

Leveraging RGB-D Data with Cross-Modal Context Mining for Glass Surface Detection

Jiaying Lin* Yuen-Hei Yeung* Shuquan Ye† Rynson W.H. Lau†

City University of Hong Kong

Abstract

Glass surfaces are becoming increasingly ubiquitous as modern buildings tend to use a lot of glass panels. This, however, poses substantial challenges to the operations of autonomous systems such as robots, self-driving cars, and drones, as these glass panels can become transparent obstacles to navigation. Existing works attempt to exploit various cues, including glass boundary context or reflections, as priors. However, they are all based on input RGB images. We observe that the transmission of 3D depth sensor light through glass surfaces often produces blank regions in the depth maps, which can offer additional insights to complement the RGB image features for glass surface detection. In this work, we first propose a large-scale RGB-D glass surface detection dataset, *RGB-D GSD*, for rigorous experiments and future research. It contains 3,009 images, paired with precise annotations, offering a wide range of real-world RGB-D glass surface categories. We then propose a novel glass surface detection framework combining RGB and depth information, with two novel modules: a cross-modal context mining (CCM) module to adaptively learn individual and mutual context features from RGB and depth information, and a depth-missing aware attention (DAA) module to explicitly exploit spatial locations where missing depths occur to help detect the presence of glass surfaces. Experimental results show that our proposed model outperforms state-of-the-art methods.

Code, Dataset and Extended version —

<https://jiaying.link/AAAI25-RBGDglass/>

Introduction

Although most objects that we see in our daily life exhibit distinctive visual characteristics, glass surfaces, however, typically do not possess any distinctive visual properties. Their contents are essentially represented by the contents behind them, making it difficult for existing computer vision systems to detect them. However, modern architects and designers tend to use glass panels frequently in their designs, which makes them prevalent both indoors and outdoors, *e.g.*, glass doors/windows, shop windows, and glass

* These authors contributed equally.

† Shuquan Ye and Rynson Lau are joint corresponding authors. Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Advantages of detecting glass surfaces with RGB-D images. These examples show that the depth map can provide a strong cue for glass surface detection. State-of-the-art methods, GSDNet (Lin, He, and Lau 2021) and EBLNet (He et al. 2021), relying only on input RGB images are not able to correctly separate the glass surfaces from the background. Through learning the cross-modal contexts and the correlation between depth-missing regions and glass surface regions, our proposed model can detect the glass surfaces accurately in all three challenging scenes. Note that red regions in the depth images represent missing depths.

walls. Hence, detecting these glass panels accurately is crucial for the safe operation of autonomous machines such as robots and drones.

The potential of deep learning models for glass surface detection (GSD) was first demonstrated by (Mei et al. 2020) and GSDNet(Lin, He, and Lau 2021). EBLNet (He et al. 2021) was proposed for GSD through enhanced boundary learning on glass surfaces. However, since EBLNet is an RGB-based method and relies heavily on boundary learning, it may fail if the boundaries of glass surfaces are ambiguous or there are some distracting glass-like regions in the input image. Fig. 1 shows three challenging scenarios where state-of-the-art RGB-based glass surface detection methods fail. In the first row, both GSDNet and EBLNet are confused by the similarity between glass regions and tiles. As a result, some tiles are mis-recognized as glass regions. The other rows show challenging scenes where some glass surfaces are difficult even for humans to identify. GSDNet and EBLNet either over- or under-detect them.

In this paper, we aim to address the limitation of RGB-based glass surface detection methods, which heavily rely on boundary learning, causing them to fail if the glass boundaries are ambiguous or there are glass-like regions in the

image. Here, we propose to address this problem from a depth-aware perspective. Our key observation is that compared with the *correct* depth of a glass surface, the depth *captured* by a depth sensor typically has two properties, due to the transmission and possibly reflection of the glass surface: 1) the captured depth is noisier, and 2) missing depths frequently appear around the surface. These two properties indicate that glass surfaces have different contextual characteristic representations in the depth map compared with those of the RGB image. Our insight from this observation is that cross-modal context and spatial information of missing depths can provide strong cues for glass surface detection. This motivates us to design a novel depth-aware glass surface detection method with two new modules: a cross-modal context mining (CCM) module to adaptively fuse multi-modal data by utilizing depth scans coupled with RGB data, and a depth-missing aware attention (DAA) module to explicitly learn the correlation between missing depths and glass regions.

In order to train our model, we need to have a RGB-D dataset for glass surfaces. Although there is an RGB-D dataset (Seib et al. 2017a) available for transparent object detection, it is not suitable for our purpose for two reasons. First, this dataset targets localizing small transparent objects with limited patterns of shapes and object types, while we focus on a more challenging task of detecting glass surfaces without well-defined shapes. Second, it contains only 440 RGB-D images captured from a single scene with four object types (*i.e.*, beer mugs, water glass, white beer glass, and wine glass), limiting real-world applicability. To address these limitations, we have constructed a new large-scale RGB-D glass surface dataset from diverse scenes with glass surfaces. It contains 3,009 RGB-D images with glass surfaces and corresponding annotated masks. A dataset of this scale allows robust model training and evaluation. Extensive experiments are conducted to evaluate our method, in comparison with the state-of-the-art methods from relevant tasks. The proposed model outperforms existing methods on our proposed dataset.

