

A Saliency-Guided Street View Image Inpainting Framework for Efficient Last-Meters Wayfinding

Chuanbo Hu^a, Shan Jia^{b,*}, Fan Zhang^c, Xin Li^a

^a*Lane Department of Computer Science and Electrical Engineering, West Virginia University, 395 Evansdale Dr, Morgantown, 26505, West Virginia, United States*

^b*University at Buffalo, State University of New York, 12 Capen Hall, Buffalo, 14260, New York, United States*

^c*Department of Civil and Environmental Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China*

Abstract

Global Positioning Systems (GPS) have played a crucial role in various navigation applications. Nevertheless, localizing the perfect destination within the last few meters remains an important but unresolved problem. Limited by the GPS positioning accuracy, navigation systems always show users a vicinity of a destination, but not its exact location. Street view images (SVI) in maps as an immersive media technology have served as an aid to provide the physical environment for human last-meters wayfinding. However, due to the large diversity of geographic context and acquisition conditions, the captured SVI always contains various distracting objects (e.g., pedestrians and vehicles), which will distract human visual attention from efficiently finding the destination in the last few meters. To address this problem, we highlight the importance of reducing visual distraction in image-based wayfinding by proposing a saliency-guided image inpainting framework. It aims at redirecting human visual attention from distracting objects to destination-related objects for more efficient and accurate wayfinding in the last meters. Specifically, a context-aware distracting object detection method driven by deep salient object detection has been designed to extract distracting objects from three semantic levels in SVI. Then we employ a large-mask inpainting method with fast Fourier convolutions to remove the detected distracting objects. Experimental results with both qualitative and quantitative analysis show that our saliency-guided inpainting method¹ can not only achieve great perceptual quality in street view images but also *redirect* the human's visual

^{*}Corresponding author (shanjia@buffalo.edu)

¹Codes are available at https://github.com/cbhu523/saliency_last_way_finding/.

attention to focus more on static location-related objects than distracting ones. The human-based evaluation also justified the effectiveness of our method in improving the efficiency of locating the target destination.

Keywords: street-level image, saliency-guided image inpainting, image-based wayfinding, image segmentation

1. Introduction

The worldwide coverage of the Global Positioning System (GPS) has made it a standard means for navigation in modern life. Ranging from map applications (e.g., Google Maps) to navigation in mobile platforms, GPS has provided accurate and reliable localization services for billions of users. However, limited by the horizontal accuracy of approximately 5 meters in smartphone-based GPS [1], even around 10 meters in an urban environment [2], identifying the exact location after being navigated to a destination is still challenging for users, especially in highly urbanized areas (e.g., urban canyons, as shown in Figure 1 (a) and (d)). This challenge is called *the last-meters wayfinding problem* [1, 3, 4].

Existing studies on last-meters wayfinding focus primarily on improving accuracy for people with visual impairments [1, 4]. However, for average users, relying on human vision to identify the destination is also not an easy task, such as in strange surroundings or highly diverse and dynamic environments [3]. To address this problem, users tend to use Street View Images (SVI) in Google Maps to virtually explore the precise destination in the last few meters before starting [5, 6]. Although SVI, as an interactive panorama from the street, has allowed users to virtually visit remote areas and intuitively query the target destination (such as stores, hotels, gas stations, and parks), there is always confusion in identifying the exact destination in the last few meters. The high diversity of geographic context and acquisition conditions of the SVI will inevitably cause visual distraction due to the fact that various large-scale objects in the SVI, such as pedestrians, vehicles, and road signs, will distract human attention from finding the destination (see the example in Figure 1 (b) and (e)).

Unlike existing studies that propose to improve last-meters route finding based on additional signals [3, 7] or fine-grained map generation [8], we highlight the importance of relying on human vision for last-meters route finding, while proposing to improve location efficiency by reducing visual distraction in street view images. Since there is little research on improving destination finding in SVI using visual attention from humans, how to facilitate the task of SVI-based wayfinding by suppressing those distracting objects remained unexplored.



Figure 1: The proposed saliency-guided street view image inpainting method can successfully remove distracting objects to redirect human visual attention to static landmarks. Note the shift of salient regions highlighted by the warm color before and after inpainting.

Humans have a tremendous ability to quickly adjust their gaze when observing a scene to obtain selected visual information [9]. This ability enables them to deploy limited processing resources to the most relevant visual information and to understand complex real-world scenes efficiently and accurately. Understanding and simulating this redirecting mechanism has scientific and practical values. One of the most important goals of vision research is to reveal the mechanisms of visual attention and the fixation behavior of human visual systems. Eye tracking, as the technological process of recording gaze movements and viewing patterns, can help researchers reveal the mechanisms of visual attention. With an eye tracking camera, eye movements such as fixations, saccades, and smooth pursuit can be captured [10]. The lengths/durations of these movements vary with the subconscious state of the brain. However, saliency modeling for SVI has remained an underresearched territory in both vision and geospatial communities.

The mechanism of human visual attention has been explored based on simulations and experiments in several large-scale data sets based on eye tracking (e.g., MIT300, MIT1003 [11], SALICON [12] and DHF1K [13]). For example, in [14], the

authors demonstrated the importance of object and semantic information in predicting human gaze, especially face and text, which have been shown to outweigh other attributes (e.g., motion, sound, touch, watchability, and operability). However, moving objects in SVIs that attract visual attention often distract noise from the task of finding your way. It is desirable to suppress distraction from moving objects and better support our navigation by redirecting visual attention.

Therefore, in this paper, we make the first attempt to improve human last-meters wayfinding by reducing visual distractions in image-based navigation. We present a hierarchical strategy for context-aware salient object detection from SVIs and design a saliency-guided inpainting method to redirect visual attention. Specifically, inspired by saliency models developed for visual perception [15, 16], we propose to construct a context-sensitive saliency detection model from three semantic levels: walking pedestrians, moving vehicles, and man-made objects. Such hierarchical decomposition of the saliency map allows us to locate and remove distracting objects at the semantic level instead of the object level based on their attraction to human visual attention. Taking into account the interaction between saliency detection (mask generation) and image inpainting (object removal), we can redirect human visual attention to the task of wayfinding.

1.1. Motivations

Given that various methods have been proposed to improve navigation performance, why is our method necessary? The differences of our study from existing methods which primarily focus on the navigation performance of a camera (e.g., in mobile devices or driving systems) are twofold: i) we aim to improve human vision-based navigation; ii) we only focus on the wayfinding in the last meters that current GPS systems can hardly cover. We justify the motivations by providing two possible applications, as follows.

- *Navigation based on SVI on maps.* A user who has never been to a destination before often relies on the SVI on maps to explore the exact location before departure. However, street view images may contain various large-scale distracting objects as a result of the high diversity of acquisition conditions. Thus, it is always difficult to find the exact destination, especially in images captured in busy commercial districts or road sections (see McDonald’s example in Figure 1 (a) or REGAL in Figure 1 (d)). In this application setup, our work can serve as a post-processing tool for the SVIs before their release in maps. We can extract and remove distracting visual objects based on our saliency-guided image inpainting framework, which will not only redirect human attention to

destination-related objects, but can also provide privacy protection for attractive but sensitive objects (as shown in Figure 1 (c) and (f)).

- *Augmented Reality (AR) navigation*, which provides directions to users on the screen overlaid on top of real environments seen through the camera of a device (e.g., a smartphone or headset). Based on GPS data, such scenarios also occur where a user is brought to the vicinity of a destination (especially in outdoor scenarios) but cannot find the exact location due to many distracting objects in the captured images. In this case, our work can serve as a post-processing tool in AR navigation systems to reduce visual distraction only in the last few meters and redirect human attention to the destination finding.

1.2. Contributions

The technical contributions² of this paper are summarized below.

- *Context-aware salient object detection*. Unlike existing methods using object-level regions for editing, we propose to construct a hierarchical decomposition of the saliency map, which allows us to detect and extract distracting objects at semantic levels in street view images.
- *Saliency-guided image inpainting to redirect visual attention*. Unlike the previous formulation of the image-inpainting model, we focus on the interaction between saliency detection and image inpainting. The context-aware saliency map facilitates the task of inpainting; in the meantime, the image inpainting achieves the objective of redirecting visual attention to facilitate the task of human wayfinding.
- *Evaluation of visual attention changes*. We design two metrics to quantitatively measure changes in human attention in distracting objects and target regions. For the first time, we experimentally verified the intuition of reducing visual distraction to accelerate the SVI-based wayfinding task. A platform is also designed to compare human search time before and after redirecting visual attention to SVIs.
- *Improved efficiency for last-meters wayfinding with good perceptual quality*. Unlike previous methods requiring well-defined or user-provided editing masks

²Codes of the designed framework are available at https://github.com/cbhu523/saliency_last_way_finding/.

as input, the proposed framework takes a single street view image as input and automatically outputs the inpainted image by removing distracting objects. Extensive evaluation demonstrates the benefit of saliency-guided image inpainting to accelerate the task of last-meters wayfinding in street view images and indoor scenarios.

The remainder of the paper is organized as follows. Section 2 provides an overview of related studies. Section 3 describes the proposed saliency-guided street view image inpainting approach. In Section 4, we present the experimental evaluation and results, including both qualitative and quantitative comparisons. We discuss research limitations and potential applications in Section 5, and draw some conclusions about future research directions in Section 6.

2. Related Work

This study aims to apply a saliency-driven image inpainting technique to reduce human visual distraction in street view images for last-meters wayfinding. In this section, we cover three related areas, including last-meters wayfinding, image-inpainting techniques, and saliency detection methods for human visual attention prediction.

Last-meters wayfinding. The last-meters wayfinding problem refers to the problem of finding the precise geographic position of a destination after being guided to an approximate destination using map data [1]. These last meters of navigation cannot be covered by current GPS systems for two main reasons: lack of granular map data and loss of absolute positioning. To bridge the gap between the user’s exact destination and the end position provided by current navigation services, several studies have been proposed in recent years. A category of existing methods explored signal/sensor-assisted approaches for highly accurate location, including WiFi signal strength [17, 18, 19] for indoor localization, audio information [20] for tourist wayfinding, and geomagnetic field [3] for indoor and outdoor navigation. Another category of solutions proposes the use of existing techniques to improve last-meters wayfinding, such as crowd-sourcing techniques [21, 8] and AI methods [1] to capture the relationship between landmarks and the person’s behavior information for people with visual impairments.

