

Self-Supervised Learning for Visual Summary Identification in Scientific Publications

Shintaro Yamamoto^a, Anne Lauscher^b, Simone Paolo Ponzetto^b, Goran Glavaš^b and Shigeo Morishima^c

^aDepartment of Pure and Applied Physics, Waseda University, Japan

^bData and Web Science Group, University of Mannheim, Germany

^cWaseda Research Institute for Science and Engineering, Japan

Abstract

Providing visual summaries of scientific publications can increase information access for readers and thereby help deal with the exponential growth in the number of scientific publications. Nonetheless, efforts in providing visual publication summaries have been few and far apart, primarily focusing on the biomedical domain. This is primarily because of the limited availability of annotated gold standards, which hampers the application of robust and high-performing supervised learning techniques. To address these problems we create a new benchmark dataset for selecting figures to serve as visual summaries of publications based on their abstracts, covering several domains in computer science. Moreover, we develop a self-supervised learning approach, based on heuristic matching of inline references to figures with figure captions. Experiments in both biomedical and computer science domains show that our model is able to outperform the state of the art despite being self-supervised and therefore not relying on any annotated training data.

Keywords

scientific publication mining, multimodal retrieval, visual summary identification

1. Introduction

Given the exponential growth in the number of scientific publications [1], providing concise summaries of scientific literature becomes increasingly important. Accordingly, previous work has focused on the automatic creation of textual summaries [2, 3, 4, 5, 6, 7]. However, specifically in the case of scientific publications (and especially in some domains), information is also conveyed in the form of figures, which allow the reader to understand the scientific contributions better, offering visual representations of data, experimental design, and results. Some scientific publishing companies (e.g., Elsevier) even require authors to submit a figure as a Graphical Abstract (GA), which is “a single, concise, pictorial and visual summary of the main findings of the article”¹. GAs, in turn, are then used to provide multi-modal online search results, following the observations that humans better remember and recall visual information [8].

BIR 2021: 11th International Workshop on Bibliometric-enhanced Information Retrieval at ECIR 2021, April 1, 2021, online

✉ s.yamamoto@fuji.waseda.jp (S. Yamamoto)



© 2021 Copyright for this paper by its authors.
Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://www.elsevier.com/authors/journal-authors/graphical-abstract>

Recently, Yang et al. [9] introduced the concept of the central figure, referring to the figure that is the best candidate for a GA of a paper. To build a dataset for automatically finding central figures in scientific publications, they asked authors of papers in PubMed² to identify one central figure in each of their scientific publications. Though a GA is not required for all publications, authors were shown to be able to identify the central figure from their publications for 87.6% of the papers. Using the obtained datasets of publications with GAs, they devised a supervised machine learning approach for identifying central figures of scientific publications. Such approaches for automatically identifying central figures can be employed to create GAs for large collections of scientific documents. As a result, researchers and students can profit from the visual support in concrete scenarios: for instance, as they obtain an impression of the discussed research at a glance without reading the text, more efficient analysis of online search results is possible. In this paper, we address two major limitations of Yang et al.'s seminal contribution. First, the dataset of Yang et al. consists of PubMed data only, limiting the applicability of the devised supervised central figure identification model to the biomedical domain. The use of figures in scientific literature, however, is a common practice in a much broader set of research fields and areas [10]. Secondly, while supervised learning is known to generally provide the best results, it critically depends on (sufficiently) large amounts of labeled data to be used for training the models: expensive and time-consuming data annotation processes impede the scalability of central figure identification across the plethora of research domains in which figures encode valuable information. Whereas in some tasks, labeled data can be acquired more economically with crowd-sourcing, this is not the case for the task at hand: identification of the central figure for a publication requires annotators to be knowledgeable in the publication domain. In other words, collecting datasets large enough to support supervised learning for central figure identification for a wide range of many domains is impractical (if not infeasible) due to the high annotation costs stemming from having to recruit expert annotators.

