SHINRA: Structuring Wikipedia by Collaborative Contribution

Anonymous authors

Abstract

We are reporting the SHINRA project, a project for structuring Wikipedia with collaborative construction scheme. The goal of the project is to create a huge and well-structured knowledge base to be used in NLP applications, such as QA, Dialogue systems and explainable NLP systems. It is created based on a scheme of "Resource by Collaborative Contribution (RbCC)". We conducted a shared task of structuring Wikipedia, and at the same, submitted results are used to construct a knowledge base.

There are machine readable knowledge bases such as CYC, DBpedia, YAGO, Freebase Wikidata and so on, but each of them has problems to be solved. CYC has a coverage problem, and others have a coherence problem due to the fact that these are based on Wikipedia and/or created by many but inherently incoherent crowd workers. In order to solve the later problem, we started a project for structuring Wikipedia using automatic knowledge base construction shared-task.

The automatic knowledge base construction shared-tasks have been popular and well studied for decades. However, these tasks are designed only to compare the performances of different systems, and to find which system ranks the best on limited test data. The results of the participated systems are not shared and the systems may be abandoned once the task is over.

We believe this situation can be improved by the following changes:

1. designing the shared-task to construct knowledge base rather than evaluating only limited test data
2. making the outputs of all the systems open to public so that we can run ensemble learning to create the better results than the best systems
3. repeating the task so that we can run the task with the larger and better training data from the output of the previous task (bootstrapping and active learning)

We conducted "SHINRA2018" with the above mentioned scheme and in this paper we report the results and the future directions of the project. The task is to extract the values of the pre-defined attributes from Wikipedia pages. We have categorized most of the entities in Japanese Wikipedia (namely 730 thousand entities) into the 200 ENE categories. Based on this data, the shared-task is to extract the values of the attributes from Wikipedia

pages. We gave out the 600 training data and the participants are required to submit the attribute-values for all remaining entities of the same category type. Then 100 data out of them for each category are used to evaluate the system output in the shared-task.

We conducted a preliminary ensemble learning on the outputs and found 15 F1 score improvement on a category and the average of 8 F1 score improvements on all 5 categories we tested over a strong baseline. Based on this promising results, we decided to conduct three tasks in 2019; multi-lingual categorization task (ML), extraction for the same 5 categories in Japanese with a larger training data (JP-5) and extraction for 34 new categories in Japanese (JP-34).

1. Introduction

Wikipedia is a great resource as a knowledge base of the names in the world. However, Wikipedia is created for human to read rather than machines to process. Our goal is to transform the current Wikipedia to a machine readable format based on a clean structure. There are some machine readable knowledge bases such as CYC[Lenat, 1995], DBpedia[Lehmann et al., 2015], YAGO[Mahdisoltani et al., 2015], Freebase[Bollacker et al., 2008], Wikidata[Vrandečić and Krötzsch, 2014] and so on, but each of them has problems to be solved. CYC has a coverage problem, and others have a coherence problem due to the fact that these are based on Wikipedia and/or created by many but inherently incoherent crowd workers. In order to solve these problems, we started a project for structuring Wikipedia using automatic knowledge base construction shared-task using a cleaner ontology definition.

The automatic knowledge base construction shared-tasks have been popular for decades. In particular, there are popular shared-tasks on Information Extraction, Knowledge Base population, attribute extraction and so on, such as KBP[U.S. National Institute of Standards and Technology (NIST)] and CoNLL. However, these tasks are designed only to compare the performances of participated systems, and to find which system ranks the best on limited test data. The results of the participated systems are not shared and the systems may be abandoned once the task is over.

We believe this situation can be improved by the following changes:

1. designing the shared-task to construct knowledge base rather than evaluating only limited test data

2. making the outputs of all the systems open to public so that we can run ensemble learning to create the better results than the best systems

3. repeating the task so that we can run the task with the larger and better training data from the output of the previous task (bootstrapping and active learning)

We conducted "SHINRA2018" with the above mentioned scheme, we call it "Resource by Collaborative Contribution", and in this paper we report the first results and the future directions of the project.

