

Fast Training Data Acquisition for Object Detection and Segmentation using Black Screen Luminance Keying

Thomas Pöllabauer

Volker Knauth

André Boller

Arjan Kuijper
Fraunhofer IGD, Fraunhoferstrasse 5, 64283, Darmstadt, Germany
TU Darmstadt, Karolinenplatz 5, 64289, Darmstadt, Germany
thomas.poellabauer@igd.fraunhofer.de

Dieter W. Fellner

ABSTRACT

Deep Neural Networks (DNNs) require large amounts of annotated training data for a good performance. Often this data is generated using manual labeling (error-prone and time-consuming) or rendering (requiring geometry and material information). Both approaches make it difficult or uneconomic to apply them to many small-scale applications. A fast and straightforward approach of acquiring the necessary training data would allow the adoption of deep learning to even the smallest of applications. Chroma keying is the process of replacing a color (usually blue or green) with another background. Instead of chroma keying, we propose luminance keying for fast and straightforward training image acquisition. We deploy a black screen with high light absorption (99.99%) to record roughly 1-minute long videos of our target objects, circumventing typical problems of chroma keying, such as color bleeding or color overlap between background color and object color. Next we automatically mask our objects using simple brightness thresholding, saving the need for manual annotation. Finally, we automatically place the objects on random backgrounds and train a 2D object detector. We do extensive evaluation of the performance on the widely-used YCB-V object set and compare favourably to other conventional techniques such as rendering, without needing 3D meshes, materials or any other information of our target objects and in a fraction of the time needed for other approaches. Our work demonstrates highly accurate training data acquisition allowing to start training state-of-the-art networks within minutes.

Keywords

Machine Learning, Object Detection, Object Segmentation, Deep Neural Networks.

1 INTRODUCTION

Modern machine learning (ML) is dominated by deep neural networks (DNNs). Training DNNs to state-of-the-art performance levels tends to require large amounts of training data to perform well. In many cases the lack of annotated data prevents the use of these networks. In the case of object segmentation and object detection a common approach - aside of costly manual labeling - is to resort to rendering to generate the required training images. To do so 3D meshes, textures, and material properties of the target objects are required, which is another barrier for many applications. In contrast, we only require the availability of the objects in question, a camera, a sufficiently large piece of special black cloth, and some lights. We completely circumvent many of the problems with traditional chroma keying such as color bleeding, same foreground/background color and similar, by utilizing a very low reflectance cloth and demonstrate its applicability in a fast, low-cost, high-quality setup, comparing favorably to much more complex data generation regimen.

Our contributions are the following:

- We propose a straightforward, easy to use setup to record high-quality training datasets for object segmentation and object detection.
- We present extensive evaluation and show the performance of our approach in comparison with other conventional data generation methods that require more information, such as meshes, textures, and materials, and/or much more processing time. To make our results more meaningful for the research community, we do all our evaluation on the common YCB-V dataset.
- We provide code, which automatically converts the recordings to datasets in COCO format for use with segmentation and 2D object detection algorithms, as well as our black screen recordings of the YCB-V objects.



Figure 1: A qualitative sample of all 21 YCB-V objects, that were recorded with a handheld smartphone and our proposed black background. It can be seen that the objects are well silhouetted against the background and can therefore be segmented in an easy way and many typical chroma key-associated problems are circumvented.

2 RELATED WORK

2.1 Data Generation for ML

Insufficient training data is a well documented problem in Machine Learning. In the domain of Computer Vision practitioners soon adopted rendering for data generation. Rendering has many suitable attributes, such as perfect ground truth, potentially unlimited amounts of data, and total control of scene composition, such as object pose and scene lighting. [TTS⁺18] demonstrates the combination of rendering on random backgrounds, as well as realistically placed objects in 3D scenes, while [HPH⁺19] shows another purely rendering-based approach. Another very common tool for data generation is BlenderProc [DSW⁺20], a pipeline extending Blender that allows for physically-based rendering. For a long time, the gap in feature representations between synthetic and real-world images was a problem [TFR⁺17, TPA⁺18, WGS⁺21]. With the use of physically-based rendering and the emergence of large foundation models such as CLIP [LLSH23, ODM⁺23, RKH⁺21], this problem is greatly reduced [SHL⁺23]. Another problem is not to be solved that easily however: In order to render you need 3D meshes of the objects in questions. There are approaches to estimate 3D shape [LGL⁺23] or additional views [VYB⁺24] based on a single image, but their quality is not on par with reality and their reconstruction is often purely probabilistic. Reconstructing 3D meshes with traditional methods such as photogrammetry [Sch05] is a very time consuming process. There is a move towards zero-shot ap-

proaches such as [NGP⁺23], which does not need images of the target objects for training, though still requires 3D meshes.

2.2 Significance of Reliable Data

The use of modern DNNs in computer vision improved solutions for a large variety of challenging tasks, provided sufficient training data. A prominent modern example would be Segment Anything (SAM), which is able to segment a very large amount of different objects in the wild [KMR⁺23]. However, ML processes come with significant risks and difficult to detect silent failures, if not handled correctly [HJS⁺20]. This is especially crucial in settings, where reliability and robustness is mandatory, such as in an industrial, clinical or dangerous contexts, which in turn can task-invalidate large unsupervised networks like SAM, due to its own limitations. One predominant challenge is, that trustworthy ML models need high-quality base data [LTH⁺22]. Many datasets are not task suited due to their size, inherent bias, dirtiness, or even just partial unfairness\incorrectness [WRSL23]. While these issues can be partly mitigated or worked around [RHW19, WRSL23], the relevant techniques introduce new layers of complexity and potential error sources. Albeit, all of these challenges can be overcome with manual labour and diligence, this can lead to the preemptive end for startups [BIRS22] and drive up cost for large companies.

