

Towards Automatic Power Battery Detection: New Challenge, Benchmark Dataset and Baseline

Xiaoqi Zhao^{1,2†}, Youwei Pang^{1,2†}, Zhenyu Chen¹, Qian Yu¹,
 Lihe Zhang^{1*}, Hanqi Liu², Jiaming Zuo^{2*}, Huchuan Lu¹

¹Dalian University of Technology ²X3000 Inspection Co., Ltd

{zqx, lartpang, dluczy, ms.yuqian}@mail.dlut.edu.cn, {jerry, klaus}@3000gy.com
 {zhanglihe, lhchuan}@dlut.edu.cn

Abstract

We conduct a comprehensive study on a new task named power battery detection (PBD), which aims to localize the dense cathode and anode plates endpoints from X-ray images to evaluate the quality of power batteries. Existing manufacturers usually rely on human eye observation to complete PBD, which makes it difficult to balance the accuracy and efficiency of detection. To address this issue and drive more attention into this meaningful task, we first elaborately collect a dataset, called X-ray PBD, which has 1,500 diverse X-ray images selected from thousands of power batteries of 5 manufacturers, with 7 different visual interference. Then, we propose a novel segmentation-based solution for PBD, termed multi-dimensional collaborative network (MDCNet). With the help of line and counting predictors, the representation of the point segmentation branch can be improved at both semantic and detail aspects. Besides, we design an effective distance-adaptive mask generation strategy, which can alleviate the visual challenge caused by the inconsistent distribution density of plates to provide MDCNet with stable supervision. Without any bells and whistles, our segmentation-based MDCNet consistently outperforms various other corner detection, crowd counting and general/tiny object detection-based solutions, making it a strong baseline that can help facilitate future research in PBD. Finally, we share some potential difficulties and works for future researches. The source code and datasets will be publicly available at [X-ray PBD](#).

1. Introduction

With the development of power battery technology, new energy vehicles are receiving more and more attention. The power battery is the only source of driving energy for bat-

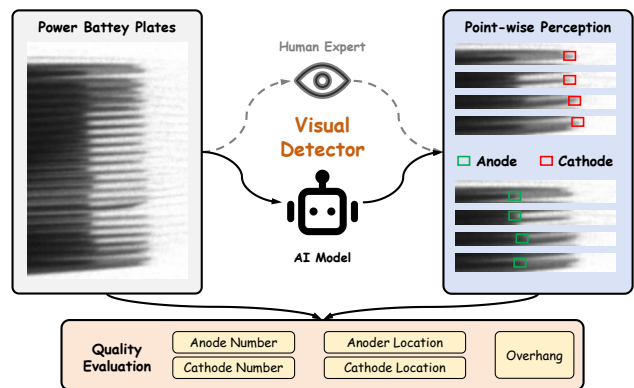


Figure 1. Illustration of the power battery detection task.

tery electric vehicle (BEV), which directly affects the power performance, endurance and safety of BEV [49]. To ensure the safety of power battery, the functional evaluation has to be done through power battery detection (PBD). As shown in Fig. 1, the PBD can provide accurate coordinate information for all anode and cathode endpoints. With the help of digital radiography (DR) device, the internal shape of the battery cell can be obtained from the X-ray image, thereby providing humans with the number and overhang¹ information to judge whether the current battery cell is OK or NG. Most factories evaluate the quality of power batteries through human eyes. However, visual fatigue will be aggravated with long-term work. Although some top battery factories have a large number of re-inspection technicians, this PBD manner is high-cost and low-efficiency. T

We believe that it is imminent to explore an intelligent PBD model. To this end, we build a X-ray PBD dataset via the high-quality industry DR device under thousands of

¹Number and overhang are two criteria for power battery manufacturers to evaluate battery quality. Number: The number of both anode and cathode plates. Overhang: The average value of the vertical distance between a single anode plate and two adjacent cathode plates.

† Equal contribution.

* Corresponding author.

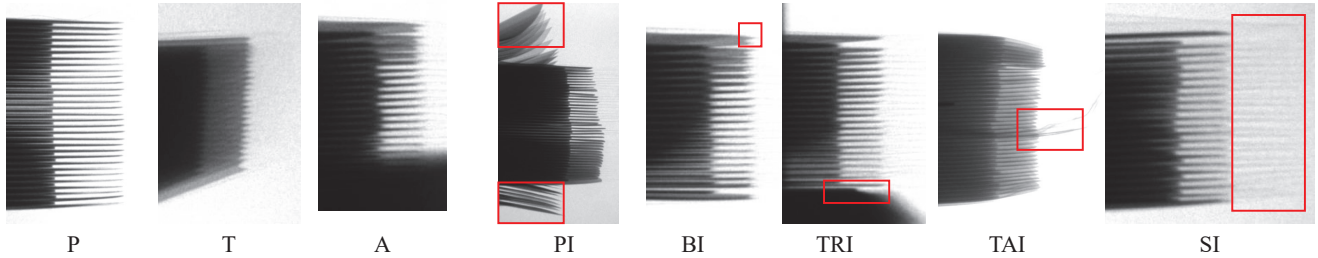


Figure 2. Examples of various attributes from our X-ray PBD dataset (best viewed zoomed in). See Tab. 1 for details.

different batteries. It has several characteristics: (1) rich shapes and complex interference (See Fig. 2); (2) vertex coordinates of plates as the annotated ground truth (GT); (3) widely distributed number of plates; (4) widely distributed overhangs between each pair of anode and cathode plates; (5) various views with long, medium and close shots. The main challenges for computer vision community are: (1) **Open vision modeling problem.** Like emerging tasks such as character stroke extraction [27], insubstantial object detection [63], open vocabulary objects detection and segmentation [9, 53], how to model a suitable AI-based solution is critical for the PBD task. (2) **Improving the discriminativeness of weak features in deep networks.** This task aims to accurately locate each independent plate endpoint (1 pixel) from the entire image (millions of pixels), which likes finding a needle in a haystack. In Fig. 2, we can see that the features of plate endpoints are weak and similar with around regions. It is difficult to capture all the endpoints in the face of interference from the plate, bifurcation and separator. Therefore, how to leverage useful cues to boost weak features is another important challenge.

