

2023

DeepPatent2: A Large-Scale Benchmarking Corpus for Technical Drawing Understanding

Kehinde Ajayi
Old Dominion University

Xin Wei
Old Dominion University

Martin Gryder
Old Dominion University

Winston Shields
Old Dominion University

Jian Wu
Old Dominion University

See next page for additional authors

Follow this and additional works at: https://digitalcommons.odu.edu/computerscience_fac_pubs



Part of the [Computer Sciences Commons](#), [Data Science Commons](#), and the [Law Librarianship Commons](#)

Original Publication Citation

Ajayi, K., Wei, X., Gryder, M., Shields, W., Wu, J., Jones, S. M., Kucer, M., & Oyen, D. (2023). DeepPatent2: A large-scale benchmarking corpus for technical drawing understanding. *Scientific Data*, *10*(1), 1-14, Article 772. <https://doi.org/10.1038/s41597-023-02653-7>

This Article is brought to you for free and open access by the Computer Science at ODU Digital Commons. It has been accepted for inclusion in Computer Science Faculty Publications by an authorized administrator of ODU Digital Commons. For more information, please contact digitalcommons@odu.edu.

Authors

Kehinde Ajayi, Xin Wei, Martin Gryder, Winston Shields, Jian Wu, Shawn M. Jones, Michal Kucer, and Diane Oyen



OPEN

DATA DESCRIPTOR

DeepPatent2: A Large-Scale Benchmarking Corpus for Technical Drawing Understanding

Kehinde Ajayi¹, Xin Wei¹, Martin Gryder¹, Winston Shields¹, Jian Wu¹, Shawn M. Jones², Michal Kucer² & Diane Oyen²

Recent advances in computer vision (CV) and natural language processing have been driven by exploiting big data on practical applications. However, these research fields are still limited by the sheer volume, versatility, and diversity of the available datasets. CV tasks, such as image captioning, which has primarily been carried out on natural images, still struggle to produce accurate and meaningful captions on sketched images often included in scientific and technical documents. The advancement of other tasks such as 3D reconstruction from 2D images requires larger datasets with multiple viewpoints. We introduce DeepPatent2, a large-scale dataset, providing more than 2.7 million technical drawings with 132,890 object names and 22,394 viewpoints extracted from 14 years of US design patent documents. We demonstrate the usefulness of DeepPatent2 with conceptual captioning. We further provide the potential usefulness of our dataset to facilitate other research areas such as 3D image reconstruction and image retrieval.

Background & Summary

Technical illustrations, sketches, and drawings are images constructed to convey information more straightforwardly to humans than using text alone^{1–3}. One goal of computer vision is to build models to understand information contained in these images. Specific tasks include recognizing objects and determining their attributes, capturing the relationships between sub-images, and understanding the context of objects⁴. Natural images such as those contained in MS COCO⁵ and ImageNet⁶ have been used to develop solutions to these tasks using deep neural networks. Different from natural images^{5,6}, technical drawings are a type of image frequently found in design patents. Although they usually do not contain many features of natural images such as various colors, gradient, and environmental detail, drawings often abstract away unnecessary distractions leaving the strokes, lines, and shading that are sufficiently detailed so that drawn objects and aspects are still recognizable by humans. Compared with natural images, technical drawings, are under-studied by the computer vision and information retrieval communities. Recently, several sketch-based datasets have been developed, e.g.^{7,8}. *QuickDraw* consists of 50 M images sketched by letting users draw designated objects within a short time⁹. Other sketch datasets are much smaller, including TU-Berlin¹⁰, Sketchy³, and sheepMarket¹¹. These free-hand sketches tend to have very few strokes and limited viewpoints. Therefore, they are not suitable for tasks that rely on drawing details, like understanding scientific and technical information. CLEF-IP 2011 provides two patent image datasets of 10k images of heterogeneous types, including flow charts and chemical structures, used for patent retrieval and classification into 9 classes¹². *ImageNet-Sketch* contains 50k images of 1000 classes, based on search results of text queries against Google Image¹³. The technical drawing data we introduce here is different from these datasets in that it contains enriched semantic information, including object names and multiple views of the same object. Our dataset also contains segmented figures. Compared with *QuickDraw*, our dataset contains more details for drawings, diverse object names (more than 132,000), and viewpoint information (Table 1). Recently, Kucer *et al.*¹⁴ introduced DEEPPATENT, a dataset of over 350 K images from design patents for image retrieval. However, the DEEPPATENT dataset does not include identification of the objects or descriptions of their viewpoints. In addition, compound technical drawings were not segmented to isolate individual figures.

In this paper, we introduce a new dataset, DEEPPATENT2, consisting of more than 2 million automatically segmented and tagged technical drawings from more than 300,000 design patents granted from 2007 to 2020.

¹Computer Science, Old Dominion University, Norfolk, Virginia, 23529, US. ²Los Alamos National Laboratory, Los Alamos, 87545, New Mexico, US. ✉e-mail: jlwu@odu.edu; doyen@lanl.gov

Dataset	Size (number of images)	#Categories	Captions	Viewpoints
ImageNet-1k ³⁸	2M photos	1,000	No	No
ImageNet-21k ⁶	14M photos	21,000	No	No
WebVision ¹⁵	2.4M photos	1,000	No	No
Flickr-8K ⁶²	8K photos	—	Yes	No
Flickr-30K ⁶³	30K photos	—	Yes	No
ImageNet-Sketch ¹³	50K drawings	1,000	No	No
TU-Berlin ¹⁰	20K sketches	250	No	No
QuickDraw ⁹	50M + sketches	345	No	No
Sketchy ³	75K sketches, 12K photos	125	No	No
Sheep 10K ¹¹	10K sheep sketches	1	No	No
Sketch Flickr15K ⁶⁴	330 sketches, 15K photos	33	No	No
CLEF-IP 2011 ¹²	10k diagrams	9	No	No
DeepPatent ¹⁴	350K + drawings	—	No	No
DeepPatent2	2M + drawings	132,890	Yes	Yes

Table 1. DEEPPATENT2 compared with major sketch-based and natural image datasets. #Categories means the number of object names.

Figure 1 shows an example of the figures segmented and tagged from a single patent. We use an approach reminiscent of a large-scale image collection for real-world images with real human tags; similar to WebVision¹⁵, which consisted of 2.4 million images crawled from the Web with associated metadata. It was demonstrated that the noisy web images like WebVision were sufficient for training a good deep learning model for visual recognition¹⁵. We propose a novel pipeline using natural language processing (NLP) models to extract object names and viewpoints from figure captions, using computer vision (CV) methods to segment compound figures, and aligning text information with corresponding figures. The pipeline is developed based on our solid results on label extraction¹⁶, visual descriptor extraction¹⁷, and image segmentation¹⁸. To demonstrate the usefulness of our dataset, we train baseline deep-learning models on conceptual captioning and demonstrate that the performance of the models benefits significantly from an increasing size of training data. We expect that DEEPPATENT2 will facilitate training robust neural models on other tasks such as 3D image reconstruction and image retrieval for technical drawings.

Methods: Building the DeepPatent2 Dataset

Overview. DEEPPATENT2 contains over 2 million technical drawings from the United States Patent and Trademark Office USPTO¹⁹ design patent documents published from 2007 to 2020. Our dataset extends the DEEPPATENT dataset¹⁴ in three aspects.

1. DEEPPATENT2 is more than 5 times larger than DEEPPATENT;
2. DEEPPATENT2 contains both *original* and *segmented* patent drawings;
3. The metadata of each drawing contains object name and viewpoint information, automatically extracted using a supervised sequence-tagging model with high accuracy.

The pipeline to create the dataset is illustrated in Fig. 2. The pipeline includes three major components, namely, data acquisition, text processing, and image processing. A patent document includes an XML file containing textual content and TIFF files containing associated images. One patent TIFF file may contain several patent figures, as shown in Fig. 1, which we refer to as a compound figure. In the following, we refer to a “figure file” as an image file containing a single or a compound figure. We refer to a “figure” as an individual figure (e.g., a figure with a unique label such as “Fig. 1”).

The text processing step aims at automatically tagging individual figures with text that a person would use to describe them. In addition to extracting the design category information included in the patent XML file, we extract human-readable object names from figure captions. One challenge to this goal is that although the XML document contains captions and inline references for individual figures, the document does not directly map figures to figure files, because a figure file may contain multiple figures. For example, Fig. 1 contains 7 individual figures in 4 figure files. To overcome this challenge, we develop modules that first segment compound figures, resolve figure labels, and link them to text processing results to associate captions to their respective figures. The final data includes JSON files containing metadata, automatically extracted descriptions (object names and viewpoints), and images of individual figures, as well as original figure files.