To alleviate these issues, we propose (1) a novel benchmark for central figure identification covering several subareas of computer science, and (2) a self-supervised learning approach for which we do not need any labeled training data. For our proposed benchmark for central figure identification, we ask two (semi-expert) annotators to rank the top three figures in a scientific paper that would be the best candidates for a graphical abstract. The papers are collected from four computer science subdomains: natural language processing (NLP), computer vision (CV), artificial intelligence (AI), and machine learning (ML). Accordingly, our newly collected dataset allows for a comparison of the performance of central figure identification models across diverse (sub)domains. Secondly, to eliminate the reliance on labeled training data, we introduce a self-supervised learning approach for automatic identification of central figures in scientific publications. The core idea of our approach is outlined as follows. In most scientific publications, a figure is mentioned in an article's body by using a direct reference (e.g., "*In Figure 3, we illustrate ...*"). This typically means that the paragraph of the direct link sentence roughly describes in text what the figure depicts visually, i.e., that the paragraph's content is clearly associated with the content of the figure. We exploit these direct links between an article's body and use these paragraph-figure pairs as training instances for a supervised central figure identification model. We then train several Transformer-based [11] models, which take pairs of

²<https://pubmed.ncbi.nlm.nih.gov/>

body text and figure captions as input, and we train the models to judge whether the paragraph text (from the body of the article) matches its paired figure. At inference (i.e., test) time, we rank the article’s figures by (1) predicting the scores for each article figure by pairing them all with the abstract and feeding them to the model and (2) ranking the figures based on the scores output by the model reflecting their degree of match with the article’s abstract. In contrast to sentence matching approaches [12, 13, 14, 15], which perform sentence-pair classification, we tackle a ranking problem, scoring and ordering all figures of an article given its abstract. Although self-supervised, our approach outperforms the existing fully supervised learning approach for central figure identification [9] in terms of top-1 accuracy. Finally, we provide an extensive analysis of performance differences across different domains.

2. Related Work

While the majority of related work in scientific paper summarization has focused on automatically creating textual summaries [2, 3, 4, 5, 6, 7], only a few have investigated the creation of visual summaries, i.e., selection of images that best reflect the publication content. Kuzi and Zhai [16] proposed Keyword-based figure retrieval: they tackle the related problem of ranking figures from multiple papers (in ACL anthology reference corpus [17]). The task that we tackle in this work differs in that we focus on selecting the best figure for a single publication, considering only the figures from that publication as candidates. Similarly, in [18, 19] the authors rank figures from a single paper based on their importance.

In this paper, we consider the problem of automatically identifying a central figure for a paper, which would then be a candidate for the paper’s visual summary, referred to as Graphical Abstract (GA) [9]. Several works have focused on analyzing GAs, e.g., their use [20] and design pattern [21].

The automatic selection of a central figure for scientific papers was first proposed by Yang et al. [9]. In their work, they built a dataset for the central figure identification from PubMed (biomedical and life science) papers. To extend the study of central figure identification, we propose a novel dataset consisting of computer science papers from several subdomains. Yang et al. proposed a supervised learning approach for central figure identification, a methodology that can hardly scale across a variety of scientific disciplines, due to the need for expert annotation of central figures. Limited sizes of existing datasets for various tasks in scientific publication mining [9, 22, 7, 23], additionally suggest that obtaining any kind of gold expert annotations on scientific text is expensive and time-consuming. To remedy for this bottleneck of annotation cost, we propose a self-supervised approach in which we make use of direct inline figure references in the article body to heuristically pair article paragraphs with figure captions and use those pairs as distant supervision.

The similarity between an abstract and a figure caption is the most important feature for the supervised model of [9]. Accordingly, we treat the task of identifying a central figure as an abstract-to-caption matching problem [13]. Approaches for sentence matching can be divided into two types: a sentence encoding-based approach and an attention-based approach. In the sentence encoding approach, sentences are encoded separately [12], which, in contrast to the attention-based approach [15, 14, 13], does not capture semantic interactions between them.

Table 1
Number of annotated papers per computer science sub-domain.