The task is to extract the values of the pre-defined attributes from Wikipedia entity pages. We used Extended Named Entity (ENE) as the definition of the category (in total 200 categories in the ontology) and the attributes (average of 20 attributes) for each category. We have categorized most of the entities in Japanese Wikipedia (namely 730 thousand

entities) into the ENE categories prior to this project. Based on this data, the shared-task is to extract values of the attributes defined for the category of each entity. We gave out the 600 training data for 5 categories at this time and the participants are supposed to submit the attribute-values for all remaining entities in Japanese Wikipedia. Then 100 data out of them are used at the evaluation of the participated systems in the shared-task. For example, there are about 200K person entities in Japanese Wikipedia, and the participants have to extract the attribute-values, such as his/her birthday, the organizations he/she have belonged, his/her mentor, the awards he/she received, from all the remaining entities (i.e. 199.4K = 200K-600 entities). Before starting the project, the participants signed the contract that all the output will be shared, so that we can conduct the ensemble learning on those outputs, and hence create a better knowledge base than the best system in the task. A very promising results of the ensemble learning is achieved and it will lead to the cleaner machine readable knowledge base construction.

2. Related Work

Structured knowledge bases have considered as one of the most important knowledge resources in the fields of Natural Language Processing. In this field, some several major projects targeted to construct structured knowledge bases in the past. From the earliest project CYC to recent Wikipedia based projects such as DBpedia, Yago, Freebase and Wikidata. Moreover, there are some shared tasks organized by knowledge base structuring project such as KBP and CoNLL. We will introduce these resources and projects and describe the points we consider as issues to be solved in those projects.

CYC ontology is a large knowledge base constructed for inferring common senses[Lenat, 1995]. As the handmade knowledge base for the general domain like as CYC, the cost of construction and maintenance has become very high, and there is a limit by hand creation also in coverage.

DBpedia is a project to construct a structured information from the semistructured data in Wikipedia such as infoboxes or categories[Lehmann et al., 2015]. DBpedia has a problem of accuracy, coverage, and coherence. For example, "Shinjuku Station" is a kind of railway station is a type of the railway company "Odakyu Electric Railway" at the DBpedia. Of course, the station is not a railway company. This problem caused by infoboxes on the "Shinjuku Station" article in Wikipedia.

Yet Another Greater Ontology (YAGO) is a ontology constructed by mapping Wikipedia articles to the WordNet synsets[Mahdisoltani et al., 2015]. YAGO has adopted attributes information extracted from infoboxes like as DBpedia because no attributes defined in WordNet synsets.

Freebase is a project to construct a structured knowledge base by crowdsourcing same as Wikipedia[Bollacker et al., 2008]. However, by the crowdsource approach, Freebase is not a well-organized ontology because it includes some noises and lack of coherence except duplication of other databases. Currently, Freebase has become integrated into Wikidata.

Wikidata is a knowledge base to provide a structured information to Wikipedia or another Wikimedia projects[Vrandečić and Krötzsch, 2014]. Wikidata have noises and lack of coherence because it has constructed by bottom-up approach same as Wikipedia and Freebase.

KBP is a shared task by NIST for establishing a technology to construct a structured knowledge base from non-structured documents[U.S. National Institute of Standards and Technology (NIST)]. KBP mainly consists of two tasks. One task is an Entity Discovery and Linking (EDL) which is to find and identify an entity defined in DB from documents. Another one is a Slot Filling which is to extract attribute information of the entity. KBP is limited entity types to Person, Location, and Organization in contrast to Wikipedia's wide coverage.

FIne Grained Entity Recognition (FIGER) is a project to identify 112 types of named entity classes that are finely defined, such as ENE, from documents[Ling and Weld, 2012]. The category of FIGER seems a bit biased, and it doesn't have attribute definitions for each category.

