

Unitail: Detecting, Reading, and Matching in Retail Scene

Fangyi Chen¹, Han Zhang¹, Zaiwang Li², Jiachen Dou¹, Shentong Mo¹, Hao Chen¹, Yongxin Zhang³, Uzair Ahmed¹, Chenchen Zhu¹, and Marios Savvides¹

¹ Carnegie Mellon University, Pittsburgh PA 15213, USA

² University of Pittsburgh, Pittsburgh PA 15213, USA

³ Tsinghua University, Beijing 100084, China

{fangyic, hanz3, jiachend, shentonm, haoc3, uzaira, marios}@andrew.cmu.edu
zal17@pitt.edu yx-zhang20@mails.tsinghua.edu.cn chenchez@alumni.cmu.edu

Abstract. To make full use of computer vision technology in stores, it is required to consider the actual needs that fit the characteristics of the retail scene. Pursuing this goal, we introduce the United Retail Datasets (Unitail), a large-scale benchmark of basic visual tasks on products that challenges algorithms for detecting, reading, and matching. With 1.8M quadrilateral-shaped instances annotated, the Unitail offers a detection dataset to align product appearance better. Furthermore, it provides a gallery-style OCR dataset containing 1454 product categories, 30k text regions, and 21k transcriptions to enable robust reading on products and motivate enhanced product matching. Besides benchmarking the datasets using various start-of-the-arts, we customize a new detector for product detection and provide a simple OCR-based matching solution that verifies its effectiveness. The Unitail and its evaluation server is publicly available at <https://unitedretail.github.io>.

Keywords: Product Detection, Product Recognition, Text Detection, Text Recognition

1 Introduction

With the rise of deep learning, numerous computer vision algorithms have been developed and have pushed many real-world applications to a satisfactory level. Currently, various visual sensors (fixed cameras, robots, drones, mobile phones, etc.) are deployed in retail stores, enabling advanced computer vision methods in shopping and restocking. Scene Product Recognition (SPR) is the foundation module in most of these frameworks, such as planogram compliance, out-of-stock managing, and automatic check-out.

SPR refers to the automatic detection and recognition of products in complex retail scenes. It comprises steps that first localize products and then recognize them via the localized appearance, analogous to many recognition tasks. However, scene products have their characteristics: *they are densely-packed, low-shot, fine-grained, and widely-categorized*. These innate characteristics result in obvious challenges and will be a continuing problem. Recent datasets in retail scenes

follow the typical setting in the common scene to initiate the momentum of research in SPR. For instance, the SKU110k dataset [17], which has recently enabled large-scale product detection, is following the MS-COCO [33] style annotation and evaluation metric. Despite their significant value, the underpaid attention to the SPR’s characteristics leads us to the question: what is the next advance towards SPR?

Firstly, traditional detection targets poorly comply with the actual needs, causing improper image alignment of the product appearances. Detection targets in common scenes [33, 57, 11, 9] are usually defined as covering the utmost visible entirety of an object with a minimal rectangle box. This format is inherited by most existing retail datasets [17, 55, 14, 52, 2]. However, because occlusion occurs more frequently between products (the densely-packed characteristic), improper alignments can easily hinder the detection performance. Detectors equipped with Non-Maximum Suppression (NMS) suffer from the overlaps among the axis-aligned rectangular bounding boxes (AABB) and rotated rectangular bounding boxes (RBOX). Moreover, poor alignment leads to inconsistent image registration of the same products, which brings extra difficulties to accurate recognition.

Secondly, even in the well-aligned cases, products from intra-classes require discriminative features due to their fine-grained characteristic. On the one hand, a slight variation in the product packaging can significantly change the product price, especially for the visually similar but textually different regions such as brand/model, flavour/version, ingredient/material, count/net weight. This requires SPR algorithms to pay attention to the particular text patterns. On the other hand, due to the labelling effort on thousands of categories per store (the widely-categorized characteristic), the available samples per category are scarce (the low-shot characteristic), which degrades the SPR robustness. These two constraints are in conjunction with our empirical observation that visual classifiers could frequently make mistakes when products look similar but vary in text information.

In this paper, we introduce the United Retail Datasets (Unitail) that responds to these issues. The Unitail is a comprehensive benchmark composed of two datasets: *Unitail-Det* and *Unitail-OCR*, and currently supports four tasks in real-world retail scene: *Product Detection*, *Text Detection*, *Text Recognition*, and *Product Matching*. (Fig.1)

Unitail-Det, as one of the largest quadrilateral object detection datasets in terms of instance number and the only existing product dataset in quadrilateral annotations by far, is designed to support well-aligned product detection. Unitail-Det enjoys two key features: 1. Bounding boxes of products are densely annotated in the quadrilateral style that cover the frontal face of products. Practically, Quadrilaterals (QUADs) adequately reflect the shapes and poses of most products regardless of the viewing angles, and efficiently cover the irregular shapes. The frontal faces of products provide distinguishable visual information and keep the appearances consistent. 2. In order to evaluate the robustness of the detectors across stores, the test set consists of two subsets to support both origin-domain and cross-domain evaluation. While one subset shares the domain

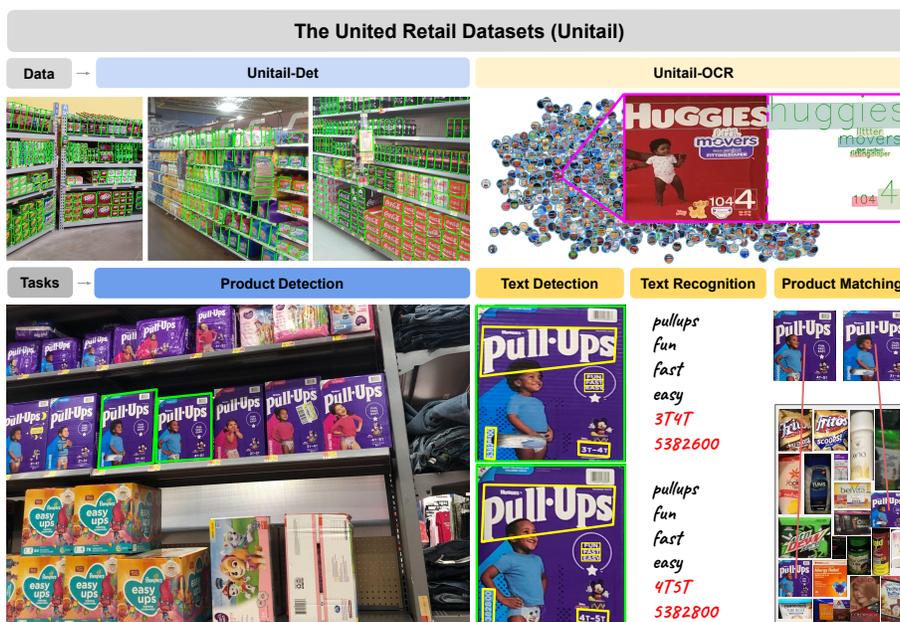


Fig. 1: The Unitail is a large-scale benchmark in retail scene that consists of two sub-datasets and supports four basic tasks.

with the training set, the other is independently collected from other different stores, with diverse optical sensors, and from various camera perspectives.

