On Human-Aware Information Seeking

Simon Schiff[1], Ralf Möller[1]

[1] University of Lübeck, Institute of Information Systems, Ratzeburger Allee 160, 23562 Lübeck, Germany

Abstract

Large parts of scientific work relies on seeking for information in very large datasets and respective metadata (e.g., document repositories on the web, databases, local image collections). Based on a search string or even sample data as a query, information retrieval systems (IR systems) return lists of ranked items that match the query, together with a short preview of the item. Using search strings or example data, it is not easy to express certain information needs, however. In this extended abstract we discuss in what way the interaction of a user with an information retrieval (IR) system can optimized with human-aware collaborative planning strategies.

Keywords

Human-Aware IR, Collaborative Planning

1. Introduction

Information seeking is a process at where humans search for information in, for instance, documents. As depicted in Figure 1, information seeking (IS) processes can be further differentiated. In case of single-step ad-hoc retrieval, every query is executed independently, which is okay if information retrieval (IR) systems return relevant documents in a single step most of the time. If this is not the case, IR systems can improve IR performance when it is known whether the user, i.e., a human, has a purposeful (telic) goal. If no, the system might return documents the human might not expect to be retrieved in the first place. For instance, in searching for news articles a human might not have a particular goal. In case there is a telic goal, the problem is to appropriately specify the respective information need to ensure decent precision values for query answers, in particular when datasets are to be accessed (and not only document repositories).

As part of our research context we consider information needs of humanities scholars, seeking for specific information, e.g., to prepare expert testimonies about certain artifacts and their materials. Not only for materials data, researchers need comprehensive information about underlying datasets to express a one-step query that enables the system to return datasets and documents with appropriate information for the task under consideration. However, detailed knowledge about datasets (and respective documents) is usually not available, such that multiple queries will be required. An IR system could improve its performance by observing the changes

Humanities-Centred AI (CHAI), Workshop at the 44th German Conference on Artificial Intelligence, September 28, 2021, Berlin, Germany

schiff@ifis.uni-luebeck.de (S. Schiff); moeller@ifis.uni-luebeck.de (R. Möller)
https://www.ifis.uni-luebeck.de/index.php (S. Schiff); https://www.ifis.uni-luebeck.de/index.php (R. Möller)
0000-0002-1986-3119 (S. Schiff); 0000-0002-1174-3323 (R. Möller)

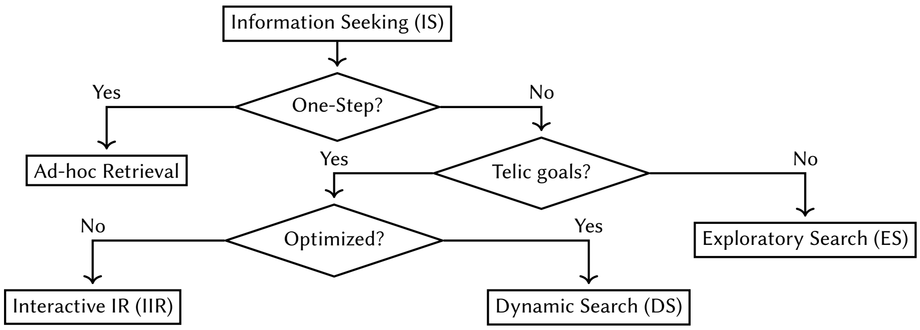

Figure 1: Information Seeking Overview [1]

of the issued queries over time and might be able to approximately represent and the scholar's (telic) long-term goal, and consider it for IR query answering tasks.

Queries issued by a human as well as the respective sets of items returned as answers can be combined to form a so-called session, ending when the information need of the human is satisfied. The IR system could compare each query with its predecessor in a session. Depending on which words are added and removed, the system could improve its performance (interactive information retrieval (IIR), see Figure 1). This still might not lead to optimal results as the system does not reflect that the human possibly has a long-term goal. Given a set of queries part of one session, an IR system has to estimate what the long-term goal of the human might be (dynamic search (DS), cf. Figure 1, see also [1]).

In this extended abstract, we discuss the basic design of an IR agent that is equipped with an IR goal, perceives its environment through various metrics (sensors), builds a set of models (online/offline) of itself as well as an approximate one of the human, and then uses these models to select (actions) sets of datasets and documents in order to fulfill the IR goals.