Domain	NLP	CV	AI	ML	Total
Conferences	ACL EMNLP	CVPR	AAAI IJCAI	ICML	
No. papers	148	158	147	144	597
Two annotators	126	127	120	123	496
Single annotator	22	31	27	21	101
<i>Figures / paper</i>					
Average	6.2±1.8	7.0±1.8	6.10±1.5	6.5±1.9	6.5±1.8
Minimum	5	5	5	5	5
Maximum	13	13	14	13	14

We employ an attention-based approach and build the model on top of pretrained Transformer networks [24, 25]. In contrast to current research on sentence matching as a classification task, we treat central figure identification as a ranking problem where all figures in a paper are scored according to their suitability to be used as a central figure.

3. Annotation Study

Data Collection. According to [10], the use of figures in scientific literature differs according to the field and the topic of research. To investigate fine-grained differences for automatic central figure identification across research domains, we collect papers published between 2017 and 2019 for four different research fields in computer science, namely natural language processing (NLP), computer vision (CV), artificial intelligence (AI) and machine learning (ML). In order to make the dataset sufficiently challenging, we keep only the publications with more than five figures. Table 1 provides the dataset statistics (number of publications and average number of figures per publication for each subdomain).

Annotation Process. Our annotation task is defined as follows: *given a paper abstract and the figures extracted from the paper, identify and rank the top 3 figures according to the degree to which they match the abstract and can therefore serve as a visual summary.* In our annotation guidelines we adopt the definition of a graphical abstract (GA) as given in the Elsevier author guidelines (cf. footnote 1). Annotations were carried out by two coders with a university degree in computer science, who were instructed to study the examples provided on the publisher page and discuss them in a group to make sure they understood the notion of a graphical abstract.

To facilitate the annotation process, we develop a web-based annotation tool with a graphical user interface displaying a paper abstract and all figures extracted from the same paper, which are randomly shuffled to avoid the bias induced by the order. We first asked our annotators to read the abstract in order to obtain an overview of the paper and then to study each figure carefully. Next, the annotators were asked to choose and rank the top 3 GA candidates. All instances are either doubly or singly annotated, and the inter-annotator agreement across the doubly annotated data amounts to .43 Krippendorff’s α (ordinal), which reflects the difficulty

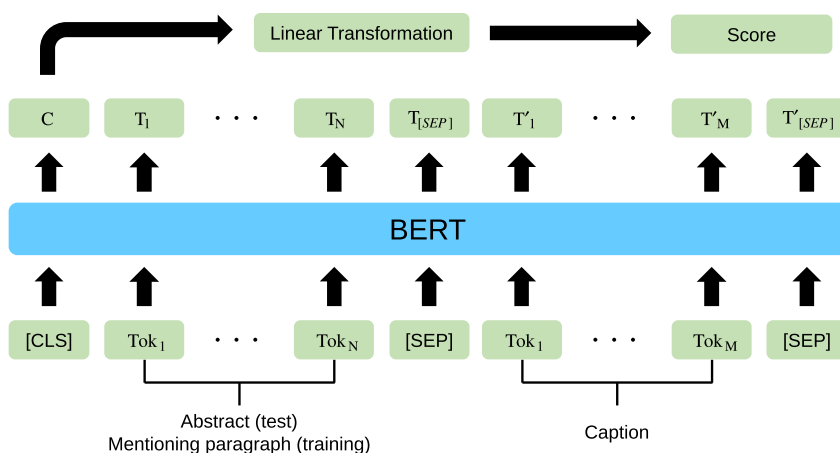


Figure 1: Our model for abstract-caption pair scoring. Paragraphs explicitly mentioning figures are paired with the figure captions during training.

and the subjective nature of the task.

4. Methodology

Problem Definition. Central figure identification can be defined in two different ways [9], namely figure-level and the paper-level: in the figure-level setting, each individual figure is classified as being a central figure or not, while in the paper-level setting, a central figure is determined from all figures in a paper. In this work, we are primarily interested in retrieving GAs as a form of summarization: hence, we opt for a document-level approach and cast it a ranking problem in which all figures from a paper are to be scored based on their suitability to provide a central figure for the publication.