Our contributions can be summarised as follows:

1. We propose a framework for glass surface detection by incorporating the depth information, with two novel modules: a CCM module to jointly learn the RGB context features and depth context features for comprehensive RGB-D context modeling, and a DAA module to explicitly exploit the locations where the depth information is missing in the depth maps for glass surface detection.
2. We have constructed a large-scale glass surface dataset of 3,009 images from diverse real-world scenes, with corresponding depth maps and ground truth labels.
3. Extensive experiments demonstrate the superior performances of our proposed method over SOTA methods.

Related work

Glass Surface Detection. (Mei et al. 2020) propose the first deep-learning model for detecting glass surfaces, given just an RGB image. This method does not require any special

hardware. As this model relies heavily on contextual contrast learning, it likely fails in complex scenes with insufficient contexts. (Lin, He, and Lau 2021) extend contextual contrast learning with modules for detecting glass boundaries and reflections. However, if the glass surfaces lack reflections or have ambiguous boundaries, the model may not be able to detect the glass surfaces correctly. Similarly, (He et al. 2021) propose EBLNet based on learning the glass boundaries. Later, RFENet is proposed (Fan et al. 2023) by further exploiting edges in a reciprocal way. More recently, Liu *et al.* propose a new dataset (Liu et al. 2024) for video glass surface detection. However, all these deep-learning based methods only consider the RGB information for glass surface detection, and can easily be confused by the distracting glass-like regions that look like glass surfaces, *e.g.*, a door frame. In contrast, we propose in this paper a multi-modal glass surface detection method, based on using RGB-D images as input. Our experimental results show that the proposed method is more robust and accurate.

Transparent Object Detection. Compared with glass surface detection, research on transparent object detection (TOD) focuses on detecting small transparent objects, such as wine glass and glass balls, which typically have classic shapes or boundary properties. Some recent TOD methods integrate more accurate measurement data, including data from light-field sensors (Xu et al. 2015) and from infrared sensors (Guo-Hua, Jun-Yi, and Ai-Jun 2019; Seib et al. 2017b). Some latest methods use auxiliary information, including polarization (Kalra et al. 2020) and specular reflection (Islam, Tahtali, and Pickering 2021), to assist the detection. Other methods use explicit boundary maps (Xie et al. 2020) for training, or a transformer architecture (Xie et al. 2021; Zhang et al. 2021b) for improved performances. However, while the former induces additional data preparation and computation costs, the latter further abstracts the object representations with a feature embedding dictionary and has a constrained set of prototype categories. Recently, some datasets (Sun et al. 2023; Chen et al. 2022; Ramirez et al. 2022, 2023) have been proposed for this task, but they are still limited to specific objects. In general, transparent objects are mostly small and with specific shapes (*i.e.*, bottle, wine glass), and tend to have different lighting properties along the boundary. However, glass surfaces generally do not possess these properties, making these TOD methods not suitable for detecting glass surfaces.

RGB-D GSD Dataset

Existing datasets (Mei et al. 2020; Lin, He, and Lau 2021) for glass surface detection do not include depth information, even though it is useful for this task. In order to train our model and to encourage further research, we propose an RGB-D Glass Surface Detection (RGB-D GSD) dataset.

Dataset construction. This dataset contains a total of 3,009 images, where 2,400 for training and 609 for testing. It is an attentively curated ensemble of three existing datasets originally developed for scene understanding, including SUN RGB-D (Armeni et al. 2017), 2D-3D-Semantics (Song, Lichtenberg, and Xiao 2015) and Matterport3D (Chang et al.

2017). Thus, our dataset offers a wide range of glass surface categories, such as window, door, wall, table, cabinet and guardrail. For these datasets to be consistent, we normalize all depth images to the range $[0, 2^{16} - 1]$, and set all invalid depth values to be the minimum value (*i.e.*, 0). In addition, as these datasets do not aim at labeling glass surfaces, they contain inaccurate glass surface labels. For example, the objects inside a glass surface instead of the glass surface itself were labeled; handles, frames or blinds were labeled as part of the glass surface. Hence, we **manually relabel** the glass surfaces in these images. Figure 2 shows some examples from our RGB-D GSD dataset.

Dataset analysis. We analyze the datasets with the following statistical metrics:

- **Glass Location.** The glass location distribution is the average of all glass surface regions in the dataset. The maps in Figure 3(a) show that glass surfaces mainly concentrate at the top region, which is consistent in the training and testing splits. This also avoids the “center bias” problem due to natural observation tendency.
- **Color Contrast.** The color contrast between glass and non-glass regions should ideally be low. Otherwise, salient color features can skew the glass surface detection task. We measure the color contrast by computing the χ^2 distance of the RGB histograms between glass and non-glass regions. Figure 3(b) compares the color contrast among GDD (Mei et al. 2020), GSD (Lin, He, and Lau 2021) and our RGB-D GSD. In general, the contrast values of our RGB-D GSD images concentrate in the lower quartile ($0 < \text{contrast} < 0.4$), which is similar to the other two datasets.
- **Area Ratio.** This metric measures the glass region size relative to the image size. This illustrates the level of semantic context that the images provide. In other words, a smaller glass region leaves more room for the surrounding environment to offer additional hints. As mentioned in GSD, GDD contains primarily close-up shots, which limits the amount of contextual information. In addition, in real-life scenario, an autonomous system is expected to be able to detect objects and perform scene understanding tasks as early as possible. Therefore, data of low area ratio is much more meaningful and beneficial for model training. In Figure 3(b), RGB-D GSD has more images with small glass areas than the other datasets.
- **Missing Depth.** Missing depths strongly correlate with the presence of glass surfaces, as glass often causes erratic depth signals. To validate this, we calculated the ratio $\frac{M \cap G}{M}$, where M denotes the set of missing depths and G denotes the set of ground truth glass regions. This ratio represents the alignment of missing depths with ground truth glass regions. The distribution is as follows: 0.0-0.2: 15.47%, 0.2-0.4: 14.28%, 0.4-0.6: 14.49%, 0.6-0.8: 18.74%, 0.8-1.0: 37.02%. These results highlight that missing depths, especially at higher ratios, strongly overlap with ground truth, confirming their association with glass surfaces.





Figure 2: Examples from our RGB-D GSD dataset. Top, middle and bottom rows show RGB images, depth maps, GT glass surface masks overlaid on the images, respectively.



(a) Location distribution (b) Color contrast and area ratio
Figure 3: Statistics of our proposed dataset.

Method

Fig. 4 shows the architecture of our proposed framework. In the RGB backbone (red), we first feed the input RGB image to the backbone network (Xie et al. 2017) to extract multi-scale RGB backbone features. Specially, the last four stages’ outputs, *i.e.*, $\{S_i\}_{i=1}^4$ is used as our RGB backbone features. In the depth backbone (yellow), the network is simpler and lighter than the RGB one due to depth maps’ sparser information and efficiency considerations. Similarly, we only use the last four stages’ outputs, denoted as $\{S_i^d\}_{i=1}^4$. We first feed the RGB and depth backbone features (*i.e.*, S_4 and S_4^d) into a CCM module to capture the multi-modal contextual features comprehensively at the latest stage (*i.e.*, *stage4*). We create a binary depth-missing map from the depth map by setting the invalid depth pixels to 1’s and valid pixels to 0’s. The multi-modal contextual features from CCM and the resized depth-missing map are then fed into the DAA module to enhance the contextual features by exploiting the spatial location information of missing depths. Finally, using the enhanced contextual features the decoder produces a coarse binary mask representing *stage4*’s detected glass surfaces. In the preceding stage (*i.e.*, *stage3*), the enhanced contextual features from *stage4* are then added to the RGB features (*i.e.*, S_3), before feeding into CCM, DAA, and decoder. We repeat this process in *stage2* and *stage1*, but omitting depth missing attention for two reasons. First, while such spatial information can be easily learned from the upper layers (*i.e.*, later stages), it is harder to learn from the lower layers (*i.e.*, earlier stages). Second, adopting additional components to the early stages will heavily decrease the efficiency of the proposed network. With this progressive refinement, the earliest stage (*i.e.*, *stage1*) outputs the finest binary mask as the final output. More details (*e.g.*, the depth backbone) are included in the *supplementary material* due to limited space.

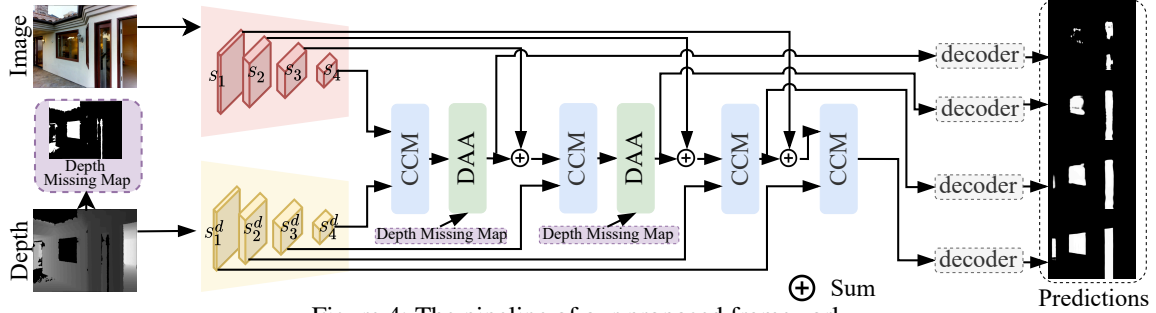


Figure 4: The pipeline of our proposed framework.