3. Extended Named Entity

In order to create structured knowledge base, we have learned that top-down design of ontology is needed. All the knowledge base created by crowds, such as Wikipedia, DBpedia, Freebase and Wikidata, have very inconsistent and imbalanced ontologies, as well as adhoc attributes. As the top-down designed ontology, we used Extended Named Entity (ENE) hierarchy. Extended Named Entity (ENE) is a named entity classification hierarchy along with the attribute definition for each category [Sekine, 2008, Sekine et al., 2002]. It includes fine-grained 200 categories of names in hierarchy of up to 4 layers. It contains not only the finer categories of "location" or "organization", but also contains new named entity types such as "products", "event", "position" and so on. As finer categories, it contains, for example, river, lake, ocean, mountain for terrain under location. Figure 1 shows the ENE definition, version 7.1.0. Attributes are defined based on the investigation of the entities in each category. For example, the attributes for "airport" categories are as follows: Reading, IATA code, ICAO code, nickname, name origine, job titl of the name origin, number of users per year, the year of the statistic, the number of airplane landing per year, longitude, latitude, location, old name,elevation, big city near by, number of runaway, length of runaway, size of the area, airport nearby, running company, running time and open year.

4. Categorization of Wikipedia Entities

In order to conduct the shared-task of the attribute-value extraction on Wikipedia entities, we first have to assign one or more categories for each entity. In other words, we have to know the Wikipedia page of "Chicago Airport" represent an airport entity. We have annotated one or more of 200 ENE categories to 782,406 entities of Japanese Wikipedia (201711 version) prior to this project. At the annotation, we have excluded the less popular 151K entities, which have less than 5 incoming links, and non-entity pages (about 53K pages) such as common nouns and simple numbers. This annotation was done by Machine learning method followed by hand checking on less reliable ones. The machine learning [Masatoshi et al., 2018] was conducted with 20K training data created by hand. Then a human check was conducted on the machine learning outputs with less reliable scores. We evaluate a sample data by multiple annotators to see the accuracy of the data and

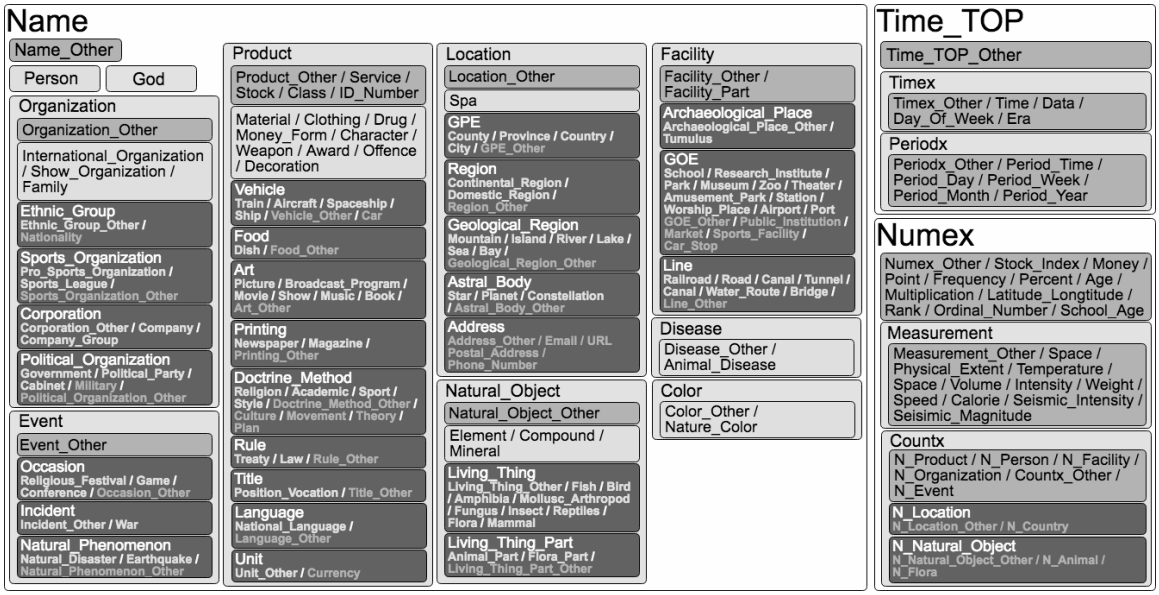

Figure 1: Definition of Extended Named Entity Hierarchy

Table 1: 14 Most Frequent Categories

Category	Frequency
Person	247,983
City	45,306
Music	41,049
Artifact other	33,453
Broadcast program	32,050
Company	26,746
School	23,609
Literature	18,515
Movie	17,901
Train station	16,901
Sports event	15,863
Road	15,360
Method other	14,766
Play group	10,181

we observed the accuracy of categorization is 98.5%. The remaining 1.5% are those which are very difficult even for the human annotators, we may need even cleaner definition of ontology for those entities. 4 show the most frequent categories in the data.