Unitail-OCR (Optical Character Recognition) aims to drive research and applications using visual texts as representations for products. This direction is partially inspired by the customers’ behavior: people can glance and recognize ice cream but need to scrutinize the flavor and calories to make a purchase. It is organized into three tasks: text detection, text recognition, and product matching. Product images in Unitail-OCR are selected from the Unitail-Det and benefit from the quadrilateral aligned annotations. Each is equipped with on-product text location and textual contents together with its category. Due to the product’s low-shot and widely-categorized characteristics, product recognition is operated by matching within an open-set gallery. To the best of our knowledge, Unitail-OCR is the first dataset to support OCR models’ training and evaluation on the retail products, and it is experimentally verified to fill in the domain blank; when evaluated on a wide variety of product texts, models trained on Unitail-OCR outperform those trained on common scene texts [24]. It is also the first dataset that enables the exploration of text-based solutions to product matching.

Based on the proposed Unitail, we design two baselines. To detect products, we analyze the limitation of applying generic object detectors in the retail scene and design RetailDet to detect quadrilateral products. To match products using

visual texts on 2D space, we encode text features with spatial positional encoding and use Hungarian Algorithm [28] that calculates optimal assignment plans between varying text sequences.

Our contributions are summarized in three folds: **(1)** we introduce the Unitail, a comprehensive benchmark for well-aligned textually enhanced SPR. **(2)** We benchmark the tasks of Unitail with various off-the-shelf methods, including [64, 54, 53, 31, 36, 56, 42, 19, 48, 47, 60, 30, 29, 12, 21, 20, 50, 39]. **(3)** we design two baselines for product detection and text-based product matching.

2 Related Work

The retail scene has drawn attention to the computer vision community for an extended period. The early evolution of product datasets [35, 27, 45] facilitates reliable training and evaluation and drives research in this challenging field. Recently, large-scale datasets [5, 17] has enabled deep learning based approaches. Datasets related to SPR can be split into two groups: product detection and product recognition. We address each in this section.

Product Detection Datasets Detection is the first step in SPR entailing the presence of products that are typically represented by rectangular bounding boxes. The GroZi-120 [35] was created using in situ and in vitro data to study the product detection and recognition with 120 grocery products varying in color, size, and shape. The D2S [13] is an instance segmentation benchmark for sparsely-placed product detection, with 21000 images and 72447 instances. The SKU110k[17] and the SKU110k-r[40] provide 11762 images with 1.7M on-shelf product instances; they are annotated with axis-aligned bounding boxes and rotated bounding boxes, respectively.

For the detection of coarsely categorized products, The RPC [55] is designed for checkout scenarios containing 0.4M instances and 200 products. The Grocery Shelves dataset [52] took 354 shelf images and 13k instances, and around 3k instances are noted in 10 brand classes. The Locount [2] simultaneously considers the detection and counting for groups of products, with an impressive number of 1.9M instances and 140 categories.

Despite the significant values of these datasets, we are still challenged by the availability of optimal bounding boxes. In the proposed Unitail-Det, we aim to shape products into quadrilaterals whose appearances are well aligned. The dataset also provides evaluation targets in the origin-domain and cross-domain, bringing algorithms closer to practical deployment.

Product Recognition Datasets Multiple product recognition datasets have been built over the past decades. Early ones have limitations in the diversity of categories and amount of images, such as the SOIL-47 [27] containing 47 categories, the Supermarket Produce Datasets [45] with 15 categories, and the

Freiburg Groceries [23] covering 25 categories. Recent collections like Products-6K[16] and Products-10K[1] focus on large-scale data, which satisfy the training of deep learning algorithms. AliProducts[5] is crawled from web sources by searching 50K product names, consequently containing 2.5 million noisy images without human annotations. The ABO[7] dataset covers 576 categories that studies the 3D object understanding. The Grozi-3.2K [14] contains 80 fine-grained grocery categories and 8350 images for multi-label classification. To the best of our knowledge, there is a lack of a dataset that encourages leveraging both visual and textual information for product recognition.

OCR on Retail Products Product recognition by texts is challenging. The lack of datasets obstructs the relevant research on this topic. The most relevant dataset is Products-6K[16] where the Google Cloud Vision API is employed to extract the textual information to enhance products’ descriptions. But the texts were not labelled by human annotators, and text location information is missing, so it is infeasible to support any advance to OCR related tasks.

There are a couple of attempts that use off-the-shelf OCR models for assisted shopping. [15] presented a system on which users search products by name and OCR models return the texts on products so that a product list ranked by word histogram is generated for the users. [38] recognize texts and then apply text embedding and language models to extract features for product verification.

Other Related Datasets Many OCR datasets [49, 37, 6, 24, 25, 34, 44, 59, 22, 18] exist prior to Unitail-OCR. Typically, an OCR dataset supports text detection and text recognition and so enables text spotting. The ICDAR2015[24] and CTW1500[34] are two widely applied benchmarks for training and evaluating OCR models in common scene. The ICDAR2015 has 1,000 training and 500 testing images causally shot indoor and outdoor with word-level annotations for each text. The CTW1500 contains 10k text (3.5k of them are curved boxes) in 1500 images collected from the internet. Compared to them, Unitail-OCR is the first that focuses on product texts and supports object-level matching task at the same time. Product texts are usually artistic words with substantial character distortions, which are hard to be localize and recognize.

