

A Quantum-based Model for Interactive Information Retrieval (extended version)

Benjamin Piwowarski
University of Glasgow
benjamin@bpiwowar.net

Mounia Lalmas
University of Glasgow
mounia@dcs.gla.ac.uk

Abstract

Even the best information retrieval model cannot always identify the most useful answers to a user query. This is in particular the case with web search systems, where it is known that users tend to minimise their effort to access relevant information. It is, however, believed that the interaction between users and a retrieval system, such as a web search engine, can be exploited to provide better answers to users. Interactive Information Retrieval (IR) systems, in which users access information through a series of interactions with the search system, are concerned with building models for IR, where interaction plays a central role. There are many possible interactions between a user and a search system, ranging from query (re)formulation to relevance feedback. However, capturing them within a single framework is difficult and previously proposed approaches have mostly focused on relevance feedback. In this paper, we propose a general framework for interactive IR that is able to capture the full interaction process in a principled way. Our approach relies upon a generalisation of the probability framework of quantum physics, whose strong geometric component can be a key towards a successful interactive IR model.

1 INTRODUCTION

In less than twenty years, search engines on the Web have revolutionised the way people search for information. The speed with which one can obtain an answer to a keyword-based query on the Web is fostering interaction between search engines and their users. Helping users to reach relevant material faster will most likely make use of such rich interaction. Another key to future search systems is the context that further defines the search, whether it be external (e.g. time of the day, location) or internal (e.g. the interests of the user).

Putting aside the problem of evaluating such contextual and interactive search (see [10] for a discussion on this topic), building models able to explicitly take into account both is of importance, especially since Information Retrieval (IR) models seem to have reached maturity and there is an obvious need to go beyond current state-of-the-art [13].

There are many reasons why we cannot assume that users will provide enough information to state an unambiguous Information Need (IN), such as a TREC topic description. First, users do not always know how to express their IN and they sometimes have only a vague knowledge of what they are looking for. Second, users knowledge and interests might evolve during the search, thereby modifying their IN. Therefore, it is important that implicit contextual and interaction “information” become integrated directly into IR models *and* experiments [6].

Beside standard relevance feedback models like the Rocchio algorithm [11] or the Okapi model [15], some recent works have attempted to capture context [8] or interaction [12]. However, there is not yet a principled framework that combines both, and that, equally importantly, tries to capture the different forms of possible interactions, namely, query (re)formulation, clicks, navigation. Those tasks are all performed frequently in web searches.

In this paper, we present a framework for interactive and contextual IR. We view search as a process with two different dynamics: (P1) The system tries to capture the user IN while (P2) the user cognitive state, and hence the user IN, is evolving and changing [19]. While the former could be modeled by standard probabilistic models, we claim that the latter can be better modeled by the generalisation of probability theory that has been developed in quantum physics (Section 4.3). Moreover, the strong geometric component of the quantum probability framework is particularly important since standard IR models rely on vector spaces and on (some variants of) the cosine similarity [14]. We show how the quantum formalism generalises these latter models (Section 4.2). In particular, we believe that one strength of the geometric models in IR is that they are intuitive. Adding a probabilistic view on this geometry opens the door for new and potentially more powerful IR models.

Our contribution is to describe how the quantum probability formalism could be used to build an interactive IR framework, first motivating the framework in Section 3 and then describing it more in details in Section 4. We shortly discuss how a working model could be built in Section 4.3.1. Finally, we detail a direct application of our framework as a principled generalisation of the Rocchio’s algorithm in Section 4.3.2.

2 Related works

Interactive IR has been actively studied by the Information Science community [6], and it has been recognised that there is still little progress regarding the development, implementation and validation of a model for interactive IR. In particular, a model that allows for the evolution of the user and the system states is yet to be developed.

An aspect of interactive IR that has been studied for a long time now is relevance feedback, where the information about the relevance of documents is exploited in order to propose a new ordering of the documents. One of the most well-known algorithm is the Rocchio algorithm [11] where the query represen-

tation in a term vector space is modified through user interaction. Each time the user deems a document relevant, the query representation drifts slightly toward the relevant document representation. Within a more principled probabilistic framework, the Okapi model can straightforwardly integrate this kind of information into the statistics used to estimate the model probabilities [15]. In contrast to Rocchio, with Okapi the last feedback has as much importance as the first ones, which is not a good property for interactive IR models.