2.3 Chroma Keying

Chroma Keying is a well established movie production technique from the 1920's, to segment objects in front of fixed mono-color backgrounds. In the beginning, a variety of different colors and shades like black/white, yellow, blue and green were popular [Fos10, Ram15]. Over time, the color green became the predominant background color for a variety of pragmatic reasons. It is easily distinguishable from human skin, rarely occurs outside of nature, does not require much lighting and is favourable for modern digital camera sensors because of higher sensitivity. However, some significant challenges remain: Green objects lead to falsely segmented foreground, which can also happen as a byproduct of color-bleeding or reflection. These effects strongly reduce segmentation quality for reflective materials, such as in metallic, transparent, and bright objects. Furthermore, object borders can become fuzzy due to lighting and merging effects with the background. While there are techniques to remedy drawbacks of conventional green screen, like color-unmixing [AAPS16], specialised capturing processes [BRS⁺22], color spill neutralization [GKTB10], threshold optimization [PJS17] and multiple background colors [SB96], they open up new problems and do not yet completely solve the inherent challenges while introducing additional capturing and/or post-processing effort.

2.4 Differentiation from similar works

LeCun et al. [LHB04] use a gray turn table to record objects and rely on chroma keying to replace foreground and background, while Dirr et al. use a white background [DBGD24]. Though a brighter background, like grey or white, would also satisfy the idea of luminance keying, they have major drawbacks in comparison to 99,99% light absorption black. First of all, they introduce far greater challenges for correct lighting conditions, arising from background shadows, background reflections and possible fuzzy edges, due to light scattering. Furthermore, it is harder to differentiate brighter object colors with bright backgrounds. In some cases, e.g. metallic or transparent objects, a non-black background is also prone to light shining partly through an object, reflections and refractions. These challenges are all naturally solved by our proposed solution, as no light is reflected or visible on the background by any relevant measure. Knauthe et al. use a capturing system, which foregoes a background for a white light source in a turn-table setup, which diminishes some of the challenges of a normal white background [KKvB⁺22]. However, this setup is very constrained in its rotation invariant use-case, due to the difficulties of building suitable shaped large light backgrounds. Agata et al. introduce dual color checker pattern backgrounds [AYK07, YAK08]. While this approach solves

foreground/background color confusion, the other issues presented earlier still persist. Additionally, the method requires more specialized backgrounds and introduces further complexity in the processing step, such as parameter tuning. Jin et al. developed a deep learning method for automatic real-time green screen keying [JLZ⁺22]. However, it is still limited by shadows and green spilling and requires a deep learning method with all inherent benefits and challenges.

3 APPROACH

3.1 Chroma and Luminance Key

We propose a simple yet effective data recording process: we place the objects on a very low-reflectance cloth, and record around 1-minute long video clips with a smartphone. These clips are processed via a cut and paste process, that uses the easily masked objects and places them on random backgrounds. Details on the processing are discussed in section 3.3.

3.1.1 Chroma Key with Green Screen

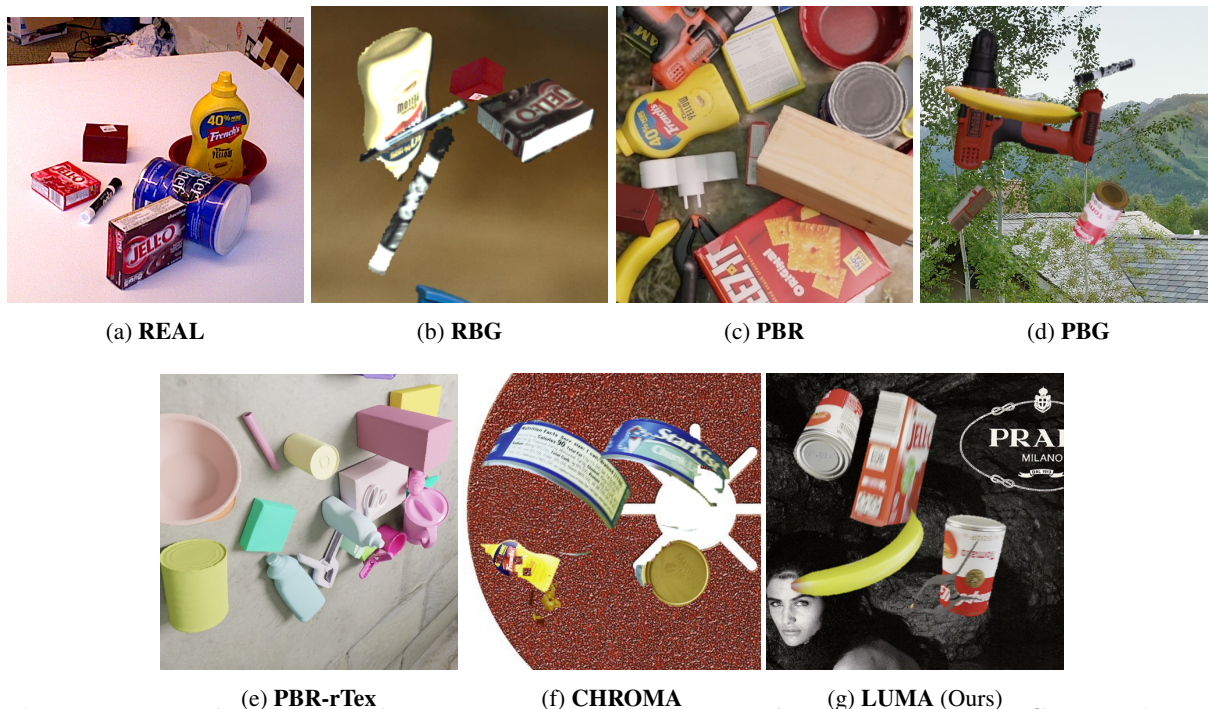
For comparison, we include recordings and experiments with the more common chroma keying approach. Arguably the most common colors are blue and green, with green being the most widely used color. Therefore we include a green screen in our evaluation.

3.1.2 Luminance Key with Black Screen

We propose Luminance Keying for fast and straightforward data acquisition for machine learning. This technique utilizes the brightness difference between an object and the background, instead of a designated color as with chroma keying. In practice, this is possible due to a textile background, which absorbs 99,99% of visible light. This leads to high contrast between object and background, even for very dark objects, allowing for high quality masking as illustrated in Figure 1. The recording process is exactly the same as with the conventional green screen described in section 3.1.1.