Despite these advances in improving the efficiency and accuracy of navigation systems for human last-meters wayfinding, the need for additional and precise signal extraction or a sustainable incentive crowd-sourcing mechanism limits its applicability. To provide a simpler, more applicable, and effective solution, we emphasize the importance of improving human visual attention for image-based last-meters wayfinding considering that street view images have been widely used in navigation

systems. Due to the process of gathering SVI in the public, it inevitably captures diverse information, including static objects (e.g., landmarks and roads) and dynamic objects (e.g., vehicles, pedestrians, and traffic signs). However, dynamic objects may not only cause occlusion both visually and pragmatically, but also collide with privacy issues with identity-related information in many SVI applications, especially the most widely used image-based localization. Therefore, several studies have proposed to clean SVIs by blurring sensitive information [22], replacing completely anonymous pedestrians [23], replacing sensitive objects with synthetic content generated by generative adversarial networks (GAN) [24], or removing dynamic objects using image-inpainting techniques [25]. Among them, advanced image-inpainting methods have shown superiority in removing dynamic objects with high visual quality.

Image inpainting. Image inpainting refers to a class of techniques that can remove unwanted objects from images. The removal of pedestrians from the SVI was first studied [26] from a privacy point of view, where pedestrian detection and pixel cloning methods were used. Rapid advances in deep learning in recent years have provided new weaponry for powerful inpainting techniques, e.g., contextual attention [27], partial convolutions [28], gated convolution [29], and recurrent feature reasoning [30]. Most recently, the knowledge distillation strategy was adapted to the problem of image inpainting [31]; a novel dynamic selection network (DSNet) was proposed to distinguish the corrupted region from the valid ones throughout the entire network architecture. In [32], robust Large Mask (LaMa) inpainting with Fourier convolutions was developed with resolution, which will serve as the baseline for our research. Note that almost all existing image-inpainting methods assume that the inpainting domain (mask) is random or known as a priori.

Saliency detection. Saliency detection [15] is a problem that has been extensively studied in the literature on visual perception and computer vision. Various saliency models, including context-aware [33], spectral residual [34], and hierarchical [35], have been proposed in the literature. Among them, hierarchical saliency detection [35] shares a motivation similar to ours. However, a significant departure from the previous work [35] is the definition of hierarchy - in the context of last-meters wayfinding, we will consider a semantically meaningful definition of hierarchy in this work. More recently, deep learning for saliency detection has also been extensively studied - from the unification of global context with local context [36] to the Holistically Nested Edge Detector (HED) [34]. For the current state-of-the-art in deep-saline models, refer to recent reviews such as [37] and [38].

Furthermore, there have been some recent studies proposed to combine the saliency with image editing techniques [39, 40, 41, 42, 43]. However, these methods cannot be applied to the street-view image based wayfinding for the following three reasons.

1) The previous methods use recoloring to emphasize the target objects or attenuate the distracted regions [39, 40, 41, 43]. They not only require accurate masks, but also cannot handle the privacy disclosure in street view images with pedestrians or vehicles. 2) They all require the input of a well-defined or user-provided mask as the manipulation region [39, 40, 41, 42, 43] and some can only handle a single or finite-size region [42, 41]. 3) In terms of the evaluation of attention retargeting, these methods compare the similarity of ground-truth masks/saliency and saliency maps of manipulated images [39, 40, 41, 42], which cannot fully represent changes in human attention in both distracting regions and target regions.

New insight. Street-view image applications face two obstacles in the real world: *occlusion* and *privacy*. The former is related to the application of last-meters wayfinding in that the region of interest might be occluded by distracting objects, which makes it more challenging to find the match. The latter generally refers to the protection of sensitive information (e.g., biometrics and vehicle license plates) in the public domain. Saliency-guided image inpainting seems to be a promising technical solution to solve both problems. However, to the best of our knowledge, saliency detection and image inpainting have been studied separately in the open literature. Their rich interaction can be justified by 1) the detected salient objects serving as the inpainting domain; and 2) removing distracting objects by inpainting could redirect visual attention to landmarks such as static buildings in street view images. How to jointly exploit the interaction between saliency detection and image inpainting remains an unexplored research topic in street view images.

3. Methodology

In this study, we propose a saliency-guided street-view image-inpainting framework to remove distracting objects for efficient last-meters wayfinding. We first formulate the problems in terms of last-meters wayfinding and SVI inpainting in this section. Then we introduce the details of the proposed framework with the following three components: 1) hierarchical salient object detection; 2) image inpainting based on Fourier convolution for distracting object removal; 3) measurement of human visual attention for last-meters wayfinding. The overview of the proposed saliency-guided inpainting framework is illustrated in Figure 2.

3.1. Problem Formulation

Given a street view image of an approximate destination guided by map data, the goal of human last meters wayfinding in our study is to reduce visual distraction in the image and guide human vision to efficiently locate the precise geographic

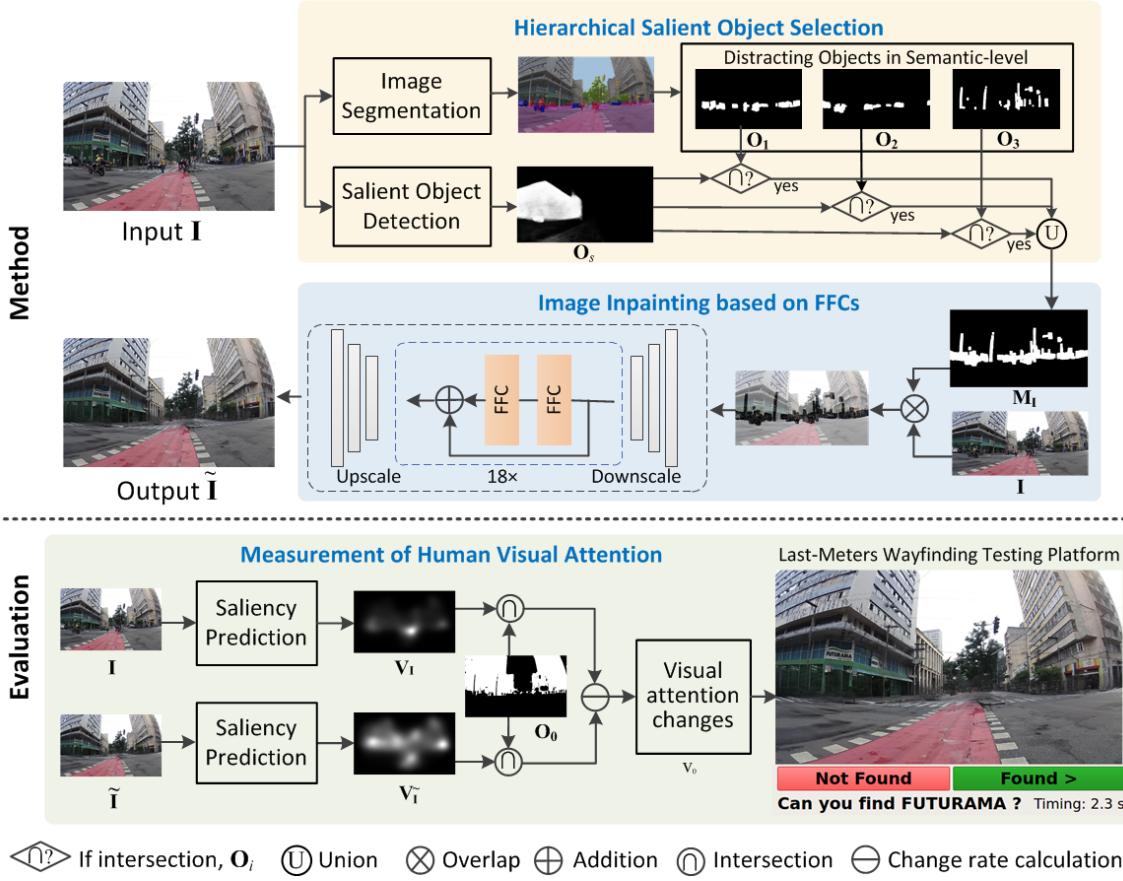


Figure 2: Overview of the proposed saliency-guided street view image inpainting framework. It consists of three building blocks: hierarchical salient object selection, saliency-guided image inpainting based on fast Fourier convolutions (FFCs), and measurement of human visual attention by visual attention changes and a self-developed last-meters wayfinding testing platform. Note that modeling the interaction between saliency detection and image inpainting leads to effective removal of distracting objects for last-meters wayfinding.

location. Let the street view image be \mathbf{I} , which could contain various landmarks and distracting objects with different levels of importance for last-meters wayfinding. Semantic image segmentation techniques [44] can be employed to partition the image into multiple segments (sets of pixels, also known as image objects), and denoted $\mathbf{O}_i \in \mathbf{I}$ ($1 \leq i \leq N$, N is the number of object categories). Based on the importance level of the objects in street view images for last-meters wayfinding, we give the following definitions to classify the objects into three categories.

Definition 1 (Object of interest). *An object of interest for last-meters wayfind-*

ing is the object that is location-related. It generally contains static buildings, such as landmarks, shopping malls, stores, and supermarkets.