5. Shared-Task Definition

In this section, we will describe the definition of the shared-task. The task is to extract the attribute-values of entities from the Wikipedia page. The category of the entity is identified and given to the participants, and the list of the attributes for the category is predefined. For this year's task (SHINRA2018), we picked 5 categories, namely "person", "city", "company", "airport" and "chemical compound" for the shared-task. This selection was done on the largest subcategories of "person", "location", "organization", which are the traditional three categories of named entity (person has no subcategories, and itself is the only category in ENE, though) We give out 600 training data for each category. In the training data, all attribute-values mentioned in the Wikipedia page are manually extracted and form the training data in JSON format. The participants also received the list of all entities, i.e. Wikipedia pages, for 5 categories, and they are required to extract attribute-values from all the remaining data. The evaluation is conducted on 100 entities for each category, but the participants will not be notified which 100 are used for evaluation even after the evaluation is over. This is for the purpose of the data construction so that the participants have to do their best to annotate all the data, and the purpose of the future comparison (if the test data is known, the participants could tune their system to the test data even unintentionally through a number of experiments). The results are reported by precision, recall and F1 scores, as usual. The systems submitted before the deadline are reported as the formal results and the results submitted after the deadline are reported as a reference results. In the ensemble learning, we use all the results regardless of the formal or reference results so that we can achieve the best results for the resource construction purpose. The participants are not required to submit the results for all 5 categories, as some participants might be interested in a particular category or may have smaller machine resources to run for all the categories.

6. Building the Data

The manual creation of the training and test data was not easy. We tried three strategies for the annotation as a preliminary experiment.

- Expert of the construction of linguistic data create the data
- Student who are supervised by experts create the data
- Workers on the crowdsourcing (Lancers) create the data

We figured out that the upper in the list, the more expensive, but at the same time the more accurate. However, we found that the crowdsourcing has relatively high coverage. The task of crowdsourcing is designed with three stages. The first stage is to identify the sections where the given attribute-value is written. In this stage, even the worker find the value in the page, they are not requested to extract the value. This identification of the sections will be repeated until two workers found no value is found, because some attributes have multiple values in one page. Then in the second stage, the values are extracted from the sections which are identified to contain the value(s). The final stage is to check if the extracted value is really the value for the attribute. Maybe this careful strategy might

produced the relatively high coverage. Based on the preliminary annotation experiments, we decided to use "expert" and "crowd" strategies. The first round annotation is done by both "expert" and "crowd" independently for the same attributes, then both results are used to create the final annotation by a different expert.

7. SHINRA2018 Shared-Task: Results

In this section, we will explain the results of the shared-task. 5 months are given to the participants to develop their systems and run their experiments from April to September 2018. 16 systems by 8 participants are submitted at SHINRA2018. The first two columns in 2 shows the participants (some in abbreviations) and their methods. Here "pattern" means that they created a hand made patterns for the attaribute-value extraction, and "DL" means some sort of "Deep Learning". "DrQA" is an open source QA system, where the participant transformed the infobox into a sentence by pattern, e.g. "The birthday of Barack Obama is August 4, 1961" and attributes to be extracted is transformed to a question, e.g. "Who is the father of Barack Obama?" in order to extract "father" attribute of "Barack Obama" entity. Then they train and run DrQA for each category. We are very happy to see there are various technologies used in the shared-task.

The results are shown in 3rd to 7th columns in the same table. The top result is shown in bold for each category. The Unisys's DrQA system performs the best in three categories, which don't have so many information in infobox. As their method use all the attributes and all types of data are treated in the same system (regardless of infobox or in the explanation sentence), the amount of training data for the system becomes relatively larger and it may receive the benefit at the situation where the training data is relatively small. TUT's pattern based system performed very well on airport category, in which the most information are described in infobox, and practically only one infobox template is used in the category. Note that the category "person" has so many different infobox templates depending on the vocation, and the company's infobox template are depending on the type of the company.