One limitation of these two IR models is that they rely on explicit relevance feedback, which might not be available or precise enough [7]. Fortunately, it is possible to adapt these approaches when feedback is implicit. Shen et al. [12] uses a term vector space representation of a user. The vector is updated when the user clicks on search result links, based on the summaries of the documents (and not on the content of the documents). However, the update still relies on the Rocchio algorithm which lacks a principled motivation. In this paper, we show how our framework can extend Rocchio’s algorithm, providing a principled way to model the IN drift.

More general and holistic models have been proposed to build an interactive retrieval system (e.g. Fuhr [4] and Shen et al. [12]) that rely on a decision-theoretic framework to determine what is the best next action the system should perform, i.e. what documents should be presented to the user. In such approaches, decisions are made based on the relevance of documents when considering past interactions. In this paper, we focus on the latter problem and do not discuss how to select the best next action the system has to perform.

Finally, some works have already explored the use of the quantum theory framework in IR. Firstly, the proposition of using quantum theory to model IR processes and the interpretation of (among others) the standard vector model within the quantum probability formalism have been detailed in [14]. Following, some works have tried to exploit the formalism, from a very abstract and generic level [2] to the formalisation of existing IR models [21]. Another work looked at how to represent documents in a space different from the standard term space [5]. Our contribution with respect to these works is to define a new space, the IN space, and to clearly map different aspects of user-system interaction to operations within this space (Section 3). In our representation, we also make use of the distinction between mixtures and superposition, two related but different concepts in quantum theory (Section 4.2), which proves to be useful in our representation.

The most related work in that field is that of Melucci’s [8], which computes the probability of having a given context $\Pr_d(C)$, where \Pr_d is the probability distribution generated by the document vector d , and the subspace C is equalled to the context and is built through user interaction. More specifically, given a set of documents deemed relevant, either using user feedback or pseudo-relevance feedback, one can compute a subspace C corresponding to the principal components of a subspace spanned by those documents. A document vector fully included in this space will be fully relevant (probability of 1), an orthogonal one will be fully irrelevant (zero probability). Melucci’s approach is dual to ours, in the sense that instead of representing users in an IN space, he considers

documents in a contextual space. Our approach, which relies on an IN space, facilitates the possibility of using the different quantum evolution mechanisms to model the interaction between the user and the retrieval system.

In this paper, we explain how our formalism can exploit implicit (i.e. clicks) and explicit feedback, like relevance judgements or query reformulation. Other forms of interaction like e.g. going back to the retrieval list, exploration of the result list, etc. could also be captured by our formalism using the same represent as for e.g. query (re)formulation, but we do not deal with those in this paper. To the best of our knowledge, there is no proper theoretical framework that allows to capture this different feedback information in a principled way and within a uniform framework.

3 An Information Need Space

We build upon the work described in [9] where a high level description of the framework was first given. Our working hypothesis is that a *pure*, in the sense that we know exactly what the user is looking for, user IN can be represented as a system in quantum physics, i.e. as a unit vector in a Hilbert space¹, and that this state evolves while the user is interacting with the system.

According to the quantum probability formalism, this (IN) vector generates a probability distribution over the different subspaces of the Hilbert space. We make the hypothesis that among other possible uses, such subspaces can be related to the relevance of documents, therefore enabling the computation of a relevance score for a document, and to user interactions (like typing a query or clicking on a document), making it possible to exploit them.

From a geometric perspective, using subspaces to describe “regions” of INs has been (sometimes implicitly) studied and motivated in some works relying on a vector space representation. First, Wang et al. [16] studied the problem of negative feedback, finding that this could be exploited by describing the IN by a set of vectors (instead of one as in Rocchio’s algorithm). Second, Widdows [17] has shown that term negation in a query was better modelled with orthogonality, which implicitly implies the use of subspaces to represent negations of INs. Finally, Zuccon [22] has shown that the cluster hypothesis still holds when representing documents as subspaces.