Definition 2 (Distracting object). *A distracting object for last-meters wayfinding is the object that distracts a human’s visual attention from searching and identifying the location-related landmarks in the street view image. It usually contains moving vehicles, pedestrians, and dynamic traffic signs.*

Definition 3 (Neutral object). *A neutral object for last-meters wayfinding is the object that is neither location-related nor eye-catching in the street view image. It usually contains a sky and vegetation.*

Given the above definitions, we formally define the problem of street view image inpainting in last-meters wayfinding. Let \mathbf{O}_o , \mathbf{O}_d be objects of interest and distracting objects, respectively, and the goal of street view image inpainting is to inpaint the distracting objects \mathbf{O}_d in the images and redirect visual attention to objects of interest \mathbf{O}_o , shown below.

$$\mathcal{M}(\mathbf{I}) = \mathbf{O}_d \xrightarrow{\tilde{\mathbf{I}} = \mathcal{I}(\mathbf{I})} \mathcal{M}(\tilde{\mathbf{I}}) = \mathbf{O}_o \quad (1)$$

where $\mathcal{M}(\cdot)$ is the visual attention prediction function (gaze), $\mathcal{I}(\cdot)$ is the image inpainting operation, and $\tilde{\mathbf{I}}$ is the image inpainted. To achieve this goal, we will first propose an approach to detect the existence of distracting objects \mathbf{O}_d in street view images. Note that these distracting objects \mathbf{O}_d are extracted from different semantic levels based on the image content. Then we will design an image inpainting method $\mathcal{I}(\cdot)$ to remove distracting objects for the reduction of visual distraction and finally introduce the saliency prediction $\mathcal{M}(\cdot)$ to predict and measure human visual attention. To facilitate our presentation, Table 1 summarizes the mathematical symbols used in this section.

3.2. Hierarchical Salient Object Detection

Unlike existing image inpainting methods [27, 28, 29, 30, 32, 45] which are mainly based on randomly or manually selected inpainting masks, we develop an automatic inpainting domain extraction method by detecting context-sensitive salient objects in street view images.

Several previous studies [46, 47, 14, 48] have observed that objects that convey more information, such as humans, faces, cars, text, and animals, tend to be more prominent and guide the human gaze in a scene. However, in street-view images, these salient objects carry no/less location-related information but will distract visual attention from last-meters wayfinding. To remove these distracting objects, we

Table 1: List of symbols used to describe the proposed method and what they mean.

Notation	Meaning
\mathbf{I}	Input street view image
$\tilde{\mathbf{I}}$	Inpainted street view image
\mathbf{O}_o	Object region of interest
\mathbf{O}_d	Distracting object region
\mathbf{O}_0	Object region in Semantic level 0 (Building)
\mathbf{O}_1	Object region in Semantic level 1 (Human)
\mathbf{O}_2	Object region in Semantic level 2 (Vehicle)
\mathbf{O}_3	Object region in Semantic level 3 (Sign)
\mathbf{O}_s	Salient object region
$\mathbf{M}_{\mathbf{I}}$	Inpainting mask
$\mathbf{V}_{\mathbf{I}}$	Visual attention map of \mathbf{I}
$\mathbf{V}_{\tilde{\mathbf{I}}}$	Visual attention map of $\tilde{\mathbf{I}}$
$\mathbf{X}, \mathbf{X}^l, \mathbf{X}^g$	Input of FFC module
$\mathbf{Y}, \mathbf{Y}^l, \mathbf{Y}^g$	Output of FFC module
f	Convolution operations of FFC
γ	Threshold to binarize the object saliency map
λ	Loss function weights
v_o	Visual attention changes on the object of interest
v_d	Visual attention changes on distracting object
\mathcal{L}	Loss function
$\mathcal{I}(\cdot)$	Image inpainting function
$\mathcal{M}(\cdot)$	Visual attention prediction function

first classify them into three levels and then design a hierarchical distracting object extraction method based on semantic segmentation and salient object detection. Specifically, this process consists of the following steps:

1) **Image segmentation.** We employ the state-of-the-art semantic image segmentation method, DeepLabv3+ [49] to assign semantic labels to every pixel in a street view image. Due to the encoder-decoder structure with atrous separable convolution, DeepLabv3+ achieves promising segmentation performance, especially on the Cityscape image dataset [50] with diverse street scenes. Given a street view image as input, the DeepLabv3+ method will capture and output the object boundaries with different semantic labels. The rich semantic information in street view images has been finely classified and labeled into 29 classes in the Cityscapes dataset, as



Figure 3: Semantic segmentation results and labels on the Cityscapes dataset using DeepLabv3+. Best view in color.

shown in Figure 3. Therefore, the input street view image will be partitioned into several segments as \mathbf{O}_i at the semantic level.

2) Salient object detection. Image segmentation alone cannot determine distracting objects in street view images. To detect visually attractive objects in images, we then use a deep network with a nested U-structure for the detection of salient objects. Inspired by the powerful multiscale feature extraction of residual blocks, we take a street view image \mathbf{I} as input to the U^2Net model [16] with 11 stages of residual U blocks and obtain the saliency map with the same spatial resolution as \mathbf{I} . Each pixel on the map is within the range of $[0, 255]$, and the pixels of the salient object \mathbf{O}_s will have higher values (close to 255). From the first two rows of Figure 4, we can observe that informative objects in street view images, such as vehicles, pedestrians, traffic signs, and landmarks, tend to be more attractive as the predicted salient object, which is consistent with previous studies.

3) Hierarchical salient object selection. If we directly take the salient object map after removing the objects of interest as the inpainting mask, one obvious problem is that other objects with the same or similar semantic labels in the image may still catch visual attention. In other words, the salient object map generated by U^2Net is at the instance level, rather than at the semantic level. To address this problem, we designed the hierarchical salient object selection method to extend

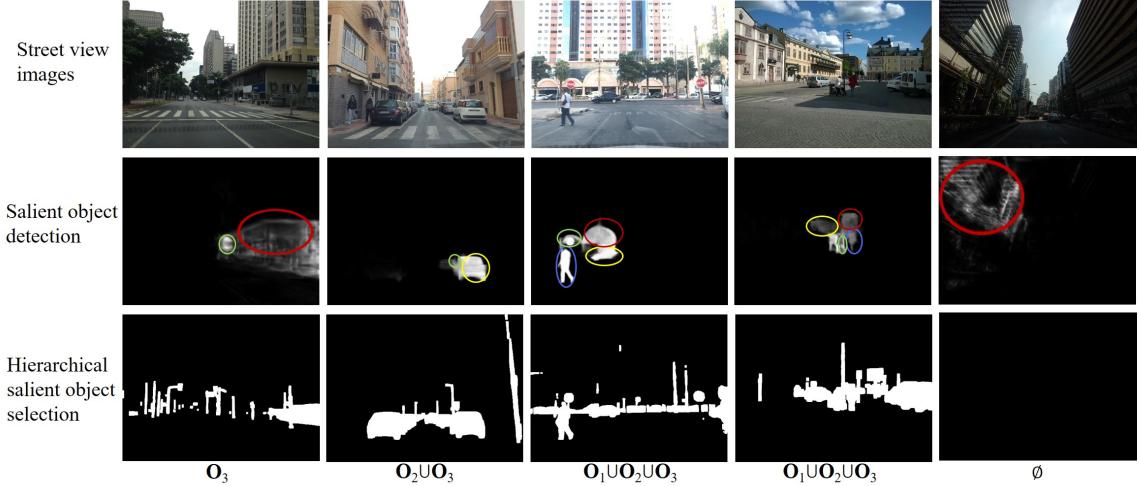


Figure 4: Examples of salient object detection and selection in street view images. Red ovals - salient objects in \mathbf{O}_0 semantic level (Building), blue - salient objects in \mathbf{O}_1 (Human), yellow - salient objects in \mathbf{O}_2 (Vehicle), and green - salient objects in \mathbf{O}_3 level (Sign). Best view in color.

Table 2: Four semantic levels for the selection of hierarchical salient objects.

Category	Segment	Semantic level (Fine class)
Object of interest	\mathbf{O}_0	Building
	\mathbf{O}_1	Human (Person, Rider)
Distracting object	\mathbf{O}_2	Vehicle (Car, Truck, Bus, Train, Motorcycle, Bicycle)
	\mathbf{O}_3	Sign (Traffic light, Traffic sign, Fence, Pole)

the salient map from the instance level to the semantic level. In street view images, we mainly focus on the object of interest and distracting objects for last-meters wayfinding. We first classify them into four categories based on the correlation in the semantic labels (see Table 2). Next, we use an empirical threshold γ to binarize the saliency map, as in previous work [51, 52, 53]. Consequently, the salient object region \mathbf{O}_s can be modified to the semantic level as follows.

$$\begin{aligned}
 \mathbf{O}_s = & \mathbf{O}_1 \text{ (if } \mathbf{O}_s \cap \mathbf{O}_1 \neq \emptyset \text{)} \\
 & \cup \mathbf{O}_2 \text{ (if } \mathbf{O}_s \cap \mathbf{O}_2 \neq \emptyset \text{)} \\
 & \cup \mathbf{O}_3 \text{ (if } \mathbf{O}_s \cap \mathbf{O}_3 \neq \emptyset \text{)} \\
 \end{aligned} \tag{2}$$

Then we can obtain the final binary inpainting mask \mathbf{M}_I with the pixels in \mathbf{O}_s as 1 and the others as 0. Figure 4 compares the saliency maps before and after the selection of hierarchical salient objects. We can see that more distracting objects can be covered with hierarchical salient object selection as the inpainting mask.