3 The Unitail

3.1 Unitail-Det

Image Collection Practically, the industry utilizes a variety of sensors under different conditions for product detection. The resolution and camera angles cover an extensive range by different sensors. For example, fixed cameras are mounted on the ceiling in most cases, and customers prefer to photograph with mobile devices. The product categories in different stores also span a great range. With these factors in mind, we collect images from two sources to support origin-domain and cross-domain detection. *In the origin domain*, training and testing



Fig. 2: Quadrilateral (in green) is a nature fit to product in real scene, removing more noisy contexts than AABB (in violet) and RBOX (in red)).

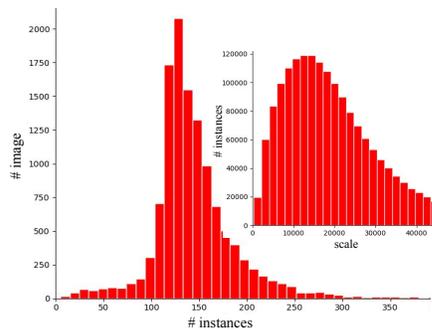


Fig. 3: Unitail-Det statistics. Left: instance density. Right: instance scale

images are supposed to share the same domain and are taken from similar perspectives in the same stores by the same sensors. As a result, we select the 11,744 images from the prior largest product dataset, SKU110k [17], to form the origin domain. *In the cross domain*, we collect 500 images in different stores through multiple sensors, covering unseen categories and camera angles.

Annotation We annotate each product with a quadrilateral style bounding box, denoted as QUAD. Fig.2 is an illustration of its advance. A QUAD refers to 4 points $p_{tl}, p_{tr}, p_{br}, p_{bl}$ with 8 degrees of freedom $(x_{tl}, y_{tl}, x_{tr}, y_{tr}, x_{br}, y_{br}, x_{bl}, y_{bl})$. For regular products shaped mainly in cuboid and cylinder, the (x_{tl}, y_{tl}) is defined as the top-left corners of their frontal faces, and the other points represent the rest 3 corners in clockwise order. For spherical, cones, and other shapes whose corners are difficult to identify, and for irregularly shaped products where so defined quadrilateral box cannot cover the entire frontal face, we first draw the minimum AABB and then adjust the four corners according to the camera perspective. As belabored, the frontal face of a product has the most representative information and is also critical for appearance consistency, but we still annotate the side face if the front face is invisible.

Totally, 1,777,108 QUADs are annotated by 13 well-trained annotators in 3 rounds of verification. The origin-domain is split to training (8,216 images, 1,215,013 QUADs), validation (588 images, 92,128 QUADs), and origin-domain testing set (2,940 images, 432,896 QUADs). The cross-domain supports a testing set (500 images, 37,071 QUADs). Their density and scale are shown in Fig.3

3.2 Unitail-OCR

Gallery and Testing Suite A product gallery setup is a common practice in the retail industry for product matching applications. All known categories are first registered in the gallery. In case of a query product, the matching algorithms find the top ranked category in the gallery. The gallery of Unitail-OCR contains

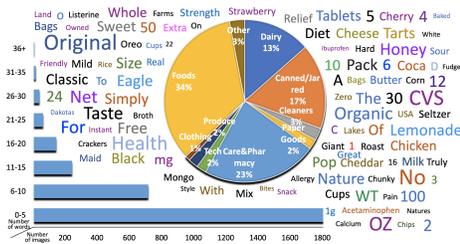


Fig. 4: Unitail-OCR statistical graphic. The pie chart reflects sections that source images were collected from. The bar chart is a histogram for the count of words on products. The font size of the words reflects the frequency of occurrence.

1454 fine-grained and one-shot product categories. Among these products, 10709 text regions and 7565 legible text transcriptions (words) are annotated. This enables the gallery to act as the training source and the matching reference.

The testing suite contains four components: 1. 3012 products labeled with 18972 text regions for text detection. 2. Among the pre-localized text regions, 13416 legible word-level transcriptions for text recognition. 3. 10k product samples from the 1454 categories for general evaluation on product matching. 4. From the 10k products, we select 2.4k fine-grained samples (visually similar for humans) for hard-example evaluation on product matching.

Image Collection and Annotation Images are gathered from the Unitail-Det cross-domain and cropped and affine transformed according to the quadrilateral bounding boxes to form an upright appearance. We remove the low-quality images with low resolution and high blurriness. Some products kept in the Unitail-OCR might exclude text regions, like those from the produce and clothes departments. We randomly select one sample from each category to form the product gallery, and the remaining is further augmented by randomly adjusting the brightness and cropping for matching purposes.

We annotate 29681 text regions from 4466 products as quadrilateral text boxes. Fig.4 shows the statistics. The bounding boxes are first classified as *legible* or *illegible*. For the 20981 legible ones, the alphanumeric transcriptions are annotated ignoring letter case and symbols. Numerical values with units are commonly seen on products such as *120mg*, and we regard them as entire words. We also provide a vocabulary that covers all words present. The usage of vocabulary is more practical in our case than in other scenes [24], because the presence of products and texts are usually known in advance by the store owner.

3.3 Tasks and Evaluation Metrics

Product Detection Task The goal is to detect products as quadrilaterals from complex backgrounds. Unitail-Det supports the training and evaluation.

We use the geometric mean of mean average precision (mAP) calculated on the origin-domain test set and cross-domain test set as the primary metric for the product detection, where the mAP is calculated in MS-COCO style[33]. Compared to arithmetic mean, the geometric mean is more sensitive when the model overfits to origin-domain but gains low performance on the cross-domain.

Text Detection Task The goal is to detect text regions from pre-localized product images. Unitail-OCR supports the training and evaluation.

We adopt the widely used precision, recall and hmean [24, 34] for evaluation.

Text Recogniton Task The goal is to recognize words over a set of pre-localized text regions. Unitail-OCR supports the training and evaluation.

We adopt the normalized edit distance (NED) [25] and word-level accuracy for evaluation. The edit distance between two words is defined by the minimum number of characters edited (insert, delete or substitute) required to change one into the other, then it is normalized by the length of the word and averaged on all ground-truths.

Product Matching Task The goal is to recognize products by matching a set of query samples to the Unitail-OCR gallery. The task is split into two tracks: **Hard Example Track**, which is evaluated on 2.5k selected hard examples; this track is designed for scenarios in which products are visually similar (for example pharmacy stores). And **General Track**, which is conducted on all 10k samples.

We adopt the top-1 accuracy as the evaluation metric.

