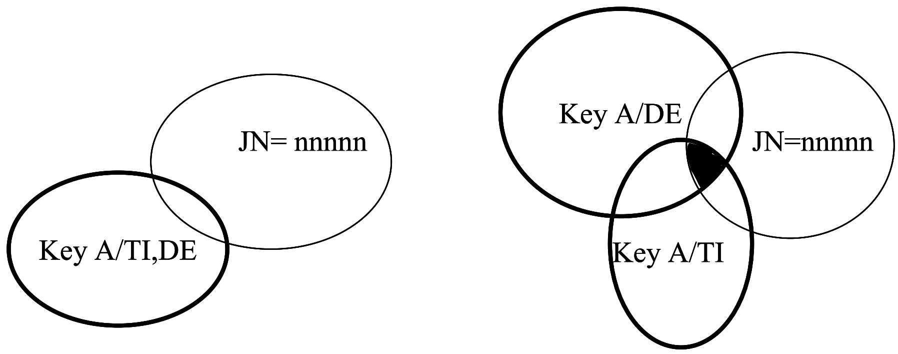

Fig. 5. Traditional Boolean query (left) retrieving a topic (A) from union of inverted index fields (/TI,DE) and metadata (JN=nnnn) for recall reasons; typical polyrepresentative query (right) intersecting the inverted separate index fields by identical search key (A) and metada. From (Ingwersen & Järvelin, 2005, p. 208).

Skov's and colleagues' (2004) tests used command structures similar to (2), but in the best match InQuery system. Here, the query language supports combining search keys in the Boolean manner, e.g. expressing relations between them, and weighting them. The examples above could be expressed as follows:

(3) #band(#field(TI food) #field(AB food) #field(DE food)

(4) #band(#field(TI #syn(food nutrition) #field(DE #syn(food nutrition))

In the literature the concept of *structured queries* typically refers to queries formulated with the Boolean operators in contrast to bag-of-words queries (e.g., Croft, Turtle & Lewis 1991; Hull 1997). A *facet structure* – see Fig. 6 under the label 'Search Keys' along the vertical dimension of the model –implies that facets must be identified, as is the case with Boolean queries in the conjunctive normal form. Partial match techniques interpret logical Boolean operations as arithmetic operations (soft

Boolean operators) with which the individual weights of search keys are combined to rank the documents. The weights are numerical values attached to keys, concepts, facets; the operators indicate logical or arithmetic operations, or the position of keys in documents.

Kekäläinen (1999) proposes a classification for queries ranging from highly structured to unstructured. Queries can be classified according to their structural features (or decreasing degree of context): (1) are concepts identified? (2) are facets identified? (3) are concepts weighted? (4) are facets weighted? (5) are search keys based on single words or phrases? (6) are search keys weighted? (7) by what operators are facets / concepts / search keys connected? – see vertical mapping on Fig. 6.

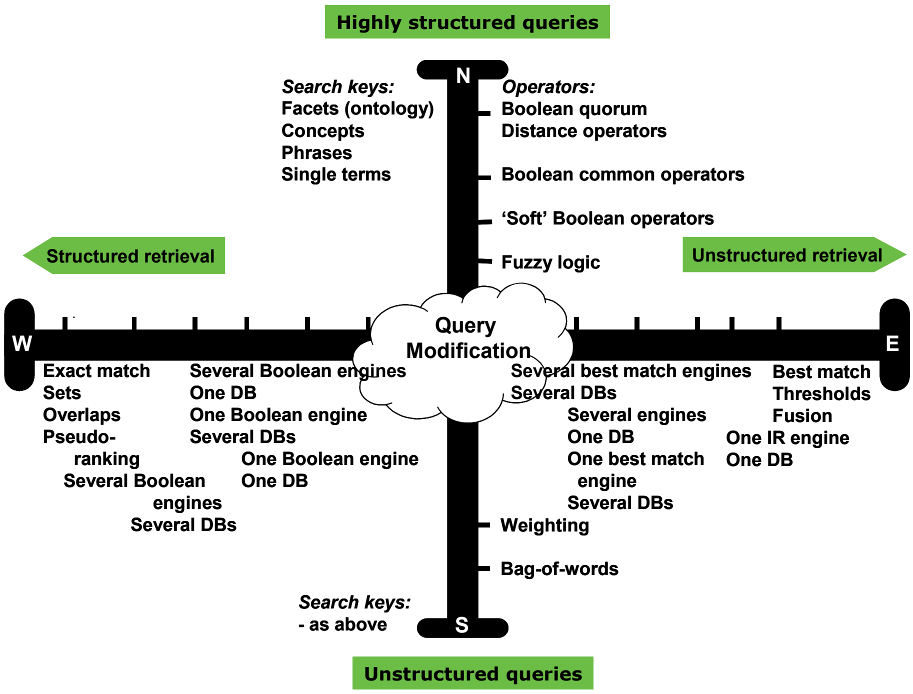

Fig. 6. The extended structural polyrepresentation continuum. From (Larsen, 2004; Larsen & Ingwersen, 2005).