2. Collaboration Process

The agent can only satisfy the information need of the human if it can anticipate his IR goal. To anticipate his IR goal, it requires to collaboratively work with him. Collaboration starts at where the human expresses his information need as a query. The human sends the query to the agent and the agent then updates its model that contains an approximation of the IR goal of the human. The agent then sends a result to the human, depending on its updated model. Depending on the result, the human sends a reformulated query and optionally feedback to the agent. The feedback helps the agent to close the gap between its model and the IR goal of the human. This process repeats until the information need of the human is satisfied. We refer to the whole process to as a session that ends if the agent has anticipated the IR goal of the human and thus is able to satisfy the information need of the human.

Many different kind of problems can occur during the session, as the IR agent and the human have to collaborate with each other via queries, feedback, and results without being able to physically observe each other. If a human expresses a query, then he has an approximate mental model of the IR agent, which has an influence on of how the human expresses his query. For

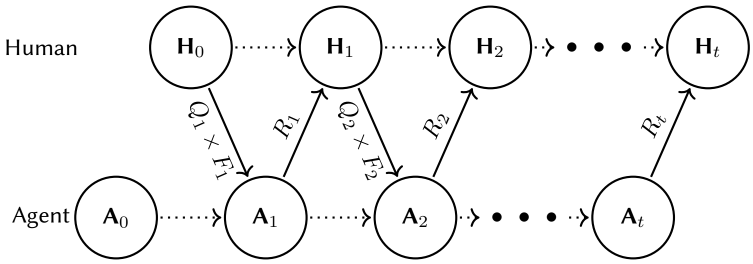

Figure 2: Collaborative Planning over a Session S_t

instance, the human expects that his query needs to be in a specific foreign language, the IR agent only compares the query with the titles of documents, or the IR agent weights some terms part of the query to high such that the human omits them. However, the human might be wrong with his expectations. The IR agent is aware of that the human has specific expectations of itself and aims to anticipate them. If the IR agent correctly anticipates them, then it might be able to adapt its behaviour. Adapting the behaviour needs to not only satisfy the expectations of the human. In addition, the behaviour of the IR agent needs to be explicable from the perspective of the human. Otherwise the human is not able to have an accurate approximate mental model of the IR agent and has difficulties in expressing a query. The IR agent cannot correctly anticipate, adapt, and act explicable for every human of who it collaboratively seeks for informations and is aware of that. In case the IR agent in uncertain whether it can collaborate with the human or the human explicitly requests an explanation it has to explain its behaviour.

More formally, an IR agent and a human interact with each other over a finite number t of time steps T in a session, with $t \in \mathbb{N}_0$. As proposed by Kambhampati et al. [2] with a slightly different notation, while the human and the agent interact with each other, both are modeled as $\mathcal{M}^{\mathbf{H}}$ and $\mathcal{M}^{\mathbf{A}}$ respectively. Model $\mathcal{M}^{\mathbf{H}}$ contains the information need of the human that can only be approximated by the agent as $\widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}$ and model $\mathcal{M}^{\mathbf{A}}$ is represented as $\widetilde{\mathcal{M}^{\mathbf{A}}}$ and approximated by the human as $\mathcal{M}^{\mathbf{A}}_{\mathbf{h}}$. In addition, the IR agent approximates $\mathcal{M}^{\mathbf{A}}_{\mathbf{h}}$ as $\widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h}}}$, in order to reflect whether itself acts explicable or not. As depicted in Figure 2, the agent updates incrementally at each time step t its models $\mathbf{A}_t = (\widetilde{\mathcal{M}^{\mathbf{A}}}, \widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}, \widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h}}})$, given a query Q_t and possibly available feedback F_t and the human updates its models $\mathbf{H}_t = (\mathcal{M}^{\mathbf{H}}, \mathcal{M}^{\mathbf{A}}_{\mathbf{h}})$ given a result R_t that the agent selects and might satisfy the information need of the human.

At each time step t during a session, the IR agent's $\mathbf{A}_t$ and human's mental models $\mathbf{H}_t$ are in a specific state that change with each time step t. Thus a session S is a sequence of session states $S = (S_t)_{t=0}^{T}$, where each session state S_t contains the current query Q_t, feedback F_t the human expresses, given results R_{t-1} the IR agent has send to the human in the previous session state, result R_t that might satisfy the information need of the human, and updated mental models $\mathbf{H}_t$ of the human and $\mathbf{A}_t$ of the agent.