In this paper, we propose a multi-dimensional collaborative network (MDCNet) for accurate PBD. **First**, we cast the PBD problem into the segmentation task and build a multi-task framework on multiple dimensions of point, line and number. Point segmentation predictor produces the position of the endpoint of each plate one by one in the form of a point map. Line segmentation and counting predictors are guided by the point prediction map to separately aggregate low-level and high-level features to predict the trend line and number of plates. Here, the point segmentation branch is the core, and the other two auxiliary tasks utilize spatial and semantic cues to simultaneously fine-tune the point branch. **Next**, we choose the clear X-ray Image with a standard battery shape as the prompt to generate the prompt filter for the point prediction branch. In this way, our PBD model can have the anti-interference ability and enhance the discrimination of plates. **Last**, we calculate the spatial distance of adjacent homopolar plates based on the annotations and use it as the reference radius of the point map to adaptively generate the ground truth for point segmentation.

Our main contributions can be summarized as follows:

Attr	Description
P	<i>Pure Plate.</i> High-quality sample without any internal and external interference.
T	<i>Tilted Plate.</i> Deformation caused by excessively dense plates.
A	<i>Aberrant Plate.</i> 1) Plates visualization is incomplete due to occlusion. 2) Plates are not arranged in the order of anode → cathode → anode.
PI	<i>Plate Interference.</i> Other batteries blend into current view.
BI	<i>Bifurcation Interference.</i> A single plate produces a bifurcation.
TRI	<i>Tray Interference.</i> A battery tray is a device used for holding large battery packs in place. Close distance between the side plate and tray may cause visual interference.
TAI	<i>Tab Interference.</i> Battery tabs are the anode and cathode connectors that carry the cells' electrical current.
SI	<i>Separator Interference.</i> A battery separator is a type of polymeric membrane that is positioned between the anode and cathode.

Table 1. Attribute descriptions (see examples in Fig. 2).

- We propose a new challenging task named power battery detection (PBD) and construct a complex PBD dataset, design an effective baseline, formulate comprehensive metrics, and explore label generation strategies to promote research on the PBD.
- We formulate the PBD problem into a segmentation task and develop a novel multi-dimensional collaborative network (MDCNet), which leverages the point, line and number cues through the multi-scale feature fusion and prompt filter techniques to achieve accurate PBD under complex and diverse background interference scenes.
- We compare our segmentation-based method with the other potential solutions, including corner detection, crowd counting and general/tiny object detection. Extensive experiments demonstrate that the proposed MDCNet performs favorably against the state-of-the-art methods under eight metrics.

2. Related Work

In this section, we introduce multiple possible modeling schemes for the PBD task and analyze related techniques.

2.1. Corner Detection

A corner point is usually defined as the intersection of two edges. Corner detection is widely used in motion detection, image matching and video tracking. It can also be called feature point detection. Harris and Stephens [11] propose one of the earliest corner detection algorithms by considering the differential of the corner score with respect to di-

rection directly. Shi–Tomasi [39] corner detector directly computes minimum euclidean distance between corners detected because under certain assumptions, the corners are more stable for tracking. Sub-pixel [4] corner detection can obtain real coordinate values to provide more accuracy for performing geometric measurement or calibration. Corner detectors are not usually very robust and often require large redundancies introduced to prevent the effect of individual errors from dominating the recognition task.

2.2. General/Tiny Object Detection

Object detection task deals with detecting instances of semantic objects of a certain class (such as humans, buildings, or cars) in digital images and videos. Object detectors usually can be divided into two types: 1) two-stage detectors: Fast R-CNN [10], Faster R-CNN [36], FPN [25], R-FCN [6]. 2) one-stage detectors: RetinaNet [26], YOLOV5 [18], FCOS [43], DETR [2]. Although most single-stage methods have fast detection speed and can learn the generalized features of objects, they are usually not as accurate as the anchor-based two-stage detectors, especially in small objects. To this end, some tiny object detectors have been proposed to achieve the balance between accuracy and speed, e.g., MRDet [34], QueryDet [50], C3Det [21] and CFINet [52].

2.3. Dense Object Counting

Dense object counting task is to infer the number of objects in dense scenarios. Most CNNs-based methods [19, 33, 35, 41, 45] predict a density map from a crowd image, where the summation of the density map is the crowd count. For the structures, some works are proposed to address scale variation [1, 19, 54], to refine the predicted density map [35, 38, 40, 42], and to encode context information [41, 48]. For the loss function, L2 loss is usually used in many models but it is sensitive to the choice of variance in the Gaussian kernel [44]. Therefore, Bayesian loss [28] is proposed with point-wise supervision.

2.4. Image Segmentation

Image segmentation aims to simplify and/or change the representation of an image into something that is more meaningful and easier to analyze. There are many kinds of image segmentation in real life, such as salient object segmentation, camouflaged object segmentation, glass segmentation and medical image segmentation. Most methods [14, 16, 17, 29, 30, 32, 55, 57–59] tend to adopt U-shape [25, 37] with the encoder and decoder as the baseline and then combine multi-level features to gradually reconstruct the high-resolution feature maps. To improve the representation ability, some methods [24, 55, 56, 62, 65] introduce the attention or gate mechanism and propose different spatial attention and channel-wise attention mod-

ules. Besides, objects with size-varying is another challenge in the segmentation filed. This motivates many efforts [5, 15, 22, 23, 31, 51, 55, 60, 61] to investigate multi-scale feature extraction module, such as the ASPP [5], Fold-ASPP [55] and DenseASPP [51].

3. X-ray PBD Dataset

3.1. Image Collection and Annotation

The proposed X-ray PBD dataset is collected from thousands of power batteries cells of 5 manufacturers captured by the same DR device, covering close, medium, and long shots. It is divided into 2 plate classes with 8 attributes according to appearance and interference types. We first deliver the initially collected 4,000+ X-ray battery images to six professional battery factory technicians for quality screening. Through careful observation, duplicate, blank, and severely defective (i.e., invalid) images are eliminated and we obtain $\sim 2,000$ battery images. Then, six technicians use the labeling software to annotate 2,000 images, respectively. The annotation information include the endpoint coordinates of all anode and cathode plates, and the numbers of plates. To avoid human judgment errors, six technicians jointly evaluate all six sets of marked data and choose the examples with consistent marked information. Finally, we obtain 1,500 high-quality power battery images.