An *unstructured query* does not indicate a facet or concept structure with operators (i.e., queries with a single or no explicit operator; differentiated relations between search keys). In a *highly structured query* search keys representing different concepts or facets are separated by operators.

Boolean structured queries have been claimed to be more effective than unstructured queries, both in a probabilistic retrieval system (e.g., InQuery) and the vector space model (Belkin et al. 1995; Hull 1997). Turtle and Croft (1991) explain the performance improvements due to structural information (phrase structure, compound nominals, and negation) captured in Boolean queries, which information is not exploited in unstructured queries. Hull (1997) confirms these conclusions. He compared a weighted Boolean model (based on probabilistic principles) with the vector space model.

It is consequently to be expected that structured Boolean-like query configurations will best support polyrepresentation in IR. At the unstructured end one can move from bag-of-words up-wards over simplistic weighting schemes towards increasingly sophisticated schemes – like weights assigned the 'soft' Boolean operators. The extended continuum, Fig. 6, attempts to demonstrate the most central points of reference considered for a range of structuring attributes to queries. Many different relevance feedback and query modification solutions can be carried out (in the cloud).

6.1 Mapping polyrepresentative research

If we map the implementations in Larsen (2002) and Christoffersen (2004) on to the continuum, they are placed at the structured retrieval axis at the point of 'One Boolean engine & Several DBs', applying single word queries in a traditional Boolean intersection structure – in the NW quadrant close to the horizontal axis. Larsen (2004), however, is located at the more unstructured retrieval pole, applying several databases (DBs) and one best match engine plus bag-of-word queries to the involved types of representations – located in the SE quadrant. Lund (2005) is located in the SE quadrant, at the horizontal axis at 'Several best match engines & one DB' applying fusion of rankings into a continuous one by simplistic document weighting based on bag-of-word queries. Skov and colleagues (2004) also represents an attempt to move towards to unstructured retrieval pole (One best match engine & one DB) – but including both unstructured (SE quadrant) and highly structured facetted queries (NE quadrant). Relevance feedback and query modification algorithms are tried out following the polyrepresentative principle by White and colleagues (2005a+b; 2006a+b) attacking the cloud of the framework. One should note, that the common baseline retrieval applied in IR evaluation tests is located at the outmost corner of the SE quadrant.

Kelly and colleagues (2005) are located also in the SE quadrant, but – in contrast to all other investigations – they focus on polyrepresentation and simplistic weighting of different representations of the *searcher's* cognitive space, operating in one database with one system and no distinction between different representations of documents.

The novel extended polyrepresentative continuum, Fig. 6, assumes *either* a) that the query derives from *one* information requirement statement, as in almost all (interactive) IR experiments *or* b) that it consists of *two or more* simultaneous representations of the same searcher's cognitive space – as demonstrated by Kelly and colleagues (2005). In the latter case, one may apply the representations as separate search key strings or – as done by Kelly and colleagues (2005) – in the form of one union of strings with weighting of single term overlaps. Search keys can be of 'higher order' (facets, concepts, phrases and non-textual semantics), and thus require pre-processing in order to function, or of 'low order' like single keys, such as, terms of pixel values.

7 Concluding Remarks

The principle of polyrepresentation is a coherent and comprehensive cognitive framework that can be applied simultaneously to the cognitive space of the user and the information space of IR systems. The principle has the potential to guide the design of interactive IR systems that take full advantage of the available document representations *and* user's context to improve retrieval performance. The polyrepresentation continuum attempts to function as a framework for such research. So far, the most promising results have been achieved at the NW and NE portions of the continuum, that is, in the more structured query quadrants of the framework. This is true for laboratory text retrieval involving two or more cognitively or functionally different representations of documents or retrieval techniques, based on one searcher statement type. However, in Kelly and colleagues' promising case (2005) we deal with the opposite research configuration of interaction located in the SE quadrant: several functionally different searcher statements of the information requirement – and one test collection processed by one single engine.

