# Modeling Information Flow in Dynamic Information Retrieval


Felipe Moraes
Department of Computer Science
Universidade Federal de Minas Gerais
Belo Horizonte, MG, Brazil
felipemoraes@dcc.ufmg.br

Mário S. Alvim
Department of Computer Science
Universidade Federal de Minas Gerais
Belo Horizonte, MG, Brazil
msalvim@dcc.ufmg.br

Rodrygo L. T. Santos
Department of Computer Science
Universidade Federal de Minas Gerais
Belo Horizonte, MG, Brazil
rodrygo@dcc.ufmg.br



## ABSTRACT

User interaction with a dynamic information retrieval (DIR) system can be seen as a cooperative effort towards finding relevant information. In this cooperation, the user provides the system with evidence of his or her information need (e.g., in the form of queries, query reformulations, relevance judgments, or clicks). In turn, the system provides the user with evidence of the available information (e.g., in the form of a set of candidate results). Throughout this conversational process, both user and system may reduce their uncertainty with respect to each other, which may ultimately help in finding the desired information. In this paper, we present an information-theoretic model to quantify the flow of information from users to DIR systems and vice versa. By employing channels with memory and feedback, we decouple the mutual information among the behavior of the user and that of the system into directed components. As a result, we are able to measure: (i) by how much the DIR system is capable of adapting to the user; and (ii) by how much the user is influenced by the results returned by the DIR system. We discuss implications of the proposed framework for the evaluation and optimization of DIR systems.



## CCS CONCEPTS

•**Information systems** →**Users and interactive retrieval;**
•**Mathematics of computing** →*Information theory;*

## KEYWORDS

Directed Information Flow; Search Effectiveness; Formal Methods


## 1 INTRODUCTION

Nowadays people spend a considerable part of their daily lives interacting with information retrieval systems across multiple devices, from web search engines to personal assistants running on a variety of smart devices. To better support information seeking, modern information retrieval systems must dynamically adapt their retrieval strategy as they interact with the user. In particular, given a user's query, a *dynamic information retrieval* (DIR) system [17, 18] may respond with an initial result set that best matches its understanding of the user's information need. This set may provide the user with an improved understanding of the information available to the system, prompting him or her to respond back (e.g., by performing a query reformulation, a relevance judgment, or a click). The system may then adapt to the user's response and retrieve a further (ideally improved) result set. This conversational process may continue until the desired information is retrieved to the user or until the user decides to abandon the interaction.

The effectiveness of a DIR system is a function of the gain it provides (in terms of the relevance of the retrieved results) and the effort it incurs (in terms of the total interaction time) to the user [13]. Quantifying such effectiveness is challenging, primarily because the behavior of both the user and the system is inherently uncertain. On the one hand, user actions cannot be fully determined by the system outputs alone, as they also depend on non-observable variables, including the user's cultural background and state of mind. On the other hand, system outputs may not be fully determined by the user actions alone, as the system may benefit from further exploration of the output space to improve its knowledge about the user's need [9]. User modeling under uncertainty has been recently investigated in the context of click models for web search [5], which are probabilistic graphical models aimed at describing and predicting user behavior when interacting with a search engine. In turn, probabilistic retrieval systems have been mostly investigated in the context of the *multi-armed bandit problem* in reinforcement learning [15], in which the system faces a trade-off between exploring new knowledge and exploiting existing knowledge [8].

To provide a unified analysis of the mutual interaction between a user and a DIR system, we propose an information-theoretic model to quantify the amount of information flowing between both agents. Building upon recent advances in information theory [16], we quantify not only the correlation between user and system behavior, but also how much they influence each other. In particular, a high flow from user to system may indicate a highly *sensitive* system, i.e., one that could effectively learn about the user's information need and adapt accordingly. In turn, a high flow from system to user may indicate a highly *influential* system, i.e., one that could effectively teach the user about the available information. While system influence may be interpreted as a measure of retrieval effectiveness—as the user behavior is highly influenced by the system's responses—, system sensitivity could measure lack of privacy—since the more tailored the system is to a particular user, the more information about this user is revealed to an observer of the system's behavior.

In the remainder of this paper, after a brief review of crucial concepts from information theory, we formalize the proposed theoretical model and discuss the challenges involved in instantiating it in a practical scenario as well as implications for evaluating and optimizing DIR systems.



