



Evidential fully convolutional network for semantic segmentation

Zheng Tong¹ · Philippe Xu¹ · Thierry Denœux^{1,2}

Accepted: 5 March 2021 / Published online: 3 April 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

We propose a hybrid architecture composed of a fully convolutional network (FCN) and a Dempster-Shafer layer for image semantic segmentation. In the so-called evidential FCN (E-FCN), an encoder-decoder architecture first extracts pixel-wise feature maps from an input image. A Dempster-Shafer layer then computes mass functions at each pixel location based on distances to prototypes. Finally, a utility layer performs semantic segmentation from mass functions and allows for imprecise classification of ambiguous pixels and outliers. We propose an end-to-end learning strategy for jointly updating the network parameters, which can make use of soft (imprecise) labels. Experiments using three databases (Pascal VOC 2011, MIT-scene Parsing and SIFT Flow) show that the proposed combination improves the accuracy and calibration of semantic segmentation by assigning confusing pixels to multi-class sets.

Keywords Evidence theory · Belief function · Fully convolutional network · Decision analysis · Semantic segmentation

1 Introduction

In the past few decades, one of the most difficult problems in computer vision has been image semantic segmentation, defined as the process of partitioning a digital image into multiple sets of pixels. The result of image segmentation is a set of segments that collectively cover the entire image, called the *segmentation mask*. This mask constitutes a simplified representation, more meaningful and easier to analyze than the original image. Semantic segmentation has been widely applied to advanced driver assistance systems [6, 16, 36], human-machine interaction [39], medical imaging [18], and so on.

In the last decade, deep learning-based models, especially fully convolutional networks (FCNs) [25] and variants [22, 28], have been developed for semantic segmentation and have achieved remarkable success. FCNs take advantage of existing deep neural networks, which have the capacity to learn reliable and robust features. An FCN transforms existing and well-known classification models such as VGG (16-layer net) [31] or ResNet [21] into fully convolutional ones by replacing the fully connected layers with convolutional ones to output spatial maps instead of classification scores. Those maps are upsampled using fractionally-strided convolutions (called *deconvolutions* [41, 42]) to produce dense per-pixel labeled outputs. This approach has allowed for significant improvements in segmentation accuracy over traditional methods on benchmark databases like Pascal VOC 2011 [17]. However, despite the power and flexibility of the FCN-based models, they still face the following three problems:

1. *How to perform novelty detection?* In many learning sets, not all classes are labeled, especially for some objects in the background. An ideal image segmentation algorithm should detect “unknown” objects belonging to classes that are not represented in the learning set. This capacity is called *novelty detection* [8]. FCN-based models generally randomly assign unknown objects to one of the known classes, though some models tend to assign unknown objects to the background class.

This article belongs to the Topical Collection: *30th Anniversary Special Issue*

✉ Thierry Denœux
thierry.denoeux@utc.fr

Zheng Tong
zheng.tong@hds.utc.fr

Philippe Xu
philippe.xu@hds.utc.fr

¹ Université de technologie de Compiègne, CNRS, Heudiasyc, Compiègne, France

² Institut universitaire de France, Paris, France

2. *How to process pixels with confusing information?* In image-segmentation training sets, all pixels are precisely labeled, even if the true label is actually uncertain. This is the case, for example, for the pixels at object borders. Pixels with precise but incorrect labels may have negative effects on the performance of learning systems [2, 27].
3. *When will the FCN-based methods fail?* In decision-making systems, a neural network should not only be as accurate as possible, but it should also have the ability to signal when it is likely to be incorrect. Neural networks developed nowadays tend not to be well calibrated [20], though they are more accurate than they were a decade ago. In other words, the accuracy of modern neural networks, including FCN-based models, does not match their confidence.

In this paper, the above problems are tackled using the Dempster-Shafer (DS) theory of belief functions [7, 30]. This mathematical framework, also referred to as *Evidence Theory*, is based on representing independent pieces of evidence by mass functions and combining them using a generic operator called Dempster's rule. It is a well-established formalism for reasoning and making decisions under uncertainty [10, 13, 38]. A mass function has more degrees of freedom than a probability distribution, which allows it to represent a wider range of belief states, from complete ignorance to full certainty.

One of the applications of the DS theory is to design *evidential classifiers* (e.g., [9, 14, 24, 33]), which compute a predicted mass function for each input vector. The output mass function can then be used for decision-making [4, 8, 19]. Over the years, two main principles have been developed to design an evidential classifier: the *model-based* and *distance-based* approaches [15]. The former uses estimated class-conditional distributions [32], while the latter constructs mass functions based on distances to prototypes [9, 14, 24]. Thanks to the generality and expressiveness of the belief-function formalism, an evidential classifier provides more informative outputs than those of conventional classifiers (e.g., a neural network with a softmax output layer). The flexibility of evidential classifiers can be exploited for uncertain data classification [40] and set-valued classification [8, 26]. Therefore, it may be advantageous to combine an FCN-based model with an evidential classifier for semantic segmentation.

The objective of this study is to take advantage of object representations generated by an FCN and use them as the input features of an evidential classifier for decision-making. The proposed model, referred to as the *evidential fully convolutional network* (E-FCN), transforms an FCN model by replacing its softmax layer by a distance-based DS layer and a utility layer. In an E-FCN, an FCN architecture

is used to extract pixel-wise high-order features from an input image. Then, the features are converted into pixel-wise mass functions by the DS layer. Finally, the mass functions are used to compute the utilities of acts assigning pixels to a set of classes for semantic segmentation in a so-called "utility layer". An end-to-end learning procedure allows us to train the E-FCN using a learning set with soft labels. The effectiveness of the E-FCN is demonstrated and discussed based on experiments using three benchmark databases (Pascal VOC 2011 [17], MIT-scene Parsing [43], and SIFT Flow [34]).

The rest of the paper is organized as follows. Section 2 starts with a brief reminder of DS theory, the DS layer for constructing mass functions, and feature representation via FCN. The E-FCN model is then introduced in Section 3. Section 4 presents numerical experiments, which demonstrate the advantages of the E-FCNs. Finally, we conclude the paper in Section 5.

2 Background

This section first recalls some necessary definitions regarding DS theory (Section 2.1), and the evidential neural network at the basis of the DS layer (Section 2.2). A brief description of feature representation via FCNs is then provided in Section 2.3.

2.1 Dempster-Shafer theory

The main concepts underlying DS theory are only briefly presented in this section, and some basic notations are introduced. Detailed information can be found in Shafer's original work [30] and in the recent review [13].

Let $\Omega = \{\omega_1, \dots, \omega_M\}$ be a set of classes, called the *frame of discernment*. A *mass function* on Ω is a mapping m from 2^Ω to $[0,1]$ such that $m(\emptyset) = 0$ and

$$\sum_{A \subseteq \Omega} m(A) = 1. \quad (1)$$

For any $A \subseteq \Omega$, each mass $m(A)$ is interpreted as a share of a unit mass of belief allocated to the hypothesis that the truth is in A , and which cannot be allocated to any strict subset of A based on the available evidence. Set A is called a *focal set* of m if $m(A) > 0$. A mass function is said to be *logical* if it has only one focal set.

Two mass functions m_1 and m_2 representing independent items of evidence can be combined conjunctively by Dempster's rule \oplus [30] as

$$(m_1 \oplus m_2)(A) = \frac{(m_1 \cap m_2)(A)}{1 - (m_1 \cap m_2)(\emptyset)} \quad (2a)$$

for all $A \neq \emptyset$, with

$$(m_1 \cap m_2)(A) = \sum_{B \cap C = A} m_1(B) m_2(C), \quad (2b)$$

and $(m_1 \oplus m_2)(\emptyset) = 0$. Mass functions m_1 and m_2 can be combined if and only if the denominator on the right-hand side of (2a) is strictly positive. The operator \oplus is commutative and associative.

For decision-making with belief functions, let $u_{ij} \in [0, 1]$ denote the utility of selecting ω_i when the true state is ω_j , and f_{ω_i} the act of selecting ω_i . We define the *pignistic expected utility* [11] of act f_{ω_i} as

$$\mathbb{E}_m(f_{\omega_i}) = \sum_{j=1}^M u_{ij} \text{Bet} P_m(\{\omega_j\}), \quad (3a)$$

where $\text{Bet} P_m$ is the *pignistic* probability measure computed from mass function m by the *pignistic transformation*, defined as

$$\text{Bet} P_m(\{\omega_j\}) = \sum_{\{A \subseteq \Omega: \omega_j \in A\}} \frac{m(A)}{|A|}, \quad (3b)$$

for all $\omega_j \in \Omega$. The act with the highest pignistic expected utility can then be selected. Other decision criteria in the belief function framework are reviewed in [11] and [26].

2.2 Evidential neural network

Denceux [9] proposed a distance-based neural-network based on DS theory, known as the *evidential neural network (ENN) classifier*. The ENN classifier summarizes the learning set by a small number of prototypes, and treats the proximity of an input vector to each prototype as a piece of evidence about its class. The different pieces of evidence are represented by mass functions, which are combined using Dempster's rule (2.1). This section provides a brief description of the ENN classifier.

We consider a training set $\mathcal{X} = \{\mathbf{x}^1, \dots, \mathbf{x}^N\} \subset \mathbb{R}^P$ of N examples represented by P -dimensional feature vectors, and n prototypes $\{\mathbf{p}^1, \dots, \mathbf{p}^n\} \subset \mathbb{R}^P$. For a test sample \mathbf{x} , the ENN classifier constructs mass functions that quantify the uncertainty about its class in $\Omega = \{\omega_1, \dots, \omega_M\}$, using a three-step procedure. This procedure can be implemented in a neural network with two hidden layers and one output layer. These three layers can be considered as a single complex layer called the "DS layer", which will be plugged into an FCN architecture as explained in Section 3.1. The three-step procedure can be described as follows.

Step 1: The similarity between \mathbf{x} and each prototype \mathbf{p}^l is computed as

$$s^l = \alpha^l \exp\left(-\left(\eta^l d^l\right)^2\right), \quad l = 1, \dots, n, \quad (4)$$

where $d^l = \|\mathbf{x} - \mathbf{p}^l\|$ is the Euclidean distance between \mathbf{x} and prototype \mathbf{p}^l , $\eta^l \in \mathbb{R}$ is a scale parameter and α^l is a parameter in $[0, 1]$. Prototypes $\mathbf{p}^1, \dots, \mathbf{p}^n$ can be considered as vectors of connection weights between the input layer and a hidden layer of n Radial Basis Function (RBF) units. The number n of prototypes is a hyper-parameter and can be tuned using a validation set or by cross-validation.