Building on the result from [9] that the similarity between an abstract and a figure caption is the most important factor for central figure identification, we use a pair of abstract and figure caption as input. Given the sets of figures extracted from a paper $X = \{x_i : i\}$ and an abstract y , we learn a scoring function $f(x, y)$ that predicts the appropriateness of the figure to act as the central figure for the abstract (and accordingly, the paper). All figures are then ranked according to the model’s prediction $S = \{s_i : s_i = f(x_i, y)\}$.

Model. Our model consists of two components, a Transformer [11] as a language encoder and a score predictor (Figure 1).

We build upon finding from recent work in NLP that has shown the benefits of an attention-based approach for sentence matching [14, 15, 13], and accordingly opt for a pre-trained BERT [25] model as the text encoder: specifically, we use in our experiments a SciBERT model [24], which is pretrained on scientific publications from Semantic Scholar [26]. We provide

the pair (x_i, d_j^k) as positive if $i = k$, while $i \neq k$ for a negative one.

Optimization. We train the model to rank the positive pairs higher than negative ones. Due to BERT’s input token sequence length restriction, we randomly sample one sentence from the paragraph. The pair of a sampled sentence and a caption is fed into the model as "[CLS] sentence [SEP] caption [SEP]". For the training objective, we formulate the following loss similar to the Triplet loss [27] as $L = \max(s_p - s_n + \alpha, 0)$, where s_p and s_n denote the predictions of the model for the positive and negative pairs, respectively. In the experiments, we set $\alpha = 1.0$. For a single training instance, we sample one positive and one negative pairs including the figure x_i as (x_i, d_j^i) and $(x_i, d_k^{i'})$ ($i \neq i'$), respectively. The training objective makes the score for a positive pair lower than that for a negative one: therefore, the figure with the lower score is ranked higher.

5. Experiments

5.1. Implementation Details

We conduct our experiments using BERT’s implementation from the Hugging Face library [28]. In all fine-tuning procedures, we use the Adam optimizer [29] with the learning rate $1e-6$, train in batches of size 32, and apply a dropout at the rate of 0.2 and a gradient clipping threshold of 5. We train the model for 1 epoch. To extract text from collected PDF versions of papers, we rely on the Science Parse library⁴. Explicit inline references of figure are identified via the keywords "Figure" or "Fig.". To extract figure captions, we employ the image-based approach from [30].

5.2. Experimental Setting

PubMed. In [9], 7,295 biomedical and life science papers from PubMed are annotated for central figure identification. We managed to obtain the PDFs from PubMed for 7,113 of those papers and divide the papers into training, validation, and test portions in the same ratio as Yang et al. (8:1:1). Using our figure mention heuristic, we create 40k paragraph-figure pairs from the training portion of the dataset. As only a single figure is annotated as the central figure in the PubMed dataset, we use the top-1 and top-3 accuracy as evaluation metrics following [9].

Computer Science (CS). We additionally evaluate our model on our new CS dataset (Section 3). Unlike the PubMed dataset, in which only a single figure is annotated as central, our annotators ranked three figures for each CS paper. Consequently, we use Mean Average Precision (MAP), Mean Reciprocal Rank (MRR), and normalized Discounted Cumulative Gain (nDCG) as our evaluation metrics on the CS dataset. For the optimization procedure, we collect papers from the same subdomains as the annotated test data, from between 2015 and 2018. We divide them into training and validation portions with ratio of 90% and 10%, respectively. Here we also obtain around 40k paragraph-figure instances for model training. For performance evaluation, we utilize our annotated data described in Section 3.

⁴<https://github.com/allenai/science-parse>

Table 2
Performance evaluation on the PubMed dataset [9].

Method	Model	Accuracy@1	Accuracy@3
Baseline	Random	0.280	0.701
	Pick first	0.301	0.733
Yang et al. [9]	Text-only	0.333	0.810
	Full	0.344	0.793
Ours	Vanilla BERT	0.331	0.770
	RoBERTa	0.347	0.741
	SciBERT	0.383	0.787

5.3. Experiments

Performance on the PubMed dataset. We first evaluate our self-supervised approach using the PubMed dataset (Table 2). We follow [9] and make use of two baselines: random and ‘select first image’. The random baseline ranks figures randomly and the ‘select first image’ ranks based on the order of the figures as they appear in the paper (i.e., figure 1 is ranked 1st). For comparison, we also provide the results of two methods from [9], a text-only model that uses cosine similarity of TF-IDF between the abstract and the figure caption as the input feature, and a full model that takes the figure type label (e.g., diagram, plot) and layout (e.g., section index, figure order) as inputs, as well as text features. We compare these against the performance of our models based on three different pretrained Transformers, namely vanilla BERT [25], RoBERTa [31] and SciBERT [24].