*ICTIR'17, October 1–4, 2017, Amsterdam, The Netherlands.*



DOI: 10.1145/3121050.3121102

## 2 PRELIMINARIES

In this section we briefly review key concepts from information theory; for more details we refer to [6]. Let $\mathcal{A} = \{a_1, \ldots, a_n\}$ and $\mathcal{B} = \{b_1, \ldots, b_m\}$ be two finite sets with associated random variables $A$ and $B$, and following probability distributions $p_A$ and $p_B$, respectively. (We shall omit the subscripts from distributions when they are clear from the context.)

The *entropy* of $A$ is defined as $H(A) = -\sum_{a \in \mathcal{A}} p(a) \log p(a)$, and it measures the uncertainty of $A$: the higher its value, the less information one has about the value of $A$. It takes its minimum value $H(A) = 0$ when $p_A$ is a point-mass distribution, and maximum value $H(A) = \log |\mathcal{A}|$ when $p_A$ is the uniform distribution. Usually the base of the logarithm is set to be 2 and the entropy is measured in *bits*. The *conditional entropy* of $A$ given $B$ is defined as $H(A|B) = -\sum_{a \in \mathcal{A}, b \in \mathcal{B}} p(a,b) \log p(a \mid b)$, and it measures the uncertainty of $A$ when $B$ is known. It takes minimum value $H(A|B) = 0$ when $A$ is completely determined by $B$, and maximum value $H(A|B) = H(A)$ when $A$ and $B$ are independent. The *mutual information* between $A$ and $B$ is defined as $I(A;B) = H(A) - H(A|B)$, and it measures the amount of information about $A$ that one gains by observing $B$. It can be shown that mutual information is symmetric, $I(A;B) = I(B;A)$, and that $0 \le I(A;B) \le H(A)$.

A (*discrete memoryless*) *channel* $C$ is a tuple $(\mathcal{A}, \mathcal{B}, p_{A|B})$, where $\mathcal{A}$ and $\mathcal{B}$ are sets of *input* and *output* symbols, respectively, and $p_{B|A}(b_j \mid a_i)$ is the probability of the channel producing output symbol $b_j$ when the input symbol is $a_i$. An input distribution $p_A$ over $\mathcal{A}$ determines, together with the channel $C$, the joint distribution $p_{AB}(a_i, b_j) = p_A(a_i) p_{B|A}(b_j \mid a_i)$, and consequently $I(A;B)$. The maximum $I(A;B)$ over all possible input distributions is the channel's *capacity*, and it represents the maximum rate by which information can be transmitted using the channel.

## 3 INFORMATION-THEORETIC MODEL OF DIR SYSTEMS

In this section we present an information-theoretic model for reasoning about the flow of information between a user and a DIR system, when both agents can behave probabilistically. We start by introducing some notation. When we have a *sequence* of symbols (ordered in time) from a set $\mathcal{A}$, we use a Greek letter $\alpha_t$ to denote the symbol at time $t$, and $\alpha^t$ to denote the sub-sequence $\alpha_1 \alpha_2 \ldots \alpha_t$. We use $A^t$ to denote the sequence of $t$ consecutive occurrences $A_1, \ldots, A_t$ of the random variable $A$. Moreover, we use $\mathcal{A}^*$ to denote the Cartesian product of set $\mathcal{A}$ with itself a non-negative number of times (so the elements of $\mathcal{A}^*$ are all finite strings whose symbols are elements of $\mathcal{A}$), and $\mathbb{D}\mathcal{A}$ to denote the set of all probability distributions on the set $\mathcal{A}$.

### 3.1 Formalization of probabilistic DIR systems

Let $\mathcal{A} = \{a_1, \ldots, a_n\}$ be the set of all possible *user's actions*, where each $a_i \in \mathcal{A}$ represents a possible input the user can provide to the DIR system (e.g., the posing of a query, or the relevance judgment of some retrieved result). Let $\mathcal{D} = \{d_1, \ldots, d_m\}$ be the set of all possible *result sets* the DIR system can return to the user (e.g., a list of documents). The user and the DIR system can interact, and a *session* of length $T \ge 1$ of the interactive process is an alternating sequence $\alpha_1 \ \delta_1 \ \alpha_2 \ \delta_2 \ \ldots \ \alpha_T \ \delta_T$ in which each $\alpha_t$ and $\delta_t$ represent the action taken by the user and the result set returned by the system, respectively, at time step $1 \le t \le T$. A *(session) history* up to time $t$ is the sub-sequence $\alpha_1 \delta_1 \ \ldots \ \alpha_t \ \delta_t$ of a session, and we may denote it simply by $\alpha^t, \delta^t$, where $\alpha^t$ is the *action history* (i.e., the projection of $h$ onto $\mathcal{A}^*$), and $\delta^t$ is the corresponding *result-set history* (i.e., the projection of $h$ onto $\mathcal{D}^*$).