Using those IN “regions”, the search process would be modelled as follows. At the very beginning of the search process, the user IN is underspecified and is a mixture of *all* possible pure INs. That is, without any information about the user, we can only know that the user is in one of all the possible IN states with a probability that depends e.g. on how popular this IN is.

We believe that using an IN space can model interactive IR since users change their point of view during a search, and relevance, contrarily to topicality, is expected to evolve within a search session [19]. More specifically, we can identify two different types of dynamics within the search process: (P1) The IN becomes

¹In brief, an inner product vector space defined over the complex field, see [14] for a formal definition

increasingly specific *from a system point of view*, e.g. when a user types some keywords or clicks on some documents, i.e. the uncertainty is reduced; and (P2) The IN changes *from a user point of view*. The IN can become more specific as the user reads some documents, or it can slightly drift as user interests do. This type of drift has been observed and analysed by [18] within search sessions.

An example of such a two-part process is described in [20] where the hypotheses are that a user has a topicality and a novelty profile, the former providing the context of the IN while the latter changes during a search session. Our model captures both processes within the same space, which we refer to as an *Information Need Space* in this paper.

Whereas the first process can be easily described within a standard probabilistic framework (we restrict the IN to subspaces of the whole space), the latter would benefit from a quantum probability formalism as the INs can drift from two overlapping subspaces. For example, consider Figure 1 and assume that the possible user INs lie within a three-dimensional space. If we first restrict the IN to subspace A , then this corresponds to process (P1): Vectors orthogonal to A would be discarded. Later, if we know based on further evidence that the IN subspace is now B , then this corresponds to process (P2): As vectors belong to A , they are not orthogonal to B and thus projecting them in B makes them drift away from their original representation. We posit that the classical probabilistic framework would address the uncertainty of the system view over the retrieval process (P1) whereas the quantum probability framework addresses the changes of the user internal state (P2). As the quantum probability framework is a generalisation of the probabilistic one, we can use the same representation and evolution operators to model both processes. We describe this framework in the following section.

4 A quantum view

In this section, we describe our framework while also introducing the framework of quantum probability.

4.1 Link with standard probability theory

Quantum probability can be thought of as an extension of classical probability theory, and relies on linear algebra in Hilbert spaces. The equivalent of a logical proposition or event is a subspace. As there is a one-to-one relationship between subspaces and projector on subspaces, we will switch between one interpretation and the other, depending on what we want to illustrate. A projector is a special case of an *observable*, more precisely a yes/no observable: The observable is binary because a state (vector) can belong or not to the associated subspace. It is possible to naturally extend standard probability theory, by defining a probability density operator over those subspaces, i.e. by associating with each subspace/projector a probability.

Next, we first present how to reconstruct a classical probability space using a Hilbert space, in order to show how the quantum formalism is a general case of probability theory. We then show how this can be extended to the quantum probability formalism, and highlight their differences and uses in our interactive IR framework.




Figure 1: Three two-dimensional subspaces (A, B, C) in a three dimensional space.

For simplicity, let us consider a probability space with a finite sample space $\Omega = \{\omega_1, \dots, \omega_N\}$ of cardinality N where each ω_i corresponds to a *pure* IN. Moreover, as each ω_i corresponds to a distinct atomic event, we require that those INs be orthogonal (i.e. a user with the IN ω_i will not be interested by a piece of information answering ω_j where $i \neq j$), a restriction that can easily be removed within the quantum formalism as discussed in Section 4.2. As usual, a probability distribution is associated with this sample space, i.e. a probability $\Pr(\omega_i)$ is associated with each *atomic* event in Ω . Then, an event is a subset \mathbf{A} of the sample space Ω whose probability is $\Pr(\mathbf{A}) = \sum_{\omega \in \mathbf{A}} \Pr(\omega)$.

In the sequel, we use the bra-ket Dirac notation as done in quantum theory. Restricted to finite vectorial spaces, a vector of the Hilbert space is denoted by $|x\rangle$ while its conjugate transpose is denoted $\langle x| \doteq |x\rangle^\dagger$.

