

HYCEDIS: HYbrid Confidence Engine for Deep Document Intelligence System

Bao-Sinh Nguyen*
Quang-Bach Tran*
simon@cinnamon.is
neath@cinnamon.is
Cinnamon AI Inc
Vietnam

Tuan-Anh Nguyen Dang
Duc Nguyen
simon@cinnamon.is
john@cinnamon.is
Cinnamon AI Inc
Vietnam

Hung Le
thai.le@deakin.edu.au
Deakin University
Australia

ABSTRACT

Measuring the confidence of AI models is critical for safely deploying AI in real-world industrial systems. One important application of confidence measurement is information extraction from scanned documents. However, there exists no solution to provide reliable confidence score for current state-of-the-art deep-learning-based information extractors. In this paper, we propose a complete and novel architecture to measure confidence of current deep learning models in document information extraction task. Our architecture consists of a Multi-modal Conformal Predictor and a Variational Cluster-oriented Anomaly Detector, trained to faithfully estimate its confidence on its outputs without the need of host models modification. We evaluate our architecture on real-world datasets, not only outperforming competing confidence estimators by a huge margin but also demonstrating generalization ability to out-of-distribution data.

CCS CONCEPTS

• **Information systems** → **Document structure.**

KEYWORDS

uncertainty, neural networks, supervised learning, information extraction

1 INTRODUCTION

Recent advances in machine learning enables creations of automatic information extractors that can read the input document in image format, locate and understand relevant text lines before organizing the information into computer-readable format for further analysis [10, 36]. Despite these successes, in critical domains such as healthcare and banking, humans still have to involve to scrutinize AI outputs as there is no room for AI errors in making important decisions that can affect human life. Confidence score estimation is one critical step towards implementing practical industrial systems wherein AI automates most of the operations yet human will intervene if necessary [39].

Unfortunately, to the best of our knowledge, there exists no holistic solution to reliably estimate the confidence score for the task of document information extraction. Current confidence score approaches are either generic methods verified only for simple image classification tasks [6] or applied only for part of the information extraction process [25].

In this paper, we introduce a novel neural architecture that can judge the result of extracted structured information from documents provided by the information extracting neural networks (hereafter referred to as the IE Networks). Our architecture is hybrid, consisting of two models, which are a Multi-modal Conformal Predictor (MCP) and a Variational Cluster-oriented Anomaly Detector (VCAD). The former aims to combine the neural signals from 3 main stages of information extraction processes including text-box localization, OCR, and key-value recognition to predict the confidence level for each extracted key-value output. The later computes anomaly scores for the raw input document image, providing the MCP with additional features to produce better confidence estimation. The VCAD works on global, low-level features and plays a critical role in lifting the burden of detecting outliers off the MCP, which focuses more on local, high-level features.

We demonstrate the capacity of our proposed architecture on real-world invoice datasets: SROIE [14], CORD [28]) and 2 in-house datasets. The experimental results demonstrate that our method outperforms various confidence estimator baselines (including Dropout [6], temperature scaling [8]). In short, we summarize our contribution as follows:

- We propose a *Multi-modal Conformal Predictor (MCP)* using a Feature Fusion module over 3 Feature Encoders to fuse signals extracted from IE Networks and compute the confidence score of the IE Networks' outputs.
- We provide a *Variational Cluster-oriented Anomaly Detector (VCAD)* to equip the MCP with an ability to handle out-of-distribution data.
- We unify the proposed MCP and VCAD in a single hybrid confidence engine, dubbed as HYCEDIS, that for the first time, can well estimate the confidence of document intelligent system.
- We conduct intensive experiments on 4 datasets with detailed ablation studies to show the effectiveness and generalization of our hybrid architecture on real-world problems.

2 BACKGROUND

A typical Document Intelligence System consists of multiple smaller steps: text detection, text recognition and information extraction (IE). Given a document image, the usual first step is to detect text lines, using segmentation [3, 9, 23] or object detection method [18, 19, 22]. The detected text line images can each go through an OCR model to transcribe into text [7, 17, 27, 33]. After all text contents are transcribed, the relevant text entities can be extracted, using entity

*Both authors contributed equally to this research.

recognition (sequence tagging) method [20, 35, 37], segmentation-based method [5, 36], or graph-based method [21, 29, 34] which formulates the document layout as a graph of text-lines/words.