3.2. Dataset Features, Statistics and Challenges

- *Categories.* As illustrated in Fig. 3 (a), the collected images have 2 categories and 8 attributes and their proportions in each shot type are summarized in detail. Different battery types, manufacturing and packaging processes produce different thicknesses of batteries. Even if the DR device is calibrated, it still has different degrees of penetration for different thicknesses, thereby resulting in “clear” and “blur” X-ray images. Attributes descriptions and visualization are provided in Tab. 1 and Fig. 2.
- *Number and Overhang.* Generally speaking, different types of batteries have different numbers of plates. Even in the same type of product, overhang is not exactly the same. As shown in Fig. 3 (b, c), the X-ray PBD dataset covers a large range of overhang and number distributions.
- *Attributes.* Fig. 3 (d, e) demonstrate co-occurrence distributions and dependencies among these attributes, respectively. We can see that more than four attribute dependencies appear in all samples, which illustrates the diversity and complexity of this dataset.
- *Dataset splits.* To provide a large amount of training data for deep learning models, X-ray PBD is split into 900 images for training and 600 for testing, randomly selected from each sub-class with different attributes. And, we further divide the test set into regular (239 images), difficult (187 images) and tough (174 images) subsets according to

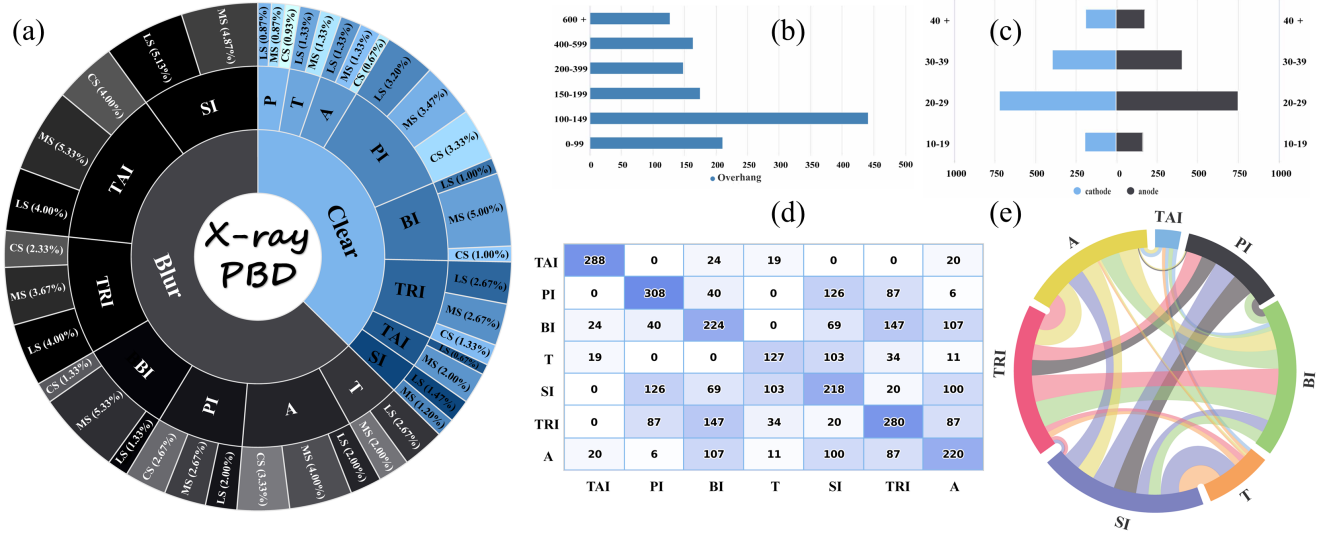


Figure 3. Statistics of the X-ray PBD dataset. (a) Taxonomic of interference and shots. (b) Overhang distributions. (c) Number distributions. (d) Co-occurrence distribution of attributes. (e) Multi-dependencies among these attributes.

the interference degree caused by different attributes.

- **Challenges.** To realize effective power battery detection, it faces multiple challenges caused by battery characteristics, photography restrictions, internal and external interference. First, there exist three types of shots in the X-ray PBD dataset. A well-performed PBD model needs to perceive the size-varying battery regions sensitively. Second, various interference and their dependencies make the discriminative information of the plates inconspicuous and easily overwhelmed by disturbances, which requires the model to have strong anti-interference ability. Third, both intra-sample and inter-sample plates have different overhangs, it brings the PBD model with strict requirements on both semantic discrimination and spatial localization. Fourth, the error tolerance rate is extremely low. In the real-life industrial production, all plates of a single battery cell must be detected with 100% accuracy, otherwise the automatic inspection is meaningless.

3.3. Evaluation Metrics

According to the number and overhang criteria adopted by the manufacturers, we design eight complementary metrics to quantitatively evaluate the performance of algorithms. Specifically, they are number mean absolute error (AN-MAE, CN-MAE), number accuracy (AN-ACC, CN-ACC, PN-ACC) and position mean absolute error (AL-MAE, CL-MAE and OH-MAE) for the anode level, cathode level, and pair level, respectively. Note that, for an image, only when its number of AN or CN is predicted with 100% accuracy, we can further evaluate the corresponding position MAE. More details are provided in the supplementary material.

4. Proposed Framework

4.1. Multi-dimensional Collaborative Network

Motivation. Imitating the coarse-to-fine human visual perception system [13, 46], the PBD model first perceives the battery area, and then gradually identifies each plate. Once the endpoint of each plate is accurately segmented, the position coordinates and overall quantity information can further be obtained. General U-shape [25, 37] structures are adopted as the point branch to generate initial point prediction. Inspired by multi-task learning [3, 7, 59] and characteristics of power battery, we extra design line segmentation and counting tasks to constrain and tune the point segmentation branch, helping it to complete the prediction from coarse to fine. The predictions of three dimensions of point, line and number cooperate with each other, dubbed as multi-dimensional collaborative network (MDCNet).

Overview. As shown in Fig. 4, the MDCNet has the encoder-decoder structure with the ResNet-50 [12] backbone. The encoder first extracts five levels of features, meanwhile, prompt and current pipelines share the same weights. Multi-scale module and the prompt filter module are embedded into the point segmentation branch. The former integrates features with receptive fields of different scales for improving the perception of multi-scale objects. The latter utilizes the prompt features to filter interference information from the current features and generate pure plate appearance features. Point predictor outputs the point segmentation map after passing through five decoder blocks. Line and counting predictors assist point prediction at low-level details and semantic level, respectively.