Let us see how our probability space is formalised within the quantum probability formalism. We associate our sample space Ω with one basis of a Hilbert space of dimension N . Let us choose an arbitrary orthonormal basis $\mathcal{B} = \{|\omega\rangle \mid \omega \in \Omega\}$ where each element of the basis $|\omega\rangle$ corresponds to the atomic event ω . These vectors are pairwise orthogonal and of norm 1. To each atomic event of the sample space is associated a one-dimensional projector \mathbb{P}_ω corresponding to the one-dimensional subspace defined by $|\omega\rangle$, and which is equal to $|\omega\rangle\langle\omega|$. Note that it corresponds to a $n \times n$ matrix where n is the dimension of the Hilbert space. The probability distribution over the Hilbert space is defined by a so-called *density operator* ρ defined by

$$\rho = \sum_{\omega \in \Omega} \Pr(\omega) \mathbb{P}_\omega \quad (1)$$

where the $\Pr(\omega)$ sum up to 1. Note that a special case of a density operator is simply \mathbb{P}_ω for a given atomic event ω when the exact state of the system is known, e.g. the IN of the user is completely specified. In this case, the whole probability mass is concentrated onto a single atomic event, i.e. $\Pr(\omega) = 1$.

The event \mathbf{A} is associated with a yes/no observable $O_{\mathbf{A}}$. It is a projector associated with the subspace spanned by the basis vectors $|\omega\rangle$ for $\omega \in \mathbf{A}$. This projector can be written $O_{\mathbf{A}} = \sum_{\omega \in \mathbf{A}} \mathbb{P}_\omega$. In the following formula where tr denotes the trace operator, we give both quantum and standard definitions to

show how they relate:

$$\underbrace{\Pr_\rho(O_A) \doteq \text{tr}(\rho O_A)}_{\text{quantum definition}} = \text{tr} \left(\sum_{\omega \in \Omega} \Pr(\omega) \mathbb{P}_\omega \sum_{\omega \in \mathbf{A}} \mathbb{P}_\omega \right) = \underbrace{\sum_{w \in \mathbf{A}} \Pr(\omega)}_{\text{standard definition}} \doteq \Pr(\mathbf{A}) \quad (2)$$

The equality is due to the fact that $\mathbb{P}_\omega \mathbb{P}_{\omega'} = 0$ whenever $\omega \neq \omega'$, as we assumed the vectors were pairwise orthogonal. This is exactly the way of computing a probability by decomposing it into atomic and incompatible events (as shown on the right side of the equation). Without giving too many details, we point out that the quantum definition is more general because O_A can be expressed in a different basis than \mathcal{B} .

We have just defined a formalism that allows us to express a classical probability distribution using linear algebra. All the information about the probability distribution is contained into a *density operator* ρ , and it can be shown that for *any* probability distribution over a Hilbert space there exists a corresponding density operator [14, p. 81]. From a practical point of view, the above description of standard probabilities with Hilbert spaces unlocks the potential of defining probabilities through geometric relationships, and permits a generalisation to a non standard probability formalism, which we describe in the next section.

We posit that at this level, we are able to model the first component of the search process, which corresponds to finding the right *subspace* of the IN, i.e. in classical terms to find the subset of the IN sample space. However, it is intuitive to think that INs are not mutually exclusive. We make the hypothesis that such a non-exclusiveness is captured by the geometry of IN space, and this can be modelled within a quantum probability formalism, as discussed in the next section.

4.2 Superposition, mixtures and information needs

Now we turn to how probability theory is generalised. We introduce the notion of superposition and mixture, and relate them to their use in our model of interactive IR. Said shortly, superposition relates to an ontologic uncertainty (the system state is perfectly known, but some events are true *only* with a given probability) whereas mixture relates to standard probabilistic uncertainty (the system is in one of the states with a given probability). Superposition is a salient characteristic of quantum probabilities and is important since it gives us a way to represent geometrically new INs while the quantum probability framework ensures we can still compute probabilities for the new INs. Mixture and superposition gives us more flexibility in the way we can represent our current state of knowledge of an IN.