In this paper, we adopt a common IE Network that consisted of 3 main modules: text detection (Layout Analysis), text recognition (CRNN) and graph-based information extraction model (Graph KV). The text detection model shares the same architecture with [3] which utilizes segmentation masks to detect text-lines in the document image. The text recognition (CRNN) uses popular CNN+Bi-LSTM+CTC-loss architecture to transcribe each text-line images into text. Finally, the GCN model [21] performs the node classification tasks from the input document graph constructed from the text-lines' location and text to extract relevant information. Here, for our problems, we classify each node into different key types that represent the categories of the text-line.

3 METHODOLOGY

3.1 Multi-modal Conformal Predictor (MCP)

Given extracted intermediate features of IE Networks, our Multi-modal Conformal Predictor aims to estimate the confidence score through predicting whether the final output is *true* or *false*. The MCP architecture (see Figure 1(a)) contains two main components which are Feature Encoding and Feature Fusion. The Feature Encoding extracts features from different layers of trained IE Networks while the Feature Fusion combine them for predicting the final output.

Feature Encoding. Motivated by designs of late-fusion multi-view sequential learning approaches [4, 16], three components of the Feature Encoding layers are independent processing streams including visual, lingual and structural feature encoders. In particular, the visual feature encoder is a many-to-one LSTM $f_{VF}(\cdot)$ that captures the visual information embedded in the CRNN of the IE Network. It takes the CRNN's logits (whose shape is $T \times F_{in_vis}$ where T is the number of timesteps) as input and outputs a vector of size F_{out_vis} , which represents the knowledge of the IE Networks on its OCR model's neural activations given the input image. In particular, for the i -th extracted text-line image I_i , we compute it as: $E_i^{vis} = f_{VF}(CRNN(I_i))$.

The lingual feature encoder, which is also implemented as a many-to-one LSTM $f_{LF}(\cdot)$, processes the predicted OCR texts of the IE Networks. Each OCR-ed character in the text is represented as an one-hot vector with the size that equals to the size of the corpus. For the i -th extracted OCR-ed text, the LSTM takes a sequence of these one-hot vectors (denoted by $text_i$, whose shape is $T \times F_{in_OCR}$) and produces an output vector of size F_{out_OCR} , representing the knowledge of the IE Networks on the linguistic meaning and the syntactic pattern of its OCR-ed outputs. We compute it as: $E_i^{OCR} = f_{LF}(text_i)$.

The structural feature encoding $f_{SF}(\cdot)$ is a feed-forward neural network that accesses the information from the final layer of the IE Networks (node classification) – the Graph KV module. Here, the logits before softmax layer of the i -th node in the graph (corresponding to the i -th text box extracted from the document), denoted by $logit_i^{(KV)}$, is the input of the structural feature encoding, and

the corresponding output is node embedding vector E_i^{node} representing the knowledge of the IE networks on its final decision (node classification): $E_i^{node} = f_{SF}(logit_i^{(KV)})$.

Feature Fusion. The Feature Fusion network f_{Fusion} takes the three outputs from the Feature Encoding module and produces the ultimate feature vector. We use simple concatenation and Bi-linear pooling [38] as two options for Feature Fusion. Bi-linear pooling use outer-product to combine inputs of different modalities. For simple concatenation, we just concatenate three vectors. For Bi-linear Pooling, we first pool the pair of E_i^{vis} and E_i^{OCR} , and then pool the resulting vector with E_i^{node} to get the pooled output F_i :

$$F_i = f_{Fusion}(E_i^{vis}, E_i^{OCR}, E_i^{node}) \quad (1)$$

3.2 Variational Cluster-oriented Anomaly Detector (VCAD)

The anomaly detector aims to detect which input image is normal or abnormal, thus bolsters the MCP by a measurement of the normality that the input has. Specifically, the input to the anomaly detector is a compressed representation of the document image, and the output is a score in the range $[0, 1]$ indicating the level of anomaly of the input. This score serves as an additional input to the confidence estimator.

Representing image data with cluster-oriented embeddings. In this section, we describe the representation learning of document images. Firstly, the training dataset was classified into some categories based on the appearance and the layout structure of the document image. Then we train a CNN-based image encoder to map each document image into a lower-dimensional vector representation. Here, the CNN architecture is MobileNet [31]. We adopt the triplet loss [13, 32] to learn the compressed representation, wherein the embeddings of images from the same category tend to form a cluster in the embedding space.