In the following subsections, we elaborate on each module in the pipeline. All computation was performed on a Dell server with 4 NVIDIA GTX 2080 Ti GPUs, 24 hyperthreaded Intel Xeon Silver 4116, 300GB RAM, and 7TB disk space. In addition, we applied AWS Rekognition’s DetectText service on all patent figure files. The total size of the dataset is approximately 380GB before compression.

According to USPTO, design patent drawings are illustrations of a manufactured object’s design, which includes detailed information about contours, shapes, material texture, properties, and proportions. Drawings and text of patent documents are public domain; and are created by writers, artists, and inventors with the knowledge that these works will become public domain²⁰.

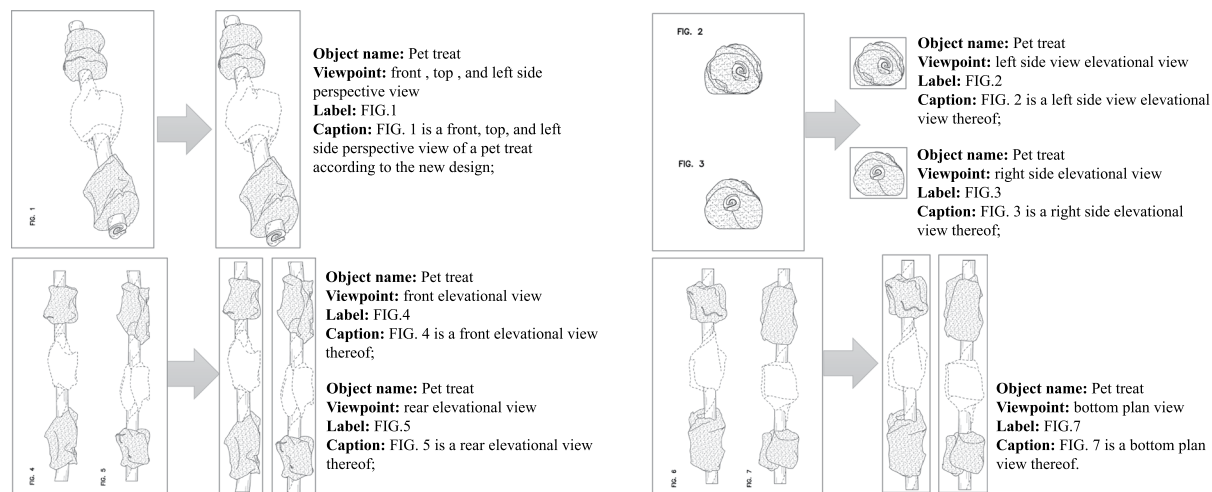


Fig. 1 An example from DEEPPATENT2 extraction results for US Design Patent #0836880.

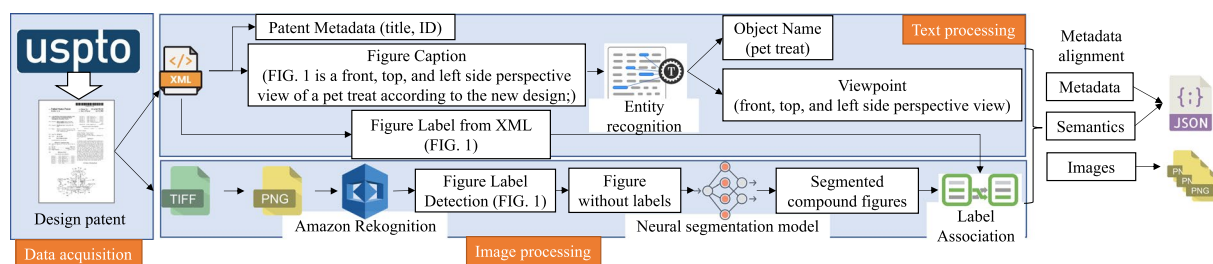


Fig. 2 The architecture of the pipeline to create DEEPPATENT2. Text in parentheses are examples. The matching module links semantic information parsed from captions to individual images segmented from original compound figures.

Patent Data Acquisition. We collected patent data ranging from the years 2007 to 2020 from the USPTO website in the form of zip and tar files. These files consist of both the full-text (in XML format) and figures (in TIFF format). Each TIFF file contains one or multiple figures. Table 2 shows statistics of the raw data. XML files before 2006 had a different schema and did not include separate XML tags for individual figures and subfigures, which may introduce further errors to figure label parsing and label-figure alignment, so we only include patents after 2007. The discrepancy between figure numbers parsed from XML (#Figure XML) and figure numbers from segmentation (#Figure Segmented) reflects the imperfection of the text and image processing methods, which will be incorporated in assessing the data quality. Figure 3 shows the average number of figure files per patent and the average number of figures per patent of our dataset, demonstrating the trends of the increasing popularity of using figures in patent applications over the last 14 years.

Text Processing: Entity Recognition. The text processing pipeline includes parsing plain text from XML documents and extracting semantic information from figure captions. In patent documents, figure captions are usually enclosed by special XML tags, so they can be accurately extracted. An inline reference is a sentence in the body text that cites one or multiple figures. We used regular expressions to match figure tags (e.g., “FIG.”, “FIGS.”, etc.) and then extracted complete sentences as inline references. Figure 4 shows an example of a figure caption and its corresponding inline reference in an XML file. Each individual figure has a caption, which we use to extract object names and viewpoints.

The detail of the text processing pipeline is elaborated in Wei *et al.*¹⁷. Here, we highlight the best model and its performance. We treated this task as an entity recognition problem. The text was first tokenized. Each token was encoded as a feature vector by a pre-trained word embedding model. Word vectors were then fed to a sequence-tagging neural network. We compared the BiLSTM-CRF (bidirectional long short-term memory or BiLSTM followed by a conditional random field or CRF) with a transformer model. Both were usually believed as the state-of-the-art sequence-tagging models²¹. We compared several word embedding models, including GloVe²², RoBERTa fine-tuned on GPT-2²³, the original RoBERTa²⁴, BERT²⁵, ALBERT²⁶, and DistilBERT²⁷. The sequence-tagging model incorporated context into the initial feature vector of each word. The CRF layer classified each token under the IOB schema. The best performance was achieved using DistilBERT²⁷ with the BiLSTM-CRF architecture. The average F1-measures for the overall entity recognition, object name, and viewpoint were 0.960, 0.927, and 0.992, respectively.

Year	#Patent	#Figure File	#Figure XML	#Figure Segmented	#Figure Mismatch	Mismatch%
2020	34,895	246,971	295,693	296,929	18,417	6.2%
2019	34,813	229,557	287,122	276,276	17,639	5.8%
2018	30,513	199,015	249,137	246,352	14,555	5.9%
2017	30,878	187,913	240,197	239,748	17,385	7.2%
2016	28,886	167,177	221,684	219,822	17,098	7.7%
2015	26,000	149,891	202,949	201,297	15,508	7.7%
2014	23,666	132,661	182,609	181,279	14,471	7.9%
2013	23,478	129,538	179,031	177,580	14,626	8.2%
2012	21,959	116,938	166,015	164,399	15,490	9.4%
2011	21,361	111,089	158,255	157,038	12,906	8.2%
2010	17,082	85,857	124,131	122,680	9,345	7.6%
2009	23,116	110,512	164,448	162,466	12,772	7.8%
2008	25,565	116,531	179,478	176,424	14,143	8.0%
2007	24,063	104,424	166,312	163,472	13,651	8.3%
Total	366,275	2,088,074	2,817,031	2,785,762	208,006	—

Table 2. The number of design patents (#Patent), figure files (#Figure file), subfigures recorded in XML (#Figure XML), segmented individual figures (#Figure segmented), mismatched figures due to OCR errors and label-figure alignment errors (#Figure mismatch), and the mismatch rate (Mismatch%), calculated as (#Figure mismatch)/(#Figure segmented), i.e., Mismatch% is defined as the number of mismatched segmented figures divided by the total number of segmented figures.