Step 2: The mass function m^l associated to reference pattern \mathbf{p}^l is computed as

$$m^l(\{\omega_j\}) = v_j^l s^l, \quad j = 1, \dots, M \quad (5a)$$

$$m^l(\Omega) = 1 - s^l, \quad (5b)$$

where $v_j^l \geq 0$ is the degree of membership of prototype \mathbf{p}^l to class ω_j with $\sum_{j=1}^M v_j^l = 1$. We denote the vector of masses induced by prototype \mathbf{p}^l as

$$\mathbf{m}^l = (m^l(\{\omega_1\}), \dots, m^l(\{\omega_M\}), m^l(\Omega))^T.$$

(5) can be regarded as computing the activation of units in a "mass functions" layer composed of n modules of $M + 1$ units each. The activations of the units in module l correspond to the belief masses assigned by m^l .

Step 3: The n mass functions \mathbf{m}^l , $l = 1, \dots, n$, are aggregated by Dempster's rule (2). The combined mass function can be computed iteratively as $\mu^1 = m^1$ and $\mu^l = \mu^{l-1} \cap m^l$ for $l = 2, \dots, n$. From (2a), we have

$$\mu^l(\{\omega_j\}) = \mu^{l-1}(\{\omega_j\})m^l(\{\omega_j\}) + \mu^{l-1}(\{\omega_j\})m^l(\{\Omega\}) + \mu^{l-1}(\Omega)m^l(\{\omega_j\}) \quad (6a)$$

for $l = 2, \dots, n$ and $j = 1, \dots, M$, and

$$\mu^l(\Omega) = \mu^{l-1}(\Omega)m^l(\Omega) \quad l = 2, \dots, n. \quad (6b)$$

The output vector $\mathbf{m} = (m(\{\omega_1\}), \dots, m(\{\omega_M\}), m(\Omega))^T$ is finally obtained by normalizing μ^n as

$$m(A) = \frac{\mu^n(A)}{\mu^n(\Omega) + \sum_{j'=1}^M \mu^n(\{\omega_{j'}\})},$$

with $A \in \{\{\omega_1\}, \dots, \{\omega_M\}, \Omega\}$.

The network parameters are the prototypes \mathbf{p}^l , the coefficients α^l and η_l , and the membership degrees v_j^l for $l = 1, \dots, n$ and $j = 1, \dots, M$. They are learnt by minimizing a loss function. To enforce the constraints $0 \leq \alpha^l \leq 1$, we introduce new variables $\xi^l \in \mathbb{R}$ such that

$$\alpha^l = \frac{1}{1 + \exp(-\xi^l)} \in (0, 1).$$

Similarly, the constraints on parameters v_j^l are enforced by introducing new parameters $\delta_j^l \in \mathbb{R}$ such that

$$v_j^l = \frac{(\delta_j^l)^2}{\sum_{j'=1}^M (\delta_{j'}^l)^2} \tag{7}$$

for $l = 1, \dots, n$ and $j = 1, \dots, M$. More details can be found in [9].

2.3 Fully convolutional network

The performance of a classifier in semantic segmentation tasks heavily depends on the information contained in its input features. Feature representation, an essential part of the machine learning workflow, consists in discovering the predictors needed for semantic segmentation from input images. In recent years, FCNs [25] and their variants [22, 28] have achieved remarkable performances thanks to their ability to construct rich pixel-wise deep feature representations.

FCNs owe their name to their architecture, which is built only from locally connected layers, such as convolution, pooling, and upsampling layers. No dense layer is used in this kind of architecture. Generally, an FCN consists of two main parts: an encoder-decoder architecture for pixel-wise object representation and a softmax layer for pixel-wise assignments. In the encoder-decoder architecture, an input image is encoded by several convolutional and pooling layers and then decoded by one or more upsampling layers. The softmax layer assigns each pixel in the input image to one of the classes based on the outputs of the encoder-decoder architecture. Therefore, the outputs of the encoder-decoder architecture, called the *pixel-wise feature maps*, are considered as a feature representation of the input image. In the study, these feature maps are used as input to a DS layer

allowing for set-valued semantic segmentation, as will be shown in Section 3.1.

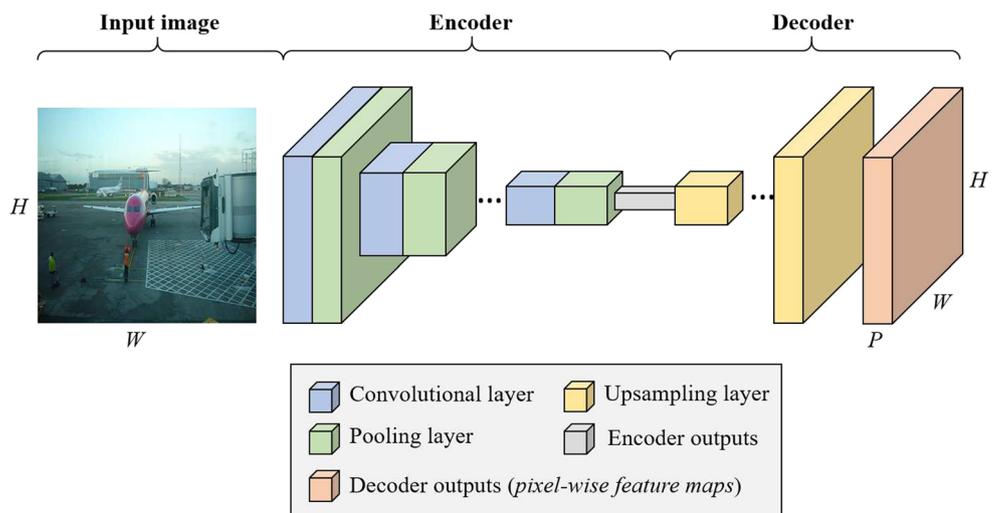
To understand the feature representation of FCNs, we briefly recall the encoder-decoder architecture illustrated in Figure 1. The encoder part consists of several convolutional and pooling layers. Each convolutional layer performs convolutions its input to produce a set of feature maps. Let $z = (z^1, \dots, z^D)$ be the input made up of D input maps or *input channels* z^i ($i = 1, \dots, D$) of size $H \times W$. The processes in a convolutional layer with input z , consisting of e convolution kernels with size $a \times b$, are expressed as

$$c^j = f(\lambda^j + \sum_i \varepsilon^{i,j} * z^i), \tag{8}$$

where $\varepsilon^{i,j}$, a matrix of size $a \times b$, is the convolution kernel between the i -th input map and the j -th output map; λ^j is the bias of kernel $\varepsilon^{i,j}$; $*$ denotes the convolution operation; c^j is the j -th output feature map, with size $\frac{h-a+1}{r} \times \frac{w-b+1}{r}$, $j = 1, \dots, e$; r is the stride with which the kernel slides over input map z^i , and f is the activation function, such as the rectified linear unit $\text{ReLU}(x) = \max(0, x)$ [23]. A pooling layer follows the convolutional layer to sub-sample feature map c^j by computing some statistics of feature values within non-overlapping $s \times s$ windows. In the case of max-pooling used in this paper, the statistic is the maximum. Thus, the outputs of the pooling layer is composed of the D feature maps sub-sampled by factor s . For example, feature map c^j with size $\frac{h-a+1}{r} \times \frac{w-b+1}{r}$ is downsized to $\frac{h-a+1}{2r} \times \frac{w-b+1}{2r}$ by a pooling layer with a 2×2 non-overlapping window.

Although the convolution and pooling operations in the encoder part contribute to feature representation by retaining only robust activations, spatial information within a receptive field is lost, which may be critical for image semantic segmentation. To address the issue, a decoder part made up of one or more upsampling layers is added at the

Fig. 1 Illustration of the encoder-decoder architecture. An encoder downsizes its input by convolution and pooling operations. The outputs of the encoder, as the sparse feature maps, are imported into a decoder. A decoder upsamples and densifies its inputs by performing the reverse operation of convolution and pooling. The final decoder outputs are the pixel-wise feature maps



output of the encoder part. The decoder performs the reverse operation of convolution and pooling for reconstructing a set of activations with the same size of the input image, as shown in Fig. 1. Thus, the outputs of the decoder part are enlarged feature maps. In this study, we use a *deconvolution layer* [28] to implement the upsampling operation.

A deconvolutional layer densifies its inputs of sparse feature maps through convolution-like operations with multiple learned kernels. However, contrary to convolutional layers, which connect multiple inputs within a kernel to a single activation, a deconvolutional layer associates a single input in a feature map to multiple outputs. Thus, the outputs of a deconvolutional layer are enlarged and dense feature maps. The processes of a deconvolution operation can also be summarized as (8), but its kernel sizes are larger than the input sizes, i.e., $a \geq H$ and $b \geq W$.

3 Evidential fully convolutional network

In this section, we describe the proposed E-FCN. Section 3.1 presents the overall architecture composed of an encoder-decoder module for feature representation, a DS layer to construct mass functions, and a utility layer for decision-making. The details of the utility layer are described in Section 3.2. Section 3.3 introduces the strategy for training E-FCN models using a learning set with soft labels.

3.1 Network architecture

The main idea of this work is to hybridize the ENN classifier presented in Section 2.2 and the FCN recalled in Section 2.3

by “plugging” a DS layer followed by a utility layer at the output of the final deconvolutional layer in the FCN. The architecture of the proposed method, called the *evidential FCN* (E-FCN), is illustrated in Fig. 2. An E-FCN classifier performs set-valued semantic segmentation and quantifies the uncertainty about the class of each pixel, taking values in $\Omega = \{\omega_1, \dots, \omega_M\}$, using a three-step procedure defined as follows.

