

LEVERAGING SEMANTIC WEB SEARCH AND BROWSE SESSIONS FOR MULTI-TURN SPOKEN DIALOG SYSTEMS

Lu Wang*, Larry Heck†, Dilek Hakkani-Tür†

*Cornell University, Department of Computer Science, Ithaca, NY, USA

†Microsoft Research, Mountain View, CA, USA

ABSTRACT

Training statistical dialog models in spoken dialog systems (SDS) requires large amounts of annotated data. The lack of scalable methods for data mining and annotation poses a significant hurdle for state-of-the-art statistical dialog managers. This paper presents an approach that directly leverage billions of web search and browse sessions to overcome this hurdle. The key insight is that task completion through web search and browse sessions is (a) predictable and (b) generalizes to spoken dialog task completion. The new method automatically mines behavioral search and browse patterns from web logs and translates them into spoken dialog models. We experiment with naturally occurring spoken dialogs and large scale web logs. Our session-based models outperform the state-of-the-art method for entity extraction task in SDS. We also achieve better performance for both entity and relation extraction on web search queries when compared with nontrivial baselines.

Index Terms— spoken dialog systems, statistical dialog management, multi-turn contextual model, dialog session, web search session, web browsing, entity extraction, relation extraction

1. INTRODUCTION

Statistical machine learning approaches for spoken dialog systems (SDS) have become popular in the last decade. However, an important problem that limits their practical application is the significant amount of annotated data that is required for training models in SDS [1]. This is particularly true for the dialog manager (DM) component, where the number of possible dialog state sequences is extremely high and difficult to enumerate. Hence scaling SDS approaches to handle even moderately complex dialog models is a key research problem.

A simple approach to solve this problem relies on a flat initialization; as the system is used, training data is obtained from real users and annotated to train better dialog models. However, such a bootstrapped statistical model is not desirable for real-world SDS as the initial user experience is poor and limited. In addition, the subsequent learning is biased towards simplified interactions, since these are the only dialogs that yield success for the user.

A more common solution is to exploit a small set of real user interactions to perform dialog simulations. With this approach, one can obtain a large set of artificial dialogs [2, 3, 4, 5]. Though simple to implement, the generated dialogs are typically not natural, and do not provide the opportunity for the system to learn the best ways to complete tasks from users. It may result in an SDS that leads a user through inefficient and occasionally non-sensical paths, causing user frustration and ultimately poor adoption rates [6].

Leveraging web search log to SDS. Nonetheless, large volumes of task-oriented and potentially relevant training data exist in web

— S₁ : Web Search and Browse Session

T₁ (Q): [the great gatsby]_{film} 2013
 T₂ (C): [www.imdb.com/title/tt1343092/] _{film}
 T₃ (Q): movies by [leonardo dicaprio]_{actor}
 T₄ (C): [www.imdb.com/name/nm0000138/] _{actor}
 T₅ (Q): [titanic]_{film}
 T₆ (C): [www.imdb.com/title/tt0120338/] _{film}
 T₇ (C): [http://www.imdb.com/character/ch0002338/] _{character}
 T₈ (Q): does [leonardo dicaprio]_{actor} have an award

— S₂ : Spoken Dialog System Session

U₁: go to bing dot com
 U₂: search for [the great gatsby]_{film}
 U₃: click that ([www.imdb.com/title/tt1343092/] _{film})
 U₄: show me movies by [leonardo dicaprio]_{actor}
 U₅: scroll down
 U₆: show me information on [titanic]_{film}
 U₇: go back
 U₈: show me information on [inception]_{film}

Fig. 1. Example fragments of a web search and browse session (S₁) and a spoken dialog system session (S₂). T_i/U_i represents the *i*th turn/utterance in web search/dialog. For the web session, Q or C in the parentheses indicates if the current turn is a query or click. For the dialog session, we also include the URL the user refers to (e.g. U₃). Each entity in the query or utterance is shown within brackets along with its entity type in subscript.

search engines. Users complete tasks everyday on the web, using a combination of search and browsing. The volume of the search data easily exceeds 100M queries per day over hundreds of millions of users. The breadth of tasks is also immense, ranging from simple (weather, finding directions, local events) to more complex (shopping, planning a trip, planning a night out).