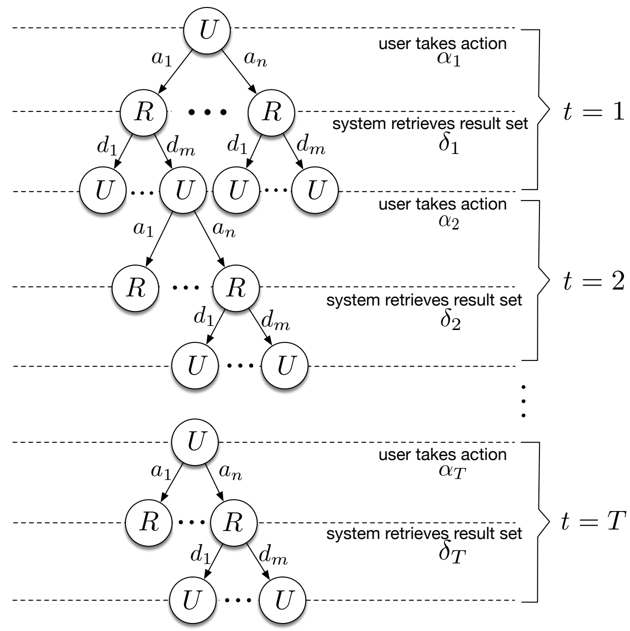


**Figure 1: A tree depicting all possible sessions between user $U$ and DIR system $R$. Each edge would also be labeled with a probability, which is omitted from the figure for simplicity.**

A *DIR system* is modeled as a probabilistic function $R : \mathcal{A}^* \times \mathcal{D}^* \to \mathbb{D}\mathcal{D}$ mapping each session history to a distribution on result sets. The function $R$ can be represented as a family $\{p(\delta_t \mid \alpha^t, \delta^{t-1})\}_{t=1}^{T}$ of conditional probability distributions, where, for each $1 \le t \le T$, $p(\delta_t \mid \alpha^t, \delta^{t-1})$ is the probability of $R$ retrieving result-set $\delta_t$, at time $t$, given that the session history so far is $\alpha^t, \delta^{t-1}$.

A *user* is modeled as a probabilistic function $U : \mathcal{A}^* \times \mathcal{D}^* \to \mathbb{D}\mathcal{A}$ mapping each history to a distribution on user actions. The function $U$ can be represented as a family $\{p(\alpha_t \mid \alpha^{t-1}, \delta^{t-1})\}_{t=1}^{T}$ of conditional probability distributions, where each $p(\alpha_t \mid \alpha^{t-1}, \delta^{t-1})$ represents the probability of the user taking action $\alpha_t$, given the history $\alpha^{t-1}, \delta^{t-1}$ so far.

All possible sessions of an interactive process between DIR system $R$ and user $U$ can be represented as a probabilistic automaton like the one in Figure 1. In such an automaton: (i) there is a node for each point in time an agent (user or DIR system) can make a move; (ii) from each node there is an emanating edge labeled with a possible move available for the agent, together with the probability of the agent choosing that move; (iii) every session (resp., history) is represented by a path from the root to a leaf (resp., some node), and the probability of the session (resp., history) can be computed by multiplying the probabilities of the edges in a path from the root to the corresponding node.

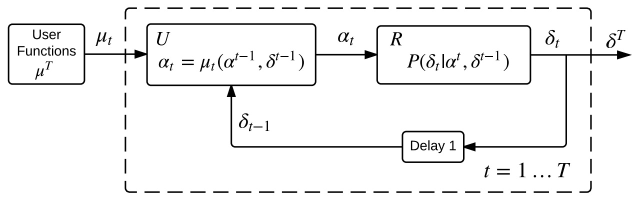


**Figure 2: Model of interaction process between user $U$ and DIR system $R$, using a channel with memory and feedback.**

## 3.2 DIR systems as channels w/ memory and feedback

In order to quantify the information flow from $U$ to $R$, and vice versa, in a session of the interactive process, we map the probabilistic automaton representing an interactive process to the model of discrete channels with memory and feedback proposed by Tatikonda and Mitter [16], and adapted to interactive systems by Alvim et al. [1].