8. Preliminary Results of Ensemble Learning

The goal of "Resource by Collaborative Contribution" is to produce better accuracy than the accuracy of the best single system. In order to see if this idea is plausible, we conducted a preliminary experiment of ensemble learning. In the past, the ensemble learning methods have been studied with various ideas; such as Bagging [Leo Breiman and Eiman, 1994], Boosting [Freund and Schapire, 1997] or stacking. These methods are generally a method to create a high accuracy system combining more than one ML systems. However, in our situation where the outputs of many systems are given and the objective is to produce the best output out of the system outputs by ensemble scheme. Because of this, the stacking method is best suitable for our purpose, but we tried two simpler methods, i.e. the simple voting method and the weighted voting method base on the accuracy on held-out data as a preliminary ensemble learning experiment.

First, we will explain the simple voting method. Assume there are n systems which outputs value v for an attribute of an entity. Then the value v receives score n. Separately,

Table 2: SHINRA2018 Results

Participants	method	ENE Category				
		Person	Company	City	Airport	Compound
TUT	Pattern	0.20	0.41	0.28	**0.72**	
OPU	(reference) Pattern	0.19				
	(reference) Pattern + Heuristics	0.16				
NUT	Pattern + LightGBM					0.42
Sansan	Pattern		0.30			
Fuji Xerox	NCRFpp	0.31			0.38	0.15
	RDFDNN					0.15
	(reference) NCRFPP		0.30	0.43	0.42	0.39
	(fefference) RDFDNN		0.28		040	0.37
TOPPAN	BRNN/LSTM		0.29		0.35	
	Pattern		0.33			
	(reference) BRNN/LSTM + Pattern		0.34		0.42	
	(reference) Pattern		0.41			
Unisys	DrQA		**0.53**		0.67	**0.47**
	(reference) DrQA	**0.44**		0.42		
AIP	(reference) RNN	0.36	0.38	**0.46**	0.71	0.46

we compute the threshold t by maximizing the accuracy on the held-out data. If n ¿ t, then we take the value v as the output of the ensemble system. In practice, we split the actual test data, which contains 100 samples, into two halves; one for the held-out data and the other for the test data for this experiment. We conducted the same experiment replacing the held-out data and test data; i.e. we conducted a 2 fold cross-validation experiment on 100 test data. The other method is the weighted ensemble method. We weighted the vote of the system by the accuracy of the system on the held-out data. Instead of a sum of the number of the systems which produce attribute-value v, we compute the sum of the accuracy as the score for the value v. The way to define the threshold and the cross-validation mechanism are the same to those of the simple voting method.

We will show the precision, recall and F1 score of the baseline and the ensemble methods in 8. Also the relative improvement of those two methods compared to the baseline method is shown in Figure 2. The baseline method is constructed by combining the best system outputs for each category, i.e. TUT system for "airport", AIP system for "city" and Unisys system for the rest, which is better than a single system, e.g. Unisys, though. We can observe from the table and the graph that the two voting methods performs better than the baseline methods in F1 score. Also the weighted voting method performs better than the simple voting method. The improvement exceed 15 F1 score on "airport" category and 3 F1 score to all the categories. The average improvement is 8 F1 score. This result of the preliminary experiment show the effectiveness of the ensemble learning methods. This

Table 3: Results of Ensemble Learning

Category	Baseline method			Simple voting			Weighted Voting		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
airport	0.790	0.658	0.718	0.802	0.854	0.828	**0.883**	**0.861**	**0.872**
city	0.378	**0.588**	0.460	0.541	0.544	0.542	**0.570**	0.575	**0.573**
company	**0.748**	0.441	0.555	0.569	0.479	0.520	0.675	**0.529**	**0.593**
person	**0.563**	0.362	0.440	0.437	0.474	0.455	0.501	**0.459**	**0.479**
compound	**0.750**	0.370	0.496	0.567	0.616	0.590	0.674	**0.653**	**0.663**
average	0.574	0.479	0.522	0.575	0.579	0.577	**0.651**	**0.602**	**0.626**