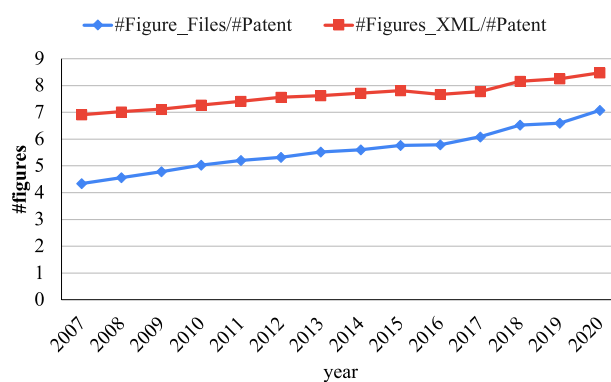


Fig. 3 Average numbers of figures and subfigures per patent each year from 2007 to 2020.

Caption: FIG. 1 is a front, top, and left side perspective view of a pet treat according to the new design.

Inline reference: The broken lines in FIGS. 1-7 illustrate portions of the pet treat that form no part of the claimed design.

Fig. 4 An example of a figure caption and its corresponding inline reference in a patent XML file. Orange highlighted text stands for object names, and blue highlighted text stands for viewpoints.

The ground truth corpus was created by randomly selecting 3300 patent figure captions from the US design patents. Each caption is manually annotated by two researchers independently using brat, a web-based text annotation tool²⁸. An example of an annotated caption is shown in Fig. 4. The ground truth was split into training, validation, and test sets, each consisting of 2700, 300, and 300 captions respectively. We then applied this model to all patent figure captions and extracted a total of 132,890 unique object names and 22,394 unique viewpoints. The result of post-extraction validation using a random sample of 100 figure captions was consistent with the evaluation based on the benchmark.

Image Processing: Figure Segmentation and Metadata Alignment. The goal of the image processing pipeline is segmenting compound figures and recognizing figure labels, which is necessary because the patent documents do not contain information that maps individual figures' captions to figure files. This process has four steps (Fig. 2). (1) Figure label detection, including text and positions; (2) Segmenting compound figure files into individual figures; (3) Associating labels with individual figures; (4) Aligning metadata with individual figures. The final dataset product is publicly available at the Harvard Dataverse repository²⁹.

OCR Tools	Recall	Precision	F1
AWS Rekognition with rectifier	96.03	97.58	96.80
AWS Rekognition without rectifier	79.37	81.97	80.65
Google Vision API	88.90	100.00	94.10
AWS Textract	80.16	99.05	88.61
Tesseract ⁶⁵	44.40	96.60	60.84
EAST deep net trained on ICDAR2015 ^{66,67}	11.90	14.29	12.99

Table 3. A comparison of OCR engines on extracting labels and bounding boxes in patent drawings.

Figure Label Detection. The original figure files are in TIFF format, which is not compatible with many optical character recognition (OCR) engines and computer vision packages. Therefore, we converted TIFF to PNG format, which is widely used in computer vision and does not introduce compression artifacts as in JPEG files. We use AWS Rekognition's DetectText service, a commercial OCR engine, to recognize figure labels, including text content and bounding boxes³⁰.

To evaluate the quality of extraction results, we extended our previous work¹⁶ and compared several top-performing OCR engines using a corpus consisting of 100 randomly selected design patent figures from the USPTO dataset in 2020 including both single and compound figures. The evaluation metrics included precision, recall, and F1. The precision was calculated as the number of correctly recognized labels divided by the total number of labels captured. The recall was calculated as the number of correctly recognized labels divided by the total number of labels shown on sampled figures. The results (Table 3) indicate that AWS Rekognition achieves the highest F1 (96.80%), compared with other OCR engines. In particular, Rekognition achieves the highest recall (96.03%). Google Vision API shows an excellent F1 (94.10%), but the recall score is relatively low (88.90%). Tesseract, an open-source OCR engine, achieves a relatively high precision (96.60%) but a poor recall (44.40%). Therefore, we adopted Rekognition into our pipeline.

The errors made by Rekognition are due to several reasons. Errors may happen when the labels are rotated by 90 degrees, resulting in reversed order of label tokens. For example, "Fig.3" in a single figure could be recognized as ["3", "Fig."], and ["Fig.7", "Fig.6"] in a compound figure could be recognized as ["7", "6", "FIG.", "FIG."]. This is likely because Rekognition scans the image from top to bottom. In addition, unlike many other OCR engines that output a single result, Rekognition outputs multiple results scored from 0 to 1. To address the two issues above, we developed a script that rectifies the Rekognition output when the confidence score output by Rekognition is above a threshold $\theta_0 = 0.5$ and then reconstructs the labels by "gluing" tokens. This rectifier improved the F1 score of figure label recognition by 16%, from 0.806 to 0.968 (Table 3).

The results of this step are figure labels and their bounding boxes for each figure file. The figure label regions in the images are whitened out before figure files are passed to the segmentation module.

Compound Figure Segmentation. A significant number of US patents use compound figures, containing more than one individual figure, so they need to be segmented before they are associated with captions. We proposed a transfer learning method, with a transformer-based neural model called MedT³¹ pre-trained on compound medical images and then fine-tuned on a small set of annotated patent figures. We compared our model against several state-of-the-art image segmentation models on a test set containing annotated patent figures. We then associated segmented figures with labels using a proximity-based method. The details of our approach can be found in Hoque *et al.*¹⁸. Here, we outline the pipeline and highlight the performance of the best models.

To our best knowledge, there was no labeled dataset that could be used for patent figure segmentation, therefore we developed an in-house ground truth dataset consisting of 500 figure files randomly selected from design patents. Here, we assume the number of figures in each figure file equals the number of labels detected by the OCR engine. If the two numbers do not equal, the segmentation is flagged as an error. The results will still be included in the final dataset as challenging cases for future models, but these cases can be easily excluded using the flags. In our ground truth of 500 figures, there were 480 compound figures each containing up to 12 individual figures. Each figure has been annotated by a graduate student using the VGG Image Annotator³² by drawing bounding boxes around individual figures. We performed an independent human validation of bounding boxes.

As a baseline, we proposed an unsupervised method called *Point-shooting*¹⁸. The idea is to randomly "shoot" open dots onto a compound figure. Dots that overlay any strokes are filled with a single solid color and the other dots are removed, so the remaining dots created a region filled with single color pixels. A contour can be drawn around the region, outlining the profile of individual figures. A bounding box can then be drawn around the contour. Using this method, we can obtain a mask of individual figures and their bounding boxes. The bounding boxes can be directly used for comparing against the ground truth. Point-shooting achieves an accuracy of 92.5%, calculated as the percentage of individual figures that are correctly segmented.

MedT was designed to train a transformer-based semantic segmentation framework³¹. The core component is a gated position-sensitive axial attention mechanism, designed to overcome the limitation of a vanilla transformer so that a more robust model can be trained using relatively small training sets. Because the point-shooting method was able to generate masks of individual figures with relatively high accuracy, we used the output of the point-shooting method as the training data for MedT. As our result would show, the MedT model achieved higher performance on the noisy training data. The training figures were first passed through a convolution block before passing through a global branch, which captures dependencies between pixels and the entire image.

The same figure is broken down into patches, which were passed through a similar convolutional block before passing through a local branch, which captures dependencies among neighboring pixels. A re-sampler aggregates the outputs from the local branch and generates the output feature maps. The outputs from both branches are finally aggregated followed by a 1×1 convolutional layer to pool these feature maps into a segmentation mask. We fine-tuned the pre-trained MedT model on our training set and achieved an accuracy of 97% with a much shorter runtime ($\approx 1/35$) on the test set compared with the point-shooting method. MedT outperformed point-shooting and other deep learning baselines including U-Net³³, HR-Net³⁴, and DETR³⁵.

Label Association and Metadata Alignment. We have two sources of metadata: labels recognized by the OCR engine and semantic information parsed from XML files. To generate the contextualized figures, two alignments were conducted.

Label Association. This step associates labels output by the OCR engine with segmented figures. This was treated as a bipartite matching problem. A general solution is the Ford-Fulkerson method³⁶. Because of the specialty of our problem, we used a simple heuristic method by matching an individual figure segmented from a compound figure with the closest label recognized by the OCR engine. The proximity was calculated as the Euclidean distance between the geometric centers between the bounding boxes of labels and figures. Despite the simplicity of the method, it achieves an accuracy of 97%, evaluated on a set of 200 randomly selected figures in the matching results. The errors in the label-figure alignment are mainly attributed to upstream errors made by the OCR engine and the segmentation model, which occurs in $\approx 7.5\%$ figures (#Figure Mismatch / #Figure Segmented) (Table 2). Here, (#Figure Mismatch) refers to the number of figures that could not be aligned with labels using the proximity method above. Another type of error is attributed to a compact and irregular arrangement of labels. The output of this step is a set of individual figures with labels. We mark cases in which the number of labels recognized does not equal the number of segmented figures. For completeness, we still include these cases in the data product. Users can easily remove these cases in downstream training/analysis or use them to design better algorithms to automatically correct errors.