The model, represented in Figure 2, can be decomposed into components as follows. The behavior $p(\delta_t \mid \alpha^t, \delta^{t-1})$ of the DIR system is modeled as the internal channel $R$ taking user actions as inputs and producing result sets as outputs. The result set produced is then fed back to the user with delay one (i.e., a result set $\delta_t$ produced at time $t$ will only become available to the user at time $t+1$). Since $R$ can be used repeatedly, its behavior after $T$ uses can be fully described by the joint probability distribution relating the sequences $\alpha^T$ of inputs and $\delta^T$ of outputs as follows:

$$p(\alpha^T, \delta^T) = \prod_{t=1}^{T} p(\alpha_t \mid \alpha^{t-1}, \delta^{t-1}) p(\delta_t \mid \alpha^t, \delta^{t-1}). \tag{1}$$

The first term $p(\alpha_t \mid \alpha^{t-1}, \delta^{t-1})$ indicates that the channel supports *feedback*, i.e., the occurrence of a current input $\alpha_t$ may depend on past outputs $\delta^{t-1}$ from the channel.[1] The second term $p(\delta_t \mid \alpha^t, \delta^{t-1})$ indicates that the channel may present *memory*, in the sense that the production of an output $\delta_t$ may not depend only on the current input $\alpha_t$, but on the whole previous history $\alpha^t, \delta^{t-1}$.

Computing the capacity of channels with memory and feedback is a significantly more complex task than it is for simple information-theoretic channels (i.e., memoryless, and without feedback). Tatikonda and Mitter's method does it by representing the channel with memory and feedback from $A^T$ to $D^T$ (the inner channel $R$ in Figure 2) as an equivalent channel without memory or feedback from user-behavior functions $U^T$ to $D^T$ (represented as everything inside the dotted-line in Figure 2). Because this outer channel is a simple channel, its capacity can be computed using usual methods. However, to do so, their model limits the production of inputs to the channel to deterministic functions. This limitation means that the term $p(\alpha_t \mid \alpha^{t-1}, \delta^{t-1})$ in Equation (1), which corresponds to the user's behavior in the interactive process, cannot be directly captured in the model. To circumvent this problem, here we follow Alvim et al.'s approach [1] and externalize the probabilistic behavior of the user to a distribution on deterministic *user-behavior functions* $\mu_t$.[2]

## 3.3 Information flow in DIR systems

To quantify the flow of information in a DIR system, we employ the concepts of directed information flow and capacity in channels with memory and feedback.

*3.3.1 A brief review of directed information.* In the presence of feedback, mutual information $I(A^T; D^T)$ does not represent the information flow from $A^T$ to $D^T$. Intuitively, this is due to the fact that mutual information expresses correlation, and therefore it is increased by feedback. The appropriate measure of transmission of information in such a channel is *directed information* [14].

*Definition 3.1.* In a channel with feedback, the *directed information from input $A^T$ to output $D^T$* is defined as

$$I(A^T \to D^T) = \sum_{t=1}^{T} I(A^t; D_t \mid D^{t-1});$$

and the *directed information from $D^T$ to $A^T$* is defined as

$$I(D^T \to A^T) = \sum_{t=1}^{T} I(A_t; D^{t-1} \mid A^{t-1}).$$

Note that directed information is not symmetric: the flow from $A^T$ to $D^T$ considers the correlation between $A^t$ and $D_t$, whereas the flow from $D^T$ to $A^T$ considers the correlation between $D^{t-1}$ and $A_t$. Intuitively, this is because $\alpha^t$ influences $\delta_t$, but, in the other direction, it is $\delta^{t-1}$ that influences $\alpha_t$. It can be proven [16] that

$$I(A^T; D^T) = I(A^T \to D^T) + I(D^T \to A^T),$$

i.e., the traditional concept of mutual information is the sum of the directed information in both directions in a channel with memory and feedback. In particular, if a channel does not have feedback, then $I(D^T \to A^T) = 0$ and $I(A^T; D^T) = I(A^T \to D^T)$, which recovers the traditional definition of information flow in such channels. In a channel with feedback, however, the correct measure of information transmitted is directed information $I(A^T \to D^T)$, and not mutual information.