Anomaly detector training. After constructing embeddings for training images, we build a Variational Auto Encoder (VAE) [15, 30] as our anomaly detector (Figure 1(b)). The VAE outlier detector is first trained on a set of normal (inlier) data to reconstruct the input it receives, with the standard VAE loss function which is the sum of KL term and reconstruction loss:

$$\mathcal{L}_{VAE}(x; \theta, \phi) = -KL(q_\phi(z|x)||p_\theta(z)) + \frac{1}{L} \sum_{l=1}^L \log p_\theta(x|z^{(l)}) \quad (2)$$

where x , z and L denote the VAE's input, latent variable, and number of samples, respectively. q_ϕ represents the encoder and p_θ the decoder of VAE.

If the input data cannot be reconstructed well, the reconstruction error (implemented as L1 loss between VAE's input and output) is high and the data can be flagged as an outlier. We apply the min-max normalization [1] to the reconstruction losses in order to get the corresponding abnormal scores in the range of $[0, 1]$.

3.3 Hybrid confidence estimation

After getting the scalar output from our VCAD, we simply concatenate this scalar with the output of the Feature Fusion module in the MCP. The resulting vector is fed to a confidence estimator (CE),

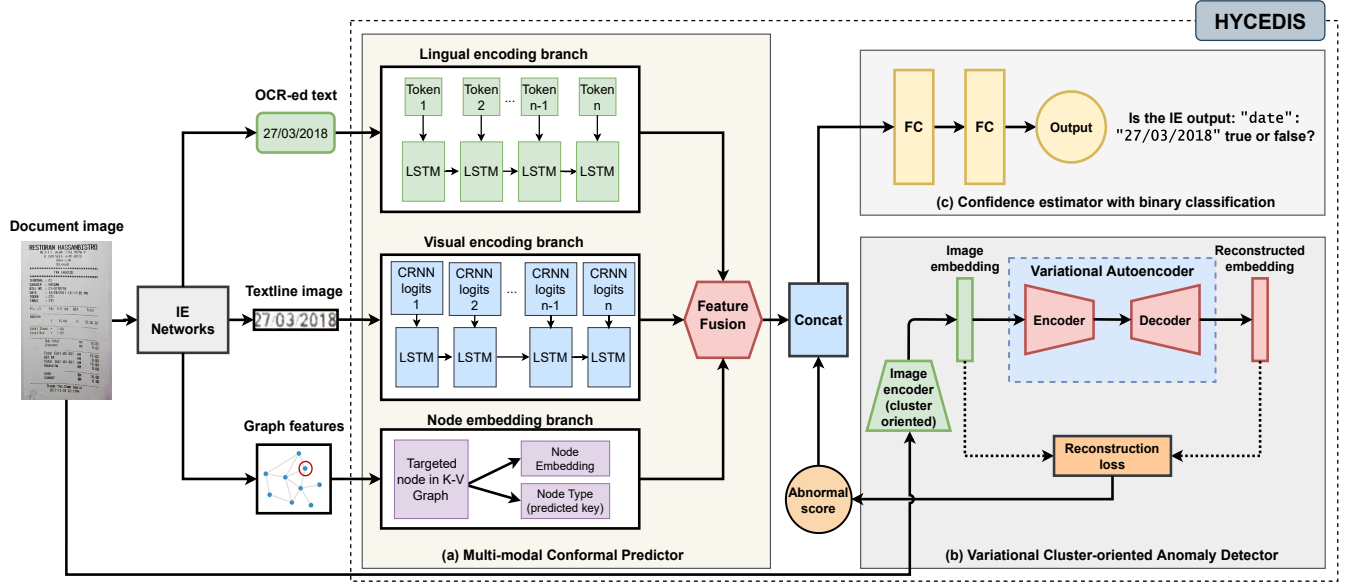


Figure 1: HYCEDIS architecture. (a) The Multi-modal Conformal Predictor (MCP). (b) The Variational Cluster-oriented Anomaly Detector (VCAD). (c) Confidence estimator (CE). MCP’s output vector, plus the VCAD’s abnormal score, is fed to fully-connected layers to produce the final output of HYCEDIS, indicating whether the extracted field true or false.

which is implemented as a 2-layer feed-forward neural network. We freeze the VCAD and train the CE and MCP end-to-end on the training data as the set of IE’s predictions in the training dataset.