3.2 Baselines and Experiments

For our evaluation we include different data sources, all visualized in Figure 2: First, we include the real-world recordings of the YCB-V dataset as used in the prominent BOP Challenge [SHL⁺23]. However, there is a problem with the YCB-V test images: one could argue, that the limited number of scenes and camera poses reduces the expressiveness in terms of generalization. Therefore we also evaluate on the physically-based rendering (pbr) dataset provided by BOP. This dataset uses the reconstructed textures, leading to a high similarity in appearance compared to the real images, but has a much wider range of camera poses, object configurations, and lighting situations. These two



(a) **REAL** (b) **RBG** (c) **PBR** (d) **PBG**
(e) **PBR-rTex** (f) **CHROMA** (g) **LUMA (Ours)**
Figure 2: Samples from a subset of our evaluated training datasets. **REAL** are real images, **RBG** are real images with replaced backgrounds, **PBR** are physical based renderings, **PBG** are physical based renderings with background replacement, **PBR-rTex** are PBRs with randomized textures and **CHROMA**, as well as **LUMA (Ours)** stand for different capturing methods.

datasets, namely **REAL** and **PBR**, are used as a baseline on which all other data sources have to be tested. Also, we test all of our approaches on in-distribution samples, that is, on disjoint splits from the dataset, i.e. when training on our luminance images, we test on **REAL**, **PBR**, as well as on a disjoint set of luminance images. This should serve as a measure on how well an experiment generalizes to out-of-distribution samples. In more detail we use the following data representations:

Real-world images (REAL). Real-world recordings as described above.

Real-world images with background replacement (RBG). Real-world recordings. To test the influence of our background replacement script, we apply it to the real images as well. This gives us an idea on how much performance we loose because of our simple crop and paste approach of data generation.

Physically-based renderings (PBR). PBR renderings with reconstructed textures, realistic object placement, and light transport, as well as a high degree of camera and object pose variation.

Physically-based renderings with background replacement (PBG). Same as above, but again using background replacement for measuring its influence on performance.

Physically-based renderings with randomized textures (PBR-rTex). To simulate cases in which geometry is available, though we lack realistic textures

and materials, we include a set of pbr images with randomized surface attributes. Namely we randomize texture and surface details such as reflectance behavior. **Chroma key images using green screen (CHROMA).** A set of green screen recordings to illustrate what results one can expect from using run-of-the-mill chroma keying methods. We use the green screen for segmentation and paste the crops on random photographs.

Luminance key images using black screen (LUMA). Our proposed method. A set of recordings utilizing the low-reflection black screen. Again, we have to crop and paste the objects on random backgrounds using our replacement script.

3.3 Training

We demonstrate the applicability of our approach on the 2D object detection use case and train YOLOX [GLW+21] networks.

Data Generation using cut and paste background replacement. To achieve technically good replacement results, we use a variety of probabilistic mechanisms that are applied during the replacement process inspired by [DMH17]. First, based on the masks provided, we filter out all objects in an image that are overlapped by others. This is trivial for luma, because there is no overlap during the acquisition, giving us near-perfect masks. In addition, we also make sure that the objects are not cut off at the edge of the image and finally do

a crop, centered around the object. Next we apply random affine transformations such as scaling, rotation and translation to the objects where the variable scaling, rotation and translation ensures variety in the appearance of the objects. After this process, the objects are placed on the background image based on randomly generated positions. We allow an overlap among objects of up to 20% in order to come closer to the real scenarios of YCB-V and make the network learn to deal with occlusion. For blending and to remove jaggies, the object masks are eroded and Gaussian noise is applied. As for the backgrounds we used random crops of images from the 50k HQ-data set [YCT⁺23]. Random cropping prevents the same background from being used more than once. **Parameterization.** To evaluate the practicability of our proposed method, we do an extensive comparison of different training sets. As shown in [HLWK18] freezing the backbone layers is useful for reducing the impact of domain shift. Since we introduce additional domain discrepancy with our background replacement, we add freezing to our evaluation of training from scratch, and using COCO-pretrained weights for initialization. As for the parameterization of YOLOX, we stick closely to the values and process as presented in their paper [GLW⁺21]. Most importantly for our comprehensive quantitative analysis we train model size "tiny" with image size 512x512, batch sizes of 64, half precision, multiscale range of 13, and for 300 epochs.

4 RESULTS

We report quantitative results on the common YCB-V object set and show samples of the achievable visual quality of the approach. In addition we show some problems of chroma keying, which luminance keying does not have in section 4.2.

4.1 Quantitative Results

In our first experiments we looked into answering the question, whether we want to train from scratch or load pre-trained weights, as well as whether to freeze the backbone. For a comprehensive list of all results, especially for unfrozen backbones, please refer to our tables presented in the appendix. Here we report results when freezing the backbone, as these led to best performance with our cut and paste background replacement. For evaluation we report both the Average Precision (AP) and the Average Recall (AR) metric in Tables 1 and 2. Since the currently available set of YCB-V objects has three modified objects (002_master_chef_can has a modified texture, for 019_pitcher_base the color changed from blue to a transparent color with a red lid, and for 035_power_drill there are some slight changes to the model), we report results on the subset of unchanged objects and call it YCB-V18, as well as

the results of the complete set.

As expected training and evaluating on REAL lead to best performance overall. This is no surprise since, among all data representations tested, features in the REAL training set are closest to the features found in the REAL test set. Physically-based rendering (PBR) comes in second place. The good performance of PBR can be explained with the photo-realistic depiction, wide range of physically-correct object placements, and lighting variations. Surprisingly the PBR set with randomized textures has very poor performance, far worse than expected. While we assumed a big drop, the extent might have to do with model size and might be less pronounced with more trainable parameters. As the 4 remaining approaches all build upon the cut and paste approach, we can directly compare the quality of the segmentation, as well as image fidelity. Here we find an outlier in the PBG (real photographs with background replacement) set: it under-performs the other approaches by a wide margin. Our assumption is that the reason is the comparatively poor masking quality. Masks are not pixel perfect and sometimes include background pixels or pixels belonging to an occluding object. Mask quality is no problem for RBG (pbr images with background replacement), having perfect masks by virtue of complete scene knowledge, but we argue RBG is held back by other problems, some similar to the ones found with chroma keying with color bleeding being the most obvious. Realistic light transport is a feature in fully rendered scenes, but becomes a detriment with cutting and pasting. Also, the textures, while being high fidelity, are not as good as real photographs. Finally, comparing LUMA with CHROMA we see an outperformance of more than 11% on both YCB-V18, as well as on the complete YCB-V set.