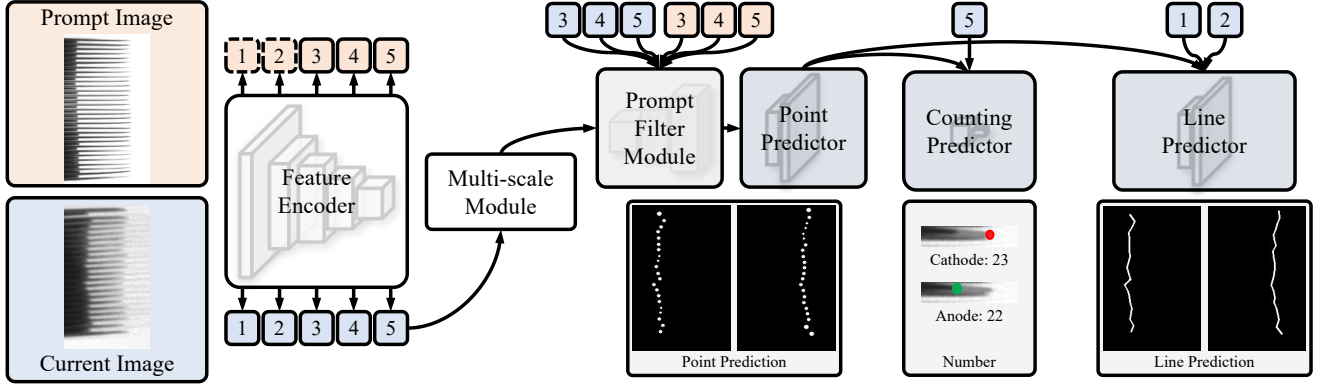


Figure 4. Overview of our MDCNet. It contains a shared encoder to extract different level features for the prompt and current images, respectively. Multi-scale module is only embedded in the high-level features. Prompt filter module are used to combine the prompt and current features to generate a series of filtered features. Point predictor include five decoder layers to produce point segmentation maps. Counting and line predictors are guided by the point prediction and fusing high-level and low-level features, respectively.

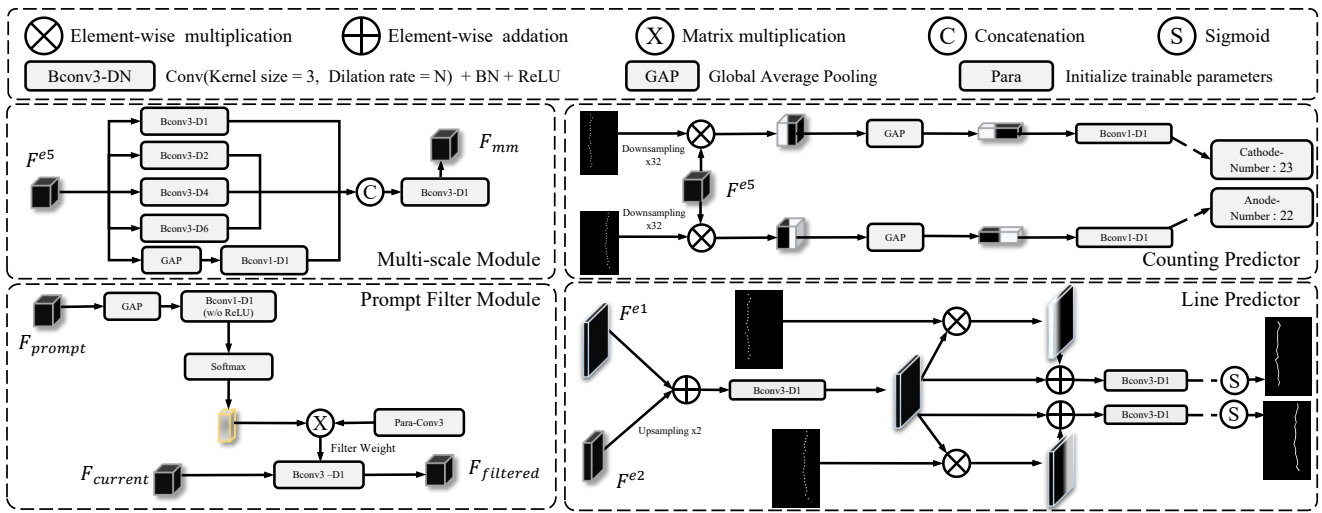


Figure 5. Detailed illustration of each component.

4.2. Components

As shown in Fig. 5, it illustrates each component in the MDCNet, including the multi-scale module, prompt filter module, counting predictor and line predictor.

- **Multi-scale module.** The single-scale features can not capture the multi-scale context for batteries with different shapes under different shots. We apply the multi-scale module to the high-level features, which have rich semantic cues for localizing objects. The multi-scale module includes four convolution branches with dilation rates of [1, 2, 4, 6] and a global average pooling branch. The features with different receptive fields are fused by cross-channel concatenation.
- **Prompt filter module.** In order to weaken external interference such as trays, bifurcation, tabs, etc., we embed the prompt filter module (PFM) in the point segmentation

branch. We randomly choose a P plate sample (see Fig 2) as the prompt image. The 3^{rd} - 5^{th} layer features of the prompt image will be separately generated the filters to edit the current image features, which progressively fuses the different levels of features through five convolution layers. We take both the 5^{th} prompt and current image features as an example to illustrate the process of prompt filtering. As shown in Fig. 5, we first conduct global average pooling and convolution for F_{prompt} . Then, the softmax function distributes F_{prompt} with the channel-wise soft attention. Next, we initialize the 3×3 convolution with trainable parameters and multiply the soft attention to generate the aggregated parameters as the weights. Finally, we obtain the enhanced feature $F_{filtered}$, which is passed to the point predictor.

- **Counting predictor.** The overall counting predictor is constructed to solve a regression problem. The inputs of the

counting predictor are the high-level features ($F_{current}^{e5}$) and the predicted point maps (M_p^{anode} , $M_p^{cathode}$) of the anode and cathode plates from the point predictor. To narrow the search range of the counting objects, we utilize two predicted point maps as spatial attention to guide $F_{current}^{e5}$. The number of anode and cathode plates are computed as:

$$\begin{cases} N^{anode} = ReLU(Conv(GAP(DS(F_{current}^{e5}) \otimes M_p^{anode}))) \\ N^{cathode} = ReLU(Conv(GAP(DS(F_{current}^{e5}) \otimes M_p^{cathode}))) \end{cases}, \quad (1)$$

where $DS(\cdot)$ is the downsampling operation, \otimes is the element-wise multiplication, $GAP(\cdot)$ is the global average pooling and $Conv(\cdot)$ refers to the convolution layer with the one channel output. As an auxiliary task for the point segmentation, counting task can improve the query ability of high-level features on the number of plates at the global level, thereby enhancing the feature representation of the point branch.

• *Line predictor.* Generally speaking, detail information of contours often exists in the low-level features. Therefore, we build the line predictor with both low-level features ($F_{current}^{e1}$, $F_{current}^{e2}$) and the predicted point maps (M_p^{anode} , $M_p^{cathode}$) as inputs. We first aggregate the $F_{current}^{e1}$, $F_{current}^{e2}$ to generate the $F_{current}^{e1,2}$:

$$F_{current}^{e1,2} = Conv(F_{current}^{e1} + US(F_{current}^{e2})), \quad (2)$$

where $US(\cdot)$ is the upsampling operation. Next, $F_{current}^{e1,2}$ is separately computed with M_p^{anode} and $M_p^{cathode}$ in the form of residual [12] and predict line segmentation maps:

$$\begin{cases} L^{anode} = S(Conv(M_p^{anode} \otimes F_{current}^{e1,2}) + F_{current}^{e1,2}) \\ L^{cathode} = S(Conv(M_p^{cathode} \otimes F_{current}^{e1,2}) + F_{current}^{e1,2}) \end{cases}, \quad (3)$$

where $S(\cdot)$ is the element-wise sigmoid function. As another auxiliary task of point segmentation, line segmentation can provide continuous segmentation cues to compensate and modify some plates that are not accurately predicted by point branch.