In particular, let x_i denote the input document image, the function $IE(\cdot)$ denote our pipeline of IE networks. The output of IE system is $\hat{v}_i = IE(x_i)$. More specifically, $\hat{v}_i = \{\hat{v}_{ik}\}_{k=1:K_i}$ is the set of K_i predictions where each \hat{v}_{ik} contains location information along with extracted text corresponding to a particular key (e.g: $\{ \text{'location': [123, 234, 184, 246]}, \text{'text': '27/03/2018'}, \text{'key': 'date'} \}$). We also have the ground truth $v_i = \{v_{ij}\}_{j=1:J_i}$ is the set of J_i elements presented in the i -th document.

Let F_{ik} denote the input of our CE corresponding to prediction \hat{v}_{ik} , the CE is represented by the function $f_{CE}(\cdot)$ yielding the softmax output $p_{ik} = f_{CE}(F_{ik})$. The label for confidence estimation task is

$$y_{ik} = \mathbb{1}\{\exists j \in \{1 : J_i\} \mid v_{ij} \stackrel{\text{match}}{=} \hat{v}_{ik}\} \quad (3)$$

The IE’s output is considered to match the ground truth element if both the text contents and the keys match and the locations’ IoU is greater than a threshold (0.3 in this paper). y_{ik} is 1 if IE’s prediction matches a ground truth element (be correct) and vice versa. Then the loss function is the standard binary cross-entropy loss with label y_{ik} and probability p_{ik} .

4 EXPERIMENTS

4.1 Datasets and evaluation metrics

4.1.1 Datasets. We collect 4 Invoice-like datasets and divide them into 2 tasks, corresponding to English and Japanese language used in the data. For each task, we use the bigger dataset as the main one, and the smaller as the out-of-distribution (OOD) dataset with respect to the main dataset.

We first use pre-trained IE Networks (see Sec. 2) to generate the intermediate features for the MCP as mentioned in Sec. 3.1. The outputs of the IE Networks and the ground-truth IE outputs are used to produce labels for the confidence estimation task (Sec. 3.3).

We only train the confidence models on the training dataset and benchmark them on the testing and corresponding OOD datasets. The evaluation on OOD data is a challenging benchmark since the OOD dataset is totally different from the main one in terms of layout, background and writing styles. Moreover, since the OOD datasets can have different type of keys from those in the main one, we only test the models on fields that share common keys with the main dataset.

(a) Public datasets (*English*)

SROIE - Main dataset. SROIE [14] is a dataset of scanned receipts. There are 4 keys: *address*, *company*, *date*, *total*. The training set has 626 files corresponding to 3859 IE’s output key-value fields. We further hold 10% of the training as the validation set. The statistics for the test set are 341 files and 1,640 fields, respectively.

CORD - OOD dataset. CORD [28] contains receipts collected from Indonesian shops and restaurants. Compared to SROIE, CORD document images are captured in the wild, thus the data is noisy and low in quality. CORD field shares only one key with SROIE, which is *total*. We use the CORD-dev set which contains 100 files corresponding to 103 IE’s output fields.

(b) In-house datasets (*Japanese*)

In-house 1 - Main dataset. In-house 1 is a dataset containing Japanese invoice documents collected from several vendors. There are 25 keys. Example keys are *issued_date*, *total_amount*, *tax*, *item_name*, *item_amount*. The training set

has 835 files corresponding to 24,697 IE’s output fields, and the test set has 338 files and 10,898 fields.

In-house 2 - OOD dataset. In-house 2 consists of 68 invoice documents from another Japanese company. The document pattern is quite different to the In-house 1 dataset. The two in-house dataset share 4 key types in common, resulting in 3,887 IE’s output fields.

4.1.2 Evaluation metrics. We use the popular *Area Under the Receiver Operating Characteristic Curve (AUC)* [2, 11, 24, 25] and *Expected Calibration Error (ECE)* [26] metrics for measuring the performance of confidence predictors.

4.2 Experimental baselines

Softmax Threshold. Our IE pipeline consists of multiple sequential models, so we adapted [12] by combining both softmax probabilities from OCR and KV models using multiplication (i.e: $p_{final} = p_{OCR} * p_{KV}$). We then specify a threshold score and considered examples with higher-than-threshold softmax probability as correctly predicted one, and vice versa. The threshold score is tuned on the training dataset.

Temperature Scaling. Temperature scaling [8] is a technique that post-processes the neural networks to make them calibrated in term of confidence. Temperature scaling divides the logits (inputs to the softmax function) by a learned scalar parameter T (temperature). We learn this parameter on a validation set, where T is chosen to minimize negative log-likelihood.