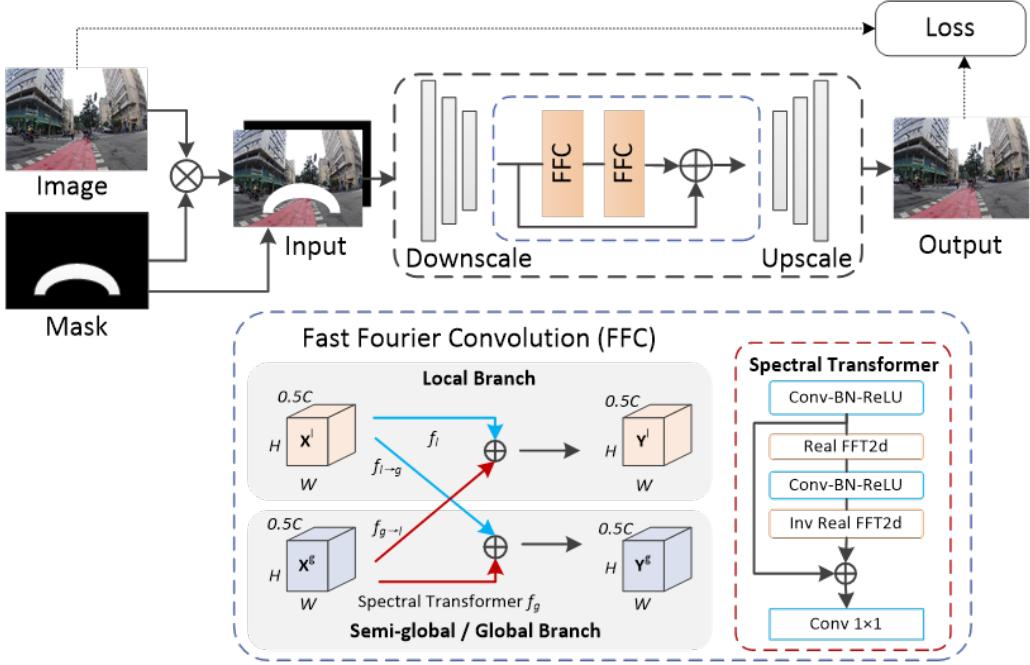


Figure 5: Overview of the fast Fourier convolutions-based image inpainting model.

3.3. Image Inpainting based on Fast Fourier Convolutions

Removing distracting objects at the semantic level means a larger inpainting domain than removing objects at the instance level. Fast Fourier convolutions (FFCs) [54] allow a large receptive field to cover an entire image even in the early layers of the network, and therefore can provide a high perceptual quality and parameter efficiency of the network. Inspired by the recent success of FFCs in the large mask inpainting model LaMa [32], we build a similar inpainting network for street view images based on the hierarchical salient distracting mask generated.

As shown in Figure 5, taking the original image \mathbf{I} stacked with the binary inpainting mask \mathbf{M}_I as input, the inpainting network first downscale the input with three downsampling blocks and then uses 18 residual FFC blocks as generator, followed by three upsampling blocks. Specifically, the FFCs split the input feature map into a local branch with conventional convolutions and a global branch with a real fast Fourier transform (FFT). Let $\mathbf{X} = \{\mathbf{X}^l, \mathbf{X}^g\}$ be the input to the FFC module ($\mathbf{X} \in \mathbf{R}^{H \times W \times C}$, \mathbf{X}^l be the local part and \mathbf{X}^g the global part), and the output of FFC $\mathbf{Y} = \{\mathbf{Y}^l, \mathbf{Y}^g\}$ can be described by the following formulas:

$$\mathbf{Y}^l = \mathbf{Y}^{l \rightarrow l} + \mathbf{Y}^{g \rightarrow l} = f_l(\mathbf{X}^l) + f_{g \rightarrow l}(\mathbf{X}^g) \quad (3)$$

$$\mathbf{Y}^g = \mathbf{Y}^{g \rightarrow g} + \mathbf{Y}^{l \rightarrow g} = f_g(\mathbf{X}^g) + f_{l \rightarrow g}(\mathbf{X}^l) \quad (4)$$

where the component $\mathbf{Y}^{l \rightarrow l}$ is obtained by applying a regular convolution f_l to \mathbf{X}^l , $\mathbf{Y}^{g \rightarrow g}$ is obtained by applying a spectral transformer f_g to \mathbf{X}^g , and the other two components $\mathbf{Y}^{g \rightarrow l}$ and $\mathbf{Y}^{l \rightarrow g}$ are obtained via an inter-path transition based on regular convolutions to take full advantage of multiscale receptive fields. Specifically, to capture the global/semi-global receptive field in the image, the spectral transform f_g first transforms the original spatial features into the spectral domain, then performs an efficient global update on the spectral data, and finally converts the data back to the spatial format. Real FFT and inverse real FFT are applied to ensure that both input and output are real-valued.

Training in the large-scale Places [55] dataset with a high diversity of scene categories and high-quality photographs, the original LaMa model achieves excellent performance in real-world image editing scenarios. Since street view images for wayfinding are more likely to contain richer information (with complex texture and content) and more diverse image quality (under different light conditions), we follow the pipeline in Figure 5 to fine-tune the whole inpainting network in street view images. The loss function to train the model is calculated as follows:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{adv} + \lambda_2 \mathcal{L}_{pl1} + \lambda_3 \mathcal{L}_{pl2} + \lambda_4 R_1 \quad (5)$$

\mathcal{L}_{adv} is the adversarial loss to ensure that the inpainting model generates natural-looking local details based on a discriminator network. Perceptual loss \mathcal{L}_{pl1} is for the high-receptive field, obtained by sequential two-stage mean operations (interlayer mean or intralayer mean) of Fourier convoluted images \mathbf{I} and $\tilde{\mathbf{I}}$. Perceptual loss \mathcal{L}_{pl2} is the loss of feature matching to stabilize training [56], and the regularization term R_1 is the discriminator gradient penalty [32] to penalize the discriminator for deviating from Nash equilibrium. $\lambda_1, \lambda_2, \lambda_3$, and λ_4 are hyperparameters.

3.4. Measurement of Human Visual Attention for Last-few-meters Wayfinding

The salient object detection method, such as the U^2Net model introduced in Section 3.2, is object-level, which cannot accurately represent human visual attention. To predict and measure how the human gaze will change in street view images, we propose to model pixel-level saliency based on the UNISAL network [57]. UNISAL integrates novel domain adaptation modules into an encoder-RNN-decoder-style network. These domain-adaptive modules are designed to shift domains between different saliency datasets and to enable high-quality shared features.

The entire network was trained in the SALICON image dataset [12] and three video datasets with eye-tracking or fixation annotations. Note that the SALICON

dataset collected 20,000 natural images with mouse tracking labeling through the Amazon Mechanical Turk (AMT) platform to obtain the ground truth of human gaze positions. Joint training in different modalities results in improved performance of the UNISAL model in visual saliency modeling. To measure human visual attention influenced by image inpainting, we first input the street view images \mathbf{I} and the inpainted image $\tilde{\mathbf{I}}$ into the pre-trained UNISAL model, and obtain the predicted visual attention map $\mathbf{V}_\mathbf{I}$ and $\mathbf{V}_{\tilde{\mathbf{I}}}$. Next, we propose to quantitatively represent the increase in visual attention in the object of interest v_o and the reduction of visual attention in distracting objects v_d by

$$v_o = \sum_{i \in \mathbf{O}_o} (\mathbf{V}_{\tilde{\mathbf{I}}}(i) - \mathbf{V}_\mathbf{I}(i)) / \sum_{i \in \mathbf{O}_o} (\mathbf{V}_\mathbf{I}(i)) \quad (6)$$

$$v_d = \sum_{i \in \mathbf{O}_d} (\mathbf{V}_\mathbf{I}(i) - \mathbf{V}_{\tilde{\mathbf{I}}}(i)) / \sum_{i \in \mathbf{O}_d} (\mathbf{V}_\mathbf{I}(i)) \quad (7)$$

4. Experiments

To evaluate our method, we performed a series of experiments on street view images with quantitative and qualitative analyses. This section first introduces the implementation details and evaluation metrics and then compares the image-inpainting performance with existing methods. Next, we demonstrate the influence of mask selection and city morphology type on inpainting results and changes in visual attention. To show the effectiveness in improving the last-meters wayfinding efficiency, we also conduct human-based evaluation. Finally, we show the applicability of the proposed framework by extending it to indoor scenario applications.

4.1. Experimental Settings

Dataset and implementation details. We used the large-scale Cityscape dataset [50] with various street scenes in our experiment. Cityscapes dataset contains a diverse set of stereo video sequences recorded from 50 different cities. For image inpainting, we fine-tuned the big LaMa Fourier model with the training data (18,000 images) in Cityscapes following the settings in the paper [32], and report all results on the testing subset (5,000 images). Loss weight λ_1 is set as 10, λ_2 as 30, λ_3 as 100, and λ_4 as 0.001. The threshold γ to binarize the object saliency map is set to 0.8, meaning that the pixels larger than the $0.8 \times$ maximum will be set as 1 and others as 0. For salient object detection, we used the pretrained model of U^2Net model in [16]. All experiments were carried out on NVIDIA TITAN RTX GPUs.

Evaluation metrics. We follow the metrics in the recent image inpainting literature [58, 30, 59, 32, 60] and report the quantitative results of the mean l_1 error,

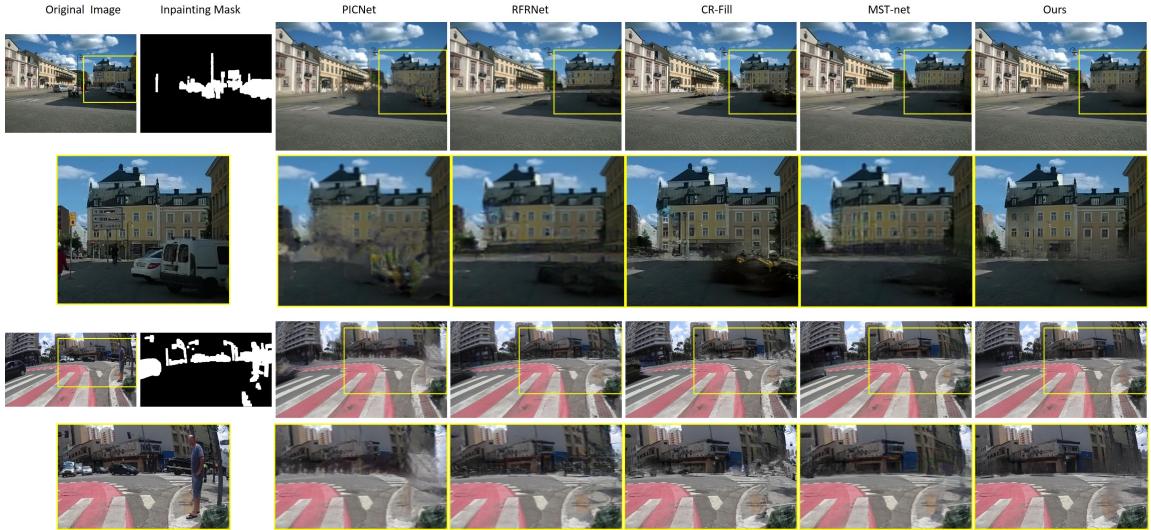


Figure 6: Visual performance of different image inpainting methods on Cityscapes dataset. Note that all images are adjusted to the same size for a clear comparison. Best view in color and zoom in.

the structural similarity index (SSIM), the peak signal-to-noise ratio (PSNR), the learned perceptual image patch similarity (LPIPS) [61] and the Frechet Inception Distance (FID) [62] based on InceptionNet [63] features. Since the calculation of l_1 , SSIM, and PSNR metrics requires the comparison of the inpainted image with ground-truth data, we inserted random objects (including vehicles, pedestrians and traffic signs) into the image in the Cityscapes dataset [64]. The processed images were input to different inpainting models with the inserted object region as the inpainting mask. Correspondingly, the l_1 , SSIM, and PSNR metrics were calculated by comparing the original Cityscapes image (as the ground truth) with the inpainted image. To evaluate changes in visual attention, we compare different methods in terms of the average v_o and v_d (as defined in Section 3.4).