A session S at time step t is successful if R_t in S_t satisfies the information need of the human. Each session starts with initial, possibly non-empty, models $\mathbf{H}_0$ and $\mathbf{A}_0$ at time step $t = 0$. Then, at the next time steps $t \geq 1$, the IR agent first updates its models $\mathbf{A}_{t-1}$, given query Q_t and optionally available feedback F_t:

$$\mathbf{A}_{t-1} \times Q_t \times F_t \longrightarrow \mathbf{A}_t$$

and then selects result R_t, depending on $\mathbf{A}_t$:

$$\mathbf{A}_t \longrightarrow R_t$$

The human receives R_t and updates his models:

$$\mathbf{H}_{t-1} \times R_t \longrightarrow \mathbf{H}_t$$

and then expresses his information need as query Q_{t+1} and feedback F_{t+1}:

$$\mathbf{H}_t \longrightarrow Q_{t+1} \times F_{t+1}$$

Session S helps the agent to close the gap between the human models $\mathbf{H}_t$ and its own models $\mathbf{A}_t$ over time and thus to anticipate the IR goal of the human.

3. Representation of Mental Models for Collaboration

As noted by Kambhampati et al., collaboration between a human and an agent works only, if both have an approximate mental model of each other with a small gap in between [3]. Thus, the human and the agent need to close the gap over time by sharing their models with each other. However, while the agent has a set of models with a concrete representation, the human has not. Both cannot directly interpret the model of the other's model and therefore need to share their models by sending/receiving queries, results, and feedback to each other respectively. The challenging task is to (i) find good representations for $\mathbf{H}_t$ and $\mathbf{A}_t$ and to incrementally update them at each time step, (ii) a language for the queries, results and feedback, and (iii) to evaluate the IR agent. For the latter, as best of our knowledge, only a few evaluation methods exist in the literature, are handcrafted by experts, and yet are not suitable for our IR agent.

The most fitting evaluation method, which we find in the literature, was part of the dynamic domain TREC 2017 conference [4]. At the conference, a session based IR agent is evaluated with a ground truth dataset, that contains a wide range of topics, handcrafted by experts. Each topic has a description and a set of subtopics. Subtopics contain sets of passages from documents part of a corpus, the IR agent has access to, each associated with a relevance score. The IR agent retrieves a description of a topic as an initial query and then has to return documents from the corpus, sorted by relevance in descending order, given the query. Documents returned by the agent are send to a human simulator that has access to the ground truth dataset. The simulator returns feedback, by comparing the ground truth dataset and the documents retrieved from the agent, whether a document is relevant or not and the IR agent decides when to stop the session. For comparing different agents, given the same ground truth dataset and corpus, the cube test (CT) is being used [5]. Given a ground truth dataset and the current iteration of a session, the human simulator sends the same feedback to the IR agent for documents the simulator has retrieved. If an IR agent approximates $\mathcal{M}^{\mathbf{H}}$ as $\widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}$, then it could improve greatly its performance. However, we argue that it is not useful, that the agent approximates $\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}$ as $\widetilde{\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}}$, as the human simulator always gives the same feedback in identical iterations, even if the complete history of two sessions differ.