Softmax Classifier. Instead of only utilizing the softmax probability of the predicted class as Softmax Threshold, Softmax Classifier is more advanced by making use of the whole softmax vector. Particularly, we build a simple classifier using a feed-forward neural network. The input for the network is the concatenation of the OCR model’s softmax vector and the KV model’s one.

Monte Carlo Dropout. MC Dropout [6] belongs to the class of Bayesian/variational approaches. By keeping the dropout enabled at test time, we can obtain the variance of the neural network’s outputs, and this variance indicates the level of uncertainty. We apply MC Dropout on our KV model, which is the final model in the pipeline.

4.3 Benchmarking results

4.3.1 Ablation study. We ablate the effect of VCAD and MCP on the whole hybrid system. Table 1 reports the results on SROIE dataset. Without VCAD, the proposed model achieves best AUC score of 86.90% using bi-linear pooling fusion strategy. Simpler concatenation method underperforms by about 3% demonstrating the importance of using outer-product to retain bit-level relationships among 3 modalities. When the VCAD is integrated, it consistently improves the performance of all fusion methods. Hence, the full hybrid HYCEDIS architecture can reach 88.12% AUC. Similar behaviors can be found with measurement using ECE metric.

4.3.2 Public English datasets result. Table 2 shows the performance of all models on public datasets. On both SROIE and its OOD CORD dataset, our full HYCEDIS is consistently the best performer regarding both ECE and AUC scores. Our MCP is the runner-up under ECE metric. The improvements of MCP in AUC and ECE suggests that the signals from intermediate features extracted from text-line

Methods	ECE	AUC
MCP (concatenation)	0.1525	83.75
MCP (bilinear pooling)	0.1175	86.90
MCP (concatenation) + VCAD	0.1385	84.37
MCP (bilinear pooling) + VCAD	0.1002	88.12

Table 1: Ablation study on SROIE dataset

Methods	SROIE		CORD	
	ECE	AUC	ECE	AUC
Softmax threshold	0.1525	83.75	0.1731	66.91
Softmax classifier	0.1400	85.50	0.3289	54.91
MC Dropout	0.1175	86.90	0.5446	43.52
Temperature scaling	0.1385	84.37	0.3787	74.58
MCP	0.1124	86.40	0.1432	75.12
HYCEDIS	0.1002	88.12	0.1259	77.45

Table 2: Performance comparison of baselines and proposed methods on SROIE and CORD datasets

Methods	In-house 1		In-house 2	
	ECE	AUC	ECE	AUC
Softmax threshold	0.1285	68.79	0.5885	53.38
Softmax classifier	0.2810	71.43	0.3945	51.22
MC Dropout	0.3733	66.14	0.3621	48.20
Temperature scaling	0.1728	64.00	0.5879	58.18
MCP	0.0782	86.32	0.3348	60.12
HYCEDIS	0.0712	90.12	0.3019	61.90

Table 3: Performance comparison of baselines and proposed methods on In-house datasets

images, OCR-ed text and graph structure help improve the accuracy of the softmax-based methods which only rely on some softmax layers of the IE Networks. In addition, when combined with VCAD, the AUC score is further increased and the ECE also downgrades. That manifests the contribution of our VCAD model. We can see a significant performance drop from baselines such as MC-Dropout when being tested on OOD CORD data. Our methods alleviate this issue, maintaining a moderate generalization to strange data.

4.3.3 In-house Japanese datasets result. We also benchmark the models on two in-house datasets. In Table 3, our model continues to show the superior performance compared with other baselines. Our MCP model improves about 14.89% and 2% AUC score and reduces 0.0503 and 0.0273 ECE score in In-house 1 and In-house 2 datasets, respectively. When adding VCAD, the performance is improved around 3.82% on In-house 1 dataset and 2.78% on In-house dataset, which again validates our hypothesis on using anomaly detector to enhance conformal predictor.

5 CONCLUSION

We have introduced a holistic confidence score architecture that aims to verify the result of IE Networks in document understanding tasks. Our architecture takes advantages of a Multi-modal Conformal Predictor and a Variational Cluster-oriented Anomaly Detector to predict whether the IE Networks’ output correct or not using features of different granularity. Our hybrid approach surpasses prior confidence estimation methods by a huge margin in benchmarks with invoice datasets. Remarkably, it demonstrates a capability of generalization to out-of-distribution datasets.