4 Two Baselines Designed for The Unitail

4.1 A Customized Detector for Product Detection

Recent studies [61, 51, 62, 26, 4] on generic object detection apply prior-art DenseBox-style head [21] to multiple feature pyramid levels. The feature pyramid is generated via feature pyramid network (FPN) [32] and contains different levels that are gradually down-sampled but semantically enhanced. An anchor-free detection head is then attached to classify each pixel on the feature pyramid and predict axis-aligned bounding boxes (AABB).

During training, assigning ground-truths to each feature pixels on the feature pyramid plays a key role. *On each pyramid level*, the centerness [51] is widely used. It is an indicator to value how far a pixel locates from the center of a ground-truth: the farther, the more likely it is to predict an inaccurate box, and the lower centerness score it gains. *Across pyramid levels*, various strategies are proposed to determine which level should be assigned, and they are grouped into scale-based and loss-based strategies. The scale-based [32, 26, 43, 51] assigns ground-truths to different levels in terms of their scales. The larger scale, the higher level is assigned so that the needs of receptive field and resolution of feature maps are balanced. The loss-based like Soft Selection [62] assigns ground-truths by calculating their losses on all levels, and trains an auxiliary network that re-weights the losses. Our design, RetailDet, adopts the DenseBox style architecture but predicts the four corners of quadrilateral by 8-channel regression head. During training, we found the prior assignment strategies unsuitable for quadrilateral products, which is specified below.

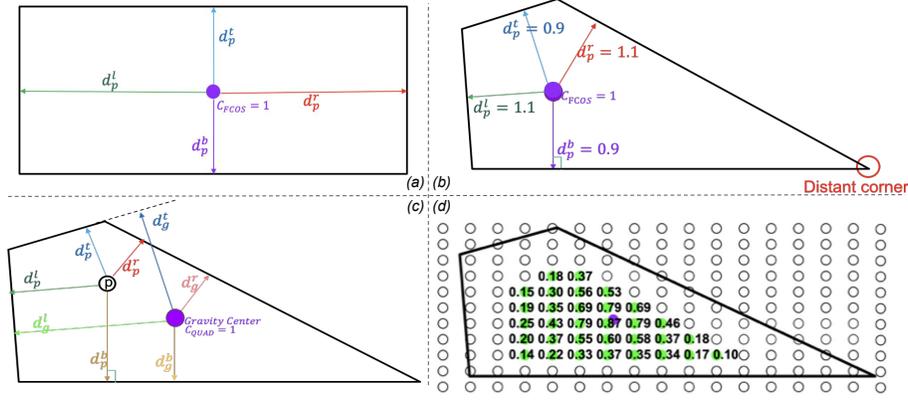


Fig. 5: (a): C_{FCOS} on AABB. (b): C_{FCOS} on QUAD. (c) and (d): C_{QUAD} .

Centerness The previous definition of the centerness [51, 62] is shown in Eq.1,

$$C_{FCOS}(p) = \left[\frac{\min(d_p^l, d_p^r)}{\max(d_p^l, d_p^r)} \cdot \frac{\min(d_p^t, d_p^b)}{\max(d_p^t, d_p^b)} \right]^{0.5} \quad (1)$$

by the Eq.1 and Fig.5(a), a location p keeps the same distance to the left/right boundaries ($d_p^l = d_p^r$) and to the top/bottom boundaries ($d_p^t = d_p^b$) will gain the highest centerness 1, and other pixels gain degraded score by Eq.1.

Limitation When adopting the same centerness to quadrilaterals, as shown in Fig.5(b), the center can be far away from a distant corner, which leads to unbalanced regression difficulty and lack of receptive field from that corner.

Our Solution We first re-define the *center* as the center of gravity (Fig.5(c)), because it is the geometric center and represents the mean position of all the points in the shape, which mitigates the unbalanced regression difficulties. We then propose Eq.2 to calculate the quad-centerness for any p ,

$$C_{QUAD}(p) = \left[\frac{\min(d_p^l, d_g^l)}{\max(d^l, d_g^l)} \cdot \frac{\min(d_p^r, d_g^r)}{\max(d^r, d_g^r)} \cdot \frac{\min(d_p^t, d_g^t)}{\max(d^t, d_g^t)} \cdot \frac{\min(d_p^b, d_g^b)}{\max(d^b, d_g^b)} \right]^{0.5} \quad (2)$$

where the $d_g^{l/r/t/b}$ denotes the distances between the gravity center g and the left/right/top/bottom boundaries. The $d_p^{l/r/t/b}$ denotes the distances between the p and the boundaries. If p locates on the gravity center, its quad-centerness gains the highest value as 1. Otherwise, it is gradually degraded (See Fig.5 (d)).

It is mentionable that when applied to AABB, Eq.2 is mathematically equivalent to Eq.1, which is proved in supplementary.

Soft Selection The loss-based Soft Selection in [62] outperforms scale-based strategies on generic objects because it assigns ground-truths to multiple levels

and re-weights their losses. This is achieved by calculating losses for each object on all levels and using the losses to train an auxiliary network that predicts the re-weighting factors.

Limitation Instances per image are numerous in densely-packed retail scene, and Soft Selection is highly inefficient (5×slower) due to the auxiliary network.

Our Solution Can we maintain the merit of Soft Selection while accelerating the assignment? We approach this issue by mimicking the loss re-weighting mechanism of the auxiliary network using scale-based calculation. This is feasible because we find the Soft Selection, in essence, follows scale-based law (detailed in supplementary). Thus, we design Soft Scale (SS) in Eq.3,4,5,6. For an arbitrary shaped object O with area $area_O$, SS assigns it to two adjacent levels l_i and l_j by Eq.3,4 and calculates the loss-reweighting factors F_{l_i} , F_{l_j} by Eq.5,6.

$$l_i = \lceil l_{org} + \log_2(\sqrt{area_O}/224) \rceil \quad (3)$$

$$l_j = \lfloor l_{org} + \log_2(\sqrt{area_O}/224) \rfloor \quad (4)$$

$$F_{l_i} = \log_2(\sqrt{area_O}/224) - \lfloor \log_2(\sqrt{area_O}/224) \rfloor \quad (5)$$

$$F_{l_j} = 1 - F_{l_i} \quad (6)$$

where 224 is the ImageNet pre-training size. Objects with exact area 224^2 is assigned to l_{org} , in which case $l_i = l_j = l_{org}$. If an object is with area 223^2 , SS assigns it to l_{org} with $F_{l_{org}} = 0.994$, and also to $(l_{org} - 1)$ with $F_{(l_{org}-1)} = 0.006$. In this work we fix l_{org} to be level 5 of feature pyramid. SS operates rapidly as scale-based strategies and keeps the loss-reweighting like Soft Selection.