Our aim is to extend the human simulator, part of the evaluation at the dynamic domain TREC 2017 conference, with a human agent, such that the IR agent can improve its performance by approximating $\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}$ as $\widetilde{\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}}$, in addition to approximating $\mathcal{M}^{\mathbf{H}}$ as $\widetilde{\mathcal{M}_{\mathbf{a}}^{\mathbf{H}}}$. The human simulator still has access to a ground truth dataset, however the dataset does only contain sets of relevant documents, each with respect to a session. Expressing queries is now the task of the human agent that does not know which documents are selected as relevant by the human simulator. From a subset of these as relevant selected documents, the human simulator generates a subjective representation in form of weighted words that simulate the information need of a human. Subjective as the human simulator generates the weighted words depending on the context of as relevant selected documents. The human agent does not know which documents are selected as being relevant by the human simulator and is not able to go through all documents in the corpus. However, the human agent can observe the weighted words that represent the information need. That is identical to a real human that is not able to compare all documents in a large corpus with a set of terms in his mind. From the set of weighted words, the human agent has to express its information need as a query. As it is for real humans the case, the human agent tries to anticipate of how the IR agent returns documents, given a query, as it only wants to retrieve documents, that are relevant. The human agent is allowed to reweight the weighted words before sending them as a query to the IR agent to retrieve a set of documents it could read. Only sending top weighted words to the IR agent, without adjusting them, is not a good strategy as the IR agent might not be able to always return relevant documents with respect to a subset of words that only subjectively represent the information need of a human. Therefore, the human agent knows that the IR agent might not be able to understand his information need and needs to find a good strategy for expressing a query. The IR agent needs to anticipate of how the human agent reweights words before sending them as a query, while it does not have access to the set of weighted words. However, it can observe the titles and contents of all documents in the corpus and the query, send by the human agent. We argue that the IR agent needs to be aware of that its actions influence the human agents actions. If the human agent reads a document, then the human simulator changes the weighted words, as the human agent has learned something new by reading a document. Collaboration between both agents only works if both aim to share and approximate the mental models of each other.

More concretely, the human agent does not compute $Q_{t+1} \times F_{t+1}$ from $\mathbf{H}_t$ and the IR agent R_t from $\mathbf{A}_t$ directly. Both $\mathbf{H}_t$ and $\mathbf{A}_t$ contain models and each model contains a function that maps, depending on the model, queries Q_t and feedback F_t or weighted words to queries Q_{t+1} and feedback F_{t+1} or weighted words, as depicted in Table 1. The human agent needs to express its information need as a query and to give optionally feedback for previously received results: $\mathbf{H}_t \to Q_{t+1} \times F_{t+1}$, with $\mathbf{H}_t = (\mathcal{M}^{\mathbf{H}}, \mathcal{M}_{\mathbf{h}}^{\mathbf{A}})$. Model $\mathcal{M}^{\mathbf{H}}$ contains a function

$$f^{\mathbf{H}\to\mathbf{A}} : \mathcal{M}_{\mathbf{h}}^{\mathbf{A}} \times W^{\mathbf{H}} \to Q_{t+1} \times F_{t+1}$$

that maps weighted words $W^{\mathbf{H}}$, generated by the human simulator to query Q_{t+1} and feedback

Table 1
Concrete Mental Models

Model	Function	From	To
$\mathcal{M}^{\mathbf{H}}$	$f^{\mathbf{H}\to\mathbf{A}}$	$\mathcal{M}^{\mathbf{A}}_{\mathbf{h}} \times W^{\mathbf{H}}$	$Q_{t+1} \times F_{t+1}$
$\mathcal{M}^{\mathbf{A}}_{\mathbf{h}}$	$f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}}$	$Q_t \times F_t$	$W^{\mathbf{A}}_{\mathbf{h}}$
$\widetilde{\mathcal{M}^{\mathbf{A}}}$	$\widetilde{f^{\mathbf{A}\to\mathbf{H}}}$	$\widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}} \times Q_t \times F_t$	$W^{\mathbf{A}}$
$\widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}$	$\widetilde{f^{\mathbf{H}\to\mathbf{A}}_{\mathbf{a}}}$	$\widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h'}}} \times W^{\mathbf{H}}_{\mathbf{a}}$	$Q_{t+1} \times F_{t+1}$
$\widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h'}}}$	$\widetilde{f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h'}}}$	$Q_t \times F_t$	$W^{\mathbf{A}}_{\mathbf{h'}}$

F_{t+1}, and $\mathcal{M}^{\mathbf{A}}_{\mathbf{h}}$ an approximated function

$$f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}} : Q_t \times F_t \to W^{\mathbf{A}}_{\mathbf{h}}$$

of the IR agent that maps Q_t and F_t to a set of weighted words $W^{\mathbf{A}}_{\mathbf{h}}$ possibly part of model $\widetilde{\mathcal{M}^{\mathbf{A}}}$ from the perspective of the human. An approximation of $f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}}$ is the humans understanding of how the IR agent interprets a query and feedback. In addition, it sends feedback F_t to the agent, containing which documents the agent has returned in the previous iteration are relevant. Thus, the human agent is aware of that the set of weighted words $W^{\mathbf{H}}$ in his mental model $\mathcal{M}^{\mathbf{H}}$ only subjectively represent the relevant documents and the IR agent is not necessarily able to return relevant documents from these words. In addition it knows, that the IR agent is aware of that and aims to reweight the words in Q_t with $f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}} : Q_t \times F_t \to W^{\mathbf{A}}_{\mathbf{h}}$ that objectively represent the information need of the human agent, from the perspective of the IR agent.