In many ways, search and browsing have elements of automated conversational interactions. Figure 1 displays one fragment of web search session, and one fragment of an spoken dialog (conversational) session. We can see that the two sessions share similar browsing patterns (in blue) in terms of the entity types visited. Web search conversations are also interactive because the system responds to what has previously been communicated. The conversations are spontaneous as the user is not constrained by domain.

The advantage of using web data is the massive volume of interactions, with both depth in the number and complexity of user interactions and breadth in the variety of user goals. Some of the user's actions (e.g., clicks) further provide enough constraints so that the labels (domains, intents, slots) can be *inferred*, enabling weakly supervised training. The observation thus inspires us to leverage the large volumes of logged user interactions that exist in centrally hosted web search engines and browsers to the SDS systems.

In this work, we take a first step to effectively exploit this source of training data for an SDS problem – entity extraction from the spoken dialogs.

Session-based model for semantic parsing in SDS. Existing work

for leveraging web search behavior to SDS mainly exploits offline information in the query-click logs. For instance, the web query-click graph has been employed for domain detection [7], slot filling [8], and intent detection [9]. Recently, the idea of combining the power of internet browser and multi-modal input (e.g. speech, gesture, etc) is discussed in [10] and implemented by [11]. Different from all of the work above, our session models are *directly trained* on the web search sessions with contextual information from the search logs. We then test the models on entity extraction problem for SDS with the acquired behavioral patterns from web and multi-turn conversational context. Authors in [12] also try to leverage contextual information for the same task, however, they do not exploit web data.

Distant supervision for query interpretation. In addition, we investigate a domain-independent framework that utilizes a structured semantic knowledge base of large scale to provide distant supervision. Modern systems of entity and relation extraction from web search queries are mostly based on supervised learning with small hand-labeled corpora, and are thus limited by the availability of training data [13]. Weakly-supervised approaches are also studied [14, 15], where seed instances are utilized for bootstrap learning. Noisy extractions, with semantic shift, are hard to avoid.

With the emergence of large scale knowledge resources (e.g. Freebase, DBPedia), *distant supervision* has been proposed for relation extraction on formal text, such as Wikipedia articles [16], and news articles [17]. It is also investigated in [18, 19] for utilizing web search results to improve the relation extraction performance in SDS. In this work, we employ distant supervision to generate training samples based on web search queries instead of web pages returned by search engine. We also test *entity extraction* and *relation extraction* tasks on both spoken utterances and web search queries.

Contributions. To the best of our knowledge, we are the first to exploit web search sessions and their contextual information and adapt it to train statistical dialog managers. Our experiments conducted in the film domain show that our session-based models trained on web search log are capable to leverage multi-turn context to produce significantly better entity extraction performance than the state-of-the-art approach [20] in SDS.

Furthermore, to address the issue of lack of training data, we propose a distantly supervised approach to semantic parsing, by employing massive unlabeled web search logs combined with large dictionaries of entities and relations gathered from knowledge bases. We train linear-chain Conditional Random Fields [21] and its variant learned on data with missing labels [22] to model web browsing behavior on session level. Our models achieve better results than nontrivial baselines on entity extraction and relation extraction for both spoken dialogs and web search query. Notice that with knowledge base as supervision, entities extracted by our models are beyond named entity slots used in conventional SDS. We are also able to find unnamed categories, such as film subject and descriptions.

2. THE QUERY INTERPRETATION MODEL

Investigations from [13] discover that three templates (or types) of queries dominate entity-based web search (i.e. more than 65% of the sampled queries fall in these three types). The three templates can be formalized as: Entity (E), which contains one entity; Type \sqcap Relation(Entity) (T-R-E), which describes one entity and another entity type with a certain relation; Entity₀ \sqcap Relation(Entity₁) (E-R-E), which consists of a pair of related entities. In spoken dialog corpora studied in this work, we also observe more than 50% of utterances can be matched to those templates, where 54% of them belong to template E. For instance, in

Figure 1, query “search for the great gatsby” (U_2) falls into template E as it has a single film entity. Query “show me movies by leonardo dicaprio” (U_4) matches template T-R-E where “movies” indicates the film type and “leonardo dicaprio” is an actor entity. A sample query for E-R-E could be “Titanic by James Cameron”.