4.2 A Simple Baseline for Product Matching

A number of directions can be explored on this new task, while in this paper, we design the simplest solution to verify the motivation: *people glance and recognize the product, and if products looks similar, they further scrutinize the text (if appears) to make decision*. To this end, we first apply a well-trained image classifier that extracts visual features $f_{g_i}^v$ from each gallery image g_i and feature f_p^v from query image p , and calculate the cosine similarity between each pair $(f_{g_i}^v, f_p^v)$ (termed as sim_i^v). If the highest ranking value sim_1^v and the second highest sim_2^v are close ($sim_1^v - sim_2^v \leq t$), we then read on products and calculate the textual similarity (termed as sim^t) to make decision by Eq.7,

$$Decision = \operatorname{argmax}_{i \in [1,2]} [w \cdot sim^t(g_i, p) + (1 - w) \cdot sim_i^v] \quad (7)$$

where threshold t and coefficient w are tuned on validation set.

Our design focuses on how to calculate sim^t . We denote the on-product texts obtained from ground-truth or OCR prediction as $S = \{s^1, s^2, \dots, s^N\}$ where N varies. People may propose to utilize sequence-to-one models (like BERT[8]) to encode S into a fixed length feature vector $f \in \mathbb{R}^d$. As shown in Fig.6 (a), a text detector is followed by a text recognizer predicting $n = 5$ words, and the

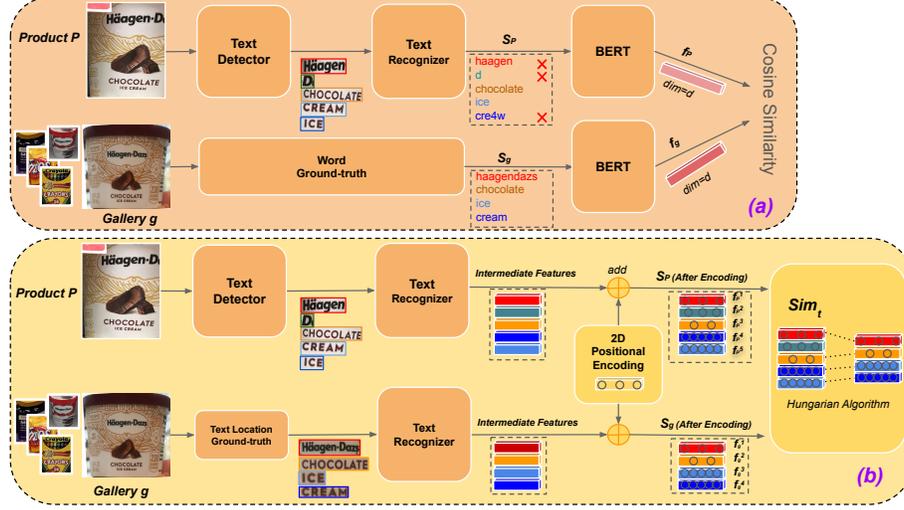


Fig. 6: (a) Pipeline with BERT encoded features, (b) Pipeline with positional encoding and Hungarian Algorithm based textual similarity.

5 words are fed into the BERT to encode a feature vector $f_p \in \mathbb{R}^d$. For each gallery image g , the same process is operated to get a feature vector $f_g \in \mathbb{R}^d$, and $sim^t(f_p, f_g)$ is calculated by the cosine similarity.

But this design does not perform well because errors from OCR models (especially from text recognizer) are propagated to the BERT causing poor feature encoding. Moreover, the positional information of text boxes is lost in the sequence. So we design a new method in Fig.6 (b). Rather than using the n recognized words, we use the n intermediate feature vectors from the text recognizer to mitigate propagated errors. For example, *CREAM* is confused as *CRE4w*, but the intermediate feature should maintain information on *A* and *M*, which is more robust than the false prediction. Each feature is then added by a 2D positional encoding[41, 3] whose calculation is based on the location of the corresponding text. It encodes the spatial information into the feature and it is predefined to keep the same dimension as the intermediate feature. Finally, we get a sequence that contains the n encoded features $f^{1 \sim n}$. As shown we get $S_p = \{f_p^1, f_p^2, f_p^3, f_p^4, f_p^5\}$ from a query product and $S_g = \{f_g^1, f_g^2, f_g^3, f_g^4\}$ from a gallery reference. Inspired by the Hungarian Algorithm [28], we design Eq.8 to directly calculate the similarity between two sequence with varying length:

$$sim^t(p, g) = sim^t(S_p, S_g) = \max_X \sum_{i=1}^n \sum_{j=1}^m \left(\frac{f_p^i \cdot f_g^j}{|f_p^i| \cdot |f_g^j|} X_{ij} \right) \quad (8)$$

where the X is a $n \times m$ boolean matrix where $\sum_j X_{ij} = 1, \sum_i X_{ij} = 1$. Eq.8 maximizes the summation of cosine similarities from assigned feature pairs, and the assignment is optimized by X .

5 Experiments

5.1 Benchmarking The Unitail

The implementation is detailed in the supplementary materials.

Quadrilateral Product Detection Along with the proposed RetailDet, we build baselines by borrowing the existing detectors capable of detecting quadrilaterals, mainly from the textual and aerial scenes[9]. These methods can be grouped into segmentation based methods [64, 54, 53, 31, 19] and regression based methods [36, 56, 42]. Segmentation methods consider QUADs as per-pixel classification masks for each region of interest, while regression methods directly regress the bounding boxes.

Table 1 shows their performances. Overall, regression based methods outperform segmentation based methods because most products can be well aligned by quadrilaterals, and the learning on segmentation mask involves extra difficulties. The RetailDet outperforms other detectors by a large margin. All detectors achieve degraded results in the cross-domain as opposed to the origin-domain, confirming that domain shift exists among stores.

Table 1: Benchmarking product detection on the Unitail. All methods are trained and tested under same setting. g-mAP is the geometric mean of mAPs.