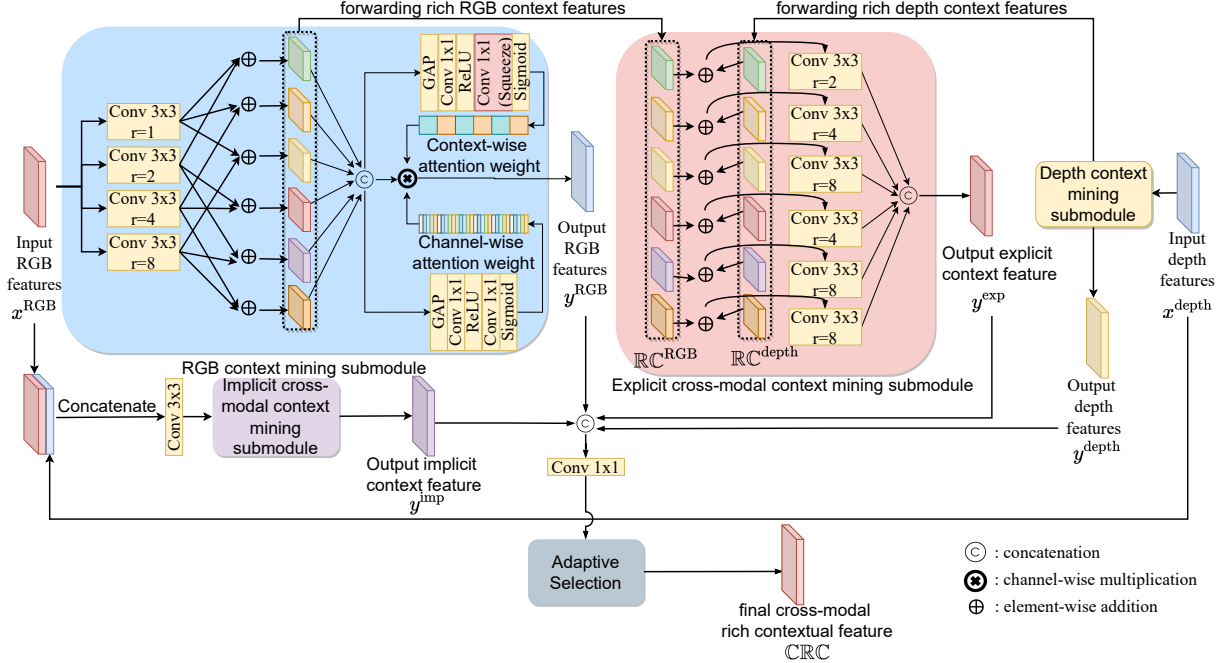


Figure 5: Illustration of our proposed cross-modal context mining (CCM) module.

Cross-modal Context Mining (CCM) Module

Previous works (Mei et al. 2020; Lin, He, and Lau 2021) focus only on single-modal glass surface detection. To exploit multi-modal data in this problem, we design the CCM module to adaptively learn the context features from both RGB and depth information. Fig. 5 illustrates the structure of the proposed CCM module. It consists of four submodules to model the multi-modal context from different aspects. The outputs of these four submodules are concatenated for adaptive selection to produce the final cross-modal feature.

RGB and Depth Context Mining Submodule. Given the input RGB features $x^{\text{RGB}} \in \mathbb{R}^{C \times H \times W}$, this submodule extracts a series of multi-scale context features $\mathbb{C}_{r_i}^{\text{RGB}}$ by atrous convolutions with dilation rate $r \in [1, 2, 4, 8]$. These multi-scale context features are then fused mutually by element-wise addition to form an interim rich representation of contextual information (\mathbb{RC}^{RGB}). After obtaining all permuted pairs of different context scales, these feature pairs are then concatenated to produce the aggregated rich contextual features for the RGB information $\mathbb{ARC}_{\text{RGB}}$. The whole process

of this context aggregation operation is:

$$\mathbb{RC}_{r_i, r_j}^{\text{RGB}} = \mathbb{C}_{r_i}^{\text{RGB}} + \mathbb{C}_{r_j}^{\text{RGB}} (r_i < r_j), \quad (1)$$

where r_i and r_j are two different dilation rates used to produce the multi-scale context features \mathbb{C} .

The aggregated rich contextual features for the RGB information $\mathbb{ARC}_{\text{RGB}}$ is then computed as:

$$\mathbb{ARC}_{\text{RGB}} = \text{Concat}(\underbrace{\mathbb{RC}_{1,2}^{\text{RGB}}, \mathbb{RC}_{1,4}^{\text{RGB}}, \dots, \mathbb{RC}_{2,8}^{\text{RGB}}, \mathbb{RC}_{4,8}^{\text{RGB}}}_{(4)}), \quad (2)$$

$\mathbb{ARC}_{\text{RGB}}$ are then forwarded to a channel-wise attention (CNA) and a context-wise attention (CXA). Unlike CNA, which computes individual weights for different channels in the input features, the attention weights for the CXA are shared across all channels within the same context by squeezing and broadcasting. We finally multiply the channel-wise attention weights, the context-wise attention weights and the input features together to form the final output y^{RGB} . Similar to the RGB context mining submodule,

we can obtain the aggregated rich contextual features for the depth information $\text{ARC}^{\text{depth}}$ by the same module design.

Implicit Multi-modal Context Mining Submodule. The RGB and depth context mining submodules only consider a single-modal input. To model cross-modal contexts, we propose the implicit multi-modal context mining submodule. Architecturally, this submodule has the same design as the RGB and the depth context mining submodules, and outputs implicit multi-modal rich context features ARC^{imp} . Like the first two submodules, we then apply channel-wise attention and context-wise attention on ARC^{imp} to obtain the final output y^{imp} of this submodule.