Comparison methods. Five state-of-the-art image-painting methods with open source codes were compared with the proposed scheme, including PICNet [58], RFR-Net [30], CR-Fill [60], MST-net [45], and the original LaMa model [32]. PICNet, CR-Fill, MST-net and the LaMa model were trained on the large-scale Places2 [55] dataset that provides 8 million training images and contains highly diverse urban and street scenes. The RFR-Net was trained on the Paris StreetView [65] dataset (with 15,000 images). Note that PICNet, RFR-Net, and MST-net model can only handle squared images and require an input image with a fixed size of 256×256 . We used the pre-trained models of these methods in our experiment.

Table 3: Quantitative evaluation of inpainting performance on Cityscapes dataset

Method	SSIM \uparrow	PSNR \uparrow	$l_1 \downarrow$	LPIPS \downarrow	FID \downarrow
PICNet [58]	0.8847	25.77	0.0189	0.8213	5.08
RFR-Net [30]	0.9089	25.00	0.0158	0.7108	5.21
CR-Fill [60]	0.9134	25.65	0.0171	0.7236	5.08
MST-net [45]	0.8416	23.79	0.0309	1.1859	5.29
LaMa[32]	0.9118	25.88	0.0171	0.7261	5.07
Ours	0.9175	26.38	0.0154	0.7296	5.04

4.2. Comparison to State-of-the-art Inpainting Methods

Quantitative comparison. For a fair comparison, we input the same inpainting masks based on the proposed hierarchical salient object detection to all methods. Table 3 shows the quantitative results of inpainting performance in five metrics. Benefiting from the FCCs for a high-receptive field in the inpainting network, the LaMa model outperforms other baselines when considering all metrics without the fine-tuning on the Cityscapes dataset. Our framework fine-tuned the model on the Cityscapes dataset (with 18,000 images), leading to the best performance in most metrics.

Qualitative comparison. Figure 6 compares the visual quality of different image-inpainting methods in street view images. We can see that the generated domain of our method is semantically more coherent with the surrounding regions, with fewer visual artifacts. Moreover, unlike RFR-Net, which has high performance in quantitative evaluation but tends to have poor inpainting performance for distracting objects in edge areas, the results of our method are more visually natural.

4.3. Comparison of Different Inpainting Masks

Quantitative comparison. To show the effectiveness of the proposed hierarchical salient object selection module in generating inpainting masks, we compare different mask schemes, including only \mathbf{O}_1 semantic level (human), only \mathbf{O}_2 semantic level (Vehicle), only \mathbf{O}_3 semantic level (Sign), all categories, and our hierarchical selection method. Table 4 presents the comparison results in terms of both inpainting performance and changes in visual attention. From the five inpainting performance metrics, we can see that only inpainting the semantic region \mathbf{O}_1 with Human achieves the best performance. This is reasonable since the \mathbf{O}_1 region with the Human tends to be much smaller than other objects in street view images, therefore resulting in a smaller inpainting domain and better image quality. However, this also leads to the smallest changes in visual attention. In contrast, only inpainting the \mathbf{O}_2 can effectively reduce visual distraction, but it still works worse in redirecting visual attention to the object of interest. Our hierarchical strategy achieves high values in

Table 4: Evaluating the performance of different inpainting masks

Inpainted domain	SSIM \uparrow	PSNR \uparrow	$l_1 \downarrow$	LPIPS \downarrow	FID \downarrow	Visual attention change	
						$v_o(\mathbf{O}_o \text{ increase})\uparrow$	$v_d(\mathbf{O}_d \text{ reduction})\uparrow$
\mathbf{O}_1	0.9923	28.62	0.0019	0.1015	4.83	0.0533	0.0018
\mathbf{O}_2	0.9574	27.26	0.0088	0.3804	5.01	0.2653	0.1469
\mathbf{O}_3	0.9606	27.88	0.0067	0.4165	4.70	0.3113	0.0361
All categories	0.9164	26.14	0.0155	0.7298	5.05	1.1155	0.1226
Ours	0.9175	26.38	0.0154	0.7296	5.04	1.0483	0.1403

both v_o and v_d , and also performs better in terms of inpainting results than using all categories as masks.

Qualitative comparison. We show the visual results of different mask generation schemes in Figure 7. Similar to the results in Table 4, small mask regions lead to high-quality images but small gaze changes. Our method can well balance the two requirements based on hierarchical salient object selection.

4.4. Comparison of Different Urban Morphologies

Quantitative comparisons. Due to the diversity of street view images in location and acquisition, the need and performance of saliency-guided image inpainting will differ greatly. For example, street view images captured in bustling cities will be more widely used for last-meters wayfinding and contain more distracting objects than those captured in *motorway*. To explore how the urban morphology in street view images will influence our method, we used the urban canyon classification method [66] to classify the street scenes in Cityscapes dataset into four categories based on the street aspect ratio α ³. Note that images with a higher α tend to be captured in the downtown area, while images with a smaller α may contain suburb scenes. From the comparison results in Table 5, we can observe that a smaller aspect ratio α leads to better inpainting performance. This is reasonable since street view images captured in the suburb tend to contain less distracting objects, meaning smaller inpainting domains. Meanwhile, images of non-street canyons with α as 0 may contain no or fewer location-related objects. As a result, our method obtains the largest v_o but the smallest v_d distracting objects. For street scenes with a large ratio α , our method shows a greater advantage in redirecting visual attention to objects of interest.

Qualitative comparisons. Figure 8 illustrates the visual results of the proposed method in different types of street view images. For images captured in non-canyon areas ($\alpha = 0$), buildings tend to be more isolated and easier to focus on after removing

³ α is the ratio of the canyon height to the canyon width.

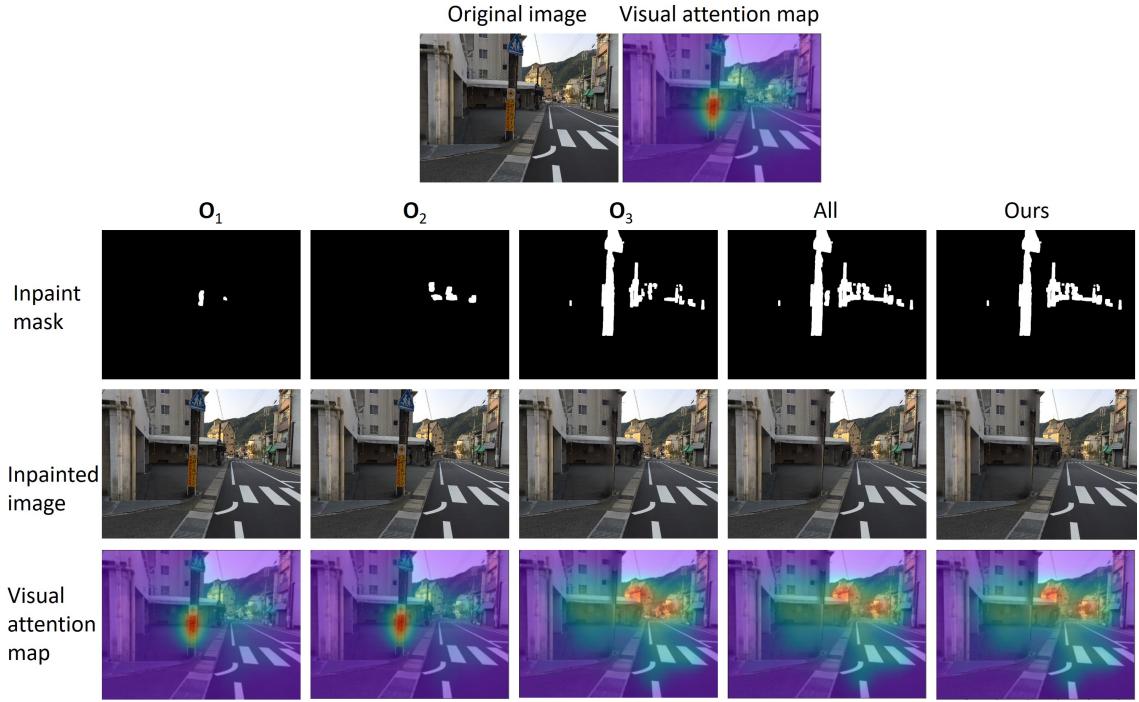


Figure 7: Visual performance of different inpainting masks on Cityscapes dataset. \mathbf{O}_1 denotes Human semantic level, \mathbf{O}_2 denotes Vehicle semantic level, and \mathbf{O}_3 denotes Sign semantic level. Best view in color and zoom in.

distracting objects and therefore lead to higher v_o (as obtained in Table 5). In contrast, our method shows more promising performance for street view images with larger street aspect ratios. Human gaze can be effectively redirected to location-related landmarks; meanwhile, inpainted images do not have obvious visual artifacts.