The IR agent needs to assign a score to all documents in the corpus and return top n of them in descending order as R_t: $\mathbf{A}_t \to R_t$ with $\mathbf{A}_t = (\widetilde{\mathcal{M}^{\mathbf{A}}}, \widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}, \widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h'}}})$. Model $\widetilde{\mathcal{M}^{\mathbf{A}}}$ contains the function

$$\widetilde{f^{\mathbf{A}\to\mathbf{H}}} : \widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}} \times Q_t \times F_t \to W^{\mathbf{A}}$$

where the set of words $W^{\mathbf{A}}$ objectively represent the information need of the human agent, $\widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}$ an approximation of $\mathcal{M}^{\mathbf{H}}$ containing the function

$$\widetilde{f^{\mathbf{H}\to\mathbf{A}}_{\mathbf{a}}} : \widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h'}}} \times W^{\mathbf{H}}_{\mathbf{a}} \to Q_{t+1} \times F_{t+1}$$

and $\widetilde{\mathcal{M}^{\mathbf{A}}_{\mathbf{h'}}}$ the function

$$\widetilde{f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h'}}} : Q_t \times F_t \to W^{\mathbf{A}}_{\mathbf{h'}}$$

The weighted words $W^{\mathbf{A}}$ objectively represent the information need of the human agent, from the perspective of the IR agent and it compares these with the documents in the corpus, by using for instance latent semantic indexing (LSI). Approximating $f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}} : Q_t \times F_t \to W^{\mathbf{A}}_{\mathbf{h}}$ as $\widetilde{f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h'}}} : Q_t \times F_t \to W^{\mathbf{A}}_{\mathbf{h'}}$ helps the agent to act according to the expectations of the human, by comparing $W^{\mathbf{A}}_{\mathbf{h'}}$ with $W^{\mathbf{A}}$. If the gab between $W^{\mathbf{A}}_{\mathbf{h'}}$ and $W^{\mathbf{A}}$ is too large, the IR agent could either adapt its behaviour to act according to the expectations of the human, even if it is a loss in its retrieval performance or if the gab is even larger, then it should explain its behaviour. Note that $W^{\mathbf{A}}_{\mathbf{h'}}$ and $W^{\mathbf{A}}$ are not necessarily identical.

Table 2
Dynamic decision network nodes

Variable	Description
$Relevant_t(D)$	the probability of a document d being relevant or not
$Weights_t(W)$	the probability of a word w being set as relevant by the human simulator at the current iteration t
$Read_t(D)$	the probability of a document d being read by the human agent at the current iteration t
$Query_t(W)$	the probability of a word w being part of the query at the current iteration t
$Relevance_t(D)$	the probability of a document d being relevant at the current iteration t from the perspective of the IR agent
$Title_t(D,W)$	the probability of a document d having word w in the title
$Content_t(D,W)$	the probability of a document d containing word w

Even if the human agent and the human simulator are not simulating a real human perfectly, we argue that it is a big step towards developing a human aware IR agent that collaboratively seeks for information together with a real human. The IR agent could be further improved in the future by evaluating it in the real world. If the IR agent performs bad in the real-world, then the human agent and human simulator need to be adapted accordingly. The effort of adapting the agents pays off, if the agents are realistic enough for the evaluation of an IR agent. Evaluating an IR agent, using a simulator is cheaper and orders of magnitudes faster, than evaluating it in the real-world.