Let us illustrate this with an example. Suppose that $|\omega_T\rangle = \begin{pmatrix} 1 & 0 \end{pmatrix}^\dagger$ and $|\omega_L\rangle = \begin{pmatrix} 0 & 1 \end{pmatrix}^\dagger$ form a basis of the IN space. Suppose the former represents the IN of a user looking for information about tigers (T) and the latter about lions (L). In

order to represent a user looking for a tigrion (the offspring of a tiger and a lion), we assume that this can be represented by the vector $|\omega_{TL}\rangle = \frac{1}{\sqrt{2}}(|\omega_T\rangle + |\omega_L\rangle)$ which is a *superposition* of two INs, where the $\frac{1}{\sqrt{2}}$ factor ensures $|\omega_{TL}\rangle$ norm is one. This is a strong assumption which we will study when experimenting with the framework. Aerts and Gabora [1] worked on how to combine concepts in a (quantum) vector space, but use spaces of increasing dimensionality to do so (through the use of a tensor product). As a final remark on superposition of INs, we would like to note that complex numbers could be used to combine INs, e.g. to distinguish tigrions (the tiger is the father) from ligers (the lion is the father), and that superposition is not restricted to topicality. For instance, assuming that we know how to represent a user searching for a paragraph and a user searching for a chapter, we could imagine representing a user looking for a paragraph as a superposition of both.

The superposed IN $|\omega_{TL}\rangle$ is quite different to a user who is equally interested by tigers or lions, which would be represented as a *mixture* of the INs $|\omega_T\rangle$ and $|\omega_L\rangle$. Formally, this IN would be associated with a density operator $\rho_{T\vee L} = \frac{1}{2}(\rho_T + \rho_L)$ where ρ_L and ρ_T are respectively the projectors associated with $|\omega_T\rangle$ and $|\omega_L\rangle$. For example, $\rho_T = |\omega_T\rangle\langle\omega_T|$. The density operator $\rho_{T\vee L}$ is to be interpreted by saying that with probability one half the IN is about tigers (or equivalently about lions).

We can see also the difference if we represent the densities by their matrices in the $(|\omega_T\rangle, |\omega_L\rangle)$ basis. We have the mixture of IN $\rho_{T\vee L} = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ which is different from the pure IN $\rho_{TL} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$. An important observation is that these different densities imply different probabilities. Let us suppose that the relevance of a document corresponds to a yes/no observable, and that the relevance of a document about lions (respectively tigers, tigrions) are represented by the projectors (yes/no observables) O_L , O_T and O_{TL} associated with the subspaces generated by $|\omega_T\rangle$, $|\omega_L\rangle$ and $|\omega_{TL}\rangle$, respectively. For example, $O_T = |\omega_T\rangle\langle\omega_T|$. According to Eq. (2), we can compute the probability of relevance of the different documents, which gives:

$$\Pr_{\rho_{TL}}(O_L) = \Pr_{\rho_{T\vee L}}(O_L) = \frac{1}{2} \text{ and } \Pr_{\rho_{TL}}(O_{TL}) = 1 \neq \Pr_{\rho_{T\vee L}}(O_{TL}) = \frac{1}{2}$$

Interestingly, we cannot distinguish the probability of relevance of the document about lions when the IN is about either tigers and lions or about tigrions (two first probabilities) but there are two reasons for this: In the former, the probability $\frac{1}{2}$ is caused by the discrepancy between the IN and the document, whereas in the second case the probability is due to the fact that the document only covers a part of the information need. Next, thanks to the quantum formalism the probabilities for the same INs are different when we evaluate the relevance of the document about tigrions (two last probabilities). We thus benefit from a two-dimensional space to distinguish different INs that would be expressed similarly in a standard vector space model. One consequence is that if we search for a set of documents that satisfy T or L , we would have two different types

of documents (about tigers and lions, assuming each document covers one IN only) whereas one document would satisfy TL .

Mixtures are also useful to represent the IN density operator ρ_0 at the very beginning of the information retrieval process, as we do not know which state the user is in. We would define the initial IN density operator as $\rho_0 = \sum_i \text{Pr}_i \mathbb{P}_i$ where i ranges over all the possible information needs and Pr_i is the probability that a random user would have the IN i when starting a search. Using the mixture is also motivated by the fact that we deal with classical undeterminism, i.e. we know the user is in a given state but we do not know which. The mixture can also be thought as a set of vectors describing all the possible INs, each vector being associated with a probability. This representation is particularly useful in the next section where we show how this initial IN ρ_0 is transformed through interactions.