Step 1: As in a probabilistic FCN (P-FCN), an image of size $W \times H \times 3$ is presented as input to the encoder-decoder architecture of an FCN to generate pixel-wise feature maps of size $W \times H \times P$, where P is the number of *output channels*. Each feature vector $1 \times 1 \times P$ from a pixel-wise feature map is a P -dimensional representation of the corresponding pixel, ready to be fed into the DS layer. This architecture generates reliable pixel-wise representations of the input image. Thanks to the representations, the E-FCN yields similar or even better performance for precise semantic segmentation than does a P-FCN with the same encoder-decoder architecture, as will be shown in Section 4.2.

Step 2: Each feature vector from the encoder-decoder architecture is fed into the DS layer, in which it is used to compute a mass function as explained in Section 2.2. The output of the DS layer for a given feature vector is an $(M + 1)$ -dimensional mass vector

$$\mathbf{m} = (m(\{\omega_1\}), \dots, m(\{\omega_M\}), m(\Omega))^T.$$

Thus, given pixel-wise feature maps of size $W \times H \times P$ from Step 1, the output of the DS layer is a tensor of size $W \times H \times (M + 1)$. Each mass vector in the tensor represents the uncertainty about the class

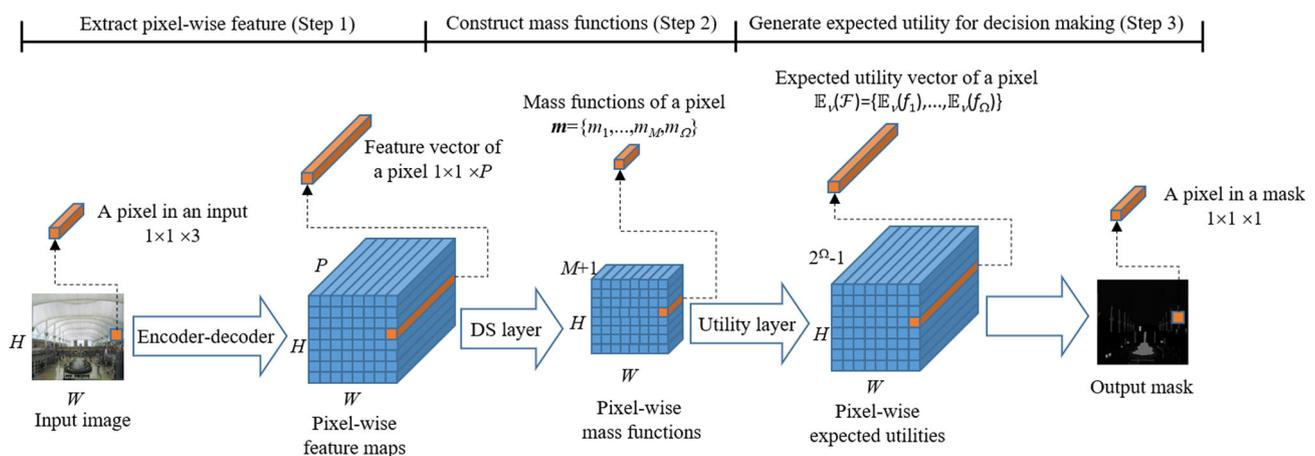


Fig. 2 Architecture of an evidential fully convolutional network (E-FCN). The E-FCN performs semantic segmentation using a three-step procedure. In the first step, an encoder-decoder architecture extracts pixel-wise feature maps from the input image. Each vector in the feature maps is fed into a DS layer to construct the pixel-wise mass

functions in the second step. These mass functions are finally fed into a utility layer to generate the pixel-wise expected utilities of all acts. Finally, the segmentation mask is computed based on the expected utilities

of the corresponding pixel. More precisely, the mass $m(\{\omega_i\})$ is the degree of belief that the true class of the corresponding pixel is ω_i . The DS layer tends to allocate uniform masses if the representations contain confusing information. The additional degree of freedom $m(\Omega)$ makes it possible to quantify the lack of evidence [12] and to verify whether the model is well trained [35]. The advantages of this uncertainty representation will be demonstrated in the performance evaluation of set-valued semantic segmentation using E-FCN in Section 4.3.

Step 3: The output pixel-wise mass vectors are fed into a utility layer, where they are used to compute the expected utility of acts. Each act is defined as the assignment of a pixel to a non-empty subset A of Ω . Therefore, the output of the layer for each feature vector from Step 2 is a vector of at most $2^\Omega - 1$ expected utilities when all of the possible acts are considered. The utility layer allows the E-FCN to perform set-valued semantic segmentation. This capability will be demonstrated by comparing the performances of the two types of FCNs in set-valued segmentation (Section 4.3) and novelty detection (Section 4.4) tasks. More details about the utility layer are given in the next section.

3.2 Utility layer

In this section, we describe in greater detail the decision-making process taking place in the utility layer. Section 3.2.1 begins by introducing the precise semantic segmentation method using mass functions and utilities. Section 3.2.2 then describes a method for computing the utility of set-valued pixel-wise classification, after which an approach to set-valued classification based on mass functions is described in Section 3.2.3. In Section 3.2.4, we summarize the workflow as a neural network layer in the E-FCN model.

3.2.1 Precise semantic segmentation

Let $\Omega = \{\omega_1, \dots, \omega_M\}$ be the set of classes. For semantic segmentation problems with precise prediction, each pixel in an image is assigned to exactly one class. An act is thus defined as the assignment of a pixel to one and only one of the M classes, and the set of acts is $\mathcal{F} = \{f_{\omega_1}, \dots, f_{\omega_M}\}$, where f_{ω_i} denotes assignment to class ω_i . To make decisions, we define a utility matrix U of size $M \times M$, whose general term $u_{ij} \in [0, 1]$ is the utility of assigning a pixel to class ω_i when the true class is ω_j .

When uncertainty about Ω is described by a DS mass function, each act f_{ω_i} induces an expected utility, such as the pignistic expected utility defined by (3). Given utility matrix U and the output of the DS layer m for a given pixel, the pignistic expected utility of assigning that pixel to class

ω_i is given by (3a), where $Bet P_m$ is the pignistic probability defined by (3b). The pixel is finally assigned to class ω_{i^*} such that

$$i^* = \arg \max_{i \in \{1, \dots, M\}} \mathbb{E}_m(f_{\omega_i}). \tag{9}$$

3.2.2 Extending the utility matrix

For semantic segmentation problems with imprecise prediction, we adopt the approach described in [26] for set-valued classification under uncertainty, which allows the assignment of a pixel to any non-empty subset A of Ω . The set of acts thus potentially becomes $\mathcal{F} = \{f_A : A \subseteq \Omega, A \neq \emptyset\}$, where f_A denotes the assignment to a subset A . (In practice, when the cardinality of Ω is very large, we may only consider acts f_A for some subsets A of Ω). In this study, f_A is referred to as an *imprecise assignment* when subset A is a *multi-class set* with $|A| \geq 2$. For decision-making with \mathcal{F} , the utility matrix U has to be extended to a matrix \tilde{U} of size $(2^M - 1) \times M$, where each element $\tilde{u}_{A,j}$ denotes the utility of assigning a pixel to set A of classes when the true label is ω_j . Following [26], this extension is performed as follows.

When the true class is ω_j , the utility of assigning a pixel to set A is defined as an *Ordered Weighted Average (OWA)* aggregation [37] of the utilities of each precise assignment in A as

$$\tilde{u}_{A,j} = \sum_{k=1}^{|A|} g_k u_{(k)j}^A, \tag{10}$$

where $u_{(k)j}^A$ is the k -th largest element in the set $\{u_{ij} : \omega_i \in A\}$ made up of the elements in the utility matrix U , and weights $\mathbf{g} = (g_1, \dots, g_{|A|})$ represent the preference to choose $u_{(k)j}(A)$ if forced to select a single value in $\{u_{ij} : \omega_i \in A\}$. The components of weight vector \mathbf{g} represent the *tolerance to imprecision* of a decision maker (DM). For example, full tolerance to imprecision is achieved when the assignment act f_A has utility 1 once set A contains the true label, no matter how large A is. In this case, only the maximum utility of elements in set $\{u_{ij}, \omega_i \in A\}$ is considered: $(g_1, g_2, \dots, g_{|A|}) = (1, 0, \dots, 0)$. At the other extreme, a DM attaching no value to imprecision would consider the act f_A as equivalent to selecting one class uniformly at random from A : this is achieved when

$$(g_1, g_2, \dots, g_{|A|}) = \left(\frac{1}{|A|}, \frac{1}{|A|}, \dots, \frac{1}{|A|} \right),$$

in which case the OWA operator becomes the average. In this study, following [26], we determine the weight vector \mathbf{g} of the OWA operator by adapting O'Hagan's method [29]. We define the tolerance to imprecision as

$$TDI(\mathbf{g}) = \sum_{k=1}^{|A|} \frac{|A| - k}{|A| - 1} g_k = \gamma, \tag{11}$$

which equals 1 for the maximum, 0 for the minimum, and 0.5 for the average. In practice, we only need to consider values of γ between 0.5 and 1 as a precise assignment is always more desirable than an imprecise one when $\gamma < 0.5$ [26]. Given a value of γ , we can compute the weights of the OWA operator by maximizing the entropy

$$ENT(\mathbf{g}) = - \sum_{k=1}^{|\mathcal{A}|} g_k \log g_k \quad (12)$$

subject to the constraints $TDI(\mathbf{g}) = \gamma$, $\sum_{k=1}^{|\mathcal{A}|} g_k = 1$, and $g_k \geq 0$.

Example 1 Table 1 shows an example of the extended utility matrix generated by an OWA operator with $\gamma = 0.8$. The first three rows constitute the original utility matrix, indicating that the utility equals 1 when assigning a sample to its true class, and 0 otherwise. The remaining rows are the matrix of the aggregated utilities. For example, we get a utility of 0.8 when assigning a sample to set $\{\omega_1, \omega_2\}$ if the true label is ω_1 .