4. Select Appropriate Actions

During seeking, the human agent has to decide which queries Q_{t+1} and feedback F_{t+1} it should send to the IR agent, while the agent has to decide which results R_t it should send to the human agent, such that the human agent is able to find every relevant document in the corpus. Finding every relevant document in the corpus is the long term goal of the human agent during a session S, however both agents need to be able to decide at each session state S_t, which action to perform. Decision making could be made possible by modelling the information seeking process as a dynamic decision network (DDN), as depicted in Figure 3. A DDN is a dynamic Bayesian network (DBN) extended with decision and utility nodes. Decision nodes are identical to chance nodes, except that the value of a decision node is determined by an action. The IR agent gets a high reward if the human gives positive feedback for a document it read and the human agent if it sends a query at where the IR agent returns documents that are relevant. The nodes of the DDN could be separated mainly into three parts, namely the (i) human simulator, (ii) human agent, (iii) human aware IR agent, and are listed in Table 2. During a session, the set of relevant documents $Relevant_t(D)$ does not change over time and is not observable by the human agent or IR agent. The human simulator generates a set of weighted words $Weights_t(W)$ from a subset of relevant documents d, which subjectively represent the information need $\mathcal{M}^{\mathbf{H}}$ of the human. If the human agent reads a document d with probability $Read_{t-1}(D)$, then the human simulator adjusts the weights $Weights_t(D)$. Adjusting the weights is similar to the

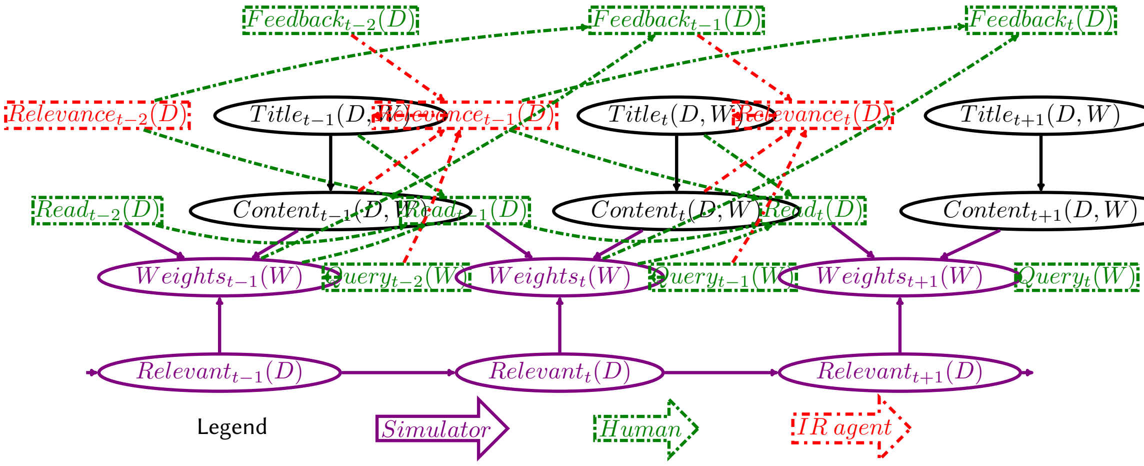

Figure 3: Collaborative seeking modelled with a dynamic decision network

behaviour of a real human, as if a human reads a document, then he learns something new. If a human learns something new, then his information need changes too. From the set of weighted words $Weights_t(W)$, the human generates a query $Query_{t-1}(W)$ that the IR agent can observe, using $f^{\mathbf{H}\to\mathbf{A}}$.

In addition to the query, the IR agent can observe the titles $Title_t(D,W)$ and $Content_t(D,W)$ of the documents d and compare them with $Weights_t(W)$ it approximates, using $\widetilde{f^{\mathbf{A}\to\mathbf{H}}}$ for computing whether a document d is relevant $Relevance_t(D)$ or not. The human agent approximates $W^{\mathbf{A}}$ as $W^{\mathbf{A}}_{\mathbf{h}}$, by approximating $\widetilde{f^{\mathbf{A}\to\mathbf{H}}}$ as $f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}}$ and the IR agent is aware of that, by approximating $f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h}}$ as $\widetilde{f^{\mathbf{A}\to\mathbf{H}}_{\mathbf{h'}}}$ for acting explicable from the perspective of the human agent. In addition, the IR agent approximates $f^{\mathbf{H}\to\mathbf{A}}$ as $\widetilde{f^{\mathbf{H}\to\mathbf{A}}_{\mathbf{a}}}$. Approximating $f^{\mathbf{H}\to\mathbf{A}}$ as $\widetilde{f^{\mathbf{H}\to\mathbf{A}}_{\mathbf{a}}}$ helps the IR agent to plan for future steps ahead. If the human agent reads a document, the human simulator reveals whether the document is relevant or not. The human agent forwards to the IR agent whether a document it read is relevant $Feedback_t(D) = true$ or not $Feedback_t(D) = false$.