# Method	Backbone	g-mAP	Origin-Domain			Cross-Domain		
			mAP	AP50	AP75	mAP	AP50	AP75
1 FCENet [64]	ResNet50	32.0	36.8	76.0	31.2	27.9	60.1	22.6
2 PANet [54]	ResNet50	35.0	40.5	72.8	41.9	30.3	53.3	31.6
3 PSENet [53]	ResNet50	39.4	45.3	77.5	49.5	34.4	58.7	36.9
4 DBNet [31]	ResNet50	45.3	51.0	86.8	55.4	40.3	71.6	42.7
5 RIDet [36]	ResNet50	45.7	51.2	82.9	58.5	40.8	70.3	43.2
6 Gliding Vertex[56]	ResNet50	46.0	52.3	89.0	56.9	40.5	76.7	38.6
7 RSDet [42]	ResNet50	46.1	51.4	83.6	58.8	41.4	71.1	44.4
8 Mask-RCNN [19]	ResNet50	52.4	57.3	91.6	66.0	48.0	77.9	53.2
9 RetailDet (ours)	ResNet50	54.7	58.7	91.6	68.4	50.9	80.6	56.7
10 RetailDet	ResNet101	57.1	60.3	92.8	70.6	54.1	83.5	60.6

Text Detection & Text Recognition We benchmark the tasks of text detection and text recognition in Table 2 and Table 3, respectively. For each of the listed algorithms, it is trained under two setting: one is following a common setting that trained on SynthText[18] and ICDAR2015(IC15)[24] for text detection and on Synth90K[22] and SynthText for text recognition, another is trained/finetuned on the Unitail. As shown all algorithms achieve better performance if trained on the Unitail, this verifies that texts in the retail domain are better handled by the proposed dataset.

Method	Training Set	R	P	hmean
DBNet[31]	SynthText+IC15	0.541	0.866	0.666
DBNet	Unitail	0.773	0.871	0.819
FCENet[64]	SynthText+IC15	0.420	0.745	0.538
FCENet	Unitail	0.795	0.857	0.825
PSENet[53]	SynthText+IC15	0.421	0.750	0.539
PSENet	Unitail	0.705	0.789	0.744

Table 2: Benchmarking text detection on Unitail. P and R stand for Precision and Recall, respectively. hmean is the harmonic mean of precision and recall.

Method	Training Set	PW	NED	Acc(%)
CRNN[48]	S90k+ST	✓	0.36	40.0
CRNN	S90k+ST+Unitail		0.25	51.4
NRTR[47]	S90k+ST	✓	0.28	55.7
NRTR	S90k+ST+Unitail		0.16	69.4
RobustScanner[60]	S90k+ST	✓	0.25	56.3
RobustScanner	S90k+ST+Unitail		0.18	65.9
SAR[30]	S90k+ST	✓	0.25	56.2
SAR	S90k+ST+Unitail		0.18	66.5
SATRN[29]	S90k+ST	✓	0.23	62.7
SATRN	S90k+ST+Unitail		0.13	74.9
ABINet[12]	S90k+ST	✓	0.17	69.2
ABINet	S90k+ST+Unitail		0.11	77.2

Table 3: Benchmarking text recognition on Unitail. S90k: Synth90k. ST: SynthText. PW: methods use public available weights. NED: Normalized Edit Distance, the lower the better. Acc: word top-1 Accuracy, the higher the better.

Product Matching The product matching results are shown in Table 4. With just texts across all 1454 categories, the method in Fig.6(b) reaches 31.71% and 47.81% on the Hard Example Track and the General Track, respectively. The result is convincing that only textual information is a strong representation for the product, but this new direction clearly requires further exploration.

Moreover, the textual information improves the regular visual classifier using the method proposed in Sec.4.2. In the hard example track, the improvement of textual information is significant (+1.76 ~ 2.75%) since the similar-looking products which are hard for regular classifier could be easier distinguished by texts. In the general track, the improvement drops (+0.56 ~ 0.94%) due to the limited ability of the textual matching algorithm.

5.2 Discussion

RetailDet Table 5 shows that the RetailDet achieves state-of-the-art results on product detection benchmark SKU110k [17]. The SKU110k and the Unitail-Det share the same training images but with different annotations (QUAD vs AABB), and the origin-domain mAP on the SKU110k is 59.0 and 61.3 on the Unitail-Det (not included in Table.1). The difference is mainly for the reason that QUAD is a natural fit for products as a smoother regression target, and it avoids densely-packed overlaps which makes post-processing easier. We also pro-

Table 4: Benchmarking on the product matching task.

Method	Acc (%)
Hard Example:	
Only text (Fig.6 (b))	31.71
EfficientNetV2 [50]	56.49
EfficientNetV2+Text	59.24 (+2.75)
ResNet101 [20, 39]	58.37
ResNet101+Text	60.52 (+2.15)
General:	
Only Text(Fig.6 (a))	30.37
Only Text(Fig.6 (b))	47.81 (+17.44)
EfficientNetV2	83.81
EfficientNetV2+Text	84.62 (+0.81)
ResNet101	85.03
ResNet101+Text	86.19 (+1.16)

Table 5: Results on SKU110k. RetailDet++ is an enhanced variant where a box refinement module is added (see supplementary for details)

Method	mAP
RetinaNet+EM[17]	49.2
FSAF [63]	56.2
CenterNet [10]	55.8
ATSS [61]	57.4
FCOS [51]	54.4
SAPD [62]	55.7
Faster-RCNN+FPN [43]	54.0
Reppoints [58]	55.6
Cascade-RCNN+Rong [46]	58.7
RetailDet++ (Ours)	59.0

vide ablation study on the quad-centerness and soft selection in supplementary materials.

Difficulty of The Text Based Product Matching Since the proposed matching pipeline is not end-to-end trained, errors are accumulated by the text detector (0.773 recall and 0.871 precision for DBNet) and text recognizer (0.772 for ABINet). Some products in the dataset do not contain texts (2%), and many words are partially or fully illegible. This requires further study on topics such as: attention mechanism on text regions, end-to-end framework for text spotting based product recognition, and one-shot learning based sequence matching.

Inference Speed The inference speed is tested on a single 2080Ti GPU, with 6.3 FPS for RetailDet, 27.6 FPS for DBNet, 41.4 FPS for ABINet, and 65.1 FPS for text+visual matching.

6 Conclusions

In this work, we introduce the United Retail Datasets (Unitail), a large-scale benchmark aims at supporting well-aligned textually enhanced scene product recognition. It involves quadrilateral product instances, on-product texts, product matching gallery, and testing suite. We also design two baselines that take advantages of the Unitail and provide comprehensive benchmark experiments on various state-of-the-art methods.