For the remaining part of the section, we describe our full query interpretation pipeline, consisting of *Entity Mention Identification*, *Entity Type Determination*, and *Relation Extraction*.

2.1. Entity Mention Identification

We first employ a linear-chain Conditional Random Fields (CRF) [21] to identify all the entity mentions in each query. We represent each query as a sequence of tokens, $\mathbf{x} = x_0x_1 \dots x_n$, and generate a sequence of labels, $\mathbf{y} = y_0y_1 \dots y_n$, that encode which tokens are part of entity mentions by using a BIOES-style label format: {B-ENT (beginning of an entity), I-ENT (inside of an entity), O (outside of any entity)}. As defined in [21], the conditional probability is given by:

$$P(\mathbf{y}|\mathbf{x}) = \frac{1}{Z(\mathbf{x})} \exp(\sum_{t=1}^T \sum_{k=1}^K \lambda_k f_k(y_{t-1}, y_t, \mathbf{x}, t))$$

where T is the length of the sequence, K is the number of feature functions f_k , and λ_k are the corresponding weights.

Feature Design. We utilize lexical features such as word in lowercase, is capitalized, has any number, and has any non-alphabetic character. For syntactic features, POS tag, NER tag, and dependency relation output by Stanford parser [23] are used. We also consider context features within the session. Figure 1 shows that “leonardo dicaprio” is mentioned twice in T_3 and T_8 . Intuitively, if an ngram appears repeatedly in the queries, it is likely to be an entity mention. Thus, we use repeat time to indicate if an ngram is repeated more than twice in recent 5 issued queries within the session.

In addition, we collect *gazetteers* from Freebase (<http://www.freebase.com/>), and use a feature to indicate if the token is part of a term in the gazetteer. Gazetteers, which are lists of entity names or relation instances, are important features for SDS to understand users’ intent or represent dialog states. We then match each entity name in the gazetteer to the query. If there is any overlapping between two mapped names, e.g. “leonardo dicaprio” v.s. “dicaprio”, the longer one is retained.

Suppose t represents the current token position and $feature_k[t]$ is the value of $feature_k$, we extract unigram, bigram and trigram features over local context, e.g. $feature_k[t-2]$, $feature_k[t-2] : feature_k[t-1]$, $feature_k[t-2] : feature_k[t-1] : feature_k[t]$.

2.2. Entity Type Determination




Fig. 2. Example of linear-chain CRF for entity type determination.

Entity type is determined after identifying entity mentions, e.g. whether the entity mention is an “actor” or a “film”. Instead of making prediction on each entity independent of the browsing history, we propose to use sequential labeling technique to predict entity types within a session simultaneously.

Suppose session S consists of turns $\{s_0, s_1, \dots, s_m\}$, where each turn s_i can be either a query or a web page click. Each entity mentioned in S represents a single state in our session-based entity type determination model. A click is mapped to an entity, if this

URL is a source URL for that entity in Freebase. Figure 2 depicts a mapping from the web session in Figure 1 to our sequential model.

CRF Training with Missing Labels. Linear-chain CRF as defined in Section 2.1 is utilized as the building block to predict entity types. Multiple entities in one query are considered as separate states. During training, it is possible that some turns do not contain any entity, which makes the labels partially observed for the session. We propose two ways to address this problem: 1) simply removing those turns in training; 2) training a CRF with missing labels.

Given a session with partial labels $\mathbf{y} = \langle y_{t_0}, y_{t_1}, \dots, y_{t_m} \rangle$, where t_i is the index of an entity with observed label, and a hidden label sequence $\mathbf{h} = \langle h_{l_0}, h_{l_1}, \dots, h_{l_m} \rangle$, where l_j is the index of an entity with latent (or unobserved) label, the conditional probability is,

$$P(\mathbf{y}|\mathbf{x}) = \sum_{\mathbf{h}} p(\mathbf{y}, \mathbf{h} | \mathbf{x}) = \frac{1}{Z(\mathbf{x})} \sum_{\mathbf{h}} \exp(\sum_{t=1}^T \sum_{k=1}^K \lambda_k f_k(y_{t-1}, y_t, \mathbf{x}, t))$$

We adopt the direct gradient ascent procedure methods from [22] to maximize $P(\mathbf{y}|\mathbf{x})$. We omit details about parameter learning due to length limit, and refer the readers to [22] for further information.