5. Conclusion and Future Work

We model the process of a human and an IR agent that collaboratively search for information during a session, by sharing their mental models, as a DDN. The IR agent is modelled such that there is a difference between its mental model $\widetilde{\mathcal{M}^{\mathbf{A}}}$ and the approximation of the human mental model $\mathcal{M}^{\mathbf{H}}$ as $\widetilde{\mathcal{M}^{\mathbf{H}}_{\mathbf{a}}}$. In a first sight, this might be counterintuitive as the goal of the IR agent should be identical to that of the human. However, the information need of a human, expressed as a query Q_{t+1} is only a tiny representation of the subjective knowledge the human has about his information need, as different humans specify very different queries for very similar information needs. The agent is aware of that and should aim to translate the query Q_t into an objective representation $W^{\mathbf{A}}$ in $\widetilde{\mathcal{M}^{\mathbf{A}}}$ of the information need of the human it could compare with the documents in the corpus. Translating the information need of the human, expressed as Q_t, needs to be done by the agent with respect to the expectations of the human, as the human approximates that as $W_{\mathbf{h}}^{\mathbf{A}}$ in $\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}$. Therefore the agent approximates $\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}$ as $\widetilde{\mathcal{M}_{\mathbf{h}}^{\mathbf{A}}}$, for estimating the subjective information need $W^{\mathbf{H}}$ of the human from query Q_t. The human either expects that the IR agent is not doing that, as he is highly certain, that his query is objective or he is uncertain and expects that the IR agent translates the query Q_t into an objective one before comparing it with each document in the corpus. Regardless of the humans expectations, the IR agent needs to act explicable, such that the human can anticipate, what the IR agent decides to do with a query Q_t. Thus, we differentiate between $\widetilde{\mathcal{M}^{\mathbf{A}}}$ and $\widetilde{\mathcal{M}_{\mathbf{a}}^{\mathbf{H}}}$.

A human aware IR agent needs to be able to act intelligent, during a session, depending on the actions of a real human. Thus, we aim to implement an IR agent, by evaluating it with a human agent in a simulated world controlled by the human simulator, that acts similar to a real human. The human agent itself could be evaluated by evaluating the IR agent that is implemented, trained, and evaluated only with the human agent, in the real world. That suffices, as this is the only and primary goal of the implementation of a human agent, even if the human agent might not act exactly as a real human. For achieving a very performant IR agent, it might be beneficial to train the IR agent offline on that corpus, it later has to answer queries on, from real humans. And even further, the humans that later send queries to the IR agent could help to train it, as we refer to machine training instead of machine learning.

In the future, we aim to implement the human simulator and the human agent for the implementation and evaluation of a human aware IR agent. The generation of a set of weighted words, that subjectively represent the information need of a human, given a subset of relevant documents, could be achieved by generating questions as proposed by Klein et al. in [6] or by Chan et al. in [7].

References

[1] Z. Tang, G. H. Yang, Dynamic Search–Optimizing the Game of Information Seeking, ACM Transactions on Information Systems (2019).

[2] T. Chakraborti, S. Sreedharan, S. Kambhampati, Explicability versus Explanations in Human-Aware Planning, in: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, 2018, pp. 2180–2182.

[3] S. Kambhampati, Challenges of Human-Aware AI Systems, arXiv preprint arXiv:1910.07089 (2019).

[4] G. Hui Yang, I. Soboroff, J. Luo, Z. Tang, TREC 2017 Dynamic Domain Track Overview, 2017. Text REtrieval Conference (TREC).

[5] J. Luo, C. Wing, H. Yang, M. Hearst, The Water Filling Model and the Cube Test: Multi-Dimensional Evaluation for Professional Search, Association for Computing Machinery, 2013. doi:10.1145/2505515.2523648.

[6] T. Klein, M. Nabi, Learning to Answer by Learning to Ask: Getting the Best of GPT-2 and BERT Worlds, 2019. arXiv:1911.02365.

[7] Y.-H. Chan, Y.-C. Fan, A Recurrent BERT-based Model for Question Generation, in: Proceedings of the 2nd Workshop on Machine Reading for Question Answering, Association for Computational Linguistics, 2019. doi:10.18653/v1/D19-5821.