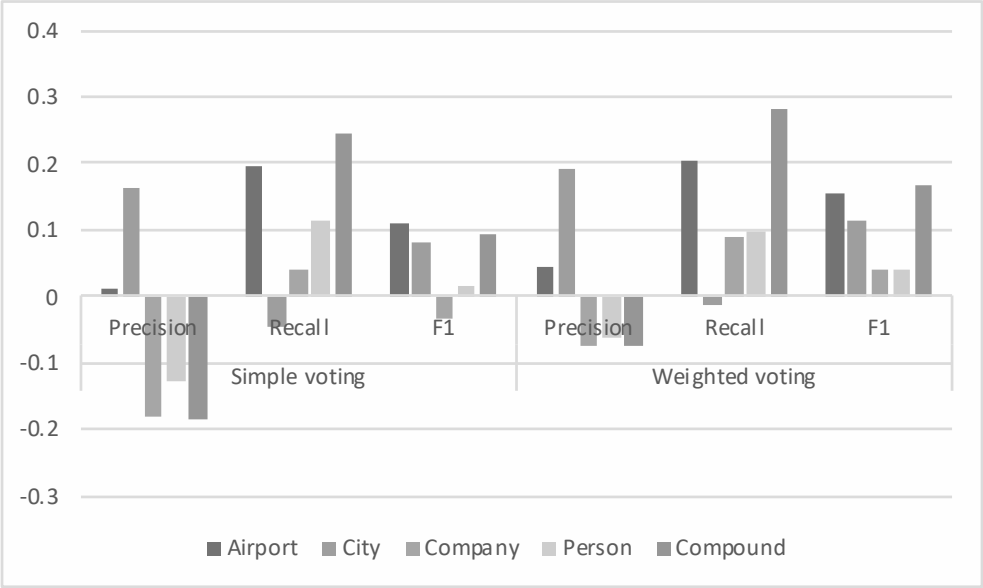

Figure 2: Relative improvements of the ensemble methods to the baseline method for each category

result is very promising and encouraging for the "Resource by Collaborative Contribution" scheme.

9. Future Projects: SHINRA2019

Based on the success of the SHINRA2018 project, we decided to continue this project as SHINRA2019. We are planning to conduct three tasks as follows:

- ML: Multi-lingual categorization task
- JP-5: Structuring task for the same 5 categories with larger training data in Japanese
- JP-34: Structuring task for 34 new categories with 100 training data in Japanese

The multi-lingual task is to expand the benefit of RbCC to the knowledge base resources in languages other than Japanese. We are planning to run it on 9 languages with the largest numbers of "users"; namely English, Spanish, French, German, Chinese, Russian, Portuguese, Italian and Arabic [Wikipedia]. Actually, Japanese is the 10th ranked language on the measure, so Wikipedias of these 9 language have more users than that of Japanese. As we don't have the category information for those 9 language Wikipedia entities, the first task is to categorize the entities. For Japanese, we annotated 20K entities as the training data for the categorization, but now we have most of the Japanese entities categorized, we can utilize this information. There are links between equivalent Wikipedia entities in different languages. For example, we observed there are about 200K entities links from Japanese to English among 720K entities already categorized in Japanese. We can use them as the training data to categorize English entities. Likewise there are language links to other 8 language Wikipedias from Japanese Wikipedia, although the number of linked entities are much smaller and some noise may exist, the participants can use much bigger training data than that for the initial Japanese categorization experiment. As there are links between the Wikipedias of other languages and possibly different types of infoboxes exist in other language, too, the participants have a lot of information to be used in the categorization task.

JP-5 is the task to extract the attribute-values for the same 5 categories in SHINRA2018; namely "person", "company", "city", "airport" and "chemical compound". At SHINRA2018, the values are prepared without contexts. In other words even there are more than one mention of a particular attribute, we didn't give out which one is the mention to that value. For example, assume the nationality for a person is "Japan", but the same string may be mentioned in the same person page but not necessarily be meant to indicate the nationality of the person, e.g. "He left Japan", we had no means to know that the context is not for the nationality. It is similar to the situation of the distant-supervision, so it is difficult to extract only the context of nationality. At SHINRA2019, we will annotate the attribute-value in the text, so that the exact context for the value can be extracted. We are also planning to expand the size of the training data from 600 to 1500, at least for the categories "person", "company" and "city".