Feature Design. We use basic features such as NER tag, is capitalized, number of tokens, URL domain, and is in gazetteer. We also encode the context (i.e. ngrams and their POS tag sequence, $n \leq 3$) of the entity mention as features. Likewise, unigram, bigram, trigram features are used as Section 2.1.

2.3. Relation Extraction

After extracting entities from previous steps, we aim to identify queries satisfying templates **T-R-E** and **E-R-E** by using relation extraction techniques. For query with more than one entity, we predict if there exists a relation between pairwise entities (**E-R-E**). For entities not classified as **E-R-E**, we then predict if the query fits **T-R-E**. Specifically, for both templates, we train two multiclass logistic regression classifiers for relations based on one v.s. all strategy.

Feature Design. Our features are based on standard lexical and shallow syntactic features (i.e. POS tags, dependency relations) from the literature [16]. We also consider the words and their POS tags between the entity pairs or entity-type pairs. A gazetteer feature is constructed to indicate if the sample is in Freebase. We also take the predicted entity types from previous module as features.

3. WEB SEARCH QUERY INTERPRETATION

Web Search Session Log (Search Log). We collected 12 weeks of web search sessions from March 18 to June 9, 2013. The sessions consisting of query-click logs based on IE browsing history. This dataset will be called “Search log” henceforth. We choose a representative domain – film, to test our semantic parsing models. We further extract sessions that users express interests in film, i.e. sessions with page views in the domain of “www.imdb.com”, “www.fandango.com”, or “www.rottentomatoes.com”. This results in 297,352 sessions in total, from which 80% are randomly selected for training, with the remaining as the test set. The same tasks for other domains will be studied in the future work.

Training via Distant Supervision. To construct training samples for entity extraction, we first collect a gazetteer consisting of entities in the form of names from Freebase along with their entity types. In the case that one entity has multiple types, the type having most instances in Freebase is chosen. To maintain high precision for the training set, we match entities with names longer than two words in the gazetteer to each query. Table 1 shows nearly half of the queries and URLs get matched with at least one entity.

A gazetteer is also constructed for relation extraction. If a query contains two entities that compose a relation instance in Freebase, we assume this query expresses that relation. Those queries make up the training set for template **E-R-E**. For template **T-R-E**, a query is a positive sample when it contains entity e (e.g. “leonardo dicaprio”) and description for entity type t (e.g. “movies”), and there exists entity e' (e.g. “Titanic”) of type t where e and e' carry a relation.

Table 2 lists the number of entity and relation instances matched in our training and test sets.

	Total	Matched	Matched %
Query	1,773,930	866,131	48.83%
URL	18,166,912	8,413,971	46.31%

Table 1. Matching statistics for Search log.

	# Ent Type	# Ent	# Rel Type	# Rel
Training	44	215,053	33	20,197
Test	44	104,657	27	5,255
Test Only	0	23,557	0	1,104

Table 2. Statistics for Training and Test dataset. Test Only denotes the instance (or type) in the test set is not observed in training.

3.1. Pilot Study: Entity-based Search Session Modeling

We first carry out experiments to validate our session assumption, i.e. current search behavior depends on the searching history. Formally, in our entity-based session modeling, a search session S is represented by a sequence of entities e_0, e_1, \dots, e_l , as described in Section 2.2. We compute the likelihood of the sessions based on Markov models with different history, where higher order model considers longer history. The likelihood of session S is expressed as the product of the entities that compose S , with each entity probability conditional on the previous $n - 1$ entities (i.e. order- n): $P(S) = \prod_{i=1}^l P(e_i | e_{i-n+1}^{i-1})$, where e_i^j denotes the sequence e_i, \dots, e_j .

Maximum likelihood estimation can be used to get $P_{ML}(e_i | e_{i-n+1}^{i-1})$ as $\frac{c(e_i^i | e_{i-n+1}^{i-1})}{c(e_{i-n+1}^{i-1})}$, where $c(\alpha)$ indicates the frequency of entity sequence α occurs. We also experiment with two smoothing techniques: *additive smoothing* and *linear interpolation* [24]. Due to space limit, we refer the readers to [24] for details. To handle unknown entities in test set, we replace the first appearance of each entity with label “<Unk>” during training. Unknown entities are also substituted with “<Unk>” during test.