Fundamentally, a number of issues need further investigation. Among these are 1) finding out which document representations that perform best in combination – depending on domains, media, genre, and presentation styles; 2) defining the optimal number and kind of search engines to combine results from and the adequate fusion techniques; 3) carrying out simulation studies that test which retrieval methods would be appropriate for matching different representations of the user's cognitive space with document representations. Such tests could apply simulated work task situations (Borlund, 2003); or they could explore exhaustively the possibilities of a number of controlled variables and thus simulate all achievable combinations Ingwersen and Järvelin, 2005). 4) Then investigations involving test persons ought to be performed to operationalize the findings from the laboratory tests. Query expansion and query adaptation to individual representations seem important and might lead to more formal IR models, which incorporate differentiated normalization and weighting for different representations. It may for instance be that the characteristics of certain representations are underrepresented by the standard *tf*idf* weighting scheme. Instead of calculating *idf* values in relation to the whole document, they might be calculated in relation to each kind of representation in the document in order to capture their individual characteristics more successfully.

Finally, one may explore further the impact of the range of retrieval approaches at the horizontal level of the continuum, in line with the models shown by Ingwersen and Järvelin (2005, p. 116-117) concerned with document and network-based retrieval models of both exact and best match nature.

Acknowledgements. We would like to thank the NORSLIS Research School for support in conducting this research.

References

Ballerini, J.P., Büchel, M., Domenig, R., Knaus, D., Mateev, B., Mittendorf, E., Schäuble, P., Sheridan, P. & Wechsler, M. SPIDER retrieval system at TREC-5. In: *Proceedings of the Fifth Text Retrieval Conference (TREC-5)*. Available at: http://trec.nist.gov/pubs/trec5/t5_proceedings.html (Cited Aug 4, 2006).

Belkin, N.J., Kantor, P., Fox, E. & Shaw, J.A. (1995). Combining the evidence of multiple query representations for information retrieval. *Information Processing & Management, 31*: 431-448.

Borlund, P. (2003): The IIR evaluation model : a framework for evaluation of interactive information retrieval systems. Information Research, 8(3), paper no. 152. [http://informationr.net/ir/8-3/paper152.html]

Buckley, C., Singhal, A. & Cormack, G.V. (1997). Using query zoning and correlation within SMART: TREC-5. In: *Proceedings of the Fifth Text Retrieval Conference (TREC-5)*: 105-118. Available at: http://trec.nist.gov/pubs/trec5/t5_proceedings.html (Cited Aug 4, 2006).

Christoffersen, M. (2004). Identifying core documents with a multiple evidence relevance filter. *Scientometrics*, *61*(3): 385-394.

Clarke, C.L.A. & Cormack, G.V. (1996). Interactive substring retrieval (Multitext retrieval for TREC-5). In: *Proceedings of the Fifth Rext Retrieval Conference (TREC-5)*: 105-118. Available at: http://trec.nist.gov/pubs/trec5/t5_proceedings.html (Cited Aug 4, 2006).

Croft, W.B, Turtle, H.R. & Lewis, D. (1991). The use of phrases and structured queries in information retrieval. In: Bookstein, A. et al. (Eds.) [ACM SIGIR 14]: 32-45.

Hull, D.A. (1997). Using structured queries for disambiguation in cross-language information retrieval. In: *AAAI Spring Symposium on Cross-Language Text and Speech Retrieval Electronic Working Notes*.

Ingwersen, P. (1996): Cognitive perspectives of information-retrieval interaction - elements of a cognitive IR theory. *Journal of Documentation*, 52(1): 3-50.

Ingwersen, P. (2002). Cognitive perspectives of document representation. In: *Emerging Frameworks and Methods. Proceedings of the 4th International Conference on Conceptions of Library and Information Science (CoLIS 4),* July 21-25, 2002, Seattle, USA. Greenwood Village, CO: Libraries Unlimited: 285-300.

Ingwersen, P. & Järvelin, K. (2005) *The Turn: Integration of Information Seeking and Retrieval in Context*. Springer.

Kekäläinen, J. (1999). *The effects of query complexity, expansion and structure on retrieval performance in probabilistic text retrieval.* Tampere, Finland: University of Tampere, Department of Information Studies, Ph.D. Thesis. [Available at: http://www.info.uta.fi/tutkimus/fire/archive/QCES.pdf. Cited Aug 4, 2006]

Kekäläinen, J. & Järvelin, K. (1998). The impact of query structure and query expansion on retrieval performance. In: *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York: ACM Press: 130-137.