4.2 Qualitative Results

We depict some of the problems of chroma keying in Figure 3 and show the quality of our approach in Figure 1. In direct comparison with chroma keying using the green screen we note a significant improvement in mask quality. At the same time the process becomes noticeably faster and less error-prone since we do not have to match a specific color when thresholding to generate our masks. This problem is illustrated in Figure 4.

4.3 Discussion

We see a noticeable drop in performance related to our background replacement. This is to be expected since we remove relevant cues for the 3D scene understanding, such as global illumination, shadows, color bleeding, and realistic occlusion. However, when considering the subset of data representations utilizing the background replacement, namely REAL, PBR, CHROMA, and LUMA, we see the overperformance of LUMA.

TRAIN SET:	PBR-rTex	PBR	PBG	REAL	RBG	CHROMA	LUMA (Ours)
002_master_chef_can*	0.19/ 1.47 /-	44.65/ 62.13 /44.65	21.74/ 53.56 /17.35	6.29/ 73.76 /73.76	2.2/ 2.84 /0.0	9.54/ 46.75 /1.56	11.74/ 46.17 /3.64
003_cracker_box	0.09/ 0.0 /-	46.86/ 20.65 /46.86	16.33/ 5.15 /14.3	12.54/ 43.34 /43.34	10.06/ 9.65 /0.0	5.31/ 10.2 /0.02	13.68/ 20.9 /0.51
004_sugar_box	0.23/ 0.02 /-	32.95/ 48.36 /32.95	11.77/ 44.75 /13.03	5.62/ 49.54 /49.54	9.37/ 9.26 /66.26	9.42/ 38.23 /0.13	14.97/ 52.64 /0.48
005_tomato_soup_can	0.08/ 0.3 /-	38.04/ 55.53 /38.04	9.59/ 59.13 /10.47	5.75/ 70.21 /70.21	9.44/ 9.43 /78.67	4.73/ 36.42 /0.46	10.27/ 29.17 /1.42
006_mustard_bottle	0.65/ 0.14 /-	38.86/ 78.2 /38.86	14.19/ 76.07 /7.38	10.68/ 88.82 /88.82	12.7/ 12.45 /9.72	16.75/ 73.42 /2.48	17.63/ 79.9 /5.3
007_tuna_fish_can	0.24/ 0.2 /-	45.98/ 70.77 /45.98	11.04/ 65.36 /19.02	12.55/ 79.35 /79.35	10.69/ 10.74 /72.85	5.8/ 53.93 /0.23	11.34/ 60.97 /1.15
008_pudding_box	0.11/ 0.0 /-	28.54/ 4.48 /28.54	17.19/ 1.2 /22.23	1.77/ 6.64 /6.64	9.88/ 9.84 /0.0	4.05/ 2.43 /1.16	5.46/ 0.92 /2.5
009_gelatin_box	0.11/ 0.0 /-	36.82/ 71.21 /36.82	15.61/ 5.73 /22.94	6.11/ 65.41 /65.41	11.61/ 11.12 /66.94	2.26/ 0.52 /0.6	1.33/ 0.55 /2.41
010_potted_meat_can	0.31/ 0.27 /-	36.28/ 37.64 /36.28	11.4/ 18.3 /11.33	4.04/ 54.23 /54.23	9.06/ 9.04 /40.73	1.29/ 0.27 /2.83	11.52/ 26.22 /0.88
011_banana	5.79/ 6.58 /-	34.84/ 54.24 /34.84	14.91/ 40.67 /20.34	15.62/ 50.18 /50.18	14.18/ 14.28 /88.81	17.17/ 51.06 /3.16	14.95/ 48.64 /0.63
019_pitcher_base*	4.27/ 3.12 /-	54.15/ 55.57 /54.15	1.08/ 29.29 /0.55	13.56/ 80.89 /80.89	0.0/ 0.0 /0.0	6.68/ 0.38 /0.67	6.28/ 0.32 /0.09
021_bleach_cleanser	0.92/ 2.56 /-	33.82/ 43.58 /33.82	8.23/ 36.19 /6.22	2.42/ 61.31 /61.31	2.09/ 2.15 /8.35	8.03/ 47.44 /1.56	10.53/ 54.61 /5.26
024_bowl	2.14/ 0.22 /-	54.84/ 46.75 /54.84	4.96/ 0.01 /5.24	9.04/ 56.9 /56.9	2.47/ 1.93 /0.0	0.06/ 0.21 /0.08	6.99/ 14.44 /1.3
025_mug	0.18/ 3.3 /-	46.37/ 58.14 /46.37	10.89/ 1.87 /1.72	4.0/ 67.35 /67.35	3.01/ 3.01 /0.0	5.45/ 31.17 /1.02	2.04/ 14.37 /0.05
035_power_drill*	0.3/ 0.0 /-	23.28/ 13.26 /23.28	3.24/ 0.76 /1.39	2.28/ 43.72 /43.72	0.0/ 0.0 /0.0	3.87/ 5.71 /0.03	4.46/ 13.23 /0.04
036_wood_block	2.1/ 0.37 /-	45.99/ 24.57 /45.99	15.57/ 8.22 /14.56	8.42/ 34.06 /34.06	5.85/ 5.8 /0.0	10.28/ 12.75 /0.64	14.1/ 17.32 /2.53
037_scissors	0.67/ 0.0 /-	19.32/ 5.48 /19.32	0.67/ 0.16 /3.42	5.98/ 1.96 /1.96	1.36/ 1.32 /0.0	1.4/ 0.33 /0.69	1.35/ 0.52 /0.01
040_large_marker	0.04/ 0.0 /-	16.1/ 47.65 /16.1	2.54/ 22.29 /5.2	1.83/ 56.73 /56.73	0.47/ 0.41 /8.58	0.87/ 31.87 /0.23	0.16/ 5.42 /0.4
051_large_clamp	0.85/ 0.0 /-	18.2/ 47.43 /18.2	0.07/ 0.25 /2.5	2.96/ 6.4 /6.4	0.07/ 0.09 /0.01	1.85/ 0.28 /2.01	2.93/ 17.75 /0.07
052_extra_large_clamp	0.38/ 0.17 /-	19.27/ 7.21 /19.27	0.01/ 0.0 /2.45	2.38/ 9.8 /9.8	0.5/ 0.08 /0.0	8.79/ 11.31 /1.13	2.95/ 0.29 /1.82
061_foam_brick	0.71/ 0.01 /-	36.94/ 55.61 /36.94	15.86/ 3.38 /24.54	8.09/ 67.82 /67.82	6.99/ 6.63 /0.0	6.43/ 1.64 /1.11	6.01/ 5.22 /0.91
AVERAGE YCB-V	0.97/ 0.89 /-	35.81/ 43.26 /35.81	9.85/ 22.49 /10.98	6.76/ 50.88 /50.88	5.81/ 5.72 /21.0	6.19/ 21.73 /1.04	8.13/ 24.27 /1.49
AVERAGE YCB-V18	0.87/ 0.78 /-	35.0/ 43.19 /35.0	10.05/ 21.59 /11.73	6.66/ 48.34 /48.34	6.65/ 6.51 /24.5	6.11/ 22.42 /1.09	8.23/ 24.99 /1.54