Given test set $T = \{S_i\}$, where S_i represents a session, we take the performance of each model as its cross-entropy [24] on T , i.e. $\frac{1}{N} \sum_{i=1}^{|T|} -\log P(S_i)$, with N as the total number of entities.

We train each model on sessions collected from May 13, 2013 to May 19, 2013, and test on sessions of June 17, 2013 to June 23, 2013. We reserve 10% of the training data as development set to tune the parameters. Results in Figure 3 show that Markov models considering longer search history produce better performance (i.e. with lower cross-entropy), which confirms our session assumption.

3.2. Results: Entity and Relation Extraction

Due to the large size of the test data, we automatically acquire labels by matching queries and entities with names longer than two words in gazetteer. The resulting labels are taken as the gold standard for evaluation. Because some predictions from our classifier are outside the gazetteer, which makes the precision underestimated, we only report recalls for measuring how many entities or relations are retrieved by the models. In the test phase, we assume our models only have access to the part of the gazetteer observed in training.

	CB Data					Netflix Data				
	Film	Director	Actor	Character	Genre	Film	Director	Actor	Character	Genre
(# utterances)	195	16	55	3	58	2,013	196	713	85	968
Baseline 1 (Gazetteer Only)	56.61	89.66	93.20	50.00	18.75	58.55	91.83	92.82	68.22	18.01
Baseline 2 (No Session)	71.71	81.25	90.57	50.00	18.75	67.98	75.29	89.23	70.23	18.01
Heck et al. [20]	-	-	-	-	-	71.72	84.62	58.61	-	29.55
Session model (CRF)	71.71	89.66	95.24†	50.00	18.75	67.98	91.83	95.12†	70.23	29.29†
Session model (CRF with missing label)	75.80	89.66	96.23†	80.00	21.54	69.39	92.01	95.12†	70.23	33.70

Table 5. Results for entity extraction on conversational corpora. Number of utterances containing each entity type is shown in the first row. F1 scores (multiplied by 100) are displayed. “-” indicates that result is unavailable. Statistically significant improvements ($p < 0.01$) over all compared systems except Heck et al. [20] are in **bold**. † denotes statistical significance ($p < 0.01$) over Baseline 1& Baseline 2.




Fig. 3. We compare models with different history – order-1 (O1), order-2 (O2), and order-3 (O3). Models with unknown entity handling are labeled with “Unk”. Two smoothing techniques are studied: additive smoothing (a.k.a Laplace smoothing) and linear interpolation (Inter). For interpolation, we also try models with entity type (“EntType”) as part of the components, i.e. adding $P_{ML}(e_i|type(e)_{i-n+1}^{i-1})$, which produce the best performance (i.e. lowest cross-entropy) combined with longer history.

	Recall on Entity Extraction
Baseline 1 (Gazetteer Only)	84.00
Baseline 2 (No Session)	86.34
CRF (mention + type)	87.33
Session model (CRF)	88.06†
Session model (CRF with missing label)	90.38

Table 3. Results for entity extraction on Search log. Recall is reported to measure the portion of entities (including entities that are unseen during training) that are identified. Our first session model (with †) achieves statistically significantly ($p < 0.01$) better result than the two baselines under a paired- t test. Our model based on CRF trained with missing labels significantly outperforms all the other compared methods.

Comparisons. We compare with 1) two baselines, and 2) a linear-CRF that predicts the entity spans and their types simultaneously [13]. Baseline 1 (Gazetteer Only) is constructed by matching the entities in gazetteer to the queries, where longest matching is selected if one token is mapped to multiple entities. For Baseline 2 (No Session), we follow [25] to construct a multiclass Maximum Entropy classifier with one v.s. all strategy.

Discussion. Table 3 shows that our two session-based models both significantly ($p < 0.01$) outperform the two strong baselines under a paired- t test. Our model based on CRF trained with missing labels also produces better performance than all the compared methods¹.

For relation extraction, we also compare with the baseline formulated by matching query to relation instances in the gazetteer. Table 4 presents the results for discovering relation instances for templates **E-R-E** and **T-R-E**, where our model outperforms the baseline.