Metadata Alignment. This step aligns labeled figures with captions extracted from XML files by matching figure labels parsed from XML files and labels associated with individual figures in the last step. The result of this step contains segmented figures, each having metadata, including the label, the caption, the bounding box coordinates, and document-level metadata, including patent ID, year, title, number of figures, etc. The method used here is based on strict integer matching, so the errors in the final data are caused by errors propagated from upstream processes.

Data Records

The final data product contains about 2 million compound PNG figures, 2.7 million segmented PNG figures, and metadata in JSON format. The dataset is organized by year (Table 2). The total size of the compressed dataset is 314 GB. The size of each year's data is roughly proportional to the number of figure files (Fig. 3).

The JSON files contain document-level and figure-level metadata (one JSON file per year). Each entry in the JSON file represents a segmented figure, which is tagged with the patent ID, original figure file, the object name, and viewpoints extracted, the figure labels, the bounding boxes of segmented figures, and their labels. The schema and a sample record of this JSON file are shown in Table 4. The patent captions contain an average of 12 words with a vocabulary of about 25,000 tokens. The semantic information extraction resulted in a total of 132,890 unique object names and 22,394 viewpoints. The distribution of the top 20 object names is shown in Fig. 5. The distribution of the top 20 viewpoints is shown in Fig. 6. The *perspective view*, *front view*, and *top plan view* are the top 3 viewpoints used in design patents. The object distribution color-coded by the top three viewpoints (Fig. 7) indicates that objects are depicted with diverse and disproportionate viewpoints, which poses challenges for 3D reconstruction from 2D sketches. The complete distribution is provided with the dataset.

We observe that one object name, such as “bottle”, may appear in multiple patents. Figure 8 shows the distribution of the number of patents per object. Similarly, Fig. 9 shows the distribution of the number of individual figures per object. The peak of Fig. 8 happens when the x -axis value is 1, meaning that most frequently, there is 1 patent for each object name. The peak of Fig. 9 happens when the x -axis value is 7, meaning that most frequently, there are 7 individual figures for each object.

The USPTO design patents contain class labels that were manually assigned using the Locarno International Classification system, which was developed by members of the Paris Convention for the Protection of Industrial Property and administered by the International Bureau of the World Intellectual Property Office WIPO³⁷. The system has 32 major categories and 16 minor categories. The fraction of patents assigned with minor categories is less than 0.5%. A full Locarno classification code is composed of two pairs of numbers separated by a hyphen with the first pair showing the top-level design classes and the second pair showing the low-level classes. Locarno classification code is available in the first section of each patent document. They follow a fixed format and thus can be accurately parsed for all patents in our dataset. Table 5 presents a list of top-level classes including the codes and descriptions. Figure 10 illustrates the distribution of patents across different classes (top panel), the average number of figures per patent (middle panel), and the average number of figures per object (bottom panel) in DeepPatent2. The analysis reveals that, on average, each patent contains approximately 5.48 figures and each object is illustrated by on average 13.83 individual figures.

The full dataset²⁹ is publicly available under the Harvard Dataverse repository at <https://doi.org/10.7910/DVN/UG4SBD>. For each year, the original and segmented figures are in separate tarball files, with a JSON file attached.

Field Name	Meaning	Example
id	A unique number representing an individual figure segmented from a compound figure	1
patentID	The patent ID followed by the patent approval date	USD0836880-20190101
patentdate	The patent approval date	2019-01-01
figid	Segmented/individual figure label	1
caption	The figure caption directly extracted from the patent XML file	Fig. 1 is a front, top, and left side perspective view of a pet treat according to the new design;
object	The object name automatically extracted from the figure caption	Pet treat
aspect	The viewpoint automatically extracted from the figure caption	front, top, and left side perspective view
figure_file	The original figure file name	USD0836880-20190101-D00001.png
subfigure_file	The segmented figure file name as shown	USD0836880-20190101-D00001_1.png
object_title	The title of the patent	Pet treat
x_figure	The x coordinate of the upper left vertex of the segmented figure's bounding box in pixels, measured from the upper-left-corner of the original figure	614
y_figure	The y coordinate of the upper left vertex of the segmented figure's bounding box in pixels, measured from the upper-left-corner of the original figure	91
w_figure	The width of the segmented figure's bounding box in pixels	969
h_figure	The height of the segmented figure's bounding box in pixels	2740
W_label	The width of the segmented figure's label's bounding box, normalized with respect to the original figure	0.0439147
H_label	The height of the segmented figure's label's bounding box, normalized with respect to the original figure	0.0502626
L_label	The normalized coordinate of the upper left vertex of the segmented figure's bounding box, measured from the upper-left corner of the original figure	0.119197
T_label	The normalized coordinate of the upper left vertex of the segmented figure's bounding box, measured from the upper-left corner of the original figure	0.861215

Table 4. The fields of an individual figure in the JSON file in the final data.

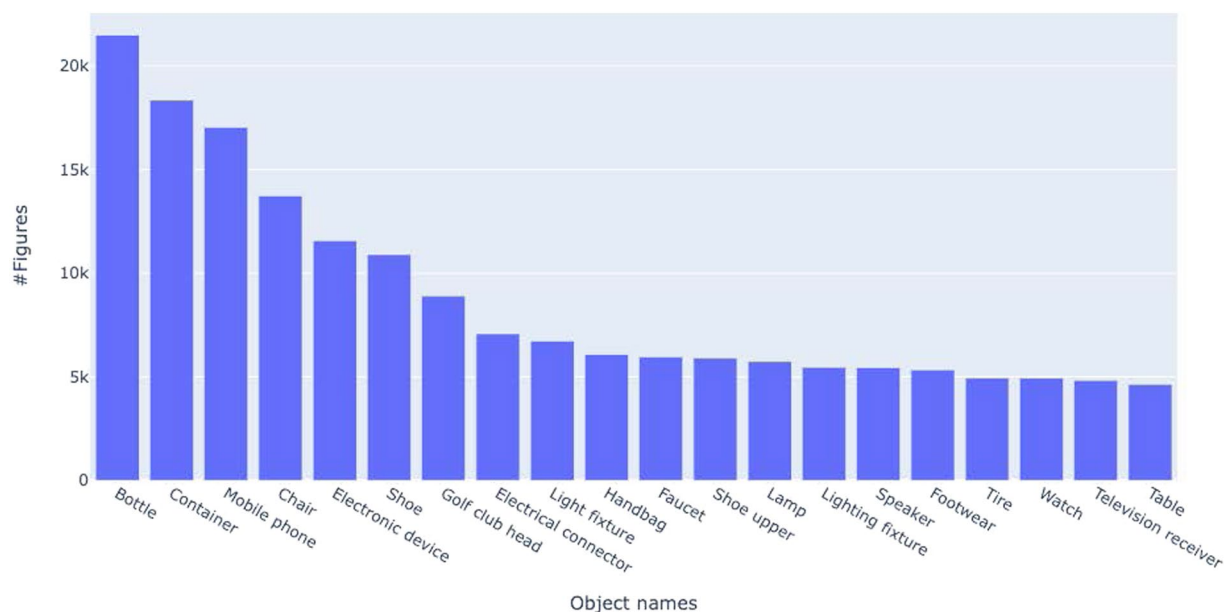


Fig. 5 Distribution of object names identified in DeepPatent2. Only the top 20 object names are shown.