3.2.3 Set-valued semantic segmentation

Based on an extended utility matrix \tilde{U} and the output of the DS layer \mathbf{m} for a given pixel, we can compute the pignistic expected utility of assigning that pixel to set A as

$$\mathbb{E}_m(f_A) = \sum_{j=1}^M \tilde{u}_{A,j} \text{Bet} P_m(\{\omega_j\}), \quad (13)$$

where $\text{Bet} P_m$ is the pignistic probability defined by (3b). The pixel is finally assigned to set A such that

$$A = \arg \max_{\emptyset \neq B \subseteq \Omega} \mathbb{E}_m(f_B). \quad (14)$$

Table 1 Utility matrix extended by an OWA operator with $\gamma = 0.8$

	Classes		
	ω_1	ω_2	ω_3
$f_{\{\omega_1\}}$	1	0	0
$f_{\{\omega_2\}}$	0	1	0
$f_{\{\omega_3\}}$	0	0	1
$f_{\{\omega_1, \omega_2\}}$	0.8	0.8	0
$f_{\{\omega_1, \omega_3\}}$	0.8	0	0.8
$f_{\{\omega_2, \omega_3\}}$	0	0.8	0.8
f_{Ω}	0.6819	0.6819	0.6819

3.2.4 Utility layer

The procedure of assigning a pixel to a set of classes using utility theory is implemented in the neural network as a *utility layer*. In this layer, the inputs and outputs are, respectively, the pixel-wise mass vectors \mathbf{m} from the preceding DS layer and the pixel-wise expected utilities of all acts in \mathcal{F} . The connection weight between unit j of the DS layer and output unit $A \subseteq \Omega$ corresponding to the assignment to set A is the utility value $\tilde{u}_{A,j}$. As coefficient γ describing the imprecision tolerance degree is predetermined, the connection weights of the expected utility layer are fixed and do not need to be updated during training.

In practice, the connections between the DS and utility layers can be designed by the user. For example, one can build a utility layer using the utility values $\tilde{u}_{A,j}$ with $|A| = 1$ to only consider precise assignments, or $0 < |A| \leq 2$ to consider assignment to sets classes of cardinality one or two. In this paper, we have only considered the acts f_A such that A is a singleton, Ω , or one of the soft labels present in the learning set (as explained in Section 3.3 below).

3.3 Learning with soft labels

In traditional learning systems for image semantic segmentation, all pixels are labeled with a single class even when their true class cannot be determined with full certainty. For example, the true class may be uncertain at object borders, but the border pixels are still given precise labels. Additionally, one cannot reliably label some small objects in an image, such as distant objects in a driving scene. Arbitrarily giving precise labels to pixels with confusing information may negatively impact the performance of learning systems in image semantic segmentation tasks. The notion of *soft label* [5, 14] is a way to address this problem.

Here, we define a soft label as a nonempty subset $A_* \in 2^{\Omega} \setminus \{\emptyset\}$ of classes a pixel may belong to, based on our current knowledge. For example, label $A_* = \{\omega_i, \omega_j\}$ indicates that the true class of a pixel is known to be either ω_i or ω_j but we cannot determine which one specifically. A strategy of end-to-end learning is proposed to train an E-FNC from an image learning set with soft labels. All parameters in the DS layer are first initialized randomly using normal distributions. For a given pixel with nonempty soft label $A_* \subseteq \Omega$, let m_l be the logical mass function with focal set A_* , i.e., such that $m_l(A_*) = 1$. The *labeling* pignistic expected utilities $\mathbb{E}_{m_l}(f_A)$ for $A \in 2^{\Omega} \setminus \{\emptyset\}$ can be computed using (13) and the pignistic belief-probability transformation (3b). Similarly, we consider the *predicted* pignistic expected utilities $\mathbb{E}_m(f_A)$ for $A \in 2^{\Omega} \setminus \{\emptyset\}$, where m is the predicted mass function from the DS layer of the E-FCN, with focal sets $\{\omega_1\}, \dots, \{\omega_M\}$ and Ω . For a given

pixel with soft label m_l and predicted mass function m , the loss $\mathcal{L}(m, m_l)$ is defined as the squared Euclidean distance between the vectors of expected utilities w.r.t. m_l and m :

$$\mathcal{L}(m, m_l) = \sum_{\emptyset \neq A \subseteq \Omega} [\mathbb{E}_{m_l}(f_A) - \mathbb{E}_m(f_A)]^2. \tag{15}$$

The derivatives of $\mathcal{L}(m, m_l)$ of the error w.r.t the output masses $m(\{\omega_k\})$ are

$$\begin{aligned} \frac{\partial \mathcal{L}(m, m_l)}{\partial m(\{\omega_k\})} &= \sum_{\emptyset \neq A \subseteq \Omega} \frac{\partial \mathcal{L}(m, m_l)}{\partial \mathbb{E}_m(f_A)} \cdot \frac{\partial \mathbb{E}_m(f_A)}{\partial m(\{\omega_k\})} \\ &= -2 \sum_{\emptyset \neq A \subseteq \Omega} [\mathbb{E}_{m_l}(f_A) - \mathbb{E}_m(f_A)] \sum_{j=1}^M \frac{\partial \mathbb{E}_m(f_A)}{\partial \text{Bet } P_m(\omega_j)} \frac{\partial \text{Bet } P_m(\omega_j)}{\partial m(\{\omega_k\})} \\ &= -2 \sum_{\emptyset \neq A \subseteq \Omega} [\mathbb{E}_{m_l}(f_A) - \mathbb{E}_m(f_A)] \sum_{j=1}^M \tilde{u}_{A,j} \left(\delta_{kj} - \frac{1}{M} \right), \end{aligned} \tag{16}$$

where $\delta_{kj} = 1$ if $k = j$ and $\delta_{kj} = 0$ otherwise.

The derivatives of $m(\{\omega_k\})$ w.r.t p_k^l , η^l , and ξ^l in the DS layer are the same as in Dencœux’s original work [9], and the gradient with respect to all network parameters can be back-propagated from the output layer to the input layer.

4 Experiments

In this section, we present numerical experiments that demonstrate the advantages of the proposed model. The databases and metrics are first introduced in Section 4.1. Precise and imprecise segmentation results are then

reported, respectively, in Sections 4.2 and 4.3. Finally, novelty detection results are presented in Section 4.4.

4.1 Databases and metrics for performance evaluation

4.1.1 Databases

Three benchmark databases were used in the study: Pascal VOC 2011 [17], MIT-scene Parsing [43], and SIFT Flow [34]. These databases were used to train and test the E-FCNs as well as probabilistic FCNs (P-FCNs) for comparison.

The Pascal VOC 2011 database contains 20 object classes in 5034 images, with segmentation masks that indicate the class of each pixel, or label it as “background” if the object does not belong to one of the twenty specified classes. The MIT-scene Parsing and SIFT Flow databases are similar to the Pascal VOC 2011 database but have, respectively, 150 categories in 20K images and 33 classes in 2688 images. The list of classes for the three databases are given in Table 2. Each of the three databases was split into 50% for training/validation and 50% for testing. The validation sets were used to determine hyper-parameters, such as the number of prototypes in each DS layer. The optimal tolerance to imprecision γ can be determined in the same way since it can also be considered as a hyper-parameter.

There is no confidence value associated with the pixel labels in any of the three databases. Thus, we defined soft labels for them. For the Pascal VOC 2011 database, we assigned each pixel in a boundary area a soft label $A \subseteq \Omega$, where A consists of the object classes around the boundary area. Some examples are shown in Fig 3a. For the MIT-scene Parsing and SIFT Flow databases with no identified

Table 2 Lists of classes for the three databases used in this study

Database	Class list
Pascal VOC 2011	background, cat, dog, horse, sheep, train, sofa, aeroplane, bicycle, bird, boat, bottle, bus, car, chair, cow, diningtable, motorbike, person, pottedplant, tv.
MIT-scene parsing	wall, floor, ceiling, bed, cabinet, earth, curtain, water, painting, shelf, house, mirror, rug, armchair, seat, desk, wardrobe, lamp, bathtub, railing, cushion, base, box, column, chest, counter, sink, skyscraper, fireplace, refrigerator, grandstand, path, stairs, runway, case, pool, pillow, screen, bookcase, blind, coffee, toilet, flower, book, hill, bench, countertop, stove, palm, kitchen, computer, swivel, bar, arcade, hovel, towel, light, truck, tower, chandelier, booth, dirt track, apparel, land, bannister, escalator, ottoman, buffet, poster, stage, van, ship, fountain, conveyer, canopy, washer, plaything, swimming, stool, barrel, basket, waterfall, tent, bag, minibike, cradle, oven, ball, food, step, tank, trade, microwave, pot, animal, lake, dishwasher, screen, blanket, sculpture, hood, sconce, vase, traffic, tray, ashcan, fan, pier, screen, plate, monitor, bulletin, shower, radiator, glass, clock, flag, sofa, airplane, building, sky, tree, road, windowpane, grass, sidewalk, person, door, table, mountain, plant, chair, car, sea, field, fence, rock, sign, sand, staircase, river, bridge, boat, bus, awning, streetlight, tv, pole, bottle, minibike, bicycle.
SIFT Flow	balcony, crosswalk, desert, moon, sun, window, awning, bird, boat, bridge, building, bus, car, cow, door, fence, field, grass, mountain, person, plant, pole, river, road, rock, sand, sea, sidewalk, sign, sky, staircase, streetlight, tree.