Kekäläinen, J. & Järvelin, K. (2000). The co-effects of query structure and expansion on retrieval performance in probabilistic text retrieval. *Information Retrieval,* 1(4): 329-344.

Kelly, D., Dollu, V.D. & Xin Fu (2005). The loquacious user: A document-independent source of terms for query expansion. In: *Proceedings of the 28th Annual ACM SIGIR Conference on Research and Development in Information retrieval*: 457-464.

Larsen, B. (2002). Exploiting citation overlaps for information retrieval: Generating a boomerang effect from the network of scientific papers. *Scientometrics*, *54*(2): 155-178.

Larsen, B. (2004). References and Citations in Automatic Indexing and Retrieval Systems: Experiments with the Boomerang Effect. Copenhagen, DK: The Royal School of Library and Information Science, Ph.D. Thesis. [Available at: http://www.db.dk/blar/dissertation . Cited Aug 4, 2006.]

Larsen, B. & Ingwersen, P. (2002). The Boomerang Effect: Retrieving Scientific Documents via the Network of References and Citations. *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York, NY: ACM Press: 397-398.

Larsen, B. & Ingwersen, P. (2005). Cognitive overlaps along the polyrepresentation continuum. In: *New Directions in Cognitive Information Retrieval*: 43-60.

Larsen, B., Lund, H., Andreasen, J.K. & Ingwersen, P. (2003). Using value-added document representations in INEX. In: *INEX 2003 Workshop Proceedings*: 67-72.

Lund, B.R. (2005). Polyrepræsentation og Datafusion: Test af teorien om polyrepræsentation gennem forsøg med fusion af TREC-5 resultater [Fusion of TREC-5 results – testing the principle of polyrepresentation]. Royal School of LIS, Copenhagen. (MSc Thesis – in Danish).

Lund, B.R., Schneider, J.W. & Ingwersen, P. (2006). Impact of relevance intensity in test topics on IR performance in polyrepresentative exploratory search systems. In: *Evaluating Exploratory Search Systems, Proceedings of the SIGIR 2006 EESS Workshop*: in press.

Shaw, W., Wood, R., Tibbo, H. (1991). The Cystic Fibrosis Database: Content and Research Opportunities. *Library and Information science Research,* 13:347-366.

Skov, M., Pedersen, H., Larsen, B. & Ingwersen, P. (2004). Testing the principle of polyrepresentation. In: *Information Retrieval in Context, Proceedings of the SIGIR 2004 IRiX Workshop*: 47-49. [Available at http://ir.dcs.gla.ac.uk/context/. Cited Aug. 4, 2006.]

Skov, M, Ingwersen, P. & Larsen, B. (2006). Inter and intra-document contexts applied in polyrepresentation. In: *Information Interaction in Context: Proceedings of the First IIiX Symposium.* Copenhagen: Royal School of Library and Information Science, 2006: in press.

Strzalkowski, T., Guthrie, L., Karlgren, J., Leistensnider, J., Lin Fang, Perez-Carballo, J. Straszheim, T., Wang Jin & Wilding, J. Natural language information retrieval: TREC-5 report. *Proceedings of the Fifth Text Retrieval Conference (TREC-5)*: 291-314. Available at: http://trec.nist.gov/pubs/trec5/t5_proceedings.html (Cited Aug 6, 2006).

Turtle, H. R. & Croft, W. B. (1991). Evaluation of an inference network-based retrieval model. *ACM transactions on Information systems,* 9(3): 187–222.

White, R.W. (2006). Using searcher simulations to redesign a polyrepresentative implicit feedback interface. *Information Processing & Management*, vol. 42(5): 1185-1202.

White, R.W., Jose, J.M. & Ruthven, I. (2006). An implicit feedback approach for interactive information retrieval. *Information Processing & Management*,vol. 42(1): 166-190.

White, R.W., Ruthven, I. & Jose, J.M. (2005a). A study of factors affecting the utility of implicit relevance feedback. In: *Proceedings of the 28th Annual ACM SIGIR Conference on Research and Development in Information retrieval*. New York, N.Y., ACM Press: 57-64.

White, R.W., Ruthven, I. & Jose, J.M. & van Rijsbergen, C.J. (2005b). Evaluating implicit feedback models using searcher simulations. *ACM Transactions on Information Systems*, 23(3): 325-361.