References

1. Bai, Y., Chen, Y., Yu, W., Wang, L., Zhang, W.: Products-10k: A large-scale product recognition dataset. CoRR **abs/2008.10545** (2020), <https://arxiv.org/abs/2008.10545>
2. Cai, Y., Wen, L., Zhang, L., Du, D., Wang, W.: Rethinking object detection in retail stores. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 947–954 (2021)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. CoRR **abs/2005.12872** (2020), <https://arxiv.org/abs/2005.12872>
4. Chen, F., Zhu, C., Shen, Z., Zhang, H., Savvides, M.: Ncms: Towards accurate anchor free object detection through l2 norm calibration and multi-feature selection. *Comput. Vis. Image Underst.* **200**, 103050 (2020)
5. Cheng, L., Zhou, X., Zhao, L., Li, D., Shang, H., Zheng, Y., Pan, P., Xu, Y.: Weakly supervised learning with side information for noisy labeled images. CoRR **abs/2008.11586** (2020), <https://arxiv.org/abs/2008.11586>
6. Chng, C.K., Chan, C.S.: Total-text: A comprehensive dataset for scene text detection and recognition. CoRR **abs/1710.10400** (2017), <http://arxiv.org/abs/1710.10400>
7. Collins, J., Goel, S., Luthra, A., Xu, L., Deng, K., Zhang, X., Vicente, T.F.Y., Arora, H., Dideriksen, T., Guillaumin, M., Malik, J.: ABO: dataset and benchmarks for real-world 3d object understanding. CoRR **abs/2110.06199** (2021), <https://arxiv.org/abs/2110.06199>
8. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. CoRR **abs/1810.04805** (2018), <http://arxiv.org/abs/1810.04805>
9. Ding, J., Xue, N., Xia, G.S., Bai, X., Yang, W., Yang, M.Y., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L.: Object detection in aerial images: A large-scale benchmark and challenges (2021)
10. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Object detection with keypoint triplets. arXiv preprint arXiv:1904.08189 (2019)
11. Everingham, M., Gool, L.V., Williams, C.K.I., Winn, J.M., Zisserman, A.: The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* **88**(2), 303–338 (2010), <https://doi.org/10.1007/s11263-009-0275-4>
12. Fang, S., Xie, H., Wang, Y., Mao, Z., Zhang, Y.: Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition (2021)
13. Follmann, P., Böttger, T., Härtinger, P., König, R., Ulrich, M.: Mvtec D2S: densely segmented supermarket dataset. CoRR **abs/1804.08292** (2018), <http://arxiv.org/abs/1804.08292>
14. George, M., Floerkemeier, C.: Recognizing products: A per-exemplar multi-label image classification approach. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *Computer Vision – ECCV 2014*. pp. 440–455. Springer International Publishing, Cham (2014)
15. George, M., Mircic, D., Sörös, G., Floerkemeier, C., Mattern, F.: Fine-grained product class recognition for assisted shopping (2015). <https://doi.org/10.48550/ARXIV.1510.04074>, <https://arxiv.org/abs/1510.04074>
16. Georgiadis, K., Kordopatis-Zilos, G., Kalaganis, F., Migkotzidis, P., Chatzilari, E., Panakidou, V., Pantouvakis, K., Tortopidis, S., Papadopoulos, S., Nikolopoulos, S., Kompatsiaris, I.: Products-6k: A large-scale groceries product recogni-

- tion dataset. In: The 14th PErvasive Technologies Related to Assistive Environments Conference. p. 1–7. PETRA 2021, Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3453892.3453894>, <https://doi.org/10.1145/3453892.3453894>
17. Goldman, E., Herzig, R., Eisenschtat, A., Goldberger, J., Hassner, T.: Precise detection in densely packed scenes. In: Proc. Conf. Comput. Vision Pattern Recognition (CVPR) (2019)
 18. Gupta, A., Vedaldi, A., Zisserman, A.: Synthetic data for text localisation in natural images. In: IEEE Conference on Computer Vision and Pattern Recognition (2016)
 19. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
 20. 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. pp. 770–778 (2016)
 21. Huang, L., Yang, Y., Deng, Y., Yu, Y.: Densebox: Unifying landmark localization with end to end object detection. arXiv preprint arXiv:1509.04874 (2015)
 22. Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Synthetic data and artificial neural networks for natural scene text recognition. In: Workshop on Deep Learning, NIPS (2014)
 23. Jund, P., Abdo, N., Eitel, A., Burgard, W.: The freiburg groceries dataset. CoRR **abs/1611.05799** (2016), <http://arxiv.org/abs/1611.05799>
 24. Karatzas, D., Gomez-Bigorda, L., Nicolaou, A., Ghosh, S., Bagdanov, A., Iwamura, M., Matas, J., Neumann, L., Chandrasekhar, V.R., Lu, S., Shafait, F., Uchida, S., Valveny, E.: Icdar 2015 competition on robust reading. In: Proceedings of the 2015 13th International Conference on Document Analysis and Recognition (ICDAR). p. 1156–1160. ICDAR '15, IEEE Computer Society, USA (2015). <https://doi.org/10.1109/ICDAR.2015.7333942>, <https://doi.org/10.1109/ICDAR.2015.7333942>
 25. Karatzas, D., Shafait, F., Uchida, S., Iwamura, M., Bigorda, L.G.i., Mestre, S.R., Mas, J., Mota, D.F., Almazàn, J.A., de las Heras, L.P.: Icdar 2013 robust reading competition. In: 2013 12th International Conference on Document Analysis and Recognition. pp. 1484–1493 (2013). <https://doi.org/10.1109/ICDAR.2013.221>
 26. Kong, T., Sun, F., Liu, H., Jiang, Y., Shi, J.: Foveabox: Beyond anchor-based object detector. arXiv preprint arXiv:1904.03797 (2019)
 27. Koubaroulis, D., Matas, J., Kittler, J.: Evaluating colour-based object recognition algorithms using the soil-47 database. In: in Asian Conference on Computer Vision. pp. 840–845 (2002)
 28. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval Research Logistics Quarterly* **2**(1-2), 83–97 (1955). <https://doi.org/https://doi.org/10.1002/nav.3800020109>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800020109>
 29. Lee, J., Park, S., Baek, J., Oh, S.J., Kim, S., Lee, H.: On recognizing texts of arbitrary shapes with 2d self-attention. CoRR **abs/1910.04396** (2019), <http://arxiv.org/abs/1910.04396>
 30. Li, H., Wang, P., Shen, C., Zhang, G.: Show, attend and read: A simple and strong baseline for irregular text recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8610–8617 (2019)
 31. Liao, M., Wan, Z., Yao, C., Chen, K., Bai, X.: Real-time scene text detection with differentiable binarization. Proceedings of the AAAI Conference on Artificial Intelligence pp. 11474–11481 (2020)

32. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017)
33. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
34. Liu, Y., Jin, L., Zhang, S., Zhang, S.: Detecting curve text in the wild: New dataset and new solution. CoRR **abs/1712.02170** (2017), <http://arxiv.org/abs/1712.02170>
35. Merler, M., Galleguillos, C., Belongie, S.: Recognizing groceries in situ using in vitro training data. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–8 (2007). <https://doi.org/10.1109/CVPR.2007.383486>
36. Ming, Q., Miao, L., Zhou, Z., Yang, X., Dong, Y.: Optimization for arbitrary-oriented object detection via representation invariance loss. IEEE Geoscience and Remote Sensing Letters pp. 1–5 (2021). <https://doi.org/10.1109/LGRS.2021.3115110>
37. Mishra, A., Alahari, K., Jawahar, C.: Scene Text Recognition using Higher Order Language Priors. In: BMVC - British Machine Vision Conference. BMVA, Surrey, United Kingdom (Sep 2012). <https://doi.org/10.5244/C.26.127>, <https://hal.inria.fr/hal-00818183>
38. Ouchekh, R., Pettersson, T., Löfström, T.: Product verification using ocr classification and mondrian conformal prediction. Expert Systems with Applications **188**, 115942 (2022). <https://doi.org/https://doi.org/10.1016/j.eswa.2021.115942>, <https://www.sciencedirect.com/science/article/pii/S0957417421012963>
39. Pan, X., Luo, P., Shi, J., Tang, X.: Two at once: Enhancing learning and generalization capacities via ibn-net. In: ECCV (2018)
40. Pan, X., Ren, Y., Sheng, K., Dong, W., Yuan, H., Guo, X., Ma, C., Xu, C.: Dynamic refinement network for oriented and densely packed object detection pp. 1–8 (2020)
41. Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A.: Image transformer. CoRR **abs/1802.05751** (2018), <http://arxiv.org/abs/1802.05751>
42. Qian, W., Yang, X., Peng, S., Yan, J., Guo, Y.: Learning modulated loss for rotated object detection. Proceedings of the AAAI Conference on Artificial Intelligence **35**(3), 2458–2466 (May 2021), <https://ojs.aaai.org/index.php/AAAI/article/view/16347>
43. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015)
44. Risnumawan, A., Shivakumara, P., Chan, C.S., Tan, C.L.: A robust arbitrary text detection system for natural scene images. Expert Systems with Applications **41**(18), 8027–8048 (2014). <https://doi.org/https://doi.org/10.1016/j.eswa.2014.07.008>, <https://www.sciencedirect.com/science/article/pii/S0957417414004060>
45. Rocha, A., Hauagge, D.C., Wainer, J., Goldenstein, S.: Automatic fruit and vegetable classification from images. Computers and Electronics in Agriculture **70**(1), 96–104 (2010). <https://doi.org/https://doi.org/10.1016/j.compag.2009.09.002>, <https://www.sciencedirect.com/science/article/pii/S016816990900180X>
46. Rong, T., Zhu, Y., Cai, H., Xiong, Y.: A solution to product detection in densely packed scenes (2021)
47. Sheng, F., Chen, Z., Xu, B.: Nrtr: A no-recurrence sequence-to-sequence model for scene text recognition. In: 2019 International Conference on Document Analysis and Recognition (ICDAR). pp. 781–786. IEEE (2019)

48. Shi, B., Bai, X., Yao, C.: An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence* (2016)
49. Singh, A., Pang, G., Toh, M., Huang, J., Galuba, W., Hassner, T.: TextOCR: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text (2021)
50. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training. *CoRR* **abs/2104.00298** (2021), <https://arxiv.org/abs/2104.00298>
51. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. *arXiv preprint arXiv:1904.01355* (2019)
52. Varol, G., Kuzu, R.: Toward retail product recognition on grocery shelves. In: *International Conference on Graphic and Image Processing* (2015)
53. Wang, W., Xie, E., Li, X., Hou, W., Lu, T., Yu, G., Shao, S.: Shape robust text detection with progressive scale expansion network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9336–9345 (2019)
54. Wang, W., Xie, E., Song, X., Zang, Y., Wang, W., Lu, T., Yu, G., Shen, C.: Efficient and accurate arbitrary-shaped text detection with pixel aggregation network. In: *ICCV*. pp. 8439–8448 (2019)
55. Wei, X., Cui, Q., Yang, L., Wang, P., Liu, L.: RPC: A large-scale retail product checkout dataset. *CoRR* **abs/1901.07249** (2019), <http://arxiv.org/abs/1901.07249>
56. Xu, Y., Fu, M., Wang, Q., Wang, Y., Chen, K., Xia, G.S., Bai, X.: Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
57. Yang, S., Luo, P., Loy, C.C., Tang, X.: Wider face: A face detection benchmark. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
58. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. *arXiv preprint arXiv:1904.11490* (2019)
59. Yao, C., Bai, X., Liu, W., Ma, Y., Tu, Z.: Detecting texts of arbitrary orientations in natural images. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1083–1090 (2012). <https://doi.org/10.1109/CVPR.2012.6247787>
60. Yue, X., Kuang, Z., Lin, C., Sun, H., Zhang, W.: Robustscanner: Dynamically enhancing positional clues for robust text recognition. In: *European Conference on Computer Vision* (2020)
61. Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection (2020)
62. Zhu, C., Chen, F., Shen, Z., Savvides, M.: Soft anchor-point object detection. In: *Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision – ECCV 2020*. pp. 91–107. Springer International Publishing, Cham (2020)
63. Zhu, C., He, Y., Savvides, M.: Feature selective anchor-free module for single-shot object detection. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
64. Zhu, Y., Chen, J., Liang, L., Kuang, Z., Jin, L., Zhang, W.: Fourier contour embedding for arbitrary-shaped text detection. In: *CVPR* (2021)