¹We do not correct the spelling errors in the queries.

	Recall on E-R-E	Recall on T-R-E
Baseline (Gazetteer Only)	79.21	74.72
Our model	80.81	78.22

Table 4. Results for relation extraction on queries from Search log.

4. ENTITY EXTRACTION IN SPOKEN DIALOG SYSTEMS

SDS Corpora. To show how our models effectively leverage web search logs for entity extraction in SDS, two corpora are collected internally from real-use scenarios of two SDS, where the users interact with a system by voice and gesture to perform tasks like browsing the internet, searching, querying, and playing. The first conversational browsing dataset is constructed for various domains, and is called **CB Data** henceforth. The other dataset is designed for a natural language film search application based on Netflix (**Netflix Data**). As our models are trained on web search log in the *film* domain, we test them on the sessions of film-related tasks with manual transcripts. Specifically, we retain sessions annotated with entities in the film domain. In total, we have 6691 utterances from 136 sessions from CB; 13851 utterances from 1895 sessions from Netflix. We thus identify the entities and their types on utterance level.

Comparisons. We use the gazetteer collected from Section 3 to construct the same baselines (i.e. Baseline 1 (Gazetteer Only), and Baseline 2 (No Session)). In addition, we compare with the state-of-the-art approach designed for the same task in dialog management from [20], where they elicit training data from Wikipedia articles instead of web queries. They employ a linear-chain CRF to predict entities spans and their types simultaneously.

Discussion. Table 5 demonstrates that our session-based models trained on web search logs outperform two nontrivial baselines and the state-of-the-art method in most entity types on spoken dialogs. This shows that our session-based models is able to effectively leverage the behavioral patterns mined from web search log and multi-turn context to improve the entity extraction performance in spoken dialog systems. The performance on genre detection is relatively low. This may due to the different context people use for genres. For example, “Disney” is a genre in “Disney movie”, but is a place name in “Disney world”. Coreference resolution is not addressed in this work, but we believe it could further improve the performance.

5. CONCLUSION

In conclusion, we present a full pipeline to leverage semantic web search and browse sessions for a semantic parsing problem in multi-turn spoken dialog systems. Our models are able to apply the learned behavioral browsing patterns from web search logs to develop spoken dialog systems. With distant supervision technique, our models achieves better performance than nontrivial baselines and the state-of-the-art methods on both web search queries and spoken dialogs on tasks of entity and relation extraction.