Classes in bold characters are included in two or three databases. Classes with close meanings, such as “minibike” and “motorbike”, are considered as identical

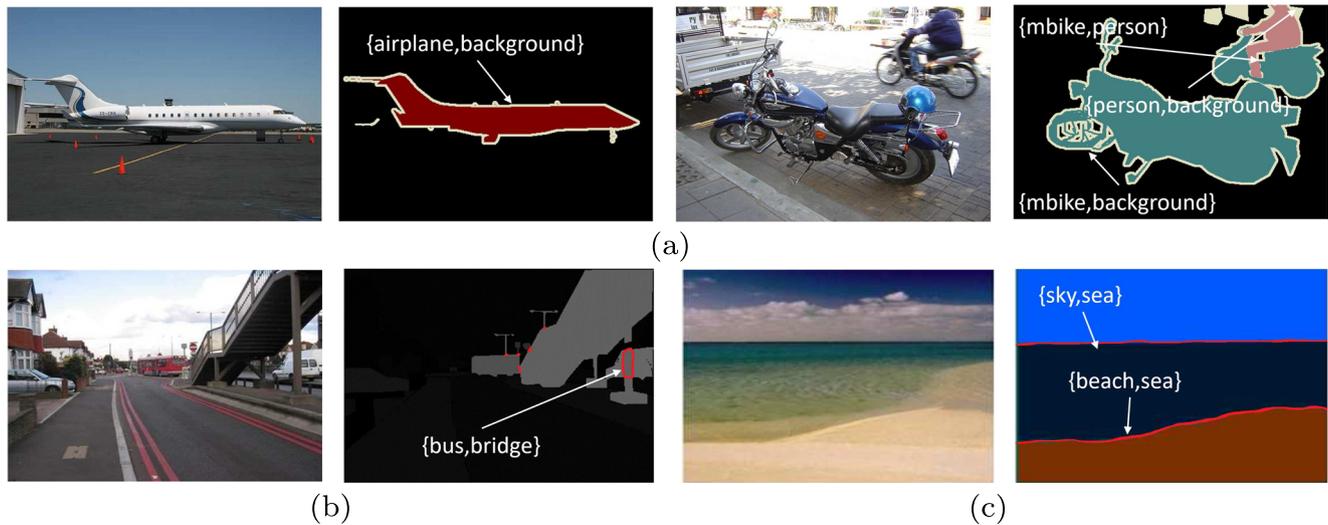


Fig. 3 Segmentation masks with soft labels: **a** Pascal VOC 2011, **b** MIT-scene Parsing, and **c** SIFT Flow

boundary areas, we assigned soft labels to the pixels situated between every two objects, as shown in Fig. 3b and c.

A semantic segmentation model should not only be accurate for the classes in the learning set, but it should also be able to detect objects belonging to classes that are not included in the learning set. To evaluate this novelty detection capacity, we mixed the three databases: for example, an FCN model trained using the Pascal VOC 2011 database was tested on the other two databases.

4.1.2 Metrics

We used three metrics for the performance evaluation of semantic segmentation: pixel utility (PU), utility of intersection over union (UIoU), and expected calibration error (ECE).

Pixel utility. For an image with T pixels, the *pixel utility* is defined as

$$PU = \frac{1}{|T|} \sum_{i=1}^{|T|} \tilde{u}_{A(i), A_*(i)} \quad (17)$$

where $A_*(i)$ is the label of pixel i , $A(i)$ is the selected set of classes for pixel i determined from (14), and using the notations introduced in Section 3.2, $\tilde{u}_{A(i), A_*(i)}$ is the utility of assigning pixel i to subset $A(i) \subseteq \Omega$ when its label is $A_*(i)$. Thus, PU is the same as pixel accuracy when only considering precise assignments and precise labels. To consider soft labels, the utility matrix \tilde{U} defined in Section 3.2 should be extended to a matrix \tilde{U}' of size $(2^M - 1) \times (2^M - 1)$ with general term \tilde{u}_{A, A_*} defined as the utility of assigning a pixel to subset $A \subseteq \Omega$ when its label is A_* , with $|A_*| \geq 1$. Soft label A_* means that we only know the true class of a pixel is in set A_* , and nothing more. To

define the utility \tilde{u}_{A, A_*} , we first compute the average of the utilities of selecting subset A when the true class is in A_* as

$$\bar{u}_{A, A_*} = \frac{1}{|A_*|} \sum_{\omega_k \in A_*} \tilde{u}_{A, k}, \quad (18a)$$

where $\tilde{u}_{A, k}$ is the utility of selecting subset A when the true class is ω_k , and we normalize this average utility to ensure that $\tilde{u}_{A_*, A_*} = 1$:

$$\tilde{u}_{A, A_*} = \frac{\bar{u}_{A, A_*}}{\bar{u}_{A_*, A_*}}. \quad (18b)$$

Example 2 Table 3 shows an example of the utility matrix considering soft labels, which is extended from Example 1. The last four columns correspond to the utility matrix for soft labels. An act achieves utility 1 only if $A = A_*$, 0 if $A \cap A_* = \emptyset$, and a value between 0 and 1 if $A \neq A_*$ and $A \cap A_* \neq \emptyset$.

Utility of intersection over union. The segmentation performance was also evaluated by the *utility of intersection over union* (UIoU) defined as

$$UIoU = \frac{1}{2^{|\Omega|} - 1} \sum_{B \subseteq \Omega} \frac{\sum_{i \in G^B \cap P^B} \tilde{u}_{A(i), B}}{|G^B \cup P^B|}, \quad (19)$$

where $P^B = \{i : A(i) \cap B \neq \emptyset\}$ is the predicted area containing pixels assigned to a set of classes that intersect B , and $G^B = \{i : A_*(i) = B\}$ is the ground truth area composed of pixels with label B . Thus, in the special case of precise segmentation with only precise labels, UIoU boils down to *intersection over union*, a widely used metric for semantic segmentation [22, 25, 28].

Table 3 Utility matrix considering soft labels with $\gamma = 0.8$

	Label							
	ω_1	ω_2	ω_3	$\{\omega_1, \omega_2\}$	$\{\omega_1, \omega_3\}$	$\{\omega_2, \omega_3\}$	Ω	
Act	$f_{\{\omega_1\}}$	1	0	0	0.625	0.625	0	0.489
	$f_{\{\omega_2\}}$	0	1	0	0.625	0	0.625	0.489
	$f_{\{\omega_3\}}$	0	0	1	0	0.625	0.625	0.489
	$f_{\{\omega_1, \omega_2\}}$	0.8	0.8	0	1	0.5	0.5	0.782
	$f_{\{\omega_1, \omega_3\}}$	0.8	0	0.8	0.5	1	0.5	0.782
	$f_{\{\omega_2, \omega_3\}}$	0	0.8	0.8	0.5	0.5	1	0.782
	f_Ω	0.682	0.682	0.682	0.853	0.853	0.853	1

Expected calibration error. In decision systems, a neural network should not only be accurate, but it should also indicate when it is likely to be incorrect. Thus, the confidence of an E-FCN should be *calibrated*. To characterize this property, we extend the *expected calibration error* (ECE) introduced in [20] as follows. We define the *prediction confidence* of pixel i as

$$co(i) = \text{Bet } P_i(A_*(i)) = \sum_{\omega_j \in A_*(i)} \text{Bet } P_i(\{\omega_j\}), \quad (20)$$

where $\text{Bet } P_i$ is the predicted pignistic probability measure for pixel i . Let I_q be the set of pixels whose prediction confidence lies in the interval $(\frac{q-1}{Q}, \frac{q}{Q}]$, $q = 1, \dots, Q$. The average utility and confidence of I_q are defined, respectively, as

$$au(I_q) = \frac{1}{|I_q|} \sum_{i \in I_q} \tilde{u}_{A(i), A_*(i)}, \quad (21a)$$

and

$$co(I_q) = \frac{1}{|I_q|} \sum_{i \in I_q} co(i). \quad (21b)$$

We consider that the classifier is well calibrated if $co(I_q) \approx au(I_q)$ for all q , and we define the ECE as

$$ECE = \frac{\sum_{q=1}^Q |I_q| \times |co(I_q) - au(I_q)|}{\sum_{q'=1}^Q |I_{q'}|} \quad (22)$$

When only considering precise acts and labels, ECE defined by (22) boils down to the original definition in [20].

4.2 Precise segmentation results

In precise segmentation, each pixel of an image is assigned to exactly one class, the set of acts being defined as $\mathcal{F} = \{f_{\omega_1}, \dots, f_{\omega_M}\}$. Three databases without soft labels mentioned in Section 4.1.1 were used to train and test the E-FCNs and probabilistic FCNs (P-FCNs). The metrics

defined in Section 4.1.2 with the utility matrix U equal to the identity matrix were used for performance assessment.

In the experiment with each database, three widely used encoder-decoder architectures were combined with the DS and utility layers, as shown in Table 4. All encoder-decoder architectures in Table 4 have the same encoder part, which consists of four stages and two convolutional

Table 4 Performance evaluation of precise segmentation: (a) Pascal VOC 2011, (b) MIT-scene Parsing, and (c) SIFT Flow

	PU	UIoU
(a)		
P-FCN-32s [25]	0.8912 ± 0.0019	0.5941 ± 0.0033
P-FCN-16s [25]	0.9001 ± 0.0015	0.6243 ± 0.0025
P-FCN-8s [25]	0.9033 ± 0.0017	0.6269 ± 0.0021
E-FCN-32s	0.8973 ± 0.0021	0.6128 ± 0.0024
E-FCN-16s	0.9045 ± 0.0014	0.6304 ± 0.0019
E-FCN-8s	0.9074 ± 0.0015	0.6337 ± 0.0020
(b)		
P-FCN-16s [25]	0.7009 ± 0.0030	0.289 ± 0.0051
P-FCN-8s [25]	0.7128 ± 0.0024	0.294 ± 0.0048
P-FCN-SegNet [1]	0.7153 ± 0.0023	0.305 ± 0.0042
E-FCN-16s	0.7090 ± 0.0026	0.292 ± 0.0048
E-FCN-8s	0.7148 ± 0.0025	0.296 ± 0.0046
E-FCN-SegNet	0.7167 ± 0.0026	0.330 ± 0.0043
(c)		
P-FCN-16s [25]	0.8489 ± 0.0034	0.3922 ± 0.0047
P-FCN-8s [25]	0.8525 ± 0.0032	0.3948 ± 0.0042
P-FCN-DilatedVGG [3]	0.8643 ± 0.0036	0.4168 ± 0.0043
E-FCN-16s	0.8521 ± 0.0030	0.3937 ± 0.0042
E-FCN-8s	0.8528 ± 0.0031	0.3961 ± 0.0040
E-FCN-DilatedVGG	0.8649 ± 0.0035	0.4182 ± 0.0038

P-FCN and E-FCN are, respectively, probabilistic and evidential FCNs. The rests of the notations, such as “-32s” and “-16s”, stand for different encoder-decoder architectures. The results are in form of “mean value ± standard deviation”. The best results for each encoder-decoder architecture are shown in bold.