JP-34 is the task to extract attribute-value for 34 new categories. As we mentioned in the previous section, creating the data is laborious, the size of the training data will be very small, namely 100. However the categories to be tested will be very close; 7 subcategories of Geographical Political Entities (GPE) such as country, prefecture/state and county, 8 subcategories of terrain such as mountain, island, river, lake and ocean, and organizational entities such as international organization, political organization, ethnic group and nationality. Although the number of training data is much smaller, we chose the very similar types and the similar attributes may exists. Some techniques of machine learning with adaptation might help creating a good result.

We hope to have many participants so that the better results can be achieved by the ensemble learning methods to all three tasks.

10. Conclusion

We proposed a scheme of knowledge base creation: "Resource by Collaborative Contribution". We conducted the Japanese Wikipedia structuring project, SHINRA2018, based on that scheme. Based on the clear definition of categories and attributed for named entities, Extended Named Entity ontology, the task is to extract the attribute-values from Japanese Wikipedia. 8 groups participated to the task, and the ensemble learning results are very promising. More than 15 F1 score improvement over the the best single system was achieved on "airport" category, and the average of 8 F1 score improvement was achieved using the weighted voting methods. We are planning to conduct SHINRA2019 based on the same scheme on 3 tasks. These are the multi-lingual categorization, the extraction of attribute-value on the same 5 categories, and the extraction of attribute-values on 34 new categories in Japanese.

We'd like to express our deep appreciation to all the participants and collaborators who helped this project. Without the participation, we couldn't even try the ensemble learning and achieve the goal. We are hoping to expand and spread the idea of RbCC scheme, not only limited to this kind of task and resource.

References

Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In *SIGMOD '08: Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-102-6. doi: http://doi.acm.org/10.1145/1376616.1376746.

Yoav Freund and Robert E Schapire. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *journal of computer and system sciences*, 55:1199139, 1997. URL http://www.face-rec.org/algorithms/Boosting-Ensemble/decision-theoretic{_}generalization.pdf.

Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sren Auer, and Christian Bizer. Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195, 2015.

Douglas Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, November 1995.

By Leo Breiman and Leo B Eiman. Bagging Predictors. Technical report, 1994. URL https://statistics.berkeley.edu/sites/default/files/tech-reports/421.pdf.

Xiao Ling and Daniel S. Weld. Fine-grained entity recognition. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, AAAI'12, pages 94–100. AAAI Press, 2012. URL http://dl.acm.org/citation.cfm?id=2900728.2900742.

Farzaneh Mahdisoltani, Joanna Biega, and Fabian M Suchanek. Yago3: A knowledge base from multilingual wikipedias. *CIDR*, 2015.

Suzuki Masatoshi, Matsuda Koji, Sekine Satoshi, Okazaki Naoaki, and Inui Kentaro. A joint neural model for fine-grained named entity classification of wikipedia articles. *IEICE Transactions on Information and Systems*, E101.D(1):73–81, 2018. doi: 10.1587/transinf.2017SWP0005.

Satoshi Sekine. Extended named entity ontology with attribute information. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, May 2008. European Language Resources Association (ELRA). ISBN 2-9517408-4-0. http://www.lrec-conf.org/proceedings/lrec2008/.

Satoshi Sekine, Kiyoshi Sudo, and Chikashi Nobata. Extended named entity hierarchy. In *LREC*, 2002.

U.S. National Institute of Standards and Technology (NIST) . TAC Knowledge Base Population (KBP) 2017. URL https://tac.nist.gov/2017/KBP/.

Denny Vrandečić and Markus Krötzsch. Wikidata: A free collaborative knowledgebase. *Commun. ACM*, 57(10):78–85, September 2014. ISSN 0001-0782. doi: 10.1145/2629489. URL http://doi.acm.org/10.1145/2629489.

Wikipedia. List of Wikipedias. https://meta.wikimedia.org/wiki/List_of_Wikipedias.