4.5. Comparison of Last-Few-Meters Wayfinding Efficiency

To show the effectiveness of our method in practical last-meters wayfinding applications, we conducted a controlled human-based evaluation to test the last-meters wayfinding efficiency before and after applying our inpainting method to street-view images. We carefully selected 50 original street view images with location-based buildings (e.g., landmarks, shopping malls, stores, and supermarkets) and then obtained their inpainted images based on our method. The 50 pairs of images were randomly divided into two subsets. Data_1 with 25 pairs and Data_2 with 25 pairs. Totally 10 volunteers (5 men and 5 women, ages 20 to 55) were asked to find a specific location using our self-developed platform (as shown in Figure 9). Volunteers were randomly divided into two groups: Group_1 with 5 persons and Group_2 with 5 per-

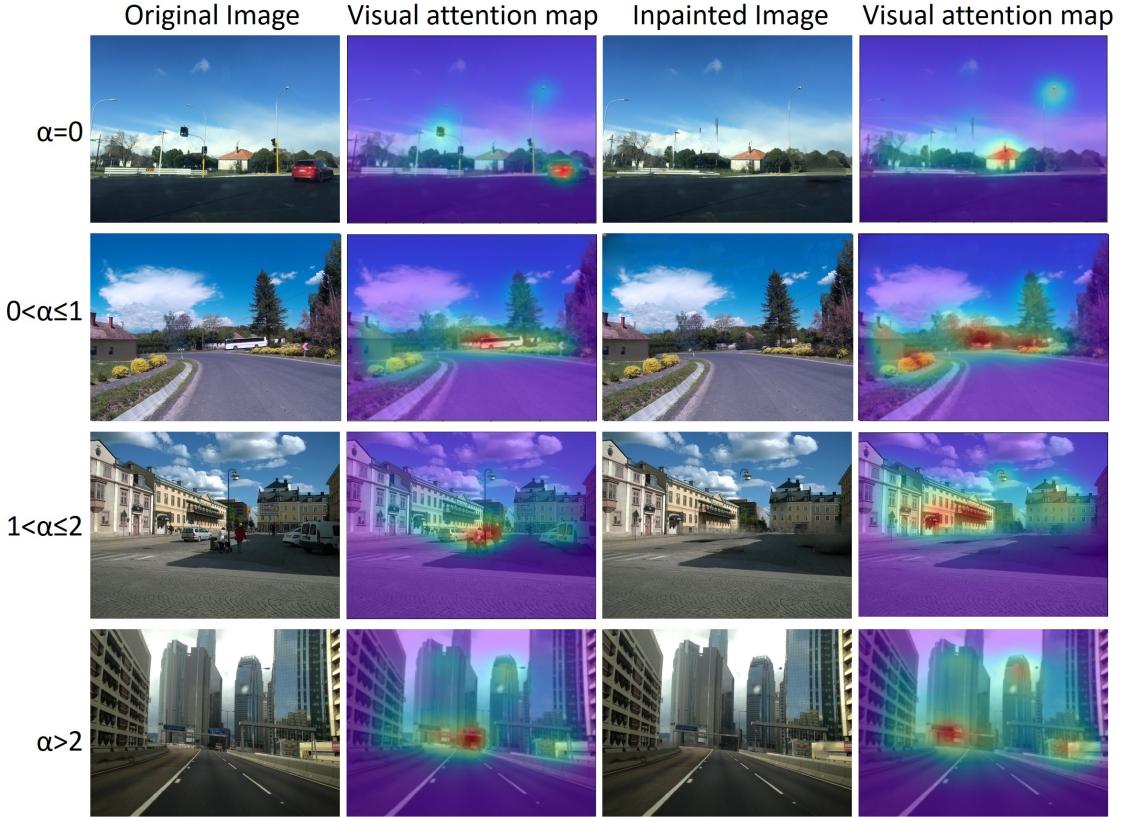


Figure 8: Visual performance of our proposed method in different types of street view images. Best view in color and zoom in.

sons, and conducted the evaluation according to the settings in Table 6. The search time between the image display and the “Found” confirmation click was recorded. To reduce individual differences in searching and matching, we normalized the time values of each volunteer between 0 and 1. Table 7 shows the average value of normalized search time, and we can see that our saliency-guided inpainting method can reduce approximately 38% search time by removing distracting objects from street view images.

4.6. Influence of different components

To explore the sensitivity of the proposed framework to different components, including semantic segmentation, salient object detection, and human visual attention detection, we performed comparison experiments to demonstrate the influence of these components on the performance of last-way finding. For the semantic segmentation module, we compared DeepLabv3+ [49] used in our framework with the

Table 5: Quantitative evaluation of the influence of α on the performance of the inpainting.

Street aspect ratio	$l_1 \downarrow$	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Visual attention change	
					$v_o(\mathbf{O}_o \text{ increase})\uparrow$	$v_d(\mathbf{O}_d \text{ reduction})\uparrow$
$\alpha = 0$	0.0130	0.9282	27.43	0.6886	1.1910	0.0732
$0 < \alpha \leq 1$	0.0175	0.9095	25.36	0.7595	0.8618	0.1903
$1 < \alpha \leq 2$	0.0180	0.9046	25.18	0.7871	0.9341	0.2326
$\alpha > 2$	0.0210	0.8923	24.29	0.8172	1.1441	0.2318



(a) last-meters wayfinding based on original image



(b) last-meters wayfinding based on inpainted image



(c) last-meters wayfinding based on original image



(d) last-meters wayfinding based on inpainted image

Figure 9: Self-developed program for human last-meters wayfinding evaluation. Paired images, for example, examples in (a) and (b), are used to report results for comparison, while unpaired images, e.g., examples in (c) and (d), are shown to each volunteer to ensure that the finding scene is not repeated.

Table 6: Human-based evaluation settings in data and volunteers

Subset	Group_1	Group_2
Data_1	Original	Inpainted
Data_2	Inpainted	Original

Table 7: Comparison of human Last-Few-Meters wayfinding efficiency in normalized search time

Subset	Original	Inpainted	Improvement
Data_1	0.4165	0.2547	38.85%
Data_2	0.4569	0.2834	37.97%

Table 8: Evaluation of the influence of different model components

Module	Method	Visual attention change	
		$v_o(\mathbf{O}_o \text{ increase}) \uparrow$	$v_d(\mathbf{O}_d \text{ reduction}) \uparrow$
Segmentation	PSPNet [67]	0.9189	0.0450
	DeepLabv3+ [49]	1.0483	0.1403
Salient object detection	BASNet [68]	1.0762	0.1362
	U^2Net model [16]	1.0483	0.1403
Visual attention detection	TransSalNet [69]	1.1112	0.1205
	UNISAL [57]	1.0483	0.1403

PSPNet [67] model, which has been proved to perform worse than the DeepLabv3+ model on several datasets [49], including the Cityscapes data [64]. Other modules were set the same as our proposed framework, i.e., with the U^2Net model [16] for salient object detection and UNISAL [57] for visual attention detection. The results in the first three rows of Table 8 show that the PSPNet [67] with worse segmentation performance achieves smaller visual attention changes than the DeepLabv3+ model. This means that the accuracy of semantic segmentation can obviously influence the final wayfinding efficiency. The more accurate the segmentation model is, the better the inpaiting method will perform, and therefore the larger improvement in last-meters wayfinding efficiency will be achieved. In Figure 10, we further present two testing examples to show the influence of segmentation performance on image inpainting and visual attention changes in wayfinding. We can also see that the worse performance of PSPNet leads to less visual distraction reduction for wayfinding than the DeepLabv3+ model.

Although the two saliency detection models, i.e., U^2Net [16] and UNISAL [57] are not specifically trained for street view images, both models are trained with a high diversity of images that includes various street view images. It is reasonable to believe that these two models can well extract the salient region and human attention in street view images on the Cityscapes dataset due to the overlapped image content with their training data. To further evaluate how saliency detection methods will influence the analysis, we compared the performance of different saliency detection models and human visual attention detection models. Specifically, for salient object detection, we compared U^2Net [16] with the widely used BASNet [68] and

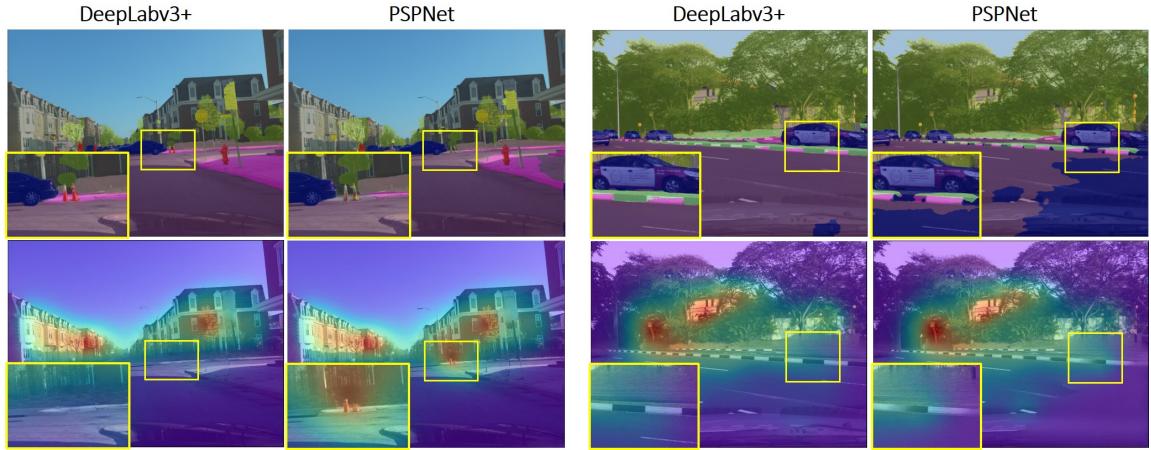


Figure 10: Influence of semantic segmentation performance on image inpainting and visual attention changes in wayfinding. Images in the first row are the semantic segmentation results, and images in the second row are the human visual attention maps on the inpainted images.

set other modules the same as our proposed framework. For human visual attention detection, we compared UNISAL [57] with the recent TranSalNet [69] for human attention detection. The results in Table 8 show that the influence of these two modules on human attention changes in last-way finding is not significant, since they achieved quite close v_o and v_d values. Not only does it show the reliability of U^2 Net and UNISAL used in our framework, but it also demonstrates the flexibility of our framework, which can be easily adjusted to other saliency detection methods.