Table 1: AP (higher is better) of all evaluated data representations, each tested on PBR, REAL, and in-distribution. In-distribution means the test set stems from the same data source (e.g. photograph, rendering, crop and paste) as the training set. Results on REAL in **bold font**. As expected, we see best results when training on the train split of the real data (REAL), followed by the physically simulated 3D scenes (PBR). Most importantly, among all 4 methods using the cut and paste approach for background replacement (PBG, RBG, CHROMA, LUMA), luminance (our proposed method) outperforms all others when tested on the real test data (REAL).

TRAIN SET:	PBR-rTex	PBR	PBG	REAL	RBG	CHROMA	LUMA (Ours)
002_master_chef_can*	8.05/ 18.8 /-	63.67/ 77.07 /63.67	39.66/ 71.67 /54.88	14.9/ 79.0 /79.0	4.59/ 4.63 /0.0	34.94/ 74.8 /9.17	33.24/ 71.97 /16.74
003_cracker_box	7.51/ 0.22 /-	62.82/ 66.44 /62.82	34.83/ 49.02 /56.82	22.93/ 62.89 /62.89	15.5/ 21.64 /0.0	13.62/ 36.22 /1.14	26.64/ 48.53 /4.63
004_sugar_box	6.75/ 2.61 /-	53.39/ 66.8 /53.39	31.07/ 74.37 /48.63	18.53/ 61.63 /61.63	27.76/ 75.68 /77.14	22.13/ 52.69 /5.68	32.84/ 69.07 /7.23
005_tomato_soup_can	4.2/ 5.96 /-	54.06/ 69.71 /54.06	22.89/ 66.91 /35.95	20.81/ 75.09 /75.09	26.39/ 68.96 /84.03	25.11/ 54.96 /7.86	30.4/ 61.8 /14.32
006_mustard_bottle	11.46/ 8.27 /-	57.22/ 87.6 /57.22	27.01/ 83.87 /34.08	23.68/ 91.27 /91.27	28.38/ 82.33 /55.56	30.54/ 83.87 /13.41	32.17/ 84.13 /21.22
007_tuna_fish_can	7.29/ 4.37 /-	57.74/ 78.83 /57.74	26.17/ 72.1 /51.98	30.23/ 83.77 /83.77	30.72/ 75.17 /84.8	24.34/ 72.77 /7.5	24.88/ 73.3 /11.4
008_pudding_box	3.31/ 0.27 /-	48.6/ 56.4 /48.6	33.25/ 23.73 /65.56	14.17/ 48.93 /48.93	24.87/ 54.27 /0.0	26.21/ 43.6 /8.24	27.76/ 27.07 /10.97
009_gelatin_box	2.96/ 0.0 /-	49.85/ 78.4 /49.85	26.41/ 50.8 /60.06	21.31/ 74.27 /74.27	28.77/ 72.4 /85.1	10.21/ 23.87 /4.89	8.33/ 16.27 /9.05
010_potted_meat_can	6.69/ 0.84 /-	53.21/ 64.98 /53.21	28.8/ 48.09 /38.76	19.73/ 59.6 /59.6	28.64/ 54.31 /78.69	7.47/ 10.4 /7.67	26.99/ 52.93 /5.43
011_banana	26.99/ 31.8 /-	49.23/ 70.87 /49.23	24.18/ 55.27 /46.25	28.16/ 61.4 /61.4	23.91/ 54.07 /90.0	25.33/ 58.27 /15.58	22.91/ 55.33 /10.21
019_pitcher_base*	29.3/ 28.53 /-	71.89/ 77.33 /71.89	4.47/ 34.76 /9.04	26.38/ 84.22 /84.22	0.0/ 0.0 /0.0	27.79/ 8.18 /5.28	25.67/ 6.44 /3.86
021_bleach_cleanser	16.59/ 17.13 /-	54.95/ 60.43 /54.95	22.5/ 58.4 /33.01	13.72/ 71.0 /71.0	2.78/ 17.27 /24.38	20.38/ 59.2 /11.14	20.94/ 63.1 /13.81
024_bowl	25.35/ 9.47 /-	66.4/ 73.07 /66.4	16.25/ 1.13 /29.93	14.96/ 63.13 /63.13	2.67/ 18.47 /0.0	0.58/ 5.0 /4.0	12.03/ 55.6 /11.33
025_mug	7.74/ 29.13 /-	60.39/ 74.13 /60.39	17.04/ 11.47 /22.32	18.81/ 71.93 /71.93	15.77/ 72.2 /0.0	19.77/ 72.53 /10.26	6.37/ 45.8 /2.7
035_power_drill*	7.57/ 0.07 /-	47.82/ 44.43 /47.82	11.34/ 7.3 /25.11	14.38/ 54.23 /54.23	0.0/ 0.0 /0.0	24.05/ 41.73 /1.92	20.71/ 54.23 /0.96
036_wood_block	15.6/ 19.87 /-	62.63/ 65.73 /62.63	32.68/ 32.93 /46.97	20.44/ 50.8 /50.8	12.61/ 28.67 /0.0	26.27/ 56.67 /8.75	28.67/ 57.07 /9.82
037_scissors	7.87/ 0.27 /-	34.22/ 20.8 /34.22	4.53/ 1.73 /35.71	12.38/ 6.4 /6.4	9.77/ 8.27 /0.0	9.97/ 6.27 /3.09	7.06/ 3.73 /1.22
040_large_marker	2.72/ 0.13 /-	30.6/ 62.0 /30.6	8.85/ 39.67 /38.86	7.12/ 63.33 /63.33	8.13/ 60.53 /33.79	6.54/ 59.8 /5.68	3.54/ 31.47 /3.1
051_large_clamp	6.62/ 0.0 /-	43.2/ 69.8 /43.2	3.06/ 12.47 /28.43	10.62/ 32.8 /32.8	0.35/ 11.47 /3.33	12.01/ 9.8 /10.44	13.17/ 36.53 /2.86
052_extra_large_clamp	8.39/ 4.2 /-	47.61/ 45.87 /47.61	0.77/ 0.0 /30.26	9.41/ 35.87 /35.87	0.07/ 0.07 /0.0	21.33/ 26.93 /6.18	12.54/ 9.07 /4.32
061_foam_brick	6.92/ 3.47 /-	52.27/ 72.93 /52.27	25.99/ 59.87 /56.78	23.88/ 76.27 /76.27	19.89/ 71.73 /0.0	22.89/ 25.47 /4.77	22.8/ 54.4 /6.92
AVERAGE YCB-V	10.47/ 8.83 /-	53.42/ 65.89 /53.42	21.04/ 40.74 /40.45	18.41/ 62.28 /62.28	14.84/ 40.58 /29.37	19.6/ 42.05 /7.27	20.93/ 46.56 /8.2
AVERAGE YCB-V18	9.72/ 7.67 /-	52.13/ 65.82 /52.13	21.46/ 41.21 /42.24	18.38/ 60.58 /60.58	17.05/ 47.08 /34.27	18.04/ 42.13 /7.57	20.0/ 46.96 /8.36