4.3 Measurement and Interaction

Beside differentiating mixture and superpositions, the quantum formalism has also consequences for computing a conditional probability. These consequences are linked to the way a measurement is performed in quantum physics. We use measurement to model interaction and describe in this section both how the measurement modifies the density operator ρ and how we link measurement to the different interactions.

For simplicity, we now use A to denote the related yes/no observable, subspace or projector. Since there is a one-to-one correspondence between them [14], they can be used to denote the same thing albeit in a different context. Given a system density operator ρ , if we observe O_A , the new density operator denoted $\rho \triangleright O_A$ is defined by

$$\rho \triangleright O_A = O_A \rho O_A / \text{tr}(\rho O_A) \quad (3)$$

This amounts to restricting ρ to the subspace defined by O_A and ensuring that $\rho \triangleright O_A$ is still a density operator. The effect of the restriction is to project every IN of the mixture ρ onto the subspace defined by O_A (with some renormalisation to ensure the probabilities still sum up to 1). One can readily verify that the probability of O_A with respect to $\rho \triangleright O_A$ is 1. It means that when A has just be measured, we know it is true at least until further interaction (or in general, evolution) modifies the density operator. Measurement can be thought as a generalisation of conditionalisation, as we can compute the conditional probability of O_A given O_B , or more precisely of measuring O_A knowing that we have measured O_B , as $\text{Pr}_\rho(O_B|O_A) = P_{\rho \triangleright O_A}(O_B)$.

In quantum theory, the order of the measurements is important, since in general the densities $\rho \triangleright O_A \triangleright O_B$ (applying two times the Eq. (3), for O_A and then for O_B) and $\rho \triangleright O_B \triangleright O_A$ are different. It is a desirable property whenever subsequent measurements of a system should yield different results, which is the case in interactive IR: The sequence of interactions represents the evolution of the user, and should be taken into account. A user drifting from an IN (e.g. hotels in Barcelona) to another (e.g. museum in Barcelona) is not the

same as the reverse, which illustrates the adequacy of the quantum formalism to handle such drifts. This is illustrated by Figure 1, where visually it can be seen that measuring O_B (hotels) then O_C (museums) is different from the reverse, since in the first case the IN vectors will lie in the subspace C whereas they would lie in B in the other case.

Starting with the initial density operator ρ_0 (section 4.2), we make the assumption that each implicit or explicit interaction between the IR system and the user corresponds to a measurement, i.e. that every interaction is associated with a yes/no observable O . After the interaction, we can recompute the IN density operator using Eq. (3). For example, a user whose internal context is associated as O_{user} , who asked a query associated with O_{q_1} and deemed a document relevant (associated with O_{d_1}), would be represented by a density operator $\rho_0 \triangleright O_{\text{user}} \triangleright O_{q_1} \triangleright O_{d_1}$. Among other users, this density operator can be used to predict the relevance of other documents.

In the following, we make our framework more concrete: In Section 4.3.1 we show how some interactions would be mapped to an observable. We argue that direct measurement can be too harsh in some cases and show how a generalisation of Rocchio’s rule can be used to overcome this problem in Section 4.3.2.

4.3.1 Mapping interactions to observables

In order to map interactions to observables, we restrict to the topical relevance and assume a vector space where dimensions are associated with terms. How to deal with more relevance dimensions is left for future work. We also assume we know how to compute the initial density operator ρ_0 – which could be approximated using the document representation described next.

Given the current IN density operator ρ_t , we can compute the probability of relevance $\text{Pr}_{\rho_t}(O_d)$ of a document d , provided O_d is the observable associated with the relevance of document d . To build such an observable, and as a first approximation, we can suppose that each paragraph p corresponds to exactly one IN $|\omega_p\rangle$, and hence that its representation is a one dimensional subspace. It is then possible to compute the subspace spanned by the vectors $\{\omega_p\}$ corresponding to the different paragraphs, and use this subspace to represent the document relevance. When a user deems a document relevant, we could use the same representation to update the current IN ρ_k . In that case, we would have the new IN density operator $\rho_{t+1} = \rho_t \triangleright O_d$.