Explicit Multi-modal Context Mining Submodule. Since the implicit multi-modal context mining submodule only takes in the fused single-scale RGB and depth features and cannot disentangle multi-modal contexts in multiple scales, it is insufficient for modeling the contextual associations between different modalities. To capture the multi-modal contextual information in multiple scales explicitly, we propose the explicit multi-modal context mining submodule. It utilizes the rich context features $\text{RC}_{r_i, r_j}^{\text{RGB}}$ and $\text{RC}_{r_i, r_j}^{\text{depth}}$ generated by the single-modal context mining submodules (*i.e.*, RGB and depth context mining submodules). Each set of rich context features from the same scale are forwarded to a 3×3 convolution with a r_j dilation rate. The aggregated rich contextual features for the explicit multi-modal rich context features ARC^{exp} are computed as:

$$\begin{aligned} \text{RC}_{r_i, r_j}^{\text{exp}} &= \psi_{r_j}(\text{RC}_{r_i, r_j}^{\text{RGB}} + \text{RC}_{r_i, r_j}^{\text{depth}}) \\ \text{ARC}^{\text{exp}} &= \text{Concat}(\underbrace{\text{RC}_{1,2}^{\text{exp}}, \text{RC}_{1,4}^{\text{exp}}, \dots, \text{RC}_{2,8}^{\text{exp}}, \text{RC}_{4,8}^{\text{exp}}}_{(4)}), \quad (3) \end{aligned}$$

Finally, we forward ARC^{exp} to our channel-wise attention and the context-wise attention to obtain the output y^{exp} .

Adaptive Selection. Simply combining single-modal contextual features and multi-modal contextual features will cause a performance drop owing to the presence of a domain gap between different modalities. Thus, we adaptively select the features from RGB context, depth context, implicit multi-modal context, and explicit multi-modal context information by dynamically adjusting the importance of different contextual features. To achieve this goal, we concatenate the outputs of these four context mining submodules as the input of our adaptive selection process, denoted as $x^{\text{sel}} \in \mathbb{R}^{4C \times H \times W}$. We feed x^{sel} to a 1×1 convolution layer to reduce its channel size to C . After that, we apply the channel-wise attention mechanism to these features in order to capture the importance of each channel. Based on the extracted channel-wise attention weights, we can adaptively select the context features by multiplying these weights to the input features to obtain the final cross-modal rich contextual features CRC as the output of the CCM module.

Depth-missing Aware Attention (DAA) Module

We observe that missing depth often appears around glass surfaces in the depth map due to light transmission, refraction and possibly reflection of the glass surface. To exploit

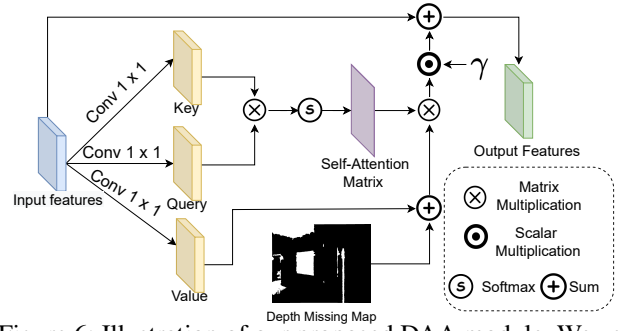


Figure 6: Illustration of our proposed DAA module. We use three individual 1×1 convolutions to generate key, query and value features from the input features. The key and query are multiplied and forwarded to Softmax to form the self-attention matrix. The depth missing map is added to the value features, multiplied by the self-attention matrix and a learnable weight γ to obtain the enhanced features, which are further added to input features to produce DAA’s output.

this cue, we propose a novel DAA module to explicitly involve the spatial information of the depth-missing regions into our framework. The proposed DAA module takes CRC , y^{RGB} , and y^{depth} as the input features. Besides, we also take the resized depth-missing map Dm as input. Fig. 6 illustrates the structure of the proposed DAA module. Formally, the DAA module is defined as:

$$\begin{aligned} f_k^i &= \varphi_v^i(x_i) + Dm, \\ f_{out}^i &= \gamma^i \sigma(\varphi_q^i(x_i) \varphi_k^i(x_i)^T) f_k^i + x_i, \end{aligned} \quad (4)$$

where φ_q^i , φ_k^i and φ_v^i are 1×1 convolution layers for the i modality features. σ is a softmax function. γ^i is a learnable weighting parameter for the i modality features. f_{out}^i are the output features for modality i , where $i \in \{\text{cm}, \text{RGB}, \text{depth}\}$ is the cross-modal, RGB and depth modality, respectively.

Loss Functions

We use a hybrid loss function $\mathcal{L} = \sum_{i=1}^N (\mathcal{L}_{BCE} + \mathcal{L}_{IoU})$, where \mathcal{L}_{BCE} and \mathcal{L}_{IoU} are the binary cross-entropy loss and the interaction-over-union loss, respectively, between the predicted glass surfaces on the i -th stage and the ground truth glass surface map. N is the number of stages.