4.7. Extended Application to Indoor Scenarios

Previous experimental results have demonstrated the effectiveness of the proposed method in improving the outdoor last-meters wayfinding. To illustrate the extensibility of our method, we show its application in indoor scenarios, where a particular brand store in a crowded mall may be difficult to identify in indoor navigation systems. To simulate this scenario, 2,000 images with interior scenes of shopping malls have been collected using the Bing Image Scraper⁴, and then the brand stores that must be found in each image are annotated. The qualitative results of our method in terms of the effects of inpainting and changes in human vision are shown in Fig. 11. It can be observed that the dense and distracting crowd can be effectively removed to redirect human visual attention to store localization. Similarly to the settings

⁴<https://pypi.org/project/bing-image-downloader/>.



Figure 11: Examples of the proposed method in indoor scene application. Best view when zooming in.

in Table 7, the quantitative results of the human search time based on 50 pairs of indoor images are shown in Table 9. We can see that our saliency-guided inpainting method can reduce approximately 28% search time by removing distracting objects in indoor scene images. The reasons why the improvement of indoor scenarios is smaller than that of outdoor images (see the results in Table 7) can be attributed to the fact that the scene diversity of indoor images is lower than outdoor street-view images, including lighting condition, image quality, and distracting region size.

5. Discussions

The proposed saliency-guided image inpainting method has demonstrated promising performance in redirecting human visual attention from distracting objects to location-related objects in street view images. As the first study to integrate saliency

Table 9: Comparison of human Last-Few-Meters wayfinding efficiency for an indoor scenario in normalized search time

Subset	Original	Inpainted	Improvement
Data_1	0.5059	0.3707	26.72%
Data_2	0.4669	0.3269	29.99%

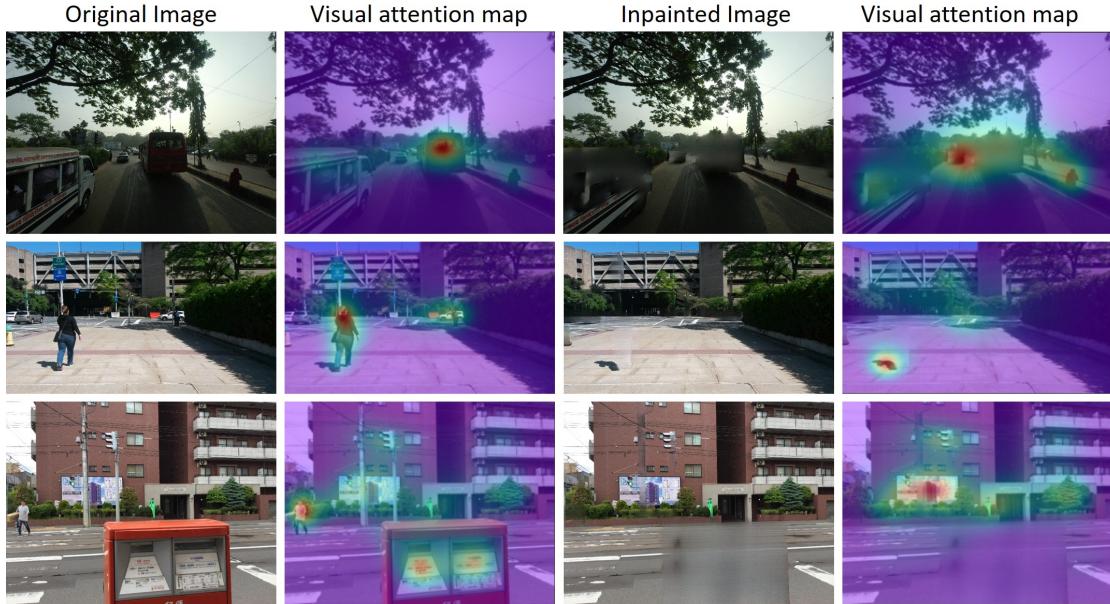


Figure 12: Examples with obvious visual inpainting artifacts in street view images under dark conditions (the first row), with shadows (the second row), and with large distracting object regions (the third row). Best view when zooming in.

detection into image inpainting for human last-meters wayfinding, there is much room for improvement of the current method. For example, inpainting performance can be further improved due to lingering visual artifacts in some challenging cases, as shown in Figure 12. These artifacts can be attributed to the following three reasons. First, street-view images have high diversity, such as lighting conditions and acquisition location. Second, the semantic image segmentation method has limitations in detecting objects with complex structures, especially those with shadows. Third, the image-inpainting method shows unstable performance in removing large domains.

Furthermore, the saliency-guided image inpainting in Figure 2 can be further optimized by end-to-end training. In the current implementation of our framework, the saliency detection and image-inpainting steps remain separate. It is possible to construct an objective function (e.g., $v_o + v_d$) related to the visual attention of

the object of interest. The goal of removing distracting objects can be formulated as a constrained optimization problem, such as maximizing the objective function with limited resources (e.g., running time) to support real-world applications [70]. From the perspective of network design, end-to-end training of joint dynamic object detection and inpainting could further improve the performance of our model. We can start from the simulated street view images, such as the CARLA dataset [71] and fine-tune the trained model on real-world street view images.

The other interesting application is the street view image processing for privacy protection [72, 26]. Instead of blurring faces or license plates, the sensitive information can be removed or replaced by virtual placeholders (e.g., faces and vehicles that do not exist). As the line between the physical world and the metaverse become fuzzy, we believe that this line of research is also applicable to privacy research, augmented reality, and intelligent transportation systems. Finally, given the rich redundancy in the temporal domain, the extension of this work to street-view video inpainting and depth-assisted image inpainting will remain as our future work.

6. Conclusions

Image-based last-meters wayfinding is easy to distract human visual attention due to the diverse and rich content in the street view images. In this paper, we make the first attempt to tackle this problem by designing a saliency-guided image inpainting framework, which aims to inpaint distracting objects on a semantic level to redirect human visual attention to location-related landmarks. Our method designs a hierarchical salient object selection module to generate inpainting masks and removes distracting objects based on the large-mask inpainting method with FFCs. Experimental results with qualitative and quantitative evaluations of street view images have validated the effectiveness of the proposed method. We believe that our framework will not only benefit image-based last-meters wayfinding but also can extend to other types of image restoration.

ACKNOWLEDGEMENTS

This work is partially supported by the DOJ under grant NIJ 2018-75-CX-0032, NSF under grant OAC-1839909, IIS-1951504 and the WV Higher Education Policy Commission Grant (HEPC.dsr.18.5).

References

- [1] M. Saha, A. J. Fiannaca, M. Kneisel, E. Cutrell, M. R. Morris, Closing the gap: Designing for the last-few-meters wayfinding problem for people with visual

impairments, in: The 21st international acm sigaccess conference on computers and accessibility, 2019, pp. 222–235.

- [2] K. Merry, P. Bettinger, Smartphone gps accuracy study in an urban environment, *PloS one* 14 (7) (2019) e0219890.
- [3] Y. Shu, K. G. Shin, T. He, J. Chen, Last-mile navigation using smartphones, in: Proceedings of the 21st annual international conference on mobile computing and networking, 2015, pp. 512–524.
- [4] J. C. Lock, G. Cielniak, N. Bellotto, A portable navigation system with an adaptive multimodal interface for the blind, in: 2017 AAAI Spring Symposium Series, 2017.
- [5] D. Anguelov, C. Dulong, D. Filip, C. Frueh, S. Lafon, R. Lyon, A. Ogale, L. Vincent, J. Weaver, Google street view: Capturing the world at street level, *Computer* 43 (6) (2010) 32–38.
- [6] R. Mahabir, R. Schuchard, A. Crooks, A. Croitoru, A. Stefanidis, Crowdsourcing street view imagery: a comparison of mapillary and openstreetcam, *ISPRS International Journal of Geo-Information* 9 (6) (2020) 341.
- [7] Q. Li, J. P. Queralta, T. N. Gia, Z. Zou, T. Westerlund, Multi-sensor fusion for navigation and mapping in autonomous vehicles: Accurate localization in urban environments, *Unmanned Systems* 8 (03) (2020) 229–237.
- [8] Q. Wang, B. Guo, Y. Liu, Q. Han, T. Xin, Z. Yu, Crowdnavi: Last-mile outdoor navigation for pedestrians using mobile crowdsensing, *Proceedings of the ACM on Human-Computer Interaction* 2 (CSCW) (2018) 1–23.
- [9] Q. Sun, A. Patney, L.-Y. Wei, O. Shapira, J. Lu, P. Asente, S. Zhu, M. McGuire, D. Luebke, A. Kaufman, Towards virtual reality infinite walking: dynamic saccadic redirection, *ACM Transactions on Graphics (TOG)* 37 (4) (2018) 1–13.
- [10] C. Lander, F. Wiehr, N. Herbig, A. Krüger, M. Löchtefeld, Inferring landmarks for pedestrian navigation from mobile eye-tracking data and google street view, in: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 2017, pp. 2721–2729.
- [11] E. Vig, M. Dorr, D. Cox, Large-scale optimization of hierarchical features for saliency prediction in natural images, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 2798–2805.