REFERENCES

- [1] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Asian conference on computer vision*, pages 622–637. Springer, 2018.
- [2] Murat Seckin Ayhan and Philipp Berens. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. 2018.
- [3] Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoon Yun, and Hwalsuk Lee. Character region awareness for text detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9365–9374, 2019.
- [4] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Senior. Lip reading sentences in the wild. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3444–3453. IEEE, 2017.
- [5] Tuan Anh Nguyen Dang and Dat Nguyen Thanh. End-to-end information extraction by character-level embedding and multi-stage attentional u-net. In *BMVC*, page 96, 2019.
- [6] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [7] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376, 2006.
- [8] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. *arXiv preprint arXiv:1706.04599*, 2017.
- [9] Pan He, Weilin Huang, Tong He, Qile Zhu, Yu Qiao, and Xiaolin Li. Single shot text detector with regional attention. In *Proceedings of the IEEE international conference on computer vision*, pages 3047–3055, 2017.
- [10] Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. An end-to-end textspotter with explicit alignment and attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5020–5029, 2018.
- [11] Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 41–50, 2019.
- [12] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- [13] Gary B Huang, Andrew Kae, Carl Doersch, and Erik Learned-Miller. Bounding the probability of error for high precision optical character recognition. *The Journal of Machine Learning Research*, 13(1):363–387, 2012.
- [14] Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and CV Jawahar. Icdar2019 competition on scanned receipt ocr and information extraction. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1516–1520. IEEE, 2019.
- [15] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [16] Hung Le, Truyen Tran, and Svetha Venkatesh. Dual memory neural computer for asynchronous two-view sequential learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1637–1645, 2018.
- [17] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2231–2239, 2016.
- [18] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [19] Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, and Xiang Bai. Rotation-sensitive regression for oriented scene text detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5909–5918, 2018.
- [20] Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. Graph convolution for multimodal information extraction from visually rich documents. *arXiv preprint arXiv:1903.11279*, 2019.
- [21] Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. Graph Convolution for Multimodal Information Extraction from Visually Rich Documents. mar 2019.
- [22] Yuliang Liu and Lianwen Jin. Deep matching prior network: Toward tighter multi-oriented text detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1962–1969, 2017.
- [23] Shangbang Long, Jiaqiang Ruan, Wenjie Zhang, Xin He, Wenhao Wu, and Cong Yao. Textsnake: A flexible representation for detecting text of arbitrary shapes. In *Proceedings of the European conference on computer vision (ECCV)*, pages 20–36, 2018.
- [24] Amit Mandelbaum and Daphna Weinshall. Distance-based confidence score for neural network classifiers. *arXiv preprint arXiv:1709.09844*, 2017.
- [25] Noam Mor and Lior Wolf. Confidence prediction for lexicon-free ocr. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 218–225. IEEE, 2018.
- [26] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- [27] Duc Nguyen, Nhan Tran, and Hung Le. Improving long handwritten text line recognition with convolutional multi-way associative memory. *arXiv preprint arXiv:1911.01577*, 2019.
- [28] Seunghyun Park, Seung Shin, Bado Lee, Junyeop Lee, Jaeheung Surh, Minjoon Seo, and Hwalsuk Lee. Cord: A consolidated receipt dataset for post-ocr parsing. 2019.
- [29] Yujie Qian, Enrico Santus, Zhijing Jin, Jiang Guo, and Regina Barzilay. GraphIE: A Graph-Based Framework for Information Extraction. 2018.
- [30] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic back-propagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [31] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [32] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [33] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(11):2298–2304, 2016.
- [34] Luca Della Vedova, Hong Yang, and Garrick Orchard. An Invoice Reading System Using a Graph Convolutional Network. 2:434–449, 2019.
- [35] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1192–1200, 2020.
- [36] Xiao Yang, Ersin Yumer, Paul Asente, Mike Kraley, Daniel Kifer, and C Lee Giles. Learning to extract semantic structure from documents using multimodal fully convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5315–5324, 2017.
- [37] Liang Yao, Chengsheng Mao, and Yuan Luo. Graph convolutional networks for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7370–7377, 2019.
- [38] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 1821–1830, 2017.
- [39] Nan-ning Zheng, Zi-yi Liu, Peng-ju Ren, Yong-qiang Ma, Shi-tao Chen, Si-yu Yu, Jian-ru Xue, Ba-dong Chen, and Fei-yue Wang. Hybrid-augmented intelligence: collaboration and cognition. *Frontiers of Information Technology & Electronic Engineering*, 18(2):153–179, 2017.