4.3. Label Generation and Supervision

For point segmentation, a direct strategy of generating the ground truth of point mask is to use circular regions of the same radius for each plate. As shown in the 1st - 3rd columns of Fig. 6, the point masks are yielded under the radius of 1, 3, and 5 pixels, respectively. It can be seen that the fixed radius often produces very dense or sparse point masks for different types of batteries, which increases the difficulty of point segmentation. To this end, we design a novel distance-adaptive label generation strategy. It computes the diameter of each point according to the distance of adjacent plates. As shown in the 4th - 6th columns of Fig. 6, the point masks are generated under the diameter of $0.1x$, $0.3x$, and $0.5x$ distance, respectively. The labels of counting and line segmentation can be separately obtained




Figure 6. Visualization of the ground-truth point masks under different strategies.

by the connected component analysis and connecting each endpoint in the point mask in sequence.

For the point and line segmentation predictors, we use the weighted IoU loss and binary cross entropy (BCE) loss, which have been widely adopted in segmentation task. We use the same definitions as in [8, 47, 58, 64]. For the counting predictor, we adopt the common L1 loss.

5. Experiments

5.1. Implementation Details

Our model is implemented based on the Pytorch and trained on one Tesla V100 GPU for 50 epochs with mini-batch size 12. The input resolutions of X-ray images are resized to 352×352 . To keep the spatial distribution pattern of the plates, we only adopt the random flipping image augmentation. We use the Adam optimizer [20]. The initial learning rate is set to 0.0001. We adopt the “step” learning rate decay policy, and set decay size as 30 and decay rate as 0.9. Our code and dataset will be made publicly available soon.

5.2. Comparison with Different Solutions

We conduct thorough comparisons with corner detection, crowd counting, general/tiny object detection methods. All these models have been retrained on our X-ray PBD dataset based on their publicly released codes. Tab. 2 shows performance comparisons in terms of eight metrics. Our method outperforms all of the competitors and obtains the best performance in terms of all metrics on the regular, diffi-

Dataset	Metrics	Corner Detection			Crowd Counting			General/Tiny Object Detection					Segmentation
		Harris [11]	Shi-tomasi [39]	Sub-pixel [4]	BL [28]	CUT [33]	IOCFORMER [42]	RetinaNet [26]	Faster RCNN [36]	YOLOv5 [18]	C3Det [21]	CFINet [52]	MDCNet (Ours)
Regular	AN-MAE↓	125.4128	81.2936	116.4404	0.3211	0.3028	0.8349	0.6330	1.2294	0.3761	3.1835	0.5505	0.0642
	CN-MAE↓	3437.5780	156.9083	121.7615	0.1468	0.8349	0.3853	0.3486	0.7798	0.5413	0.5138	0.2477	0.0183
	AN-ACC↑	0.0000	0.0000	0.0000	0.6881	0.6972	0.3211	0.6330	0.4220	0.6881	0.6330	0.7339	0.9633
	CN-ACC↑	0.0000	0.0000	0.0000	0.8532	0.1743	0.6514	0.7615	0.6972	0.6972	0.8532	0.8991	0.9908
	PN-ACC↑	0.0000	0.0000	0.0000	0.5872	0.1193	0.2294	0.4587	0.3761	0.4954	0.6330	0.6881	0.9541
	AL-MAE↓	—	—	—	—	—	—	6.6782	16.4639	5.1243	4.9344	4.0220	2.3365
	CL-MAE↓	—	—	—	—	—	—	6.0135	14.8421	4.6982	3.6087	3.8072	1.8411
	OH-MAE↓	—	—	—	—	—	—	3.2207	6.4742	5.1286	4.7642	3.9496	2.0422
Difficult	AN-MAE↓	122.1703	73.5899	106.4858	1.3211	2.4216	0.9117	1.1546	1.5047	1.5047	2.2650	1.0883	0.4259
	CN-MAE↓	3208.6183	146.3344	113.3249	1.1468	3.3215	0.5079	0.7476	1.1546	1.7192	1.1703	0.8139	0.2050
	AN-ACC↑	0.0000	0.0000	0.0000	0.7382	0.5584	0.3659	0.3975	0.4038	0.4416	0.5616	0.5741	0.7855
	CN-ACC↑	0.0000	0.0000	0.0032	0.8170	0.4479	0.5552	0.5804	0.5584	0.4069	0.7035	0.8991	0.9306
	PN-ACC↑	0.0000	0.0000	0.0000	0.6278	0.2965	0.2461	0.2587	0.3091	0.2177	0.5016	0.5426	0.7603
	AL-MAE↓	—	—	—	—	—	—	6.9021	28.0690	4.1363	4.9417	4.9602	2.4397
	CL-MAE↓	—	—	44.5243	—	—	—	6.4768	23.7668	4.0355	4.7920	4.9878	2.0978
	OH-MAE↓	—	—	—	—	—	—	3.4025	6.9621	3.1810	3.7909	3.9769	2.1092
Tough	AN-MAE↓	59.7989	16.8621	20.0287	2.5989	2.8212	2.3036	4.1552	5.0057	6.5632	3.7356	3.2989	2.0920
	CN-MAE↓	1132.0632	22.0862	18.9080	2.1494	3.2275	2.1044	3.4253	4.2011	6.1437	2.5747	3.1494	1.8966
	AN-ACC↑	0.0000	0.0000	0.0172	0.3448	0.3103	0.5057	0.2931	0.4425	0.3448	0.3621	0.3793	0.5632
	CN-ACC↑	0.0000	0.0115	0.0172	0.4023	0.2872	0.5230	0.5230	0.6494	0.2126	0.4195	0.4310	0.6839
	PN-ACC↑	0.0000	0.0000	0.0000	0.3161	0.2586	0.3218	0.1954	0.3678	0.0690	0.3046	0.3276	0.5115
	AL-MAE↓	—	—	85.5647	—	—	—	4.9547	9.9379	4.0577	4.0700	4.9448	2.0004
	CL-MAE↓	—	53.0062	51.1617	—	—	—	4.9750	7.2092	4.4330	4.6340	4.6617	1.4654
	OH-MAE↓	—	—	—	—	—	—	3.8719	1.9481	4.9554	3.6690	3.6988	1.6291

Table 2. Quantitative comparison of different methods. \uparrow and \downarrow indicate that the larger scores and the smaller ones are better, respectively. The best scores are highlighted in red. “—” represents that the results are not available because these methods can not provide coordinate information or their prediction accuracy of the number of plates is zero.