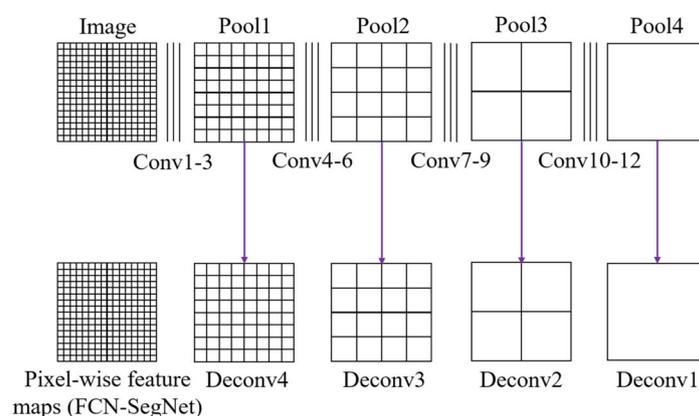
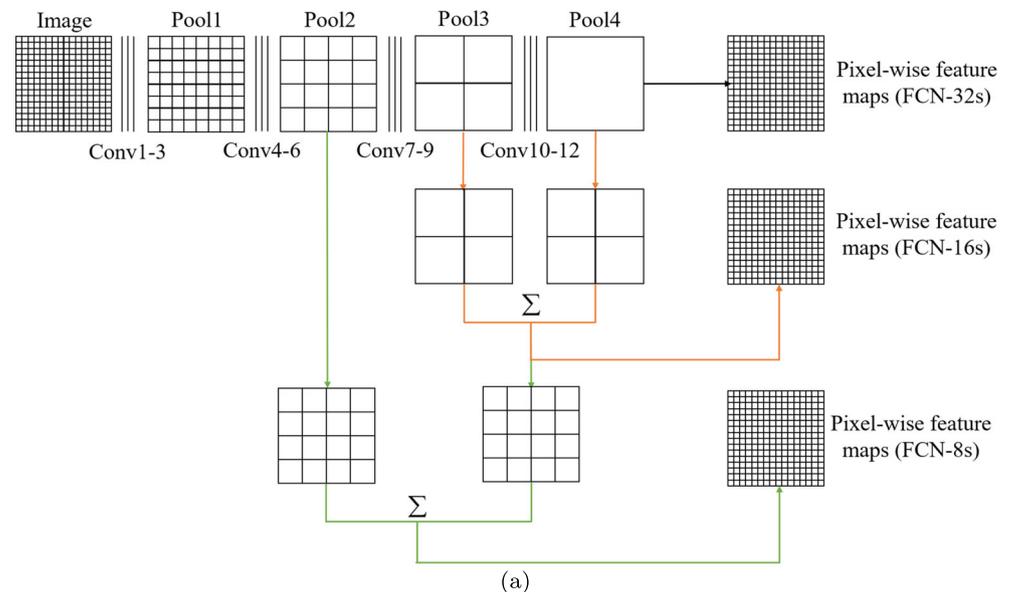
layers with 3×3 kernels. Each stage is made up of three convolutional layers with 3×3 kernels and a max-pooling layer with a 2×2 non-overlapping window. Figure 4a illustrates the differences between the FCN-32s, FCN-16s, and FCN-8s architectures in their decoder parts with a deconvolutional layer. The FCN-SegNet architecture uses four deconvolutional layers to upsample the sparse feature maps from the end of the encoder part, as well as the feature maps from the corresponding pooling layers based on pooling indices [1], as shown in Figure 4b. The FCN-DilatedVGG architecture is the same as FCN-SegNet except that it adds a fully connected conditional random field at the end of the last deconvolutional layer [3]. The numbers P of feature maps for the Pascal, MIT and SIFT databases were, respectively, 31, 128 and 64. The numbers n of prototypes in the DS layer for these three databases were set, respectively, to 75, 300 and 95.

The DS and utility layers slightly improve the accuracy of precise assignments performed by FCN models, even

though the performance of FCN models on precise segmentation mainly depends on the encoder-decoder architecture. Table 4a presents the results of PU and UIoU for the Pascal VOC database. E-FCNs achieved higher PU and UIoU than P-FCNs with the same encoder-decoder architecture, which shows the E-FCNs outperform the P-FCNs for precise segmentation. Similar improvements can also be found in the MIT-scene Parsing and SIFT Flow databases as shown, respectively, in Table 4 band c.

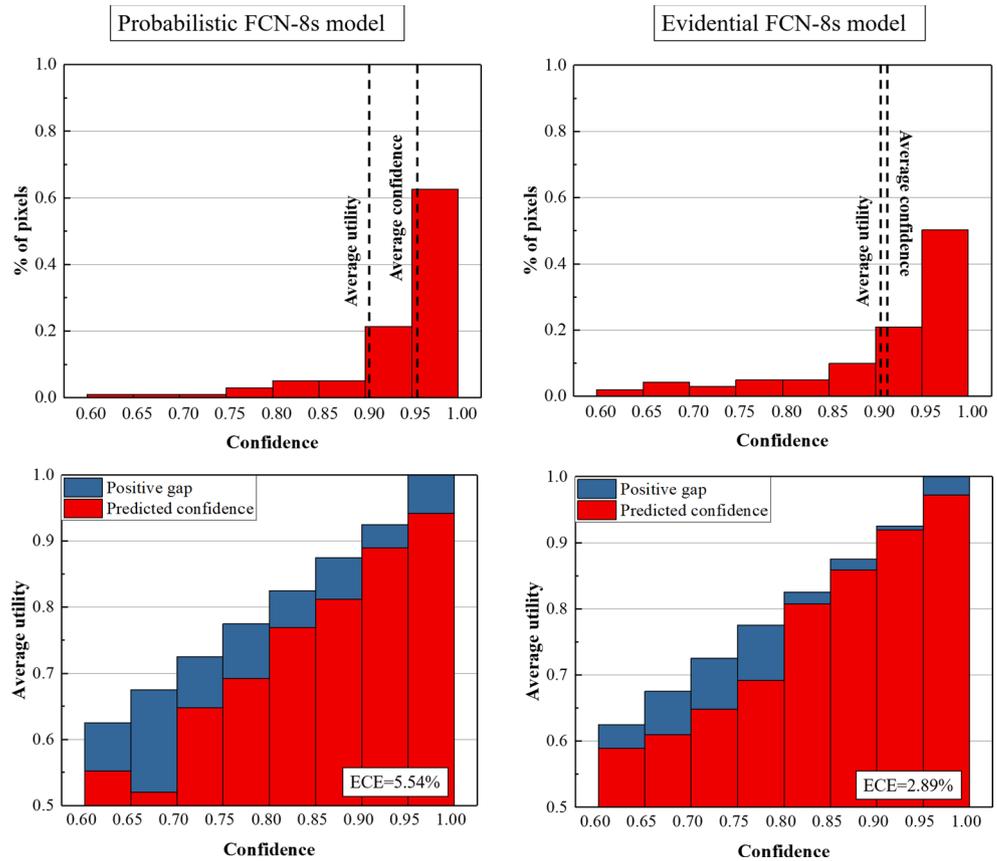
The use of DS and utility layers also makes the FCN models better calibrated. Figure 5 presents a visual calibration representation of the FCN-8s models in the Pascal VOC database. The top row shows the pixel distribution of prediction confidence (21b) as histograms. The average confidence of the E-FCN-8s model closely matches its average pixel utility, while the average confidence of the P-FCN-8s model is substantially higher than its average pixel utility. This is further illustrated in the pixel utility diagrams (bottom row of Fig. 5), which

Fig. 4 Illustration of the encoder-decoder architectures used in this paper. Pooling layers are represented as grids that show relatively sparse information. Intermediate convolution layers are omitted. **a** The FCN-32s, FCN-16s, FCN-8s architectures are used to combine sparse and high-layer information with dense and low-layer information for upsampling. Black arrow: the deconvolutional layer in FCN-32s directly upsamples the outputs of Pool 4 to pixel-wise feature maps; orange arrows: the deconvolutional layer in FCN-16s combines outputs from Pool 3 and 4, lets the net predict finer details, while retaining high-level semantic information; green arrows: the deconvolutional layer in FCN-8s acquire additional feature maps from Pool 2 to provide further precision; **b** The FCN-SegNet architecture uses four deconvolutional layers to upsample the sparse feature maps from the end of the encoder part, as well as the feature maps from the corresponding pooling layers based on pooling indices (purple arrows)



(b)

Fig. 5 Pixel confidence distributions (top) and pixel utility histograms (bottom) for P-FCN-8s (left) and E-FCN-8s (right) on the Pascal VOC database



show pixel utility as a function of confidence. The E-FCN-8s model is well calibrated since its confidence in each bin approximates the expected average utility, whereas the predicted utility of the P-FCN-8s model does not match its confidence. As a consequence, the E-FCN-8s model achieves a smaller ECE than the probabilistic one. The effect of the DS and utility layers on the calibration can also be found in the FCN-SegNet and FCN-DialtedVGG models on the MIT-scene Parsing and SIFT Flow databases as shown, respectively, in Figs. 6 and 7.

4.3 Imprecise segmentation results

In imprecise segmentation, each pixel of an image is assigned to a non-empty subset A of Ω ; the set of acts is $\mathcal{F} = \{f_A, A \in 2^\Omega \setminus \{\emptyset\}\}$, or a subset thereof. Here we only considered acts f_A such that A is a singleton, Ω or one of the soft labels in the training set. For performance evaluation, we used the metrics and the three databases described in Section 4.1. For each database, the segmentation masks with and without soft labels were used to train different FCN models. The same encoder-decoder

architectures used for precise segmentation in Section 4.2 were combined with the DS and utility layers.