Table 2: AR (higher is better) of all evaluated data representation, corresponding to Table 1. Results on REAL in **bold font**. LUMA (Ours) compares favourably, outperforming PBG, RBG, and CHROMA.



Figure 3: Some of the problems with chroma keying. Color bleeding leads to part of the object appearing greenish, which leads to imperfect masking (top left). Luminance keying (top right) in contrast gives much improved masking. Other problems with chroma are the high reflectivity of conventional backgrounds, that lead to a "halo" effect at the edges (bottom left), and the cutting out of object parts close to the background color (bottom right).



Figure 4: Another problem of chroma keying. While the lighting did not change, a slight change in camera settings between two video clips lead to very different tones of green, making the thresholding much harder compared to luminance keying.

At the same time this relative gain in performance of REAL and PBR comes at a cost: creating a suitable big dataset to be able to train on REAL requires (usually manual) labeling, while the creation of a PBR dataset presumes the existence of 3D meshes or requires the scanning of the objects. Also, we have a throughput of roughly 50 images rendered per minute on a single A100 GPU, making the PBR dataset costly in terms of resources and time.

5 CONCLUSION

In conclusion, we propose a luminance keying method using a black screen with 99.99% light absorption

for efficient training data acquisition, significantly simplifying the process of training deep neural networks for object segmentation and detection. Our technique overcomes the limitations of manual annotation and rendering, traditionally required for creating annotated datasets, by employing a high-absorption black background to facilitate quick video recording and brightness-based automatic masking of objects. This approach not only expedites the data preparation process but also eliminates common issues associated with chroma keying, such as color bleeding and color overlap. We did extensive evaluation and find that our method compares favourably to much more involved data generation approaches on the YCB-V object set. Overall, this work enables the rapid deployment of deep learning applications across various scales, democratizing access to state-of-the-art object detection and segmentation technologies. For reproducibility and further research we publish our processing code, as well as our black screen YCB-V recordings at <https://huggingface.co/datasets/tpoellabauer/YCB-V-LUMA>.

6 REFERENCES

- [AAPS16] Yağiz Aksoy, Tunç Ozan Aydin, Marc Pollefeys, and Aljoša Smolić. Interactive high-quality green-screen keying via color unmixing. *ACM Transactions on Graphics (TOG)*, 36(4):1, 2016.
- [AYK07] Hiroki Agata, Atsushi Yamashita, and Toru Kaneko. Chroma key using a checker pattern background. *IEICE TRANSACTIONS on Information and Systems*, 90(1):242–249, 2007.
- [BIRS22] James Bessen, Stephen Michael Impink, Lydia Reichensperger, and Robert Seamans. The role of data for ai startup growth. *Research Policy*, 51(5):104513, 2022.
- [BRS⁺22] Lukas Block, Adrian Raiser, Lena Schön, Franziska Braun, and Oliver Riedel. Image-bot: generating synthetic object detection datasets for small and medium-sized manufacturing companies. *Procedia CIRP*, 107:434–439, 2022.
- [DBGD24] Jonas Dirr, Johannes C Bauer, Daniel Gebauer, and Rüdiger Daub. Cut-paste image generation for instance segmentation for robotic picking of industrial parts. *The International Journal of Advanced Manufacturing Technology*, 130(1):191–201, 2024.
- [DMH17] Debidatta Dwibedi, Ishan Misra, and Martial Hebert. Cut, paste and learn: Surprisingly easy synthesis for instance