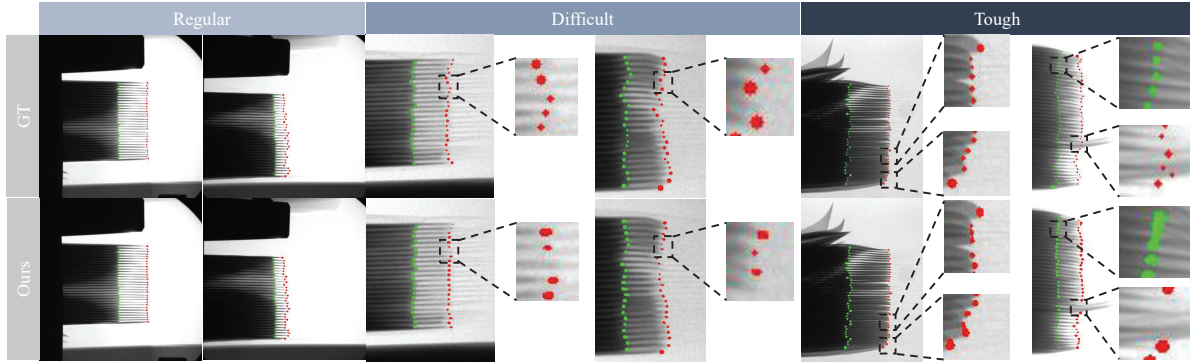


Figure 7. Qualitative results on the regular, difficult and tough examples with different shots and attributes. Best viewed zoomed in.




Figure 8. Visual comparison with other general/tiny object detection-based [18, 21, 26, 36, 52], counting-based [28, 33, 42], corner detection-based [4, 11, 39] solutions. We directly visualize the predicted results (Ours: Segmentation map, General/Tiny object detection methods: Bounding box, Counting methods: Density map, Corner detection methods: Corner map) without any post-processing operations.

cult and tough splits. It is worth noting that the MDCNet achieves dominant performance in terms of the PN-ACC compared to the second best method CFINet [52] on the regular (0.9541 vs 0.6881), difficult (0.7603 vs 0.5426) and tough (0.5115 vs 0.3276) subsets. Fig. 7 illustrates visual results of the proposed algorithm on the challenge X-ray images with different complex interference. For regular ex-

amples, our model accurately captures the body of each battery, and completes fine semantic segmentation for all the endpoints of anode and cathode plates. For difficult examples, we can overcome most of the separator interference but still lose the sporadic plate that blends well with the separator. For tough examples, the plates are very dense and the roots overlap, which will result in multiple adhesions in

Method	AN-MAE↓	CN-MAE↓	AN-ACC↑	CN-ACC↑	PN-ACC↑	AL-MAE↓	CL-MAE↓	OH-MAE↓
Baseline	1.5424	1.2732	0.6697	0.7064	0.6239	3.5317	2.8830	3.1929
+ PFM	0.8211	0.4324	0.8257	0.8440	0.7706	3.3298	2.6425	3.0231
+ CP	0.3892	0.1782	0.9174	0.9266	0.8716	3.2343	2.5322	2.9371
+ LP	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422

Table 3. Ablation experiments of each component in the MDCNet.

Method	AN-MAE↓	CN-MAE↓	AN-ACC↑	CN-ACC↑	PN-ACC↑	AL-MAE↓	CL-MAE↓	OH-MAE↓
Const-1	3.3234	2.0231	0.4954	0.5505	0.4587	2.1146	1.7322	2.0018
Const-3	1.7312	1.1254	0.7523	0.8257	0.7523	2.4215	1.8743	2.1797
Const-5	1.7622	1.0817	0.7523	0.8073	0.7339	2.4347	1.8922	2.1922
Ada-0.1	2.0236	1.4576	0.6330	0.6881	0.6055	2.4617	1.9020	2.1178
Ada-0.3	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422
Ada-0.5	0.3452	0.2971	0.9174	0.9450	0.8991	2.5301	1.9420	2.1201

Table 4. Quantitative comparison of different label generation strategies and settings for the point segmentation branch.

the predicted point map. Fig. 8 illustrates the visual comparison of different solutions with our segmentation-based MDCNet. It can be seen that general/tiny object detection methods often suffer from missed and repeated detection failures due to diverse interference of the bifurcation, tray and separator. For counting methods, the density maps provide rough and overall location of the plates but can not explicitly show each plate one by one. For corner detection methods, the corner maps introduce a lot of redundancy and can not distinguish plate endpoints from general intersection points even if we have performed edge extraction to pre-filter a large amount of background irrelevant to plate endpoints. More visual results are provided in the supplementary material.

5.3. Ablation Study

Multi-dimensional Collaborative Network. We analyze the contribution of each component to the overall network in Tab. 3, which shows the results of ablation study on the regular split. The results on the other two splits can be found in the supplementary material. First, our baseline only has the point segmentation branch, which is installed on the FPN [25] with ResNet-50 backbone. We can see that this baseline has been able to suppress most competitors except the CFNet and C3Det. Hence, the performance gain in subsequent experiments is more convincing. Next, we embed the prompt filter module (PFM) into the point segmentation branch. It can be seen that the performance is significantly improved with the gain of $\sim 15\%$ in terms of the AN-ACC, CN-ACC and PN-ACC. Finally, we add the counting predictor (CP) and line predictor (LP) to further refine the point segmentation at the high-level and low-level features. The former improves the number accuracy and the latter helps modify the failure positioning. Thus, multi-dimensional features complement each other.

Label Generation Strategies. For point segmentation, we conduct a series of experiments about label generation

No.	AN-MAE↓	CN-MAE↓	AN-ACC↑	CN-ACC↑	PN-ACC↑	AL-MAE↓	CL-MAE↓	OH-MAE↓
1	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422
2	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422
3	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422
4	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422
5	0.0642	0.0183	0.9633	0.9908	0.9541	2.3365	1.8411	2.0422

Table 5. Evaluation of the PFM with different prompt inputs.

strategies as discussed in Sec. 4.3. As shown in Tab. 4, the model trained under the Ada-0.3 performs best in terms of number accuracy and localization error among all the six settings. The distance-adaptive strategy consistently outperforms the constant strategy.