Technical Validation

Data Validation, Evaluation, and Copyright. Our data was automatically generated by high-performance machine learning and deep learning methods^{17,18} (read Section “Methods” for a detailed description). However, errors can still occur, propagate, and accumulate leading to errors in the final data product. There are four possible error sources, namely, figure label detection, compound image segmentation, label association, and entity recognition (ER). Label association is dependent on image segmentation and OCR, but label association

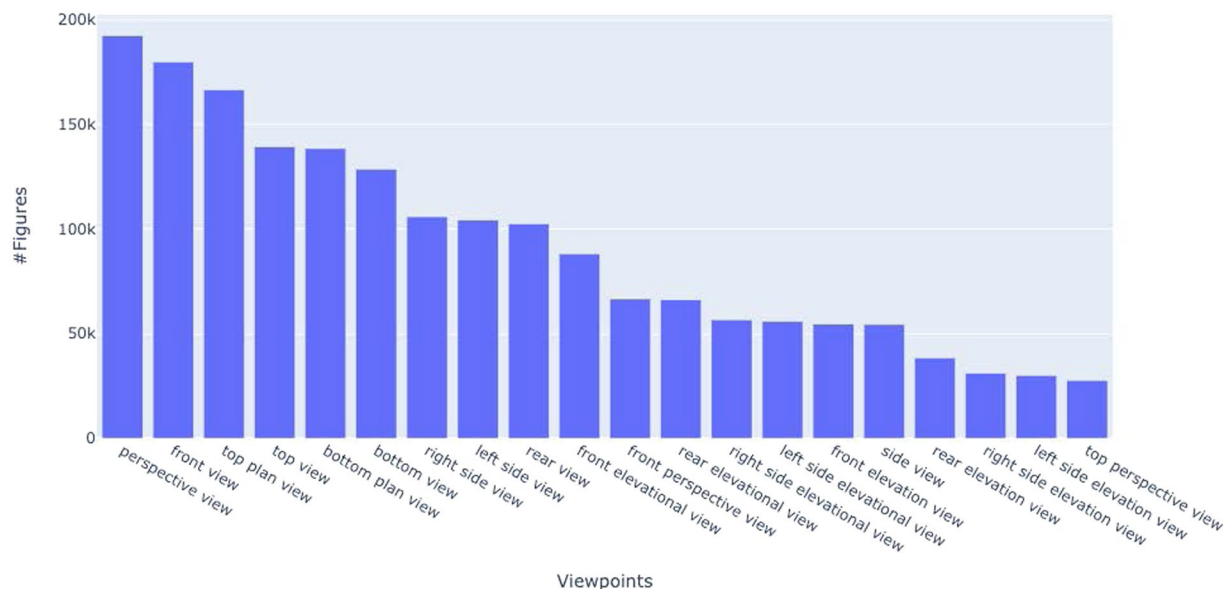


Fig. 6 Distribution of viewpoints identified in DeepPatent2. Only the top 20 viewpoint are shown.

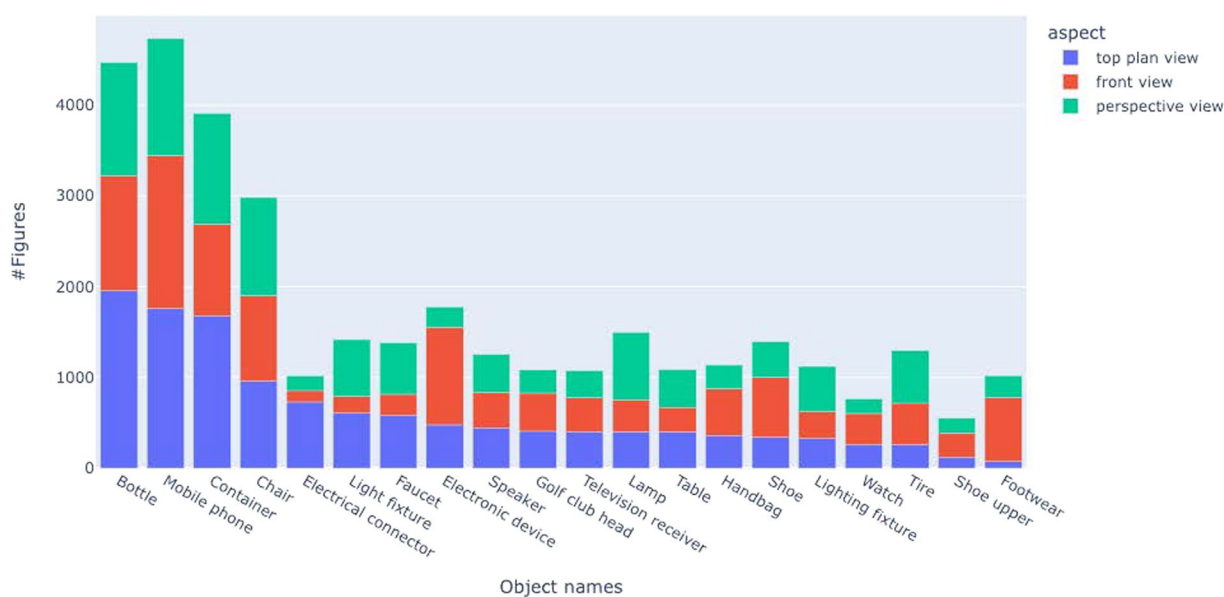


Fig. 7 Distribution of top 20 objects with the top three viewpoints color-coded. Note that the bar heights are different from Fig. 5 because not all aspects were included. The object names along the abscissa are sorted by the total number of figures.

errors could occur even when image segmentation and OCR are correct. Assuming the label association (LA) can be approximated as the mismatch errors (Mismatch% in Table 2), which is 7.5% on average (so the precision $P_{LA} = 92.5\%$), the overall error rate is calculated as

$$E \approx 1 - P_{LA} \times P_{ER} = 1 - 92.5\% \times 96.0\% \approx 11.2\%. \quad (1)$$

All figures are preserved in the final dataset, but mismatched figures are marked in their file names, so they can be used if needed.

To validate the estimated error rates, we build a verification dataset by randomly sampling 100 compound figures each year from 2007 to 2020 for a total of 1400 compound figures. We manually inspected the quality of the final data product. Specifically, we inspected whether a compound figure is correctly segmented and whether the label, object name, and viewpoint of an individual figure are correctly extracted. We calculated the error rates by dividing the number of individual figures containing *any* errors (404) by the total number of individual figures (3464). Figure 11 (Left) shows the error rates calculated for each year and the average error rate.

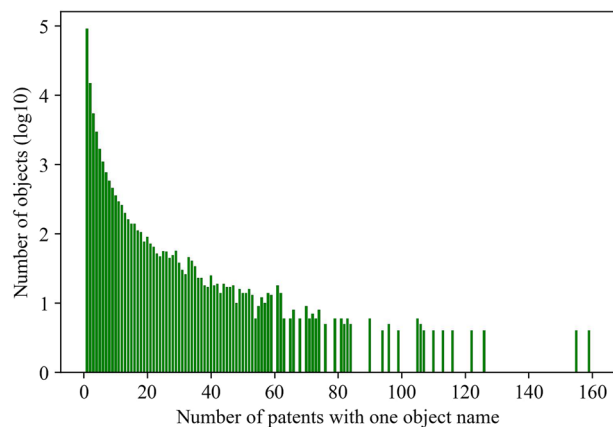


Fig. 8 One object name may be described by multiple patents. This figure shows a truncated distribution of the number of patents in one object category. For example, when the x -axis value is 7, the y -axis value is 774, meaning each of the 774 objects is described by 7 patents. The x -axis is truncated because data points beyond the truncation point are sparse.

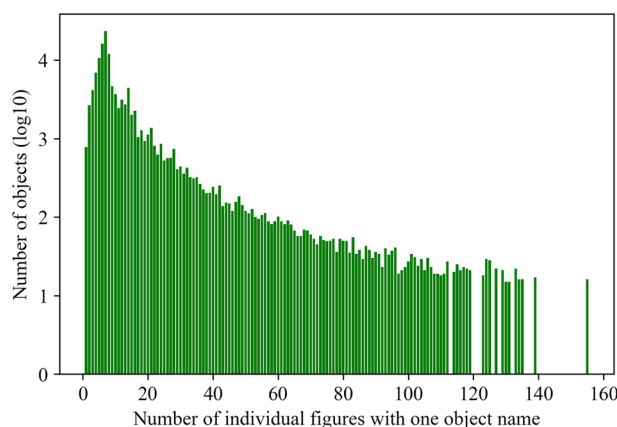


Fig. 9 One object name may appear in multiple individual figures. This figure shows the distribution of the number of individual figures with one object name. For example, when the x -axis value is 7, the y -axis value is 23367, meaning each of the 23367 objects appears in 7 individual figures. The x -axis is truncated because data points beyond the truncation point are sparse.

The verified error is consistent with the estimated error by Eq. (1). Table 6 shows the breakdown of the verified errors by source. Note that the Verified Overall row is calculated using direct counts and should not be calculated from previous rows using Eq.(1), because a fraction of figures have more than one type of errors. We also calculate error rates for each class. The verification dataset contains 23 classes, of which 10 classes contain more than 40 individual figures. Figure 11 (Right) shows the error rates across classes which have more than 40 individual figures in the verification dataset.