- [12] M. Jiang, S. Huang, J. Duan, Q. Zhao, Salicon: Saliency in context, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 1072–1080.
- [13] W. Wang, J. Shen, F. Guo, M.-M. Cheng, A. Borji, Revisiting video saliency: A large-scale benchmark and a new model, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 4894–4903.
- [14] J. Xu, M. Jiang, S. Wang, M. S. Kankanhalli, Q. Zhao, Predicting human gaze beyond pixels, *Journal of vision* 14 (1) (2014) 28–28.
- [15] L. Itti, C. Koch, E. Niebur, A model of saliency-based visual attention for rapid scene analysis, *IEEE Transactions on pattern analysis and machine intelligence* 20 (11) (1998) 1254–1259.
- [16] X. Qin, Z. Zhang, C. Huang, M. Dehghan, O. R. Zaiane, M. Jagersand, U2-net: Going deeper with nested u-structure for salient object detection, *Pattern Recognition* 106 (2020) 107404.
- [17] D. Turner, S. Savage, A. C. Snoeren, On the empirical performance of self-calibrating wifi location systems, in: 2011 IEEE 36th Conference on Local Computer Networks, IEEE, 2011, pp. 76–84.
- [18] Z. Yang, C. Wu, Y. Liu, Locating in fingerprint space: wireless indoor localization with little human intervention, in: Proceedings of the 18th annual international conference on Mobile computing and networking, 2012, pp. 269–280.
- [19] P. Roy, C. Chowdhury, A survey of machine learning techniques for indoor localization and navigation systems, *Journal of Intelligent & Robotic Systems* 101 (3) (2021) 1–34.
- [20] D. McGookin, S. Brewster, P. Priego, Audio bubbles: Employing non-speech audio to support tourist wayfinding, in: International Conference on Haptic and Audio Interaction Design, Springer, 2009, pp. 41–50.
- [21] G. Shen, Z. Chen, P. Zhang, T. Moscibroda, Y. Zhang, {Walkie-Markie}: Indoor pathway mapping made easy, in: 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13), 2013, pp. 85–98.
- [22] C. Sebastian, B. Boom, E. Bondarev, et al., Lidar assisted large-scale privacy protection in street view cycloramas, *Electronic Imaging* 2019 (11) (2019) 281–1.

- [23] A. Nodari, M. Vanetti, I. Gallo, Digital privacy: Replacing pedestrians from google street view images, in: Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012), IEEE, 2012, pp. 2889–2893.
- [24] J. Yu, H. Xue, B. Liu, Y. Wang, S. Zhu, M. Ding, Gan-based differential private image privacy protection framework for the internet of multimedia things, *Sensors* 21 (1) (2021) 58.
- [25] R. Uittenbogaard, C. Sebastian, J. Vijverberg, B. Boom, D. M. Gavrila, et al., Privacy protection in street-view panoramas using depth and multi-view imagery, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 10581–10590.
- [26] A. Flores, S. Belongie, Removing pedestrians from google street view images, in: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops, 2010, pp. 53–58.
- [27] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, T. S. Huang, Generative image inpainting with contextual attention, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 5505–5514.
- [28] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, B. Catanzaro, Image inpainting for irregular holes using partial convolutions, in: Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 85–100.
- [29] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, T. S. Huang, Free-form image inpainting with gated convolution, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 4471–4480.
- [30] J. Li, N. Wang, L. Zhang, B. Du, D. Tao, Recurrent feature reasoning for image inpainting, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 7760–7768.
- [31] M. Suin, K. Purohit, A. Rajagopalan, Distillation-guided image inpainting, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 2481–2490.
- [32] R. Suvorov, E. Logacheva, A. Mashikhin, A. Remizova, A. Ashukha, A. Silvestrov, N. Kong, H. Goka, K. Park, V. Lempitsky, Resolution-robust large mask inpainting with fourier convolutions, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2022, pp. 2149–2159.

- [33] S. Goferman, L. Zelnik-Manor, A. Tal, Context-aware saliency detection, *IEEE transactions on pattern analysis and machine intelligence* 34 (10) (2011) 1915–1926.
- [34] X. Hou, L. Zhang, Saliency detection: A spectral residual approach, in: 2007 IEEE Conference on computer vision and pattern recognition, 2007, pp. 1–8.
- [35] Q. Yan, L. Xu, J. Shi, J. Jia, Hierarchical saliency detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2013, pp. 1155–1162.
- [36] R. Zhao, W. Ouyang, H. Li, X. Wang, Saliency detection by multi-context deep learning, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 1265–1274.
- [37] R. Cong, J. Lei, H. Fu, M.-M. Cheng, W. Lin, Q. Huang, Review of visual saliency detection with comprehensive information, *IEEE Transactions on circuits and Systems for Video Technology* 29 (10) (2018) 2941–2959.
- [38] D. Zhang, H. Fu, J. Han, A. Borji, X. Li, A review of co-saliency detection algorithms: Fundamentals, applications, and challenges, *ACM Transactions on Intelligent Systems and Technology (TIST)* 9 (4) (2018) 1–31.
- [39] Y.-C. Chen, K.-J. Chang, Y.-H. Tsai, Y.-C. F. Wang, W.-C. Chiu, Guide your eyes: Learning image manipulation under saliency guidance., in: BMVC, Vol. 2, 2019, p. 3.
- [40] R. Mechrez, E. Shechtman, L. Zelnik-Manor, Saliency driven image manipulation, *Machine Vision and Applications* 30 (2) (2019) 189–202.
- [41] Y. A. Mejjati, C. F. Gomez, K. I. Kim, E. Shechtman, Z. Bylinskii, Look here! a parametric learning based approach to redirect visual attention, in: European Conference on Computer Vision, Springer, 2020, pp. 343–361.
- [42] L. Jiang, M. Xu, X. Wang, L. Sigal, Saliency-guided image translation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 16509–16518.
- [43] K. Aberman, J. He, Y. Gandelsman, I. Mosseri, D. E. Jacobs, K. Kohlhoff, Y. Pritch, M. Rubinstein, Deep saliency prior for reducing visual distraction, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 19851–19860.

- [44] X. Liu, Z. Deng, Y. Yang, Recent progress in semantic image segmentation, *Artificial Intelligence Review* 52 (2) (2019) 1089–1106.
- [45] C. Cao, Y. Fu, Learning a sketch tensor space for image inpainting of man-made scenes, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2021.
- [46] A. Borji, Boosting bottom-up and top-down visual features for saliency estimation, in: *2012 ieee conference on computer vision and pattern recognition*, IEEE, 2012, pp. 438–445.
- [47] T. V. Nguyen, M. Xu, G. Gao, M. Kankanhalli, Q. Tian, S. Yan, Static saliency vs. dynamic saliency: a comparative study, in: *Proceedings of the 21st ACM international conference on Multimedia*, 2013, pp. 987–996.
- [48] S. He, H. R. Tavakoli, A. Borji, Y. Mi, N. Pugeault, Understanding and visualizing deep visual saliency models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10206–10215.
- [49] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in: *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 801–818.
- [50] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, B. Schiele, The cityscapes dataset for semantic urban scene understanding, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [51] F. Perazzi, P. Krähenbühl, Y. Pritch, A. Hornung, Saliency filters: Contrast based filtering for salient region detection, in: *2012 IEEE conference on computer vision and pattern recognition*, IEEE, 2012, pp. 733–740.
- [52] H. Yu, K. Zheng, J. Fang, H. Guo, W. Feng, S. Wang, Co-saliency detection within a single image, in: *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [53] X. Wang, C. Qi, Detecting action-relevant regions for action recognition using a three-stage saliency detection technique, *Multimedia Tools and Applications* 79 (11) (2020) 7413–7433.
- [54] L. Chi, B. Jiang, Y. Mu, Fast fourier convolution, *Advances in Neural Information Processing Systems* 33 (2020).

- [55] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, A. Torralba, Places: A 10 million image database for scene recognition, *IEEE transactions on pattern analysis and machine intelligence* 40 (6) (2017) 1452–1464.
- [56] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, B. Catanzaro, High-resolution image synthesis and semantic manipulation with conditional gans, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8798–8807.
- [57] R. Droste, J. Jiao, J. A. Noble, Unified Image and Video Saliency Modeling, in: *Proceedings of the 16th European Conference on Computer Vision (ECCV)*, 2020.
- [58] C. Zheng, T.-J. Cham, J. Cai, Pluralistic image completion, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1438–1447.
- [59] Z. Yi, Q. Tang, S. Azizi, D. Jang, Z. Xu, Contextual residual aggregation for ultra high-resolution image inpainting, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7508–7517.
- [60] Y. Zeng, Z. Lin, H. Lu, V. M. Patel, Cr-fill: Generative image inpainting with auxiliary contextual reconstruction, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14164–14173.
- [61] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, O. Wang, The unreasonable effectiveness of deep features as a perceptual metric, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.
- [62] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two time-scale update rule converge to a local nash equilibrium, *Advances in neural information processing systems* 30 (2017).
- [63] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [64] H. Imam, B. A. Abdullah, H. E. Abd El Munim, Semantic segmentation under severe imaging conditions, in: *2019 Digital Image Computing: Techniques and Applications (DICTA)*, IEEE, 2019, pp. 1–7.

- [65] C. Doersch, S. Singh, A. Gupta, J. Sivic, A. Efros, What makes paris look like paris?, *ACM Transactions on Graphics* 31 (4) (2012).
- [66] C.-B. Hu, F. Zhang, F.-Y. Gong, C. Ratti, X. Li, Classification and mapping of urban canyon geometry using google street view images and deep multitask learning, *Building and Environment* 167 (2020) 106424.
- [67] H. Zhao, J. Shi, X. Qi, X. Wang, J. Jia, Pyramid scene parsing network, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [68] X. Qin, Z. Zhang, C. Huang, C. Gao, M. Dehghan, M. Jagersand, Basnet: Boundary-aware salient object detection, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 7479–7489.
- [69] J. Lou, H. Lin, D. Marshall, D. Saupe, H. Liu, Transalnet: Towards perceptually relevant visual saliency prediction, *Neurocomputing* 494 (2022) 455–467.
- [70] B. Bescos, J. Neira, R. Siegwart, C. Cadena, Empty cities: Image inpainting for a dynamic-object-invariant space, in: *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 5460–5466.
- [71] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, V. Koltun, CARLA: An open urban driving simulator, in: *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16.
- [72] A. Frome, G. Cheung, A. Abdulkader, M. Zennaro, B. Wu, A. Bissacco, H. Adam, H. Neven, L. Vincent, Large-scale privacy protection in google street view, in: *2009 IEEE 12th international conference on computer vision*, 2009, pp. 2373–2380.