Prompt Filter Module. In Tab. 5, we list five groups of results using different P attribute images as the anchor of prompt filter module. It can be seen that PFM shows strong stability and generalization without dependence on a specific prompt.

6. Future Works

Because the performance on difficult and tough splits is bad, there remains a long way to go to solve the PBD task. Some potential research are of benefit to the development of the PBD community. First, how to better model PBD is still an open problem. Second, the annotation of PBD relies on experienced workers and the amount of data is scarce. Semi/self-supervised and few-shot learning techniques may be necessary. Third, image enhancement techniques such as image super-resolution, restoration and deblurring can be jointly modeled in order to clean the battery plates. Last, the generality and specificity among multiple scenarios and interference are worth exploring in depth in order to design a robust unified model.

7. Conclusion

In this work, we present a new task named power battery detection (PBD), which is a crucial course in the new energy industry field. We build a complex X-ray PBD dataset, formulate evaluation metrics and propose a segmentation-based solution (MDCNet). Compared with other solutions based on corner detection, crowd counting and general/tiny object detection, MDCNet achieves a dominant performance with over 95% detection accuracy on regular samples. In the future, more effort will be devoted to handling hard samples. We also plan to extend the PBD dataset to a 3D form with the help of CT device, which can provide richer internal slices information.

Acknowledgements. This work was supported by the National Natural Science Foundation of China under Grant 62276046 and by Dalian Science and Technology Innovation Foundation under Grant 2023JJ12GX015.

References

- [1] Deepak Babu Sam, Shiv Surya, and R Venkatesh Babu. Switching convolutional neural network for crowd counting. In *CVPR*, pages 5744–5752, 2017. 3
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020. 3
- [3] Rich Caruana. Multitask learning: A knowledge-based source of inductive bias. In *ICML*, pages 41–48, 1993. 4
- [4] Dazhi Chen and Guangjun Zhang. A new sub-pixel detector for x-corners in camera calibration targets. In *International Conference in Central Europe on Computer Graphics and Visualization*, 2005. 3, 7
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40:834–848, 2017. 3
- [6] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NeurIPS*, 2016. 3
- [7] Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser. In *ACL*, pages 845–850, 2015. 4
- [8] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Pranet: Parallel reverse attention network for polyp segmentation. In *MICCAI*, pages 263–273, 2020. 6
- [9] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *ECCV*, pages 540–557. Springer, 2022. 2
- [10] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. 3
- [11] Chris Harris, Mike Stephens, et al. A combined corner and edge detector. In *Alvey vision conference*, number 50, pages 10–5244, 1988. 2, 7
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 4, 6
- [13] Jay Hegd . Time course of visual perception: coarse-to-fine processing and beyond. *Progress in neurobiology*, 84(4):405–439, 2008. 4
- [14] Wei Ji, Shuang Yu, Junde Wu, Kai Ma, Cheng Bian, Qi Bi, Jingjing Li, Hanruo Liu, Li Cheng, and Yefeng Zheng. Learning calibrated medical image segmentation via multi-rater agreement modeling. In *CVPR*, pages 12341–12351, 2021. 3
- [15] Wei Ji, Jingjing Li, Qi Bi, Chuan Guo, Jie Liu, and Li Cheng. Promoting saliency from depth: Deep unsupervised rgb-d saliency detection. *ICLR*, 2022. 3
- [16] Wei Ji, Jingjing Li, Cheng Bian, Zhicheng Zhang, and Li Cheng. Semantict: A large-scale dataset and method for robust semantic segmentation in multispectral images. In *ACM MM*, page 3307–3316, 2023. 3
- [17] Wei Ji, Jingjing Li, Cheng Bian, Zongwei Zhou, Jiaying Zhao, Alan L Yuille, and Li Cheng. Multispectral video semantic segmentation: A benchmark dataset and baseline. In *CVPR*, pages 1094–1104, 2023. 3
- [18] Glenn Jocher, Alex Stoken, Jirka Borovec, Ayush Chaurasia, Liu Changyu, A Laughing, A Hogan, J Hajek, L Diaconu, Y Marc, et al. ultralytics/yolov5: v5. 0-yolov5-p6 1280 models aws supervise. ly and youtube integrations. *Zenodo*, 11, 2021. 3, 7
- [19] Di Kang and Antoni Chan. Crowd counting by adaptively fusing predictions from an image pyramid. *arXiv preprint arXiv:1805.06115*, 2018. 3
- [20] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [21] Chunggi Lee, Seonwook Park, Heon Song, Jeongun Ryu, Sanghoon Kim, Haejoon Kim, S rgio Pereira, and Donggeun Yoo. Interactive multi-class tiny-object detection. In *CVPR*, pages 14136–14145, 2022. 3, 7
- [22] Jingjing Li, Wei Ji, Qi Bi, Cheng Yan, Miao Zhang, Yongri Piao, Huchuan Lu, and Li Cheng. Joint semantic mining for weakly supervised rgb-d salient object detection. *NeurIPS*, pages 11945–11959, 2021. 3
- [23] Jingjing Li, Wei Ji, Size Wang, Wenbo Li, and Li Cheng. Dvsod: Rgb-d video salient object detection. In *NeurIPS*, pages 8774–8787, 2023. 3
- [24] Jingjing Li, Wei Ji, Miao Zhang, Yongri Piao, Huchuan Lu, and Li Cheng. Delving into calibrated depth for accurate rgb-d salient object detection. *IJCV*, 131(4):855–876, 2023. 3
- [25] Tsung-Yi Lin, Piotr Doll r, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 3, 4, 8
- [26] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Doll r. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 3, 7
- [27] Lizhao Liu, Kunyang Lin, Shangxin Huang, Zhongli Li, Chao Li, Yunbo Cao, and Qingyu Zhou. Instance segmentation for chinese character stroke extraction, datasets and benchmarks. *arXiv preprint arXiv:2210.13826*, 2022. 2
- [28] Zhiheng Ma, Xing Wei, Xiaopeng Hong, and Yihong Gong. Bayesian loss for crowd count estimation with point supervision. In *ICCV*, pages 6142–6151, 2019. 3, 7
- [29] Youwei Pang, Lihe Zhang, Xiaoqi Zhao, and Huchuan Lu. Hierarchical dynamic filtering network for rgb-d salient object detection. In *ECCV*, pages 235–252, 2020. 3
- [30] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Multi-scale interactive network for salient object detection. In *CVPR*, pages 9413–9422, 2020. 3
- [31] Youwei Pang, Xiaoqi Zhao, Tian-Zhu Xiang, Lihe Zhang, and Huchuan Lu. Zoom in and out: A mixed-scale triplet network for camouflaged object detection. In *CVPR*, pages 2160–2170, 2022. 3
- [32] Yongri Piao, Wei Ji, Jingjing Li, Miao Zhang, and Huchuan Lu. Depth-induced multi-scale recurrent attention network for saliency detection. In *ICCV*, pages 7254–7263, 2019. 3