The concept of capacity is generalized for channels with feedback as follows. Let $\mathbb{D}U^T = \{\{p(\alpha_t \mid \alpha^{t-1}, \delta^{t-1})\}_{t=1}^T\}$ be the set of all input distributions. For finite $T$, the *capacity* of a channel $\{p(\delta_t \mid \alpha^t, \delta^{t-1})\}_{t=1}^T$ is

$$C_T = \sup_{\mathbb{D}U^T} \frac{1}{T} I(A^T \to D^T).$$

[1] In this paper we use the term "feedback" only to denote its usual meaning in information theory, as opposed to its usual meaning in retrieval systems.

[2] To formalize the transformation, let us call $\mathcal{U}_t$ the set of all measurable maps $\mu_t : \mathcal{D}^{t-1} \to \mathcal{A}$ endowed with a probability distribution, and let $U_t$ be the corresponding random variable. Let $\mathcal{U}^T$, $U^T$ denote the Cartesian product on the domain and the random variable, respectively. A *user-behavior function* is an element $\mu^T = (\mu_1, \ldots, \mu_T) \in \mathcal{U}^T$. Note that, by probability laws, $p(\mu^T) = \prod_{t=1}^T p(\mu_t \mid \mu^{t-1})$. Hence the distribution on $\mathcal{U}^T$ is uniquely determined by a sequence $\{p(\mu_t \mid \mu^{t-1})\}_{t=1}^T$. We will use the notation $\mu_t(\delta^{t-1})$ to represent the $\mathcal{A}$-valued $t$-tuple $(\mu_1, \mu_2(\delta^1), \ldots, \mu_t(\delta^{t-1}))$. We, then, formally capture a user in our model as a family $\{\mu_t \mid \mu^{t-1}\}_{t=1}^T$ of distributions on user-behavior functions defined as $p(\mu_1) = p(\alpha_1 \mid \alpha^0, \delta^0) = p(\alpha_1)$, and $p(\mu_t \mid \mu^{t-1}) = \prod_{\delta^{t-1}} p(\mu_t \mid \mu^{t-1}(\delta^{t-2}), \delta^{t-1})$ for $2 \le t \le T$.

*3.3.2 Interpretation in terms of DIR systems.* Taking advantage of the disentanglement of the information flow from user to DIR system from the information flow in the opposite direction, we can provide a refined analysis of an interactive process as follows.

$I(A^T \to D^T)$ can be used as a measure of how *sensitive* to the user the DIR system is, as it quantifies how much the DIR system's outputs are tailored to the user's actions. The higher the value of $I(A^T \to D^T)$, the more the result sets retrieved by the DIR system will differ from user to user. This has an interesting application at enforcing privacy in the user interaction. For instance, the DIR system could be parameterized to restrict itself to a certain (perhaps user-defined) level of sensitivity, so that no user information beyond the target level can flow to the system. Moreover, since $I(A^T \to D^T) = \sum_{t=1}^{T} H(A_t|A^{t-1}, D^{t-1}) - H(A^T|D^T)$, the term $\sum_{t=1}^{T} H(A_t|A^{t-1}, D^{t-1})$ can be interpreted as the *entropy of the user* and used to measure how unpredictable the user's behavior is.

On the other hand, $I(D^T \to A^T)$ is a measure of how *influential* the retrieval system is on the user. The higher its value, the more the user's actions differ from what they would be in case the result sets were just being randomly produced. This can be interpreted as a measure of the retrieval effectiveness of the system as interaction progresses, providing an alternative to current DIR evaluation efforts. Moreover, it could serve as a criterion for optimizing DIR systems, perhaps constrained to a particular sensitivity level, as previously discussed.

Finally, since the capacity $C_T$ is the maximum information flow over all possible distributions on user-behavior functions, it is an accurate theoretical limit for the retrieval system's ability to produce results tailored to users.

## 4 CONCLUSIONS

A dynamic information retrieval (DIR) system aims to identify relevant information through a series of interactions with the user. An effective DIR system should be able to satisfy the user's information need with as few interactions as possible, by dynamically adapting itself in response to user actions. In this paper, we proposed an information-theoretic model for quantifying the flow of information between a user and a DIR system throughout the interactive process. Our proposed model allows for decoupling the mutual information between the user and the system into separate flows of information from the user to the system and from the system to the user. We discussed a possible interpretation for the information flow in each direction and pointed out potential applications for the evaluation and optimization of DIR systems.