The first possible type of interaction would be the (re)formulation of a query by a user. We would associate to a given query a subspace/observable O_q , and update the current probability density operator ρ_t to $\rho_{t+1} = \rho_t \triangleright O_q$. A representation of the query could for example be computed through pseudo-relevance feedback provided we know how to represent the documents: The subspace associated with O_q would then be the subspace spanned by the observables representing the top-ranked documents (by a standard IR algorithm). For example, in Figure 1, if A and B correspond to two different top-ranked documents for a given query, then O_q would correspond to the whole three dimensional space (i.e. the join of subspaces A and B). Another way to compute the query observable

O_q , without relying on an external model, would be the union of the subspaces representing the paragraphs where each term of the query appears.

Here, we give one illustration of the usefulness of the quantum formalism for an interactive IR framework. The query observable O_q (or the document observable O_d) can be used to detect if a user’s change of mind is too important to be a simple drift, an important feature an interactive IR system should have [18]. Within the quantum framework, we use the same geometric representation to both update the density operator knowing an event and to compute the probability of this event. Indeed, when at time t the user types a new query q' , we can compute the probability of the query according to the current IN density operator ρ_t , i.e. compute $\Pr_{\rho_t}(O_{q'})$. Based on this value, our IR system would decide that the user switched to a new IN, and react accordingly.

4.3.2 A probabilistic generalisation of Rocchio’s update rule

Sometimes measurement can be too harsh, either because we are not sure we made the observation (as in relevance feedback) or because we do not expect the user IN to completely drift to the INs covered by e.g. a clicked document. Said otherwise, we want to give some inertia to the user and this is exactly what Rocchio’s algorithm does. We show how this algorithm can be extended in a principled way with our formalism.

Consider a user looking for information about tigers or lions, which can be represented, respectively, as the subspaces A and C of Figure 1. If the user deems relevant a document about tigers, the updated density operator would only be made of INs belonging to the subspace A , which is orthogonal to C . Documents about lions would be henceforth considered as non-relevant since $\Pr_{\rho \triangleright O_A}(C) = 0$ whatever the density operator ρ is. In order to give some inertia to the underlying dynamics, we could as in Rocchio’s algorithm use a parameter α such that the new density operator combines both the last IN density operator and the new one, i.e. $\rho_{t+1} = \alpha(\rho_t \triangleright O_A) + (1 - \alpha)\rho_t$. This formula is interpreted as the fact that with probability α the user IN has been projected onto the subspace defined by document A , and with probability $1 - \alpha$ it remained the same. It can be said that the above formula is a principled generalisation of the Rocchio update rule. However, it is not equivalent to it, since each IN vector composing the density operator ρ_t is (in general, i.e. when it does not belong to subspace A) decomposed into two vectors, one projected onto A while the other remains the same whereas with Rocchio, we only deal with one vector, which has limitations e.g. when using negative feedback [16].

5 Conclusion

We proposed a new interactive IR framework, which exploits the strong connection between geometry and probabilities present in the quantum probability formalism. Our framework allows for a principled and geometric mapping of user interactions into an IR model. In particular, we show how to handle

click/relevance feedback and query reformulation. How to use the latter information has not been explored in IR so far, beside providing query recommendation. Other forms of interaction (e.g. navigation) would fit our framework, through the definition of associated subspaces. Beside measurement, the quantum framework is powerful enough to provide other types of evolution of the IN density operator. This would provide a way to predict how a user might evolve, e.g. in order to predict that users looking for hotels might look for museums in a town.

Diversity and novelty are two important components of interactive IR system [20] that we did not discuss. We believe that the non-standard logic associated with the subspaces would be a useful property for such purposes, since we could represent a set of documents as the union subspace spanned by the different document subspaces. The relationship between the union subspace and the possible information needs provide enough information to compute how the possible INs are addressed and how novel the documents are.

From a more practical perspective, future work includes experimenting with different possible ways of mapping interactions to observables, dealing with complexity issues, in order to build an interactive IR model based on the ideas we developed in this paper.

Acknowledgments

This research was supported by an Engineering and Physical Sciences Research Council grant (Grant Number EP/F015984/2). Mounia Lalmas is currently funded by Microsoft Research/Royal Academy of Engineering.