Figure 8 displays the test results according to PU and UIoU for imprecise segmentation of the Pascal VOC database. For a wide range of imprecision tolerance degree γ , the E-FCN models reach higher PU and UIoU values than those obtained by the P-FCN models; this is due to the fact that the E-FCN models tend to assign ambiguous pixels to multi-class sets, instead of making precise decisions. Such imprecise assignments avoid pixel misclassification in case of high uncertainty, especially when feature vectors from an encoder-decoder architecture do not contain sufficient information to predict a precise class, and multiple classes have similar probabilities. Figure 9 shows the pixel confidence distributions for the FCN models with $\gamma = 0.8$. We can see that the average confidences of the E-FCN models are smaller than those of the P-FCN models. This observation suggests that the E-FCN models make cautious decisions for ambiguous pixels by assigning them to multi-class sets, rather than classifying them arbitrarily into a single class. The E-FCN models are thus better calibrated than those based on P-FCN, which can be over-

Fig. 6 Pixel confidence distributions (top) and pixel utility histograms (bottom) for P-FCN-SegNet (left) and E-FCN-SegNet (right) on the MIT-scene Parsing database

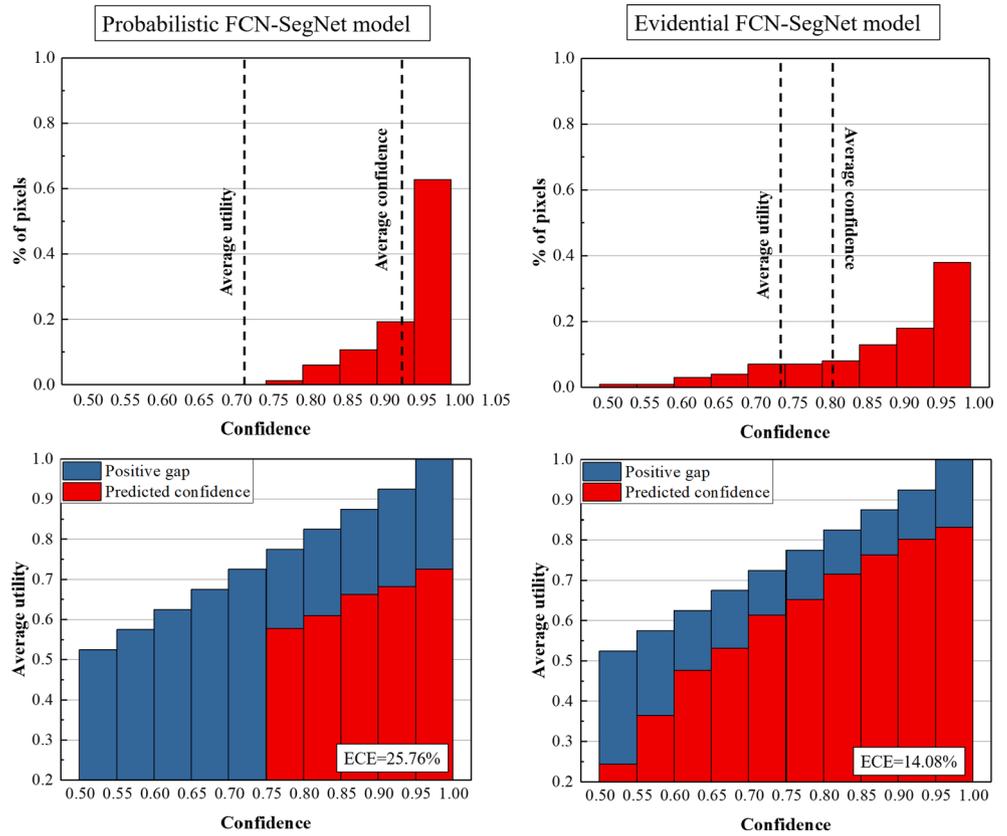
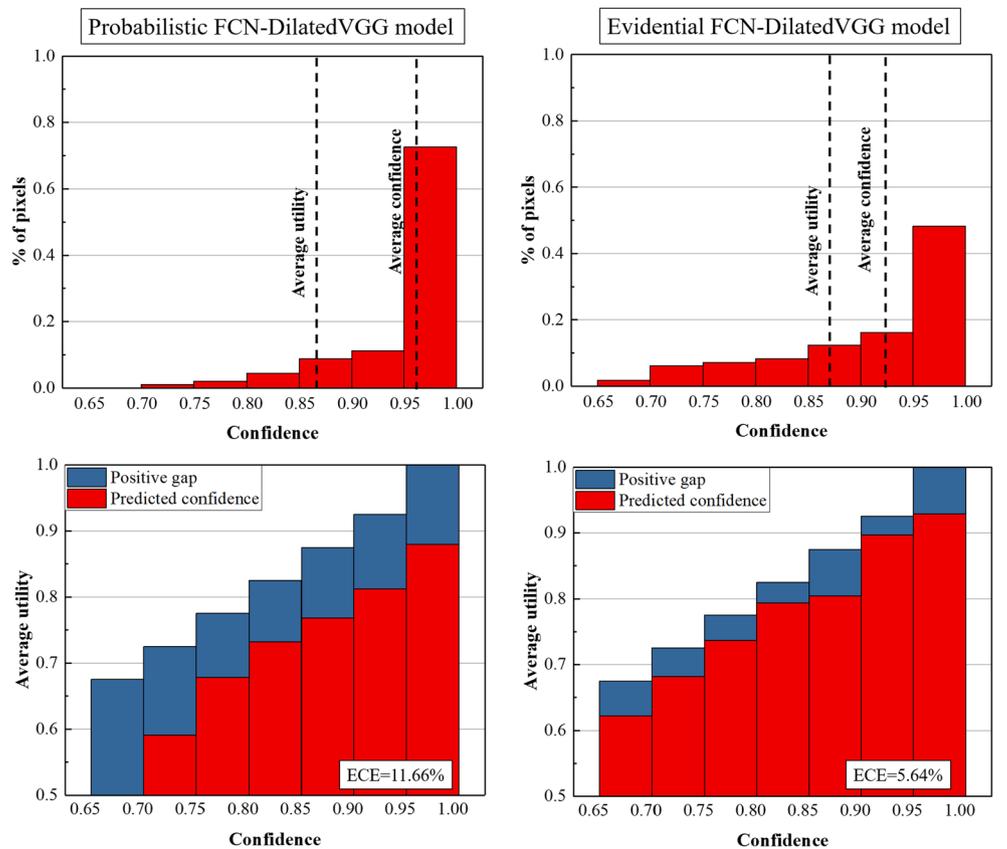


Fig. 7 Pixel confidence distributions (top) and pixel utility histograms (bottom) for P-FCN-DilatedVGG (left) and E-FCN-DilatedVGG (right) on the SIFT Flow database



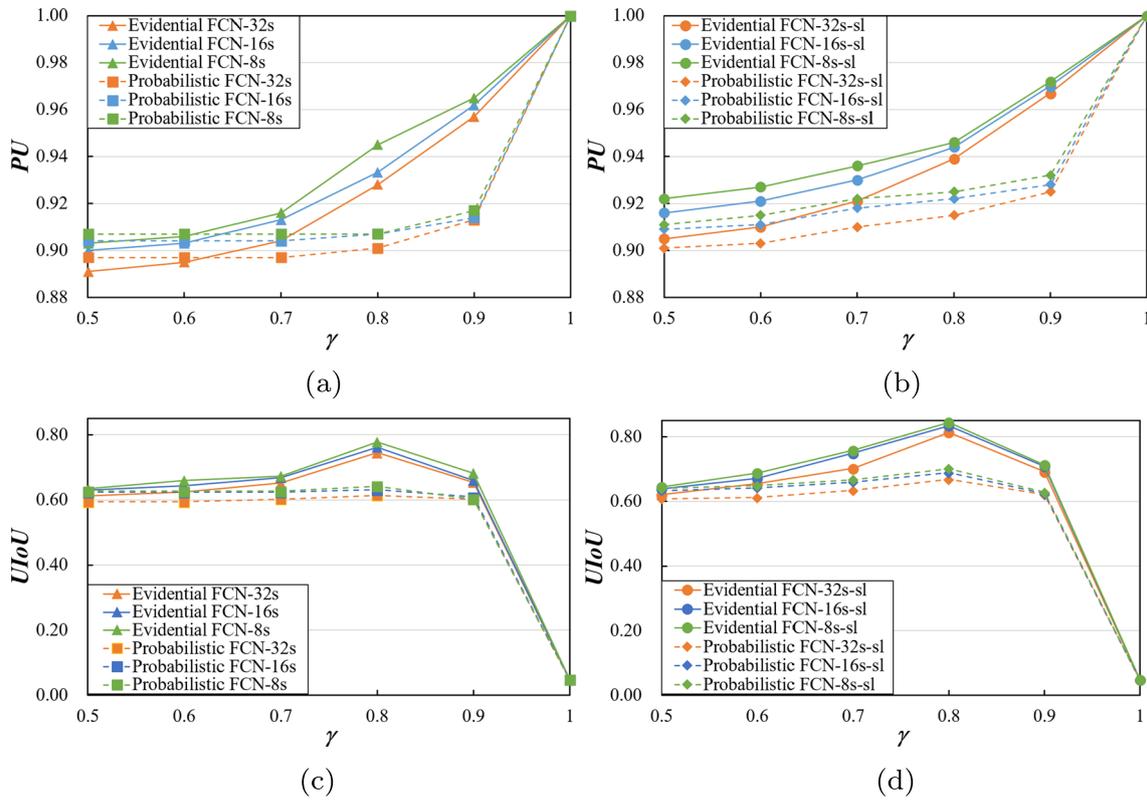
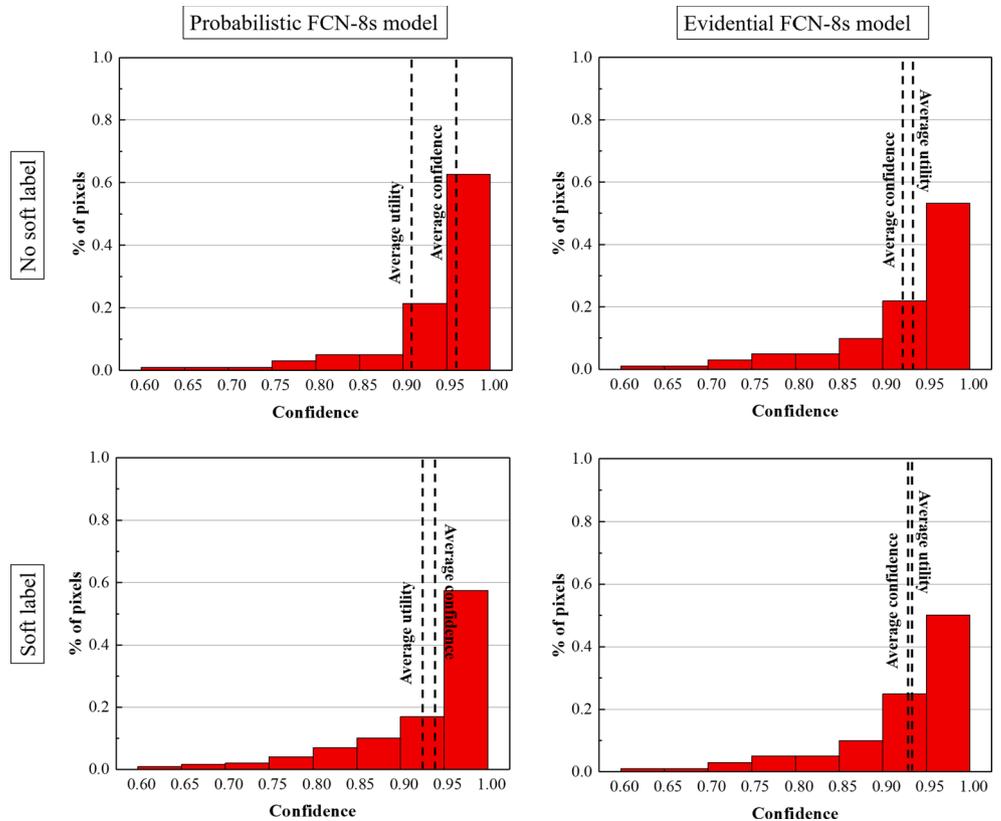