Our estimated error rate is on par with other computer vision datasets, especially considering that our methods produce automated tags. Other auto-tagged computer vision datasets like WebVision estimate up to 20% error rates;¹⁵ while curated datasets like ImageNet-1k³⁸ have been estimated to have error rates up to 6%³⁹.

USPTO advocates open data, which can be freely used, reused, and redistributed⁴⁰. The text and drawings of a patent are typically not subject to copyright restrictions²⁰. Our dataset is under the Creative Commons Attribution (CC-BY 4.0) Generic International License.

Usage Notes

In this section, we demonstrate the usefulness and value of the dataset on a *Conceptual Captioning* task, which generates short descriptive text that captures main objects and their viewpoints of a given image. This task was carried out on a dataset (hereafter CC18) consisting of natural images collected from the Web⁴¹. In this work, we carry out this task on the technical drawings presented in our dataset. We demonstrate how to use the dataset and advance the visual-to-text conversion task by running a state-of-the-art model on data in the technical drawing domain and potentially opening up research opportunities in other topics on patent drawings. In addition, we provide other tasks, in which our dataset could be potentially adopted. Here, we focus on demonstrating the potential benefit of improving the model performance by scaling up the training set using our dataset, instead of developing a new method.

Class Code	Names
1	Foodstuffs
2	Articles of clothing and haberdashery
3	Travel goods, cases, parasols and personal belongings, not elsewhere specified
4	Brushware
5	Textile piece goods, artificial and natural sheet material
6	Furnishing
7	Household goods, not elsewhere specified
8	Tools and hardware
9	Packaging and containers for the transport or handling of goods
10	Clocks and watches and other measuring instruments, checking and signaling instruments
11	Articles of adornment
12	Means of transport or hoisting
13	Equipment for production, distribution or transformation of electricity
14	Recording, telecommunication or data processing equipment
15	Machines, not elsewhere specified
16	Photographic, cinematographic, and optical apparatus
17	Musical instruments
18	Printing and office machinery
19	Stationery and office equipment, artists' and teaching materials
20	Sales and advertising equipment, signs
21	Games, toys, tents and sports goods
22	Arms, pyrotechnic articles, articles for hunting, fishing, and pest killing
23	Fluid distribution equipment, sanitary, heating, ventilation and air-conditioning equipment, solid fuel
24	Medical and laboratory equipment
25	Building units and construction elements
26	Lighting apparatus
27	Tobacco and smokers' supplies
28	Pharmaceutical and cosmetic products, toilet articles, and apparatus
29	Devices and equipment against fire hazards, for accident prevention and for rescue
30	Articles for the care and handling of animals
31	Machines and appliances for preparing food or drink, not elsewhere specified
32	Graphic symbols and logos, surface patterns, ornamentation, arrangement of interiors and exteriors

Table 5. The top-level class codes and descriptions of the Locarno International Classification designation adopted by USPTO for design patents.

Conceptual Captioning. The task of conceptual captioning involves generating a short textual description of an image. The state-of-the-art models usually employ the encoder-decoder architecture⁴¹. We employed ResNet-152⁴², a CNN-based model pre-trained on ImageNet³⁸, and fine-tuned on technical drawings selected from DeepPatent2. We varied the training data size from 500 to 1000 to 63,000 and tested the performance of each model on a corpus consisting of 600 figures with manually validated object names and viewpoints. The training sets were randomly selected from the parent sample. We selected the test set so it did not overlap with any samples from the training set. The input to the CNN encoder was an individual figure. The decoder generated a template-based caption including the object name and the viewpoint. The dimension of the encoder output was 2048 and we reduced it to 512 to match the input dimension of the LSTM decoder. We trained the image captioning network with the cross-entropy loss for 30 epochs with a mini-batch size of 32. We adopted an ADAM⁴³ optimizer and a learning rate of 0.001.

The models were evaluated using standard metrics including METEOR⁴⁴, NIST⁴⁵, Translation Error Rate (TER⁴⁶), ROUGE⁴⁷, and accuracy. The accuracy was defined as a fraction of the ground-truth patent object names and viewpoints that were correctly predicted by the model. The results in Table 7 indicate that increasing the training size for automatically tagged data improved the performance of image captioning models.

The computing environment includes a Linux server with Intel Silver CPU, Nvidia GTX 2080 Ti. Because of the large size of data, we recommend a disk capacity of at least 500GB for all data. To obtain the subfigures from compound images, we used *Pytorch* > = 1.4.0, *torchvision* > = 0.5.0, and *Python* 3.8. The *OpenCV* package was used to load PNG images. To load large JSON files, we recommend using *ijson* > 3.1 package, which helps to conserve computer memory space.

Potential Usages of the Dataset.

- **Technical drawing image retrieval and semantic understanding.** Investigation of effective and efficient retrieval methods for technical drawings has attracted the attention of the computer vision community. The

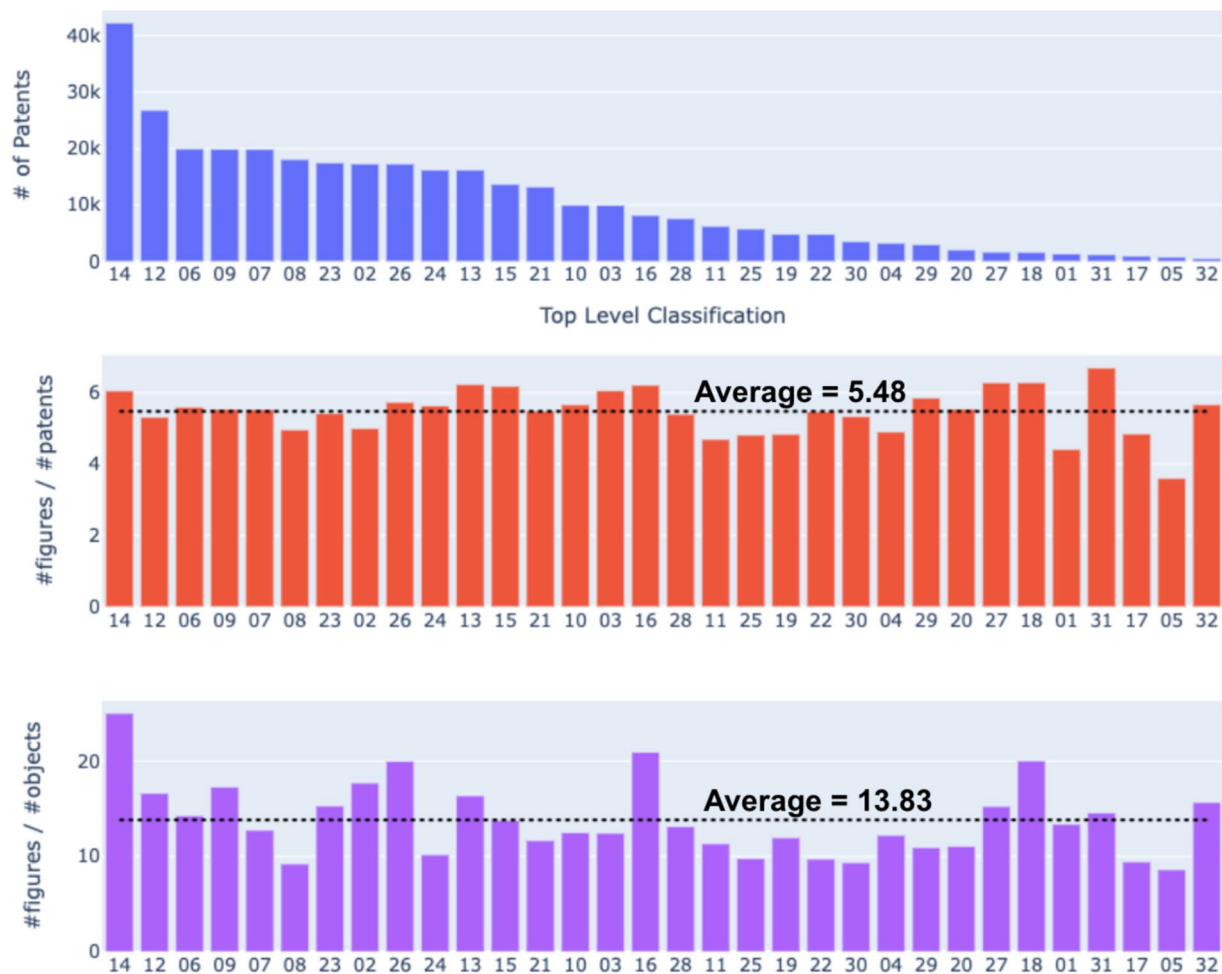


Fig. 10 Distribution of patents, figures per patent, and figures per object identified over the patent classes. Only the top 32 classifications are shown. The remaining classifications are not listed because they contain less than 5 figures. The class codes in the abscissa correspond to the Locarno International Classification described in Table 5.