References

- [1] David Suendermann and Roberto Pieraccini, *Spoken Language Understanding - Systems for Extracting Semantic Information from Speech*, chapter SLU in Commercial and Research Spoken Dialog Systems, Wiley and Sons, 2011.
- [2] Konrad Scheffler and Steve Young, “Automatic learning of dialogue strategy using dialogue simulation and reinforcement learning,” in *Proceedings of the second international conference on Human Language Technology Research*, San Francisco, CA, USA, 2002, HLT '02, pp. 12–19, Morgan Kaufmann Publishers Inc.
- [3] Jason D. Williams and Steve Young, “Partially observable markov decision processes for spoken dialog systems,” *Comput. Speech Lang.*, vol. 21, no. 2, pp. 393–422, Apr. 2007.
- [4] Oliver Lemon and Olivier Pietquin, “Machine learning for spoken dialogue systems,” in *In Proceedings of the European Conference on Speech Communication and Technologies (Interspeech07)*, Anvers, 2007.
- [5] David Griol, Giuseppe Riccardi, and Emilio Sanchis, “A statistical dialog manager for the luna project,” in *INTERSPEECH*, 2009, pp. 272–275, ISCA.
- [6] Matthew Frampton and Oliver Lemon, “Recent research advances in reinforcement learning in spoken dialogue systems,” *Knowledge Eng. Review*, vol. 24, no. 4, pp. 375–408, 2009.
- [7] Dilek Hakkani-Tur, Gokhan Tur, Larry Heck, Asli Celikyilmaz, Ashley Fidler, Dustin Hillard, Rukmini Iyer, and Sarangarajan Parthasarathy, “Employing web search query click logs for multi-domain spoken language understanding,” in *IEEE Automatic Speech Recognition and Understanding Workshop*, 2011.
- [8] Gukhan Tur, Dilek Z. Hakkani-Tur, Dustin Hillard, and Asli elikyilmaz, “Towards unsupervised spoken language understanding: Exploiting query click logs for slot filling,” in *INTERSPEECH*, 2011, pp. 1293–1296, ISCA.
- [9] Larry Heck and Dilek Hakkani-Tur, “Exploiting the semantic web for unsupervised spoken language understanding,” in *IEEE Workshop on Spoken Language Technology*, 2012.
- [10] Larry Heck, “The conversational web,” in *IEEE Workshop on Spoken Language Technology*, 2012.
- [11] Larry Heck, Dilek Hakkani-Tur, Madhu Chinthakunta, Gokhan Tur, Rukmini Iyer, Partha Parthasarathy, Lisa Stifelman, Elizabeth Shriberg, and Ashley Fidler, “Multimodal conversational search and browse,” in *IEEE Workshop on Speech, Language and Audio in Multimedia*, 2013.
- [12] Aditya Bhargava, Asli Celikyilmaz, Dilek Hakkani-Tur, and Ruhi Sarikaya, “Easy contextual intent prediction and slot detection,” in *ICASSP*, 2013, IEEE.
- [13] Jeffrey Pound, Alexander K. Hudek, Ihab F. Ilyas, and Grant Weddell, “Interpreting keyword queries over web knowledge bases,” in *Proceedings of the 21st ACM international conference on Information and knowledge management*, New York, NY, USA, 2012, CIKM '12, pp. 305–314, ACM.
- [14] Ellen Riloff and Rosie Jones, “Learning dictionaries for information extraction by multi-level bootstrapping,” Menlo Park, CA, USA, 1999, AAAI '99/IAAI '99, pp. 474–479, American Association for Artificial Intelligence.
- [15] Marius Paşca, “Weakly-supervised discovery of named entities using web search queries,” in *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, New York, NY, USA, 2007, CIKM '07, pp. 683–690, ACM.
- [16] Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky, “Distant supervision for relation extraction without labeled data,” Stroudsburg, PA, USA, 2009, ACL '09, pp. 1003–1011, Association for Computational Linguistics.
- [17] Limin Yao, Sebastian Riedel, and Andrew McCallum, “Collective cross-document relation extraction without labelled data,” Stroudsburg, PA, USA, 2010, EMNLP '10, pp. 1013–1023, Association for Computational Linguistics.
- [18] Gokhan Tur, Minwoo Jeong, Ye-Yi Wang, Dilek Hakkani-Tur, and Larry Heck, “Exploiting the semantic web for unsupervised natural language semantic parsing,” in *Proceedings of Interspeech*, 2012.
- [19] Dilek Hakkani-Tur, Larry P. Heck, and Gokhan Tur, “Using a knowledge graph and query click logs for unsupervised learning of relation detection,” in *ICASSP*, 2013, IEEE.
- [20] Larry Heck, Dilek Hakkani-Tur, and Gokhan Tur, “Leveraging knowledge graphs for web-scale unsupervised semantic parsing,” in *Proceedings of Interspeech*, 2013.
- [21] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” in *Proceedings of the Eighteenth International Conference on Machine Learning*, San Francisco, CA, USA, 2001, ICML '01, pp. 282–289, Morgan Kaufmann Publishers Inc.
- [22] Kedar Bellare and Andrew Mccallum, “Learning Extractors from Unlabeled Text using Relevant Databases,” in *Sixth International Workshop on Information Integration on the Web (IIWeb)*, 2007.
- [23] Dan Klein and Christopher D. Manning, “Accurate unlexicalized parsing,” in *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics - Volume 1*, Stroudsburg, PA, USA, 2003, ACL '03, pp. 423–430, Association for Computational Linguistics.
- [24] Stanley F. Chen and Joshua Goodman, “An empirical study of smoothing techniques for language modeling,” Stroudsburg, PA, USA, 1996, ACL '96, pp. 310–318, Association for Computational Linguistics.
- [25] James R. Curran and Stephen Clark, “Language independent ner using a maximum entropy tagger,” Stroudsburg, PA, USA, 2003, CONLL '03, pp. 164–167, Association for Computational Linguistics.