Regardless of the text encoder, our approach outperforms the baselines in terms of both top-1 and top-3 accuracy. This result indicates that our method for generating training data creation is effective for central figure identification. Among the text encoders, SciBERT performs the best for both metrics, arguably because it has been trained on a corpus of scientific papers, thus minimizing problems related to domain transfer. Despite not requiring manual annotation for training, our approach with SciBERT also outperforms the supervised approach of Yang et al. [9] in terms of top-1 accuracy.

Performance on the CS dataset. We also evaluate the performance of the model on our CS dataset (Table 3). We follow the same setting as for the PubMed data and use a random and ‘choose first image’ methods as baselines. Here, we compare SciBERT with vanilla BERT and RoBERTa, so as to additionally verify the effectiveness of SciBERT in the CS domain – since over 80% of the papers in the corpus for SciBERT pre-training are from the biomedical domain and the ratio of CS papers account for only 18% of SciBERT’s pretraining corpus [24].

Our approach outperforms the random baseline in terms of MAP, MRR, and nDCG. Among base models, SciBERT outperforms vanilla BERT and RoBERTa as in PubMed papers. Though most of the corpus for SciBERT pre-training is from the biomedical domain, a certain number of CS papers seen in pretraining still contributes to the downstream performance on central figure identification.

However, as opposed to the case of PubMed papers, the ‘pick first’ baseline here is much stronger and hard-to-beat, even for our Transformer-based approach. Note that in the proposed

Table 3
Performance evaluation on CS papers.

Method	Model	MAP	MRR	nDCG
Baseline	Random	0.616	0.693	0.732
	Pick first	0.754	0.827	0.809
Ours	Vanilla BERT	0.694	0.773	0.767
	RoBERTa	0.702	0.793	0.775
	SciBERT	0.731	0.822	0.794

Table 4
Performance comparison of training on different research fields.

(a) Test on PubMed papers.			(b) Test on CS papers.			
Training data	ACC@1	ACC@3	Training data	MAP	MRR	nDCG
– (Random b.)	0.280	0.701	– (Random b.)	0.616	0.693	0.732
PubMed	0.383	0.787	PubMed	0.728	0.822	0.789
CS	0.368	0.777	CS	0.731	0.822	0.794

self-supervised learning approach, the order of the figures is not taken into account. Consequently, our models do not consider the order in which the figures appear. This result indicates that CS papers tend to use Graphical Abstract (GA) in the beginning, and empirically highlights that our new dataset is more challenging than the PubMed-based dataset, as the ‘pick first’ baseline is hard to beat.

Cross-domain Experiments. Image usage in scientific publications is known to be different across scientific fields [10]. Accordingly, we set next to empirically evaluate the robustness of our approach in a domain transfer setup. Due to the different granularity of our PubMed and CS datasets, the latter including papers from four different research areas of computer science (AI, NLP, ML, and CV), we are able to perform two sets of domain transfer experiments, namely biomedical vs. computer science as well as across different CS subdomains.

We first compare model performance by training and testing on datasets from different domains – i.e., biomedical papers from PubMed vs. computer science publications from our CS dataset – using SciBERT as a base model (Table 4).

In the test with PubMed papers, training on the same domain performs better both in terms of top-1 and top-3 accuracy. Despite the slightly lower performance, training on the CS domain also outperforms the random baseline. On the other hand, training on different domains, somewhat surprisingly, does not degrade the performance. This would imply that papers from different domains exhibit similar text-figure (caption) matching properties.

Next, we examine the results of domain transfer for different CS subdomains dataset belonging to different areas of computer science, due to the fact that image usage and volume may potentially vary among fields like, e.g., natural language processing and computer vision, with papers from the latter containing typically more images. We train models on four different areas (NLP, CV, AI, and ML) and test them on all others. Domain comparison within several research topics in CS is summarized in Table 5.