- detection. In *Proceedings of the IEEE international conference on computer vision*, pages 1301–1310, 2017.
- [DSW⁺20] Maximilian Denninger, Martin Sundermeyer, Dominik Winkelbauer, Dmitry Olefir, Tomas Hodan, Youssef Zidan, Mohamad Elbadrawy, Markus Knauer, Harinandan Katam, and Ahsan Lodhi. Blenderproc: Reducing the reality gap with photorealistic rendering. In *International Conference on Robotics: Science and Systems, RSS 2020*, 2020.
- [Fos10] Jeff Foster. *The Green Screen Handbook*. Sybex (Wiley), 2010.
- [GKTB10] Anselm Grundhöfer, Daniel Kurz, Sebastian Thiele, and Oliver Bimber. Color invariant chroma keying and color spill neutralization for dynamic scenes and cameras. *The Visual Computer*, 26:1167–1176, 2010.
- [GLW⁺21] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. YOLO: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021.
- [HJS⁺20] Ronan Hamon, Henrik Junklewitz, Ignacio Sanchez, et al. Robustness and explainability of artificial intelligence. *Publications Office of the European Union*, 207, 2020.
- [HLWK18] Stefan Hinterstoisser, Vincent Lepetit, Paul Wohlhart, and Kurt Konolige. On pre-trained image features and synthetic images for deep learning. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018.
- [HPH⁺19] Stefan Hinterstoisser, Olivier Pauly, Hauke Heibel, Marek Martina, and Martin Bokeloh. An annotation saved is an annotation earned: Using fully synthetic training for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, pages 0–0, 2019.
- [JLZ⁺22] Yue Jin, Zhaoxin Li, Dengming Zhu, Min Shi, and Zhaoqi Wang. Automatic and real-time green screen keying. *The Visual Computer*, 38(9):3135–3147, 2022.
- [KKvB⁺22] Volker Knauth, Maurice Kraus, Max von Buelow, Tristan Wirth, Arne Rak, Laurenz Merth, Alexander Erbe, Christian Kontermann, Stefan Guthe, Arjan Kuijper, et al. Alignment and reassembly of broken specimens for creep ductility measurements. In *VMV*, pages 33–40, 2022.
- [KMR⁺23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- [LGL⁺23] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023.
- [LHB04] Yann LeCun, Fu Jie Huang, and Leon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pages II–104. IEEE, 2004.
- [LLSH23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023.
- [LTH⁺22] Weixin Liang, Girmaw Abebe Tadesse, Daniel Ho, Li Fei-Fei, Matei Zaharia, Ce Zhang, and James Zou. Advances, challenges and opportunities in creating data for trustworthy ai. *Nature Machine Intelligence*, 4(8):669–677, 2022.
- [NGP⁺23] Van Nguyen Nguyen, Thibault Groueix, Georgy Ponimatkin, Vincent Lepetit, and Tomas Hodan. Cnos: A strong baseline for cad-based novel object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2134–2140, 2023.
- [ODM⁺23] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [PJS17] Thanathorn Phoka, Warayu Jariyawattanarat, and Attawith Sudsang. Fine tuning for green screen matting. In *2017 9th*

- International Conference on Knowledge and Smart Technology (KST)*, pages 317–322. IEEE, 2017.
- [Ram15] Kathryn Ramey. *Experimental Filmmaking : Break the Machine*. Taylor Francis Ltd, 2015.
- [RHW19] Yuji Roh, Geon Heo, and Steven Euijong Whang. A survey on data collection for machine learning: a big data-ai integration perspective. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1328–1347, 2019.
- [RKH⁺21] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [SB96] Alvy Ray Smith and James F Blinn. Blue screen matting. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 259–268, 1996.
- [Sch05] Toni Schenk. Introduction to photogrammetry. *The Ohio State University, Columbus*, 106(1), 2005.
- [SHL⁺23] Martin Sundermeyer, Tomáš Hodaň, Yann Labbe, Gu Wang, Eric Brachmann, Bertram Drost, Carsten Rother, and Jiří Matas. Bop challenge 2022 on detection, segmentation and pose estimation of specific rigid objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2784–2793, 2023.
- [TFR⁺17] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017.
- [TPA⁺18] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Bochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018.
- [TTS⁺18] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. *arXiv preprint arXiv:1809.10790*, 2018.
- [VYB⁺24] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. *arXiv preprint arXiv:2403.12008*, 2024.
- [WGS⁺21] Olivia Wiles, Sven Goyal, Florian Stimberg, Sylvestre-Alvise Rebuffi, Ira Ktena, Krishnamurthy Dj Dvijotham, and Ali Taylan Cemgil. A fine-grained analysis on distribution shift. In *International Conference on Learning Representations*, 2021.
- [WRSL23] Steven Euijong Whang, Yuji Roh, Hwanjun Song, and Jae-Gil Lee. Data collection and quality challenges in deep learning: A data-centric ai perspective. *The VLDB Journal*, 32(4):791–813, 2023.
- [YAK08] Atsushi Yamashita, Hiroki Agata, and Toru Kaneko. Every color chromakey. In *2008 19th International Conference on Pattern Recognition*, pages 1–4. IEEE, 2008.
- [YCT⁺23] Qinhong Yang, Dongdong Chen, Zhentao Tan, Qiankun Liu, Qi Chu, Jianmin Bao, Lu Yuan, Gang Hua, and Nenghai Yu. Hq-50k: A large-scale, high-quality dataset for image restoration. *arXiv preprint arXiv:2306.05390*, 2023.

Appendix

Table 3: In our extensive evaluation we address the problem of domain gap / distribution shift between image sources. First we analyse the effect of freezing the backbone (first table), training from scratch (second table), and continuing training from weights pre-trained on COCO (third table). Next we train on our 7 proposed data representations. To emphasize the influence of data similarity, we next use a range of different validation sets, corresponding to different data representations and sources and determine the best checkpoint for each. Finally, we test the best checkpoint against relevant test sets. All sets and splits are disjoint. NoBGC .. Chroma data without background replacement, NoBGL .. Luma w/o bg. repl., 50 .. a set consisting of 50 percent images with and 50% images w/o bg. repl.