- [33] Yifei Qian, Liangfei Zhang, Xiaopeng Hong, Carl R Donovan, Ognjen Arandjelovic, UK Fife, and PR Harbin. Segmentation assisted u-shaped multi-scale transformer for crowd counting. 2022. [3](#), [7](#)
- [34] Ran Qin, Qingjie Liu, Guangshuai Gao, Di Huang, and Yunhong Wang. Mrdet: A multihead network for accurate rotated object detection in aerial images. *IEEE TGRS*, 60:1–12, 2021. [3](#)
- [35] Viresh Ranjan, Hieu Le, and Minh Hoai. Iterative crowd counting. In *ECCV*, pages 270–285, 2018. [3](#)
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. [3](#), [7](#)
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. [3](#), [4](#)
- [38] Deepak Babu Sam and R Venkatesh Babu. Top-down feedback for crowd counting convolutional neural network. In *AAAI*, number 1, 2018. [3](#)
- [39] Jianbo Shi et al. Good features to track. In *CVPR*, pages 593–600, 1994. [3](#), [7](#)
- [40] Zenglin Shi, Le Zhang, Yun Liu, Xiaofeng Cao, Yangdong Ye, Ming-Ming Cheng, and Guoyan Zheng. Crowd counting with deep negative correlation learning. In *CVPR*, pages 5382–5390, 2018. [3](#)
- [41] Vishwanath A Sindagi and Vishal M Patel. Generating high-quality crowd density maps using contextual pyramid cnns. In *ICCV*, pages 1861–1870, 2017. [3](#)
- [42] Guolei Sun, Zhaochong An, Yun Liu, Ce Liu, Christos Sakaridis, Deng-Ping Fan, and Luc Van Gool. Indiscernible object counting in underwater scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13791–13801, 2023. [3](#), [7](#)
- [43] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. [3](#)
- [44] Jia Wan and Antoni Chan. Adaptive density map generation for crowd counting. In *ICCV*, pages 1130–1139, 2019. [3](#)
- [45] Boyu Wang, Huidong Liu, Dimitris Samaras, and Minh Hoai Nguyen. Distribution matching for crowd counting. In *NeurIPS*, pages 1595–1607, 2020. [3](#)
- [46] RJ Watt. Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *JOSA A*, 4(10):2006–2021, 1987. [4](#)
- [47] Jun Wei, Shuhui Wang, and Qingming Huang. F³net: Fusion, feedback and focus for salient object detection. In *AAAI*, pages 12321–12328, 2020. [6](#)
- [48] Feng Xiong, Xingjian Shi, and Dit-Yan Yeung. Spatiotemporal modeling for crowd counting in videos. In *ICCV*, pages 5151–5159, 2017. [3](#)
- [49] JiYang Xu, Jian Ma, Xuan Zhao, Hao Chen, Bin Xu, and XueQin Wu. Detection technology for battery safety in electric vehicles: A review. *Energies*, 13(18):4636, 2020. [1](#)
- [50] Chenhongyi Yang, Zehao Huang, and Naiyan Wang. Querydet: Cascaded sparse query for accelerating high-resolution small object detection. In *CVPR*, pages 13668–13677, 2022. [3](#)
- [51] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *CVPR*, pages 3684–3692, 2018. [3](#)
- [52] Xiang Yuan, Gong Cheng, Kebing Yan, Qinghua Zeng, and Junwei Han. Small object detection via coarse-to-fine proposal generation and imitation learning. In *ICCV*, pages 6317–6327, 2023. [3](#), [7](#)
- [53] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *CVPR*, pages 14393–14402, 2021. [2](#)
- [54] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *CVPR*, pages 589–597, 2016. [3](#)
- [55] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Suppress and balance: A simple gated network for salient object detection. In *ECCV*, pages 35–51, 2020. [3](#)
- [56] Xiaoqi Zhao, Lihe Zhang, Youwei Pang, Huchuan Lu, and Lei Zhang. A single stream network for robust and real-time rgb-d salient object detection. In *ECCV*, pages 646–662, 2020. [3](#)
- [57] Xiaoqi Zhao, Youwei Pang, Jiaying Yang, Lihe Zhang, and Huchuan Lu. Multi-source fusion and automatic predictor selection for zero-shot video object segmentation. In *ACM MM*, pages 2645–2653, 2021. [3](#)
- [58] Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Automatic polyp segmentation via multi-scale subtraction network. In *MICCAI*, pages 120–130, 2021. [6](#)
- [59] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, and Huchuan Lu. Joint learning of salient object detection, depth estimation and contour extraction. *IEEE TIP*, 31:7350–7362, 2022. [3](#), [4](#)
- [60] Xiaoqi Zhao, Hongpeng Jia, Youwei Pang, Long Lv, Feng Tian, Lihe Zhang, Weibing Sun, and Huchuan Lu. M2snet: Multi-scale in multi-scale subtraction network for medical image segmentation. *arXiv preprint arXiv:2303.10894*, 2023. [3](#)
- [61] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Towards diverse binary segmentation via a simple yet general gated network. *arXiv preprint arXiv:2303.10396*, 2023. [3](#)
- [62] Xiaoqi Zhao, Shijie Chang, Youwei Pang, Jiaying Yang, Lihe Zhang, and Huchuan Lu. Adaptive multi-source predictor for zero-shot video object segmentation. *International Journal of Computer Vision*, pages 1–19, 2024. [3](#)
- [63] Kailai Zhou, Yibo Wang, Tao Lv, Yunqian Li, Linsen Chen, Qiu Shen, and Xun Cao. Explore spatio-temporal aggregation for insubstantial object detection: benchmark dataset and baseline. In *CVPR*, pages 3104–3115, 2022. [2](#)
- [64] Tao Zhou, Huazhu Fu, Geng Chen, Yi Zhou, Deng-Ping Fan, and Ling Shao. Specificity-preserving rgb-d saliency detection. In *CVPR*, pages 4681–4691, 2021. [6](#)
- [65] Lei Zhu, Zijun Deng, Xiaowei Hu, Chi-Wing Fu, Xuemiao Xu, Jing Qin, and Pheng-Ann Heng. Bidirectional feature pyramid network with recurrent attention residual modules for shadow detection. In *ECCV*, pages 121–136, 2018. [3](#)