References

- [1] D. Aerts and L. Gabora. A theory of concepts and their combinations II: A Hilbert space representation. *Kybernetes*, 34, 2005.
- [2] S. Arafat. *Foundations research in information retrieval inspired by quantum theory*. PhD thesis, University of Glasgow, December 2007.
- [3] P. Bruza, D. Sofge, W. Lawless, C. J. van Rijsbergen, and M. Klusch, editors. *Proceedings of the Third Quantum Interaction Symposium*, volume 5494 of *Lecture Notes in Artificial Intelligence*. Springer, March 2009.
- [4] N. Fuhr. A probability ranking principle for interactive information retrieval. *Information Retrieval*, 12, 2008.
- [5] A. F. Huertas-Rosero, L. Azzopardi, and C. J. van Rijsbergen. Eraser lattices and semantic contents. In Bruza et al. [3], pages 266–275.
- [6] P. Ingwersen and K. Järvelin. *The Turn: Integration of Information Seeking and Retrieval in Context (The Information Retrieval Series)*. Springer-Verlag, Secaucus, NJ, USA, 2005.

- [7] M. Magennis and C. J. van Rijsbergen. The potential and actual effectiveness of interactive query expansion. In N. J. Belkin, A. D. Narasimhalu, P. Willett, and W. Hersh, editors, *Proceedings of the 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 324–332, New York, NY, USA, July 1997. ACM.
- [8] M. Melucci. A basis for information retrieval in context. *ACM Transactions On Information Systems*, 26(3):1–41, June 2008.
- [9] B. Piwowarski and M. Lalmas. Structured information retrieval and quantum theory. In Bruza et al. [3].
- [10] S. E. Robertson. On the history of evaluation in IR. *Journal of Information Science*, 34:439–456, 2008.
- [11] J. J. Rocchio. Relevance feedback in information retrieval. In G. Salton, editor, *The SMART retrieval system: experiments in automatic document processing*, pages 313–323. Prentice-Hall, Englewood Cliffs, USA, 1971.
- [12] X. Shen, B. Tan, and C. Zhai. Implicit user modeling for personalized search. In *Proceedings of the 14th ACM international conference on Information and knowledge management*, pages 824–831, Bremen, Germany, 2005. ACM.
- [13] K. Sparck Jones. What’s the value of TREC: is there a gap to jump or a chasm to bridge? *SIGIR Forum*, 40(1):10–20, 2006.
- [14] C. J. van Rijsbergen. *The Geometry of Information Retrieval*. Cambridge University Press, New York, NY, USA, 2004.
- [15] S. Walker and S. E. Robertson. Okapi/keenbow at TREC-8. In E. M. Voorhees and D. K. Harman, editors, *NIST Special Publication 500-246: The Eighth Text REtrieval Conference (TREC-8)*, Gaithersburg, Maryland, USA, November 1999.
- [16] X. Wang, H. Fang, and C. Zhai. A study of methods for negative relevance feedback. In S.-H. Myaeng, D. W. Oard, F. Sebastiani, T.-S. Chua, and M.-K. Leong, editors, *Proceedings of the 31st Annual International ACM SIGIR*, pages 219–226, New York, NY, USA, July 2008. ACM.
- [17] D. Widdows. Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In *Proceedings of the 41st ACL conference*, pages 136–143, Morristown, NJ, USA, 2003. Association for Computational Linguistics.
- [18] H. I. Xie. Shifts of interactive intentions and information-seeking strategies in interactive information retrieval. *Journal of the American Society for Information Science*, 51(9):841–857, 2000.

- [19] Y. Xu. The dynamics of interactive information retrieval behavior, part i: An activity theory perspective. *Journal of the American Society for Information Science and Technology*, 58(7):958–970, 2007.
- [20] Y. Xu and H. Yin. Novelty and topicality in interactive information retrieval. *Journal of the American Society for Information Science and Technology*, 59(2):201–215, 2008.
- [21] G. Zuccon, L. Azzopardi, and C. J. van Rijsbergen. A formalization of logical imaging for information retrieval using quantum theory. In *IEEE proceedings of the 5th International Workshop on Text-based Information Retrieval*. IEEE, 2008.
- [22] G. Zuccon, L. Azzopardi, and C. J. van Rijsbergen. Semantic spaces: Measuring the distance between different subspaces. In Bruza et al. [3].