		BB Freezing, Pretrained on COCO																																	
		PBR rand. Tex		PBR		PBG		Real		RBG		Chroma		Luma																					
Train	Val	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma										
Test	PBR	1.0	0.6	1.0	1.0	35.8	35.8	35.7	35.7	9.9	9.8	9.2	9.8	6.8	6.4	6.3	6.4	5.8	5.7	5.7	5.7	6.2	6.1	6.2	6.1	6.1	6.1	6.1	8.1	8.1	8.1	8.1	8.1	7.6	
	PBG	0.2	0.2	0.2	0.2	4.6	4.6	4.6	4.6	10.8	11.0	10.2	10.9	2.1	2.1	2.1	2.1	6.5	6.6	6.6	6.6	4.9	4.9	4.9	5.0	5.0	5.0	5.0	5.5	5.5	5.6	5.5	5.6	5.3	
	Real	0.8	0.2	0.9	0.8	43.3	43.3	43.2	43.1	22.1	21.1	22.5	21.1	48.1	50.7	50.9	50.7	26.2	25.6	25.6	25.3	21.7	21.4	21.7	21.3	21.3	21.3	24.6	24.2	24.3	24.3	24.3	24.1	23.9	
	RBG	0.0	0.0	0.0	0.0	7.1	7.1	7.3	7.0	15.5	15.3	15.6	16.2	4.4	6.3	6.1	6.3	20.7	21.2	21.1	21.0	13.4	12.9	13.6	13.7	13.7	13.7	12.3	16.0	16.3	16.3	16.3	16.3	14.5	
	NoBGC	0.0	0.0	0.0	0.0	3.7	3.6	3.6	3.5	0.1	0.0	0.1	0.1	0.6	0.9	1.0	0.8	0.5	0.6	0.5	0.6	0.7	1.0	0.6	1.0	1.0	1.0	1.4	0.5	0.9	0.9	0.9	0.9	1.0	1.7
NoBGL	0.0	0.0	0.0	0.0	4.9	5.0	4.9	5.1	0.1	0.1	0.1	0.1	1.0	1.1	1.2	1.1	0.3	0.4	0.4	0.4	1.0	1.3	0.7	1.2	1.2	1.2	2.1	1.0	1.5	1.5	1.5	1.5	1.5	1.7	3.4
Ckpt		246	16	166	288	288	249	263	246	190	289	95	243	30	173	249	179	214	267	277	285	198	170	270	288	288	288	94	200	248	246	246	288	249	45
		Without BB Freezing, From Scratch																																	
		PBR rand. Tex		PBR		PBG		Real		RBG		Chroma		Luma																					
Train	Val	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma										
Test	PBR	1.8	0.8	0.1	1.6	73.7	72.9	68.7	72.3	10.4	2.4	7.6	6.8	1.6	1.6	1.6	1.6	0.7	0.2	0.5	0.1	3.5	3.1	2.9	3.1	1.6	1.6	3.1	6.0	5.1	5.5	5.1	2.8	4.2	5.9
	PBG	0.2	0.2	0.1	0.1	6.1	6.5	6.4	6.6	15.0	9.3	11.4	25.0	0.9	0.8	0.8	0.9	5.6	7.4	5.8	6.3	6.3	6.4	6.3	6.5	6.0	6.0	4.9	7.2	7.8	7.4	7.8	6.5	7.4	7.4
	Real	0.0	0.0	0.0	0.0	8.7	15.0	19.5	15.6	0.4	0.0	0.6	0.2	44.3	46.4	46.2	46.3	24.6	18.4	22.8	15.1	2.7	2.3	2.0	2.3	0.5	0.5	3.3	0.8	0.9	1.7	0.6	0.1	0.4	0.7
	RBG	0.0	0.0	0.0	0.0	0.9	1.0	2.2	1.2	4.8	1.3	1.2	4.7	2.7	2.6	1.8	2.1	35.1	43.3	36.5	39.0	8.5	9.6	7.6	9.2	4.2	4.2	6.9	10.0	10.6	11.0	10.8	4.0	9.9	9.6
	NoBGC	0.0	0.0	0.0	0.0	3.5	2.8	4.2	3.1	0.0	0.0	0.1	0.0	0.5	0.6	0.5	0.6	0.0	0.0	0.0	0.0	14.2	13.7	10.3	12.7	11.0	11.0	24.2	3.9	1.0	5.0	1.2	1.8	2.2	5.3
NoBGL	0.0	0.0	0.0	0.0	7.2	7.4	8.9	9.8	0.0	0.0	0.2	0.0	0.7	0.9	1.0	1.2	0.0	0.0	0.0	0.0	4.6	5.1	4.0	4.5	6.3	6.3	6.1	21.8	16.7	27.5	23.2	23.8	28.4	31.5	
Ckpt		48	6	2	101	169	80	43	63	18	289	6	39	103	161	182	165	29	165	32	271	48	63	85	64	250	250	22	54	188	75	172	288	219	60
		Without BB Freezing, Pretrained on COCO																																	
		PBR rand. Tex		PBR		PBG		Real		RBG		Chroma		Luma																					
Train	Val	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma	PBR	PBG	Real	RBG	Chroma	Luma										
Test	PBR	3.2	3.2	3.2	3.2	74.7	52.8	52.8	52.8	9.6	2.2	8.7	9.6	10.5	10.5	10.1	10.1	4.9	4.8	4.8	0.0	5.8	5.8	5.3	5.8	2.0	2.8	2.8	8.3	6.3	7.1	7.7	3.4	4.3	7.1
	PBG	0.6	0.6	0.6	0.6	6.2	6.5	6.5	6.5	19.8	9.4	16.6	19.8	3.3	3.3	3.4	3.4	8.9	9.5	9.5	7.5	6.6	6.6	5.4	6.6	6.3	6.7	6.7	7.2	5.6	6.4	6.3	6.7	5.6	
	Real	0.5	0.5	0.5	0.5	11.4	26.2	26.2	26.2	0.0	0.0	0.1	0.0	64.2	64.2	64.3	64.3	21.8	24.8	24.8	14.7	2.6	2.6	10.0	2.6	0.5	1.0	1.0	2.3	0.5	7.2	1.7	0.1	0.1	7.2
	RBG	0.0	0.0	0.0	0.0	2.8	6.7	6.7	6.7	15.2	0.6	13.0	15.2	11.3	11.3	10.6	10.6	31.9	34.1	34.1	41.4	9.6	9.6	14.9	9.6	6.4	6.9	6.9	10.2	7.5	12.6	13.0	4.3	6.1	12.6
	NoBGC	0.0	0.0	0.0	0.0	8.5	6.6	6.6	6.6	0.0	0.0	0.0	0.0	5.9	5.9	5.9	5.9	0.3	0.0	0.0	0.0	11.9	11.9	11.5	11.9	11.6	19.0	19.0	4.2	0.5	12.4	2.5	1.1	1.4	12.4
NoBGL	0.0	0.0	0.0	0.0	11.0	8.5	8.5	8.5	0.1	0.0	0.0	0.1	3.2	3.2	3.3	3.3	0.1	0.1	0.1	0.0	4.1	4.1	4.7	4.1	2.8	6.3	6.3	20.0	6.3	24.9	16.8	12.5	14.2	24.9	
Ckpt		1	1	1	1	140	3	3	3	13	279	7	13	9	9	6	6	5	7	7	261	8	8	2	8	246	166	166	11	158	2	20	288	249	2