Fig. 8 Testing PU and UIoU vs. γ on the Pascal VOC database. The first and second columns are the models trained with/without soft labels, respectively

Fig. 9 Pixel confidence distributions for the P-FCN-8s (left) and E-FCN-8s (right) models on the Pascal VOC 2011 database without (top)/with (bottom) soft labels



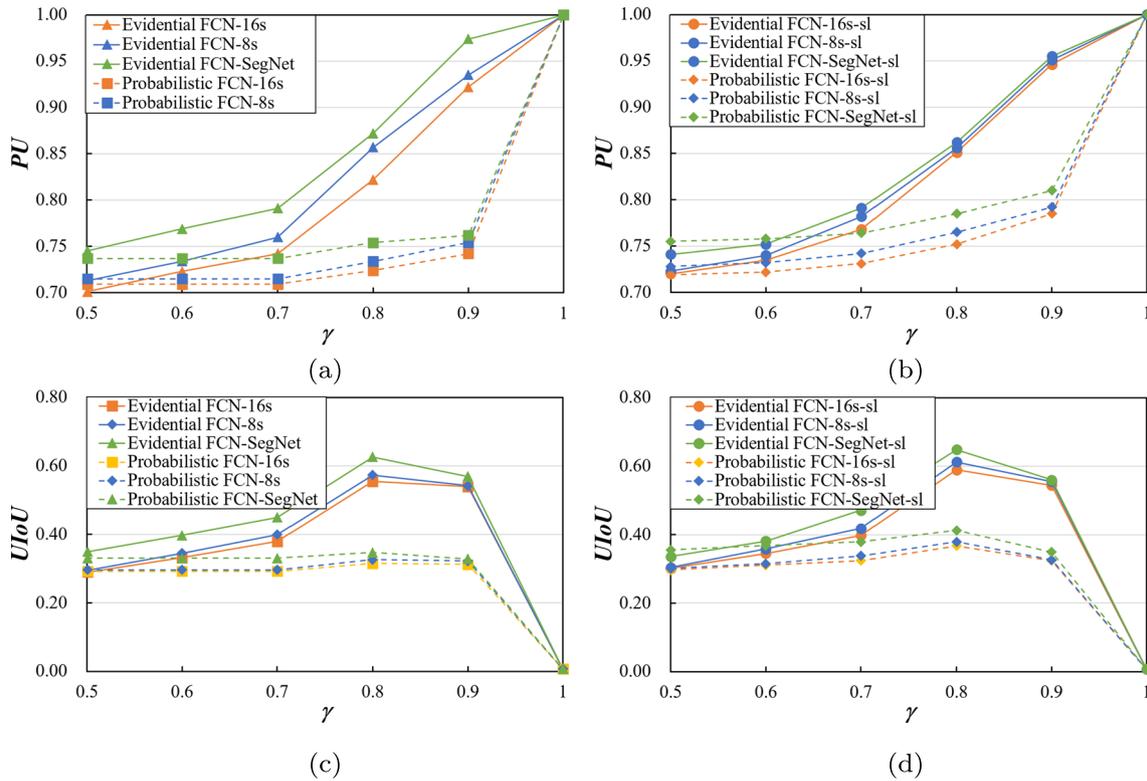
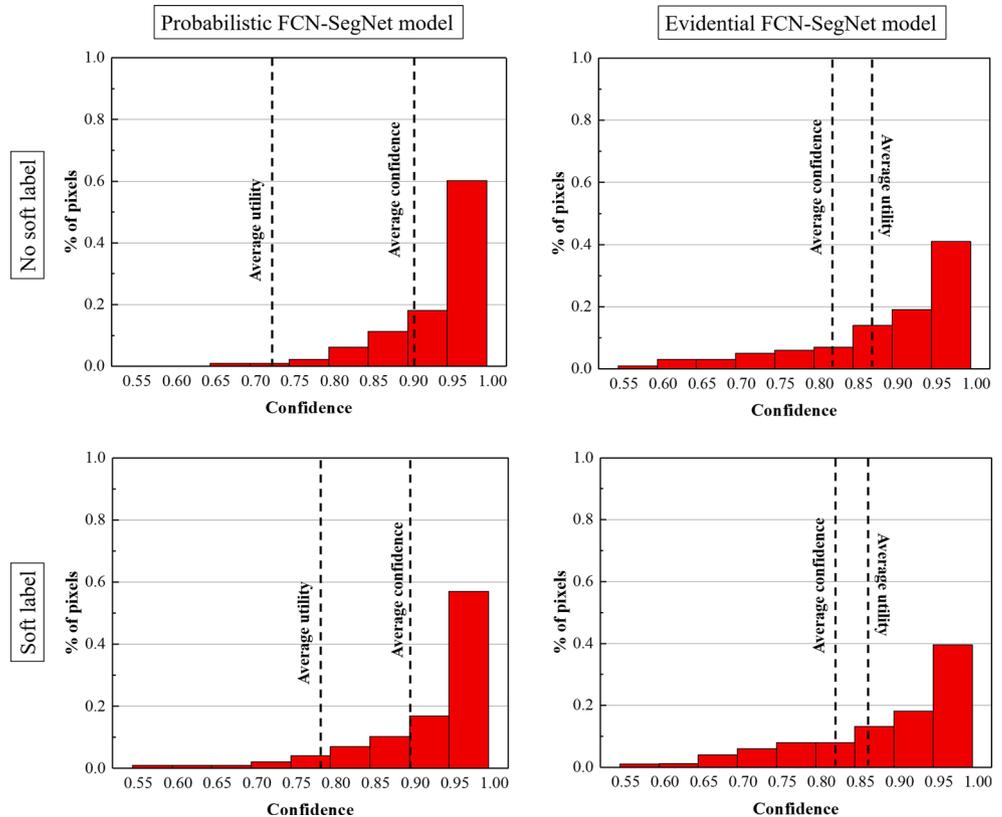


Fig. 10 Testing PU and UIoU vs. γ on the MIT-scene Parsing database. The first and second columns are the models trained with/without soft labels, respectively

Fig. 11 Pixel rate histograms for the P-FCN-SegNet (left) and E-FCN-SegNet (right) models on the MIT-scene Parsing database without (top)/with (bottom) soft labels



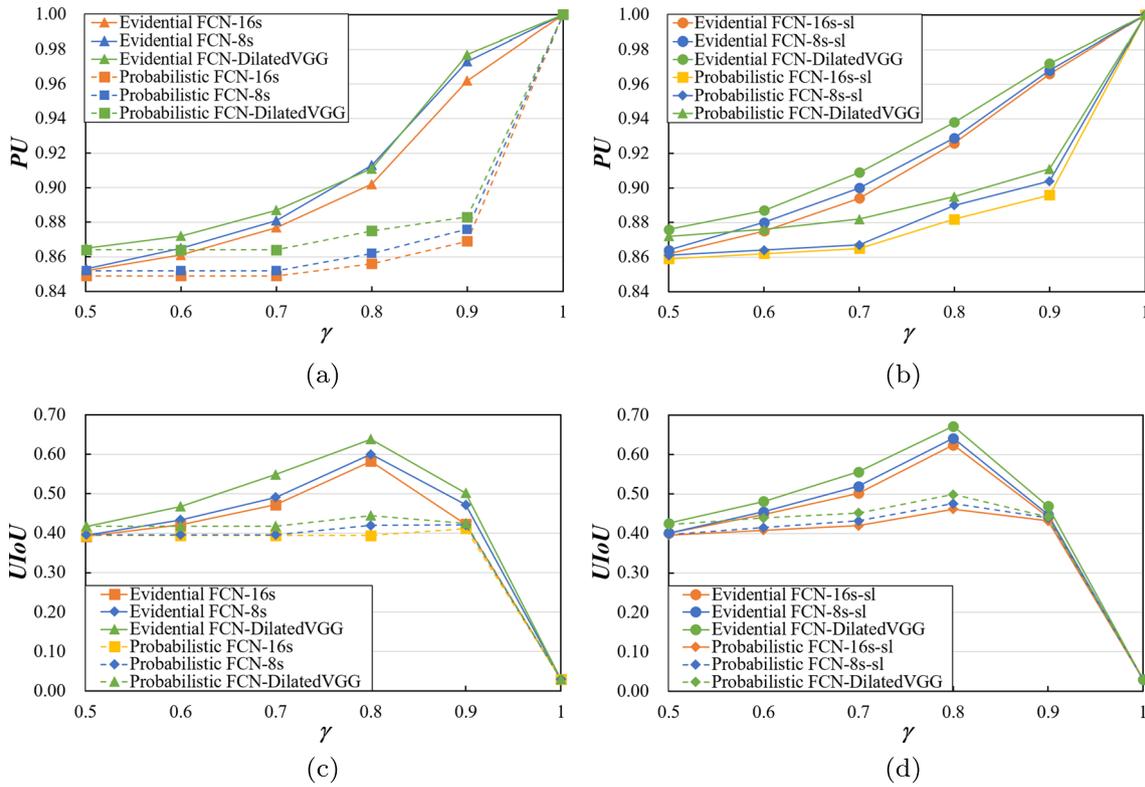