Experiments

Datasets and Evaluation Metrics. We evaluate our method in two sets of experiments: one using our RGB-D dataset and the other using existing RGB glass datasets such as GDD and GSD. We use four metrics for evaluation, including intersection-over-union (IoU), F-measure (F_β), mean absolute error (MAE), and balance error rate (BER).

Implementation Details We use ResNext-101 (Xie et al. 2017) pretrained on ImageNet as our backbone for the RGB image. We use the Adam optimizer (Kingma and Ba 2015) with an initial learning rate of $1e - 4$, training epoch 130 and batch size 14. Note that we do not apply any post-processing technique (*e.g.*, conditional random field (CRF) (Krähenbühl and Koltun 2011)) to our predicted maps

Methods	Venue	RGB-D GSD				GDD				GSD			
		IoU \uparrow	F β \uparrow	MAE \downarrow	BER \downarrow	IoU \uparrow	F β \uparrow	MAE \downarrow	BER \downarrow	IoU \uparrow	F β \uparrow	MAE \downarrow	BER \downarrow
SPNet	ICCV 2021	0.706	0.831	0.050	11.41	-	-	-	-	-	-	-	-
CLNet	ICCV 2021	0.707	0.829	0.051	11.00	-	-	-	-	-	-	-	-
GDNet	CVPR 2020	0.468	0.631	0.119	19.25	0.814	0.909	0.097	8.83	0.790	0.869	0.069	7.72
EBLNet	ICCV 2021	0.707	0.819	0.048	10.91	0.870	0.922	0.064	6.08	0.817	0.878	0.059	6.75
GSDNet	CVPR 2021	0.714	0.822	0.048	9.73	0.881	0.932	0.059	5.71	0.836	0.903	0.055	6.12
RFENet	IJCAI 2023	0.699	0.825	0.046	11.42	0.874	0.929	0.062	5.79	0.836	0.904	0.049	6.24
Ours		0.742	0.853	0.043	9.33	0.883	0.933	0.059	5.65	0.849	0.912	0.050	6.02

Table 1: Quantitative results on the RGB-D GSD, GDD, and GSD datasets.




Figure 7: Visual comparison of our method with state-of-the-art methods on images from our RGB-D GSD dataset.

for final output. Our model takes about 14 hours to converge, and 0.10s per image for inference on a single RTX2080Ti.

Quantitative Evaluation. We evaluate the performance of the proposed method quantitatively on both RGB-D and RGB glass surface detection on the proposed RGB-D GSD dataset and the existing glass dataset GDD and GSD. We compare the proposed method with SPNet (Zhou et al. 2021) and CLNet (Zhang et al. 2021a) for RGB-D salient object detection; GDNet (Mei et al. 2020), GSDNet (Lin, He, and Lau 2021), EBLNet (He et al. 2021) and RFENet (Fan et al. 2023) for RGB glass surface detection. Additional results, including comparisons with more relevant baseline methods (e.g., MirrorNet (Yang et al. 2019) and PMD (Lin, Wang, and Lau 2020) for mirror detection), are provided in the Supplemental. For RGB-based methods, we use their publicly available codes with default configurations and we only use the RGB images and the ground truth masks for training and testing on our RGB-D GSD dataset. For RGB-D-based methods (e.g., SPNet and CLNet), we do not evaluate them on the RGB datasets (i.e., GDD and GSD) due to the absence of GT depth maps in these datasets.

To adopt our method to RGB glass surface detection, we keep the RGB backbone and the RGB context mining submodule used in the CCM module, with the depth backbone network and the DAA modules removed from our original framework. Table 1 shows the experimental results. We can see that our proposed method significantly outperforms these baseline methods on all four metrics on our RGB-D GSD dataset. In particular, our method shows substantial im-

provement on MAE, with a performance increase by 10.42% over the second-best method GSDNet. Besides, our method outperforms all compared methods on both RGB glass surface detection datasets, GDD and GSD, even though it does not use any auxiliary information as some glass surface detection methods do (e.g., boundary labels used in EBLNet and the reflection maps used in GSDNet). Specifically, our method achieves a significant performance improvement on the more challenging GSD dataset. The reason for our method to obtain a relatively minor performance gain on GDD dataset is that GDD contains images mostly captured from limited scenes and existing methods can perform very well on them. These results show that our method with RGB context mining submodules is particularly effective for glass surface detection in complex real-world scenes.

Qualitative Evaluation. We further demonstrate our method visually in Fig. 7, comparing it with four state-of-the-art glass surface detection methods. The 1st row of Fig. 7 shows that state-of-the-art methods may struggle with visually ambiguous glass regions due to uncommon context. In these cases, using only RGB information to predict the glass surfaces can be challenging. In the 2nd row, our method correctly detects over-exposed glass regions, by integrating both RGB and depth contextual information. The last row shows a challenging example with extremely tiny glass surface regions. Our method still significantly outperforms the baselines. We attribute it to the multi-scale context features incorporated in our proposed CCM module.