An exact computation of information flows is intractable, as it would require computing the probability of observing a system input or output given all possible interaction histories.[3] Hence, an important direction for future investigation involves methods for approximate estimation of information flows. Promising candidates here include several techniques developed for the analysis of quantitative information flow in the context of security, which involve a mixture of model-checking techniques and/or statistical analysis, such as [4, 7, 11]. Another practical limitation of the proposed model is the unsigned nature of the computed information flows. In fact, directed information (as Shannon's mutual information) is agnostic to the meaning of the information flowing: it can measure by how much one agent's state is affected about the other agent's state (e.g., user and DIR system), but it cannot reflect whether the information flowing is helpful or unhelpful. This limitation restricts the interpretability of system sensitivity and influence. Future extensions should consider advanced information-theoretic metrics, which can take into consideration the meaning of information, such as the $g$-leakage framework [3] and its generalized capacities [2]. Lastly, in addition to the potential applications discussed in Section 3.3, we plan to investigate the applicability of the proposed model in scenarios such as statistical translation models [10] as well as conversational machines using neural networks [12]. In particular, the latter achieved state-of-the-art performance by using a mutual information-based optimization criterion. As future work, we plan to extend our measure of influence in this scenario, where a dual optimization criterion could improve the precision of the quality of messages produced by a dialogue system.

[3]It can be shown that the computational complexity is $O(|\mathcal{A}|^T |\mathcal{D}|^T \times k)$, where $|\mathcal{A}|$ is the number of user actions, $|\mathcal{D}|$ is the number of documents, $k$ is a constant to produce $|\mathcal{D}|$, and $T$ is the number of interactions.

## ACKNOWLEDGMENTS

This work was partially funded by projects InWeb (MCT/ CNPq 573871/2008-6) and MASWeb (FAPEMIG/PRONEX APQ-01400-14), and by the authors' grants from CNPq, CAPES and FAPEMIG.

## REFERENCES

[1] Mário S. Alvim, Miguel E. Andres, and Catuscia Palamidessi. 2012. Quantitative Information Flow in Interactive Systems. *JCS* (2012).
[2] Mário S. Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. 2014. Additive and Multiplicative Notions of Leakage, and Their Capacities. In *Proc. of CSF.*
[3] Mário S. Alvim, Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Geoffrey Smith. 2012. Measuring Information Leakage Using Generalized Gain Functions. In *Proc. of CSF.*
[4] Rohit Chadha, Umang Mathur, and Stefan Schwoon. 2014. Computing Information Flow Using Symbolic Model-Checking. In *Proc. of FSTTCS.*
[5] Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. 2015. *Click Models for Web Search.* Morgan & Claypool.
[6] Thomas M. Cover and Joy A. Thomas. 1991. *Elements of Information Theory.* J. Wiley & Sons, Inc.
[7] Deniz Gencaga, Kevin H. Knuth, and William B. Rossow. 2015. A Recipe for the Estimation of Information Flow in a Dynamical System. *Entropy* (2015).
[8] Artem Grotov and Maarten de Rijke. 2016. Online learning to rank for information retrieval: SIGIR 2016 Tutorial. In *Proc. of SIGIR.* 1215–1218.
[9] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. 1996. Reinforcement Learning: A Survey. *JAIR* (1996).
[10] Maryam Karimzadehgan and ChengXiang Zhai. 2010. Estimation of Statistical Translation Models Based on Mutual Information for Ad Hoc Information Retrieval. In *Proc. of SIGIR.*
[11] Yusuke Kawamoto, Fabrizio Biondi, and Axel Legay. 2016. *Hybrid Statistical Estimation of Mutual Information for Quantifying Information Flow.*
[12] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A Diversity-Promoting Objective Function for Neural Conversation Models. In *NAACL.*
[13] Jiyun Luo, Christopher Wing, Hui Yang, and Marti Hearst. 2013. The Water Filling Model and the Cube Test: Multi-dimensional Evaluation for Professional Search. In *Proc. of CIKM.*
[14] James L. Massey. 1990. Causality, Feedback and Directed Information. In *Proc. of SITA.*
[15] Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction.* MIT press Cambridge.
[16] Sekhar Tatikonda and Sanjoy K. Mitter. 2009. The Capacity of Channels With Feedback. *IEEE Trans.Inf.Theory* (2009).
[17] Ryen W. White. 2016. *Interactions with Search Systems.* Cambridge University Press.
[18] Grace Hui Yang, Marc Sloan, and Jun Wang. 2016. *Dynamic Information Retrieval Modeling.* Morgan & Claypool Publishers.