Diagram and Abstract Imagery competition uses the DeepPatent dataset⁴⁸. The rich semantic information of DeepPatent2 can potentially help build a new multimodal (text + image) ground truth and enable tasks such as semantic understanding of abstract drawings and technical document classification⁴⁹.

- **Summarization of scholarly and technical corpora.** Many existing summarization methods^{50–52} use only text. However, it has been shown that combining multimodal content improves a reader's understanding of each document's content while reducing the amount of content they must consume^{53,54}. For example, the automatic selection of relevant images⁵⁵ has focused primarily on web resources and natural images. DeepPatent2, with its segmented compound figures, can be used to evaluate summary image selection in technical corpora.
- **3D image reconstruction.** Although humans are good at perceiving 3D objects from technical drawings, the task is challenging for computers. Delanoy *et al.* developed neural methods to reconstruct 3D images from 2D sketches using training datasets corresponding to procedural, vases, and chairs⁵⁶. Our dataset contains more diverse object types with multiple viewpoints and can potentially be used for training high-fidelity models.
- **Figure segmentation.** Segmenting compound figures is a common preprocessing step before individual figures are used for analysis, retrieval, or machine learning tasks. Using 4000 DeepPatent2 figures, we achieved an accuracy of 99.5% at both $IoU = 0.7$ and $IoU = 0.9$ evaluated on the technical drawings in¹⁸. Our data can potentially be used for training a large-scale base model that is fine-tunable for scientific or medical figure segmentation problems.
- **Technical Drawing Classification.** Sketch classification methods^{57–59} have been proposed to recognize sketch images from the Web. However, these datasets contain a limited number of object categories and viewpoints. Our dataset contains more diverse object types and viewpoints and can potentially be used for training robust technical drawing classification models.
- **Create Generative and Multimodal Design Models for Innovation.** The Generative Adversarial Networks (GAN) and diffusion models perform remarkably well on text-to-image synthesis tasks⁶⁰ and large language models (LLMs) have achieved superior performance on many generative NLP tasks⁶¹. Whether it is possible

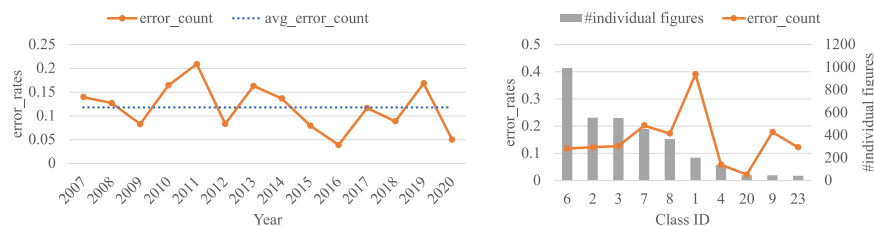


Fig. 11 Left: Error rates calculated using the verification dataset over years; Right: Error rates across classes. The classes are ordered by the number of individual figures in each class. Class 6 (Furnishing) has the highest number of individual figures. Class 1 (Foodstuffs) has the highest error rate (0.39).

Error Type	Error
Estimated - Label (Segmentation + OCR)	7.5%
Estimated - ER	4%
Estimated - Overall	11.2%
Verified - Segmentation	5.52%
Verified - Label (excluding segmentation)	1.87%
Verified - ER (Object extraction)	7.61%
Verified - ER (Aspect extraction)	0.91%
Verified - Overall	11.7%

Table 6. A breakdown analysis of the errors estimated using Eq. (1) and the errors obtained using the verification dataset in the Technical Validation Section. The results show that they are consistent with a discrepancy of 0.5%.

6–11 Train Size	ACC.	METEOR	NIST	TER	ROUGE-1			ROUGE-2		
					P	R	F1	P	R	F1
63000	0.647	0.627	1.62	0.49	0.645	0.641	0.634	0.545	0.544	0.534
1000	0.524	0.463	1.18	0.60	0.620	0.472	0.529	0.411	0.365	0.379
500	0.518	0.451	1.11	0.60	0.620	0.447	0.513	0.405	0.343	0.365

Table 7. A comparison of image captioning models with different training sizes. The best performance metrics are highlighted in bold.

to combine LLMs with the diffusion models to automatically generate multimodal design models for innovation is an open question. Diffusion models and LLMs are known to be inaccurate in their details, and so the object name and viewpoint information of DeepPatent2 could provide the needed detailed technical drawings for training multimodal generative models that are accurate enough for design innovation.

Code availability

The code used for preprocessing and segmenting figures is publicly available on GitHub: <https://github.com/lamps-lab/Patent-figure-segmentor> and <https://github.com/GoFigure-LANL/figure-segmentation>. Similarly, the software used for extracting semantic information, including object names and viewpoints from patent captions, is publicly available at <https://github.com/lamps-lab/Visual-Descriptor>.

Received: 31 January 2023; Accepted: 17 October 2023;

Published online: 07 November 2023

References

1. Carney, R. N. & Levin, J. R. Pictorial illustrations still improve students' learning from text. *Educational Psychology Review* **14**, 5–26 (2002).
2. Mayer, R. E. Illustrations that instruct. *Advances in Instructional Psychology* **10**, 253–284 (2019).
3. Sangkloy, P., Burnell, N., Ham, C. & Hays, J. The sketchy database: Learning to retrieve badly drawn bunnies. *ACM Trans. Graph.* **35**, <https://doi.org/10.1145/2897824.2925954> (2016).
4. Nadeem, U., Shah, S. A. A., Sohel, F., Togneri, R. & Bennamoun, M. Deep learning for scene understanding. In *Handbook of Deep Learning Applications*, 21–51, https://doi.org/10.1007/978-3-030-11479-4_2 (Springer International Publishing, 2019).
5. Lin, T.-Y. *et al.* Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, 740–755 (2014).
6. Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (2009).
7. Vrochidis, S., Mourtzidou, A. & Kompatsiaris, I. Concept-based patent image retrieval. *World Patent Information* **34**, 292–303 (2012).
8. Gryaditskaya, Y. *et al.* Opensketch: A richly-annotated dataset of product design sketches. *ACM Trans. Graph.* **38** (2019).