Ablation Study. Finally, we perform an ablation study to

Methods	IoU \uparrow	F_{β} \uparrow	MAE \downarrow	BER \downarrow
B	0.703	0.814	0.046	11.14
B + CCM	0.727	0.836	0.045	9.59
B + DAA	0.706	0.819	0.046	10.48
Ours w/ RGB only	0.686	0.802	0.052	11.47
Ours w/ RGB + Gray	0.325	0.646	0.112	33.02
Ours w/ RGB + Black	0.352	0.639	0.112	32.87
Ours	0.742	0.853	0.043	9.33

Table 2: Ablation studies. ‘‘B’’ is the base model. The upper section shows the ablation study on the effectiveness of CCM and DAA modules, while the lower section shows the study on the benefits of the depth cue.

evaluate each of the proposed components of our model. Table 2 shows the results of our ablation study on our proposed dataset. We can see that adding either the CCM module or the DAA module (by replacing CMM with a single convolution layer to integrate multi-modal inputs) helps improve the model performances, but our final model with both CCM and DAA modules performs the best on all four metrics. Note that adding the CCM modules to the base network (*i.e.*, RGB and depth backbone networks with decoders) to form an ablated model (‘‘B + CCM’’) significantly outperforms the ablated model (‘‘B + DAA’’) that adds the DAA module to the base network. We attribute this to the success of our cross-modal contextual mining process conducted by the CCM module, which benefits the RGB-D glass surface detection task from a global view. Figure 8 shows a visual example of the component analysis. We can see that the proposed CCM and DAA modules can help improve performance by exploiting cross-modal and depth-missing information to remove the over-predicted regions. To verify the benefits of the depth cue in RGB-D glass surface detection, we conduct the following experiments: 1) remove all depth-related modules (*i.e.*, only keep the RGB backbone and the RGB context mining submodule in CCM) and then retrain the proposed network (‘‘Ours w/ RGB only’’); 2) replacing the depth map with the grayscale image of the input RGB image during inference (‘‘Ours w/ RGB + Gray’’); 3) replacing the depth map with a black (empty) image during inference (‘‘Ours w/ RGB + Black’’). We can see that removing depth-related modules and using only RGB information 1) produces unsatisfying results on our RGB-D GSD dataset. In addition, replacing the depth map with the grayscale version of the input image or with an empty map significantly decreases the model performance. These experimental results show the effectiveness and importance of the depth cue for RGB-D glass surface detection.

Effectiveness of the DAA Module. Table 3 shows the ablation study on our proposed DAA module. ‘‘DAA w/o D_m ’’ refers to the DAA module without taking the depth-missing map as input. We design three other ablated models: ‘‘DAA on RGB’’, ‘‘DAA on Depth’’, and ‘‘DAA on CM’’ as adopting



Figure 8: A visual example of the ablation study.

Methods	IoU \uparrow	F_{β} \uparrow	MAE \downarrow	BER \downarrow
DAA w/o D_m	0.733	0.835	0.045	9.92
DAA on RGB	0.738	0.838	0.044	9.62
DAA on Depth	0.729	0.831	0.048	9.85
DAA on CM	0.739	0.846	0.045	9.30
Ours	0.742	0.853	0.043	9.33

Table 3: Ablation study of the DAA module.

the DAA module only on the RGB, depth, and cross-modal contextual features from the CCM modules, to test the effectiveness of the DAA module for extracting features in different modalities. Our final model adopts the DAA module with all three modalities. Experimental results show that the depth-missing information plays a key role in the DAA module, and our DAA module can effectively enhance the feature representation from different modalities.

Conclusion

In this paper, we have investigated the glass surface detection problem by considering the depth information. To this end, we first construct a new large-scale RGB-D glass surface detection dataset containing 3,009 images with the corresponding depth maps and annotations. This dataset covers diverse scenes with glass surfaces and can facilitate research on glass surface detection. We then propose a novel RGB-D framework for glass surface detection. Our framework consists of two novel modules: (1) a cross-modal context mining (CCM) module for mining the context information among different modalities, and (2) a depth-missing aware attention (DAA) for exploiting the depth missing information around glass surfaces. Experimental results show the superior performances of our proposed framework, compared with state-of-the-art methods from relevant fields.

Despite its success, our method struggled under extreme conditions, such as low-light scenarios. In cases of perfect transmission, where depth signals fail to return and thus appear absent in depth scans, the depth context becomes irrelevant. Additionally, reflections from mirrors, which produce an almost identical image of the scene, continue to pose a significant challenge.




Figure 9: Failure cases. Our method may fail with low light scene, perfect glass transmission or reflection.

Acknowledgments

This work is in part supported by two GRF grants from the Research Grants Council of Hong Kong (RGC No.: 11211223 and 11220724).