Table 5

Performance comparison of the model trained on papers from different research topics in CS.

(a) MAP						
Test data	Training data				Baseline	
	NLP	CV	AI	ML	Random	Pick first
NLP	0.727	0.728	0.728	0.730	0.631	0.751
CV	0.716	0.721	0.716	0.719	0.585	0.758
AI	0.727	0.729	0.728	0.730	0.637	0.776
ML	0.676	0.682	0.679	0.681	0.617	0.732

(b) MRR						
Test data	Training data				Baseline	
	NLP	CV	AI	ML	Random	Pick first
NLP	0.791	0.795	0.795	0.798	0.705	0.816
CV	0.826	0.833	0.826	0.831	0.664	0.831
AI	0.828	0.834	0.830	0.828	0.711	0.847
ML	0.759	0.769	0.763	0.769	0.686	0.814

(c) nDCG						
Test data	Training data				Baseline	
	NLP	CV	AI	ML	Random	Pick first
NLP	0.777	0.778	0.776	0.779	0.743	0.817
CV	0.785	0.790	0.785	0.787	0.708	0.803
AI	0.799	0.802	0.800	0.802	0.763	0.828
ML	0.762	0.763	0.760	0.761	0.745	0.791

Overall, the results are rather consistent across areas and indicate that, within computer science, the research topics of papers do not affect the model performance. Among the four topics, the performance is the lowest on machine learning (ML) papers. We therefore manually analyzed ML papers with poor model performance and observed that these papers tend to have figures that look rather similar. This makes the identification of the central figure – even with manual effort – difficult. The ‘pick first’ scores higher for CV papers than for NLP and ML papers, whereas the random baseline naturally performs worse on CV papers, which contain more figures.

Model Analysis. To understand the model behavior, we analyze the attention in SciBERT. We visualize the attention in the Transformer model. We find that most attention maps are consistent with typical classes reported in [32], such as vertical or diagonal attention patterns. In some attention heads, the model attends to the lexical overlap between abstract and caption. The examples of attention matrices produced by heads attending over the same or semantically similar tokens, are shown in Figure 3. In this example, instances of tokens like ‘tracking’ and ‘when’ appearing in both in abstract and caption have mutually high attention weights. Additionally, pairs of tokens with similar meaning like ‘restore’ and ‘recovering’ also receive high mutual attention weights.

We also compare attention patterns among the model trained with different topics (NLP, CV,




Figure 3: Examples of attentions from SciBERT which attends to semantically similar tokens (training and test data are from the CS domain).

Table 6

Cosine similarities of attention weight maps obtained from models trained for different CS domains. Attention maps from all layers are flattened into a single vector.

	NLP	CV	AI
CV	0.9997	0.9998	0.9998
AI	0.9998	0.9998	-
ML	0.9997	-	-

AI, and ML). Following [32], we calculate the cosine similarity of attention maps. We show the mean cosine similarity of flattened attention map for randomly selected 100 samples in Table 6. Cosine similarity is high for all combinations; this means that the attention patterns across the different CS domains are virtually identical, confirming empirically our previous assumption that there are no relevant differences between CS domains when it comes to text-figure matching (see Table 5).

Kovaleva et al. reported that after fine-tuning attention maps change the most in the last two transformer layers [32]. We therefore analyze the change in attention patterns after our task-specific fine-tuning. We compare the standard fine-tuning, in which we update all SciBERT’s parameters (and which we used in all our previous experiments), and the feature-based training, in which we freeze SciBERT’s parameters and train only the regressor’s parameters. The comparison of attention patterns between fine-tuned and frozen SciBERT is summarized in Table 7. On the one hand, if we freeze SciBERT’s parameters, we observe a major drop in performance (6 MAP points). On the other hand, high cosine similarity of attention maps between the fine-tuned and frozen SciBERT that the two transformers still exhibit similar attention patterns. This suggests that only slight changes in the parameters of Transformer’s attention heads have the potential to substantially change the predictions of the regressor.