9. Google. The quick, draw! dataset. <https://github.com/googlecreativelab/quickdraw-dataset/> (2020).
10. Eitz, M., Hays, J. & Alexa, M. How do humans sketch objects? *ACM Transactions on Graphics* **31** (2012).
11. Koblin, A. M. The sheep market. In *Proceedings of the Seventh ACM Conference on Creativity and Cognition*, 451–452, <https://doi.org/10.1145/1640233.1640348> (2009).
12. Piroi, F., Lupu, M., Hanbury, A. & Zenz, V. CLEF-IP 2011: Retrieval in the intellectual property domain. In *Conference and Labs of the Evaluation Forum* (2011).
13. Wang, H., Ge, S., Lipton, Z. C. & Xing, E. P. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems* (2019).
14. Kucer, M., Oyen, D., Castorena, J. & Wu, J. DeepPatent: Large scale patent drawing recognition and retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2309–2318 (2022).
15. Li, W., Wang, L., Li, W., Agustsson, E. & Van Gool, L. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862* (2017).
16. Gong, M. *et al.* Recognizing figure labels in patents. In *Proceedings of the Workshop on Scientific Document Understanding at AAAI Conference on Artificial Intelligence* (2021).
17. Wei, X., Wu, J., Ajayi, K. & Oyen, D. Visual descriptor extraction from patent figure captions: A case study of data efficiency between BiLSTM and transformer. In *Proceedings of The ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, <https://doi.org/10.1145/3529372.3533299> (2022).
18. Hoque, M. R. U. *et al.* Segmenting technical drawing figures in US patents. In Veyseh, A. P. B., Deroncourt, F., Nguyen, T. H., Chang, W. & Lai, V. D. (eds.) *Proceedings of the Workshop on Scientific Document Understanding co-located with 36th AAAI Conference on Artificial Intelligence, SDU@AAAI 2022, Virtual Event, March 1, 2022*, vol. 3164 of *CEUR Workshop Proceedings* (CEUR-WS.org, 2022).
19. USPTO. United States Patent and Trademark Office. <https://www.uspto.gov/>. Accessed: October 10, 2023.
20. United States Patent and Trademark Office. Terms of use for USPTO websites. <https://www.uspto.gov/terms-use-uspto-websites> (2022).
21. Li, J., Sun, A., Han, J. & Li, C. A survey on deep learning for named entity recognition. *IEEE Trans. Knowl. Data Eng.* **34**, 50–70, <https://doi.org/10.1109/TKDE.2020.2981314> (2022).
22. Pennington, J., Socher, R. & Manning, C. D. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, (EMNLP)* (ACL, 2014).
23. Radford, A. *et al.* Language models are unsupervised multitask learners. *OpenAI blog* (2019).
24. Liu, Y. *et al.* RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692* (2019).
25. Devlin, J., Chang, M., Lee, K. & Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-HLT)*, <https://doi.org/10.18653/v1/n19-1423> (2019).
26. Lan, Z. *et al.* ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations, (ICLR)* (2020).
27. Sanh, V., Debut, L., Chaumond, J. & Wolf, T. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
28. Stenetorp, P. *et al.* brat: a web-based tool for nlp-assisted text annotation. In Daelemans, W., Lapata, M. & Márquez, L. (eds.) *EACL 2012, 13th Conference of the European Chapter of the Association for Computational Linguistics, Avignon, France, April 23-27, 2012*, 102–107 (The Association for Computer Linguistics, 2012).
29. Wu, J. *et al.* Replication Data for DeepPatent2: A Large-Scale Benchmarking Corpus for Technical Drawing Understanding. *Harvard Dataverse* <https://doi.org/10.7910/DVN/UG4SBD> (2023).
30. Vijayarani, S. & Sakila, A. Performance comparison of OCR tools. *International Journal of UbiComp (IJU)* **6**, 19–30 (2015).
31. Dosovitskiy, A. *et al.* An image is worth 16 × 16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR* (2021).
32. Dutta, A. & Zisserman, A. The via annotation software for images, audio and video. In *Proceedings of the 27th ACM International Conference on Multimedia, MM '19*, 2276–2279, <https://doi.org/10.1145/3343031.3350535> (Association for Computing Machinery, New York, NY, USA, 2019).
33. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241 (Springer, 2015).
34. Wang, J. *et al.* Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence* (2020).
35. Carion, N. *et al.* End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229 (Springer, 2020).
36. Ford, L. R. & Fulkerson, D. R. Maximal flow through a network. *Canadian Journal of Mathematics* **8**, 399–404, <https://doi.org/10.4153/CJM-1956-045-5> (1956).
37. World Intellectual Property Office. Locarno Classification. <https://www.wipo.int/classifications/locarno/>. Accessed: October 10, 2023.
38. Russakovsky, O. *et al.* ImageNet large scale visual recognition challenge. *International Journal of Computer Vision* **115**, 211–252 (2015).
39. Northcutt, C. G., Athalye, A. & Mueller, J. Pervasive label errors in test sets destabilize machine learning benchmarks. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track* (2021).
40. United States Patent and Trademark Office. Why USPTO open data? <https://developer.uspto.gov/about-open-data> (2022).
41. Sharma, P., Ding, N., Goodman, S. & Soricut, R. Conceptual captions: A cleaned, hypemymed, image alt-text dataset for automatic image captioning. In Gurevych, I. & Miyao, Y. (eds.) *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, 2556–2565, <https://doi.org/10.18653/v1/P18-1238> (Association for Computational Linguistics, 2018).
42. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778 (2016).
43. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
44. Banerjee, S. & Lavie, A. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 65–72 (2005).
45. Doddington, G. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proceedings of the Second International Conference on Human Language Technology Research, HLT '02*, 138–145 (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002).
46. Snover, M. G., Dorr, B. J., Schwartz, R. M., Micciulla, L. & Makhoul, J. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers, AMTA 2006, Cambridge, Massachusetts, USA, August 8-12, 2006*, 223–231 (Association for Machine Translation in the Americas, 2006).
47. Lin, C.-Y. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81 (2004).
48. Kucer, M. ECCV 2022 DIRA workshop image retrieval challenge. <https://codalab.lisn.upsaclay.fr/competitions/5885> (2022).

49. Jiang, S., Luo, J., Hu, J. & Magee, C. L. Deep learning for technical document classification. *CoRR* **abs/2106.14269** (2021).
50. See, A., Liu, P. J. & Manning, C. D. Get To The Point: Summarization with Pointer-Generator Networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 1073–1083, <https://doi.org/10.18653/v1/P17-1099> (Vancouver, British Columbia, Canada, 2017).
51. Moirangthem, D. S. & Lee, M. Abstractive summarization of long texts by representing multiple compositionality with temporal hierarchical pointer generator network. *Neural Networks* **124**, 1–11, <https://doi.org/10.1016/j.neunet.2019.12.022> (2020).
52. Zhang, J., Zhao, Y., Saleh, M. & Liu, P. PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning* **119**, 11328–11339 (2020).
53. Capra, R., Arguello, J. & Scholer, F. Augmenting web search surrogates with images. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, 399–408, <https://doi.org/10.1145/2505515.2505714> (San Francisco, California, USA, 2013).
54. Jones, S. M., Weigle, M. C. & Nelson, M. L. Social Cards Probably Provide For Better Understanding Of Web Archive Collections. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2023–2032, <https://doi.org/10.1145/3357384.3358039> (Beijing, China, 2019).
55. Jones, S. M., Weigle, M. C., Klein, M. & Nelson, M. L. Automatically selecting striking images for social cards. In *13th ACM Web Science Conference 2021*, WebSci '21, 36–45, <https://doi.org/10.1145/3447535.3462505> (Virtual Event, United Kingdom, 2021).
56. Delanoy, J., Aubry, M., Isola, P., Efros, A. A. & Bousseau, A. 3D Sketching using Multi-View Deep Volumetric Prediction. *Proc. ACM Comput. Graph. Interact. Tech.* **1**, 21:1–21:22, <https://doi.org/10.1145/3203197> (2018).
57. Zhang, H. *et al.* Sketchnet: Sketch classification with web images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1105–1113 (2016).
58. Jearasuwan, S. & Wangsiripitak, S. Sketch image classification using component based k-NN. In *2019 IEEE 4th International Conference on Computer and Communication Systems (ICCCS)*, 267–271 (IEEE, 2019).
59. Jiang, S., Luo, J., Ruiz-Pava, G., Hu, J. & Magee, C. L. Deriving Design Feature Vectors for Patent Images Using Convolutional Neural Networks. *Journal of Mechanical Design* **143**, 061405, <https://doi.org/10.1115/1.4049214> (2021).
60. Dhariwal, P. & Nichol, A. Diffusion Models Beat GANs on Image Synthesis. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P. & Vaughan, J. W. (eds.) *Advances in Neural Information Processing Systems*, vol. 34, 8780–8794 (Curran Associates, Inc., 2021).
61. Zhao, W. X. *et al.* A survey of large language models. *CoRR* **abs/2303.18223**, <https://doi.org/10.48550/arXiv.2303.18223> (2023).
62. Hodosh, M., Young, P. & Hockenmaier, J. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research* **47**, 853–899 (2013).
63. Young, P., Lai, A., Hodosh, M. & Hockenmaier, J. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* **2**, 67–78 (2014).
64. Hu, R. & Collomosse, J. A performance evaluation of gradient field hog descriptor for sketch based image retrieval. *Computer Vision and Image Understanding* **117**, 790–806 (2013).
65. Smith, R. An overview of the Tesseract OCR engine. In *IEEE International Conference on Document Analysis and Recognition (ICDAR)* (2007).
66. Zhou, X. *et al.* East: an efficient and accurate scene text detector. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5551–5560 (2017).
67. Karatzas, D. *et al.* ICDAR 2015 competition on robust reading. In *13th International Conference on Document Analysis and Recognition (ICDAR)* (IEEE, 2015).

Acknowledgements

Research conducted by KA, XW, MG, WS, and JW presented in this paper was supported by Los Alamos National Laboratory subcontract BA601958 awarded to Old Dominion University. Research conducted by SJ, MK, and DO presented in this paper was supported by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20200041ER.

Author contributions

K.A. assembled different modules as a pipeline and drafted the paper, X.W. built the text extraction module and populated semantic data into the database, M.G. collected the raw data, W.S. developed the web user interface to verify the data, J.W. managed the project, proposed ideas, coordinate human resources, and edited the paper, S.J. extracted figure labels using the A.W.S. service, M.K. and D.O. participated in the discussion and edited the draft.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to J.W. or D.O.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023