References

- Armeni, I.; Sax, A.; Zamir, A. R.; and Savarese, S. 2017. Joint 2D-3D-Semantic Data for Indoor Scene Understanding. *ArXiv e-prints*.
- Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Niessner, M.; Savva, M.; Song, S.; Zeng, A.; and Zhang, Y. 2017. Matterport3D: Learning from RGB-D Data in Indoor Environments. *International Conference on 3D Vision (3DV)*.
- Chen, X.; Zhang, H.; Yu, Z.; Opipari, A.; and Chadwicke Jenkins, O. 2022. Clearpose: Large-scale transparent object dataset and benchmark. In *ECCV*, 381–396. Springer.
- Fan, K.; Wang, C.; Wang, Y.; Wang, C.; Yi, R.; and Ma, L. 2023. RFENet: Towards Reciprocal Feature Evolution for Glass Segmentation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, 717–725. ijcai.org.
- Guo-Hua, C.; Jun-Yi, W.; and Ai-Jun, Z. 2019. Transparent object detection and location based on RGB-D camera. *Journal of Physics: Conference Series*, 1183: 012011.
- He, H.; Li, X.; Cheng, G.; Shi, J.; Tong, Y.; Meng, G.; Prinnet, V.; and Weng, L. 2021. Enhanced Boundary Learning for Glass-Like Object Segmentation. In *ICCV*, 15859–15868.
- Islam, M. N.; Tahtali, M.; and Pickering, M. 2021. Specular Reflection Detection and Inpainting in Transparent Object through MSPLFI. *Remote Sensing*, 13(3).
- Kalra, A.; Taamazyan, V.; Rao, S. K.; Venkataraman, K.; Raskar, R.; and Kadambi, A. 2020. Deep Polarization Cues for Transparent Object Segmentation. In *CVPR*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Krähenbühl, P.; and Koltun, V. 2011. Efficient inference in fully connected crfs with gaussian edge potentials. In *NeurIPS*.
- Lin, J.; He, Z.; and Lau, R. W. 2021. Rich Context Aggregation with Reflection Prior for Glass Surface Detection. In *CVPR*.
- Lin, J.; Wang, G.; and Lau, R. W. 2020. Progressive Mirror Detection. In *CVPR*.
- Liu, F.; Liu, Y.; Lin, J.; Xu, K.; and Lau, R. W. 2024. Multi-View Dynamic Reflection Prior for Video Glass Surface Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3594–3602.
- Mei, H.; Yang, X.; Wang, Y.; Liu, Y.; He, S.; Zhang, Q.; Wei, X.; and Lau, R. W. 2020. Don’t Hit Me! Glass Detection in Real-World Scenes. In *CVPR*.
- Ramirez, P. Z.; Costanzino, A.; Tosi, F.; Poggi, M.; Salti, S.; Mattoccia, S.; and Di Stefano, L. 2023. Booster: a benchmark for depth from images of specular and transparent surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ramirez, P. Z.; Tosi, F.; Poggi, M.; Salti, S.; Mattoccia, S.; and Di Stefano, L. 2022. Open challenges in deep stereo: the booster dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21168–21178.
- Seib, V.; Barthen, A.; Marohn, P.; and Paulus, D. 2017a. Friend or foe: exploiting sensor failures for transparent object localization and classification. In *2016 International Conference on Robotics and Machine Vision*, volume 10253, 94–98. SPIE.
- Seib, V.; Barthen, A.; Marohn, P.; and Paulus, D. 2017b. Friend or foe: exploiting sensor failures for transparent object localization and classification. In Bernstein, A. V.; Olaru, A.; and Zhou, J., eds., *2016 International Conference on Robotics and Machine Vision*, volume 10253, 94 – 98. International Society for Optics and Photonics, SPIE.
- Song, S.; Lichtenberg, S. P.; and Xiao, J. 2015. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *CVPR*, 567–576.
- Sun, T.; Zhang, G.; Yang, W.; Xue, J.-H.; and Wang, G. 2023. Trosd: A new rgb-d dataset for transparent and reflective object segmentation in practice. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(10): 5721–5733.
- Xie, E.; Wang, W.; Wang, W.; Ding, M.; Shen, C.; and Luo, P. 2020. Segmenting Transparent Objects in the Wild. In *ECCV*.
- Xie, E.; Wang, W.; Wang, W.; Sun, P.; Xu, H.; Liang, D.; and Luo, P. 2021. Segmenting Transparent Object in the Wild with Transformer. In *IJCAI*.
- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *CVPR*.
- Xu, Y.; Nagahara, H.; Shimada, A.; and Taniguchi, R.-i. 2015. TransCut: Transparent Object Segmentation From a Light-Field Image. In *ICCV*.
- Yang, X.; Mei, H.; Xu, K.; Wei, X.; Yin, B.; and Lau, R. W. 2019. Where is my mirror? In *ICCV*.
- Zhang, J.; Fan, D.-P.; Dai, Y.; Yu, X.; Zhong, Y.; Barnes, N.; and Shao, L. 2021a. RGB-D Saliency Detection via Cascaded Mutual Information Minimization. In *ICCV*.
- Zhang, J.; Yang, K.; Constantinescu, A.; Peng, K.; Müller, K.; and Stiefelhagen, R. 2021b. Trans4Trans: Efficient Transformer for Transparent Object and Semantic Scene Segmentation in Real-World Navigation Assistance. *CoRR*, abs/2108.09174.
- Zhou, T.; Fu, H.; Chen, G.; Zhou, Y.; Fan, D.-P.; and Shao, L. 2021. Specificity-preserving RGB-D Saliency Detection. In *ICCV*.