Table 7

Comparison between fine-tuned and frozen SciBERT trained on CS papers.

(a) Model performance.

SciBERT	MAP	MRR	nDCG
fine-tune	0.733	0.827	0.794
freeze	0.677	0.752	0.754

(b) Cosine similarity of attention map for randomly sampled 100 sentence-caption pairs in each layer.

Layer	1	2	3	4	5	6
Similarity	0.9999	0.9993	0.9983	0.9982	0.9983	0.9980
Layer	7	8	9	10	11	12
Similarity	0.9977	0.9971	0.9964	0.9951	0.9948	0.9945

6. Conclusion

While research efforts have been mostly spent on increasing information access to scientific literature by creating textual summaries, it is known that a large amount of information is often conveyed visually in the form of figures. In this work, we have addressed the problem of central figure identification from scientific publications, the task of identifying a candidate for a visual summary. Starting from previous work, which has introduced a dataset enabling supervised learning for central figure identification, we identified and addressed two main issues: (1) the only existing data set is limited to the biomedical domain, and (2) large-scale annotations for new domains are impractical and costly. To alleviate these issues, we first presented a new benchmark collection of scientific publications annotated for central figures in the computer science domain covering four different subfields. Secondly, we proposed a self-supervised approach to central figure identification. Our method exploits the link between portions of text explicitly referencing figures and figure captions, thereby bypassing the need for large manually annotated training data. We have experimentally demonstrated the effectiveness of our approach, outperforming the supervised approach in terms of rank-1 accuracy. Finally, a follow-up analysis of cross-domain performance differences and models' attention scores revealed only slight differences across the individual CS subdomains, but interestingly, our findings also indicate that the positioning of the central figure differs between the CS and the biomedical domain. We hope that our results fuel further research on cost-effective visual summary creation for increased information access in light of the exponentially growing body of scientific literature.

Acknowledgments

This work was supported by the Program for Leading Graduate Schools, "Graduate Program for Embodiment Informatics" of the Ministry of Education, Culture, Sports, Science and Technology (MEXT) of Japan, and JST ACCEL (JPMJAC1602). Computational resource of AI Bridging Cloud Infrastructure (ABCI) provided by National Institute of Advanced Industrial Science and Technology (AIST) was used.

References

- [1] L. Bornmann, R. Mutz, Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references, *Journal of the Association for Information Science and Technology* 66 (2015) 2215–2222.
- [2] A. Cohan, F. Deroncourt, D. S. Kim, T. Bui, S. Kim, W. Chang, N. Goharian, A discourse-aware attention model for abstractive summarization of long documents, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 2018, pp. 615–621.
- [3] A. Cohan, N. Goharian, Scientific article summarization using citation-context and article’s discourse structure, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 390–400.
- [4] Q. Mei, C. Zhai, Generating impact-based summaries for scientific literature, in: *Proceedings of 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 2008, pp. 816–824.
- [5] V. Qazvinian, D. R. Radev, Scientific paper summarization using citation summary networks, in: *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1*, 2008, pp. 689–696.
- [6] A. Lauscher, G. Glavaš, K. Eckert, University of mannheim@ clscisumm-17: Citation-based summarization of scientific articles using semantic textual similarity, in: *CEUR workshop proceedings, volume 2002*, RWTH, 2017, pp. 33–42.
- [7] M. Yasunaga, J. Kasai, R. Zhang, A. R. Fabbri, I. Li, D. Friedman, D. R. Radev, Scisummnet: A large annotated corpus and content-impact models for scientific paper summarization with citation networks, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2019, pp. 7386–7393.
- [8] D. L. Nelson, V. S. Reed, J. R. Walling, Pictorial superiority effect., *Journal of experimental psychology: Human learning and memory* 2 (1976) 523–528.
- [9] S. T. Yang, P.-S. Lee, L. Kazakova, A. Joshi, B. M. Oh, J. D. West, B. Howe, Identifying the central figure of a scientific paper, in: *2019 International Conference on Document Analysis and Recognition (ICDAR)*, 2019, pp. 1063–1070.
- [10] P.-S. Lee, J. D. West, B. Howe, Viziometrics: Analyzing visual information in the scientific literature, *IEEE Transactions on Big Data* 4 (2018) 117–129.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, I. Polosukhin, Attention is all you need, in: *Advances in Neural Information Processing Systems* 30, 2017, pp. 5998–6008.
- [12] S. R. Bowman, G. Angeli, C. Potts, C. D. Manning, A large annotated corpus for learning natural language inference, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 632–642.
- [13] Z. Wang, W. Hamza, R. Florian, Bilateral multi-perspective matching for natural language sentences, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2017, pp. 4144–4150.
- [14] M. Liu, Y. Zhang, J. Xu, Y. Chen, Original semantics-oriented attention and deep fusion network for sentence matching, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on*

- Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 2652–2661.
- [15] C. Duan, L. Cui, X. Chen, F. Wei, C. Zhu, T. Zhao, Attention-fused deep matching network for natural language inference, in: Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI'18, 2018, pp. 4033–4040.
 - [16] S. Kuzi, C. Zhai, Figure retrieval from collections of research articles, in: European Conference on Information Retrieval, Springer, 2019, pp. 696–710.
 - [17] S. Bird, R. Dale, B. Dorr, B. Gibson, M. Joseph, M.-Y. Kan, D. Lee, B. Powley, D. Radev, Y. F. Tan, The ACL anthology reference corpus: A reference dataset for bibliographic research in computational linguistics, in: Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08), 2008, pp. 1755–1759.
 - [18] F. Liu, H. Yu, Learning to rank figures within a biomedical article, PLOS ONE 9 (2014) 1–14.
 - [19] H. Yu, F. Liu, B. P. Ramesh, Automatic figure ranking and user interfacing for intelligent figure search, PLOS ONE 5 (2010) 1–12.
 - [20] J. Yoon, E. Chung, An investigation on graphical abstracts use in scholarly articles, International Journal of Information Management 37 (2017) 1371–1379.
 - [21] J. Hullman, B. Bach, Picturing science: Design patterns in graphical abstracts, in: P. Chapman, G. Stapleton, A. Moktefi, S. Perez-Kriz, F. Bellucci (Eds.), Diagrammatic Representation and Inference, 2018, pp. 183–200.
 - [22] A. Lauscher, G. Glavaš, S. P. Ponzetto, K. Eckert, Investigating the role of argumentation in the rhetorical analysis of scientific publications with neural multi-task learning models, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, 2018, pp. 3326–3338.
 - [23] X. Hua, M. Nikolov, N. Badugu, L. Wang, Argument mining for understanding peer reviews, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), 2019, pp. 2131–2137.
 - [24] I. Beltagy, K. Lo, A. Cohan, SciBERT: A pretrained language model for scientific text, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 3615–3620.
 - [25] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), 2019, pp. 4171–4186.
 - [26] W. Ammar, D. Groeneveld, C. Bhagavatula, I. Beltagy, M. Crawford, D. Downey, J. Dunkelberger, A. Elgohary, S. Feldman, V. Ha, R. Kinney, S. Kohlmeier, K. Lo, T. Murray, H.-H. Ooi, M. Peters, J. Power, S. Skjonsberg, L. Wang, C. Wilhelm, Z. Yuan, M. van Zuylen, O. Etzioni, Construction of the literature graph in semantic scholar, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers), 2018, pp. 84–91.
 - [27] E. Hoffer, N. Ailon, Deep metric learning using triplet network, in: A. Feragen, M. Pelillo, M. Loog (Eds.), Similarity-Based Pattern Recognition, 2015, pp. 84–92.
 - [28] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf,

- M. Funtowicz, J. Brew, Huggingface's transformers: State-of-the-art natural language processing, ArXiv abs/1910.03771 (2019).
- [29] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- [30] N. Siegel, N. Lourie, R. Power, W. Ammar, Extracting scientific figures with distantly supervised neural networks, in: Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries, 2018, pp. 223–232.
- [31] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, ArXiv abs/1907.11692 (2019).
- [32] O. Kovaleva, A. Romanov, A. Rogers, A. Rumshisky, Revealing the dark secrets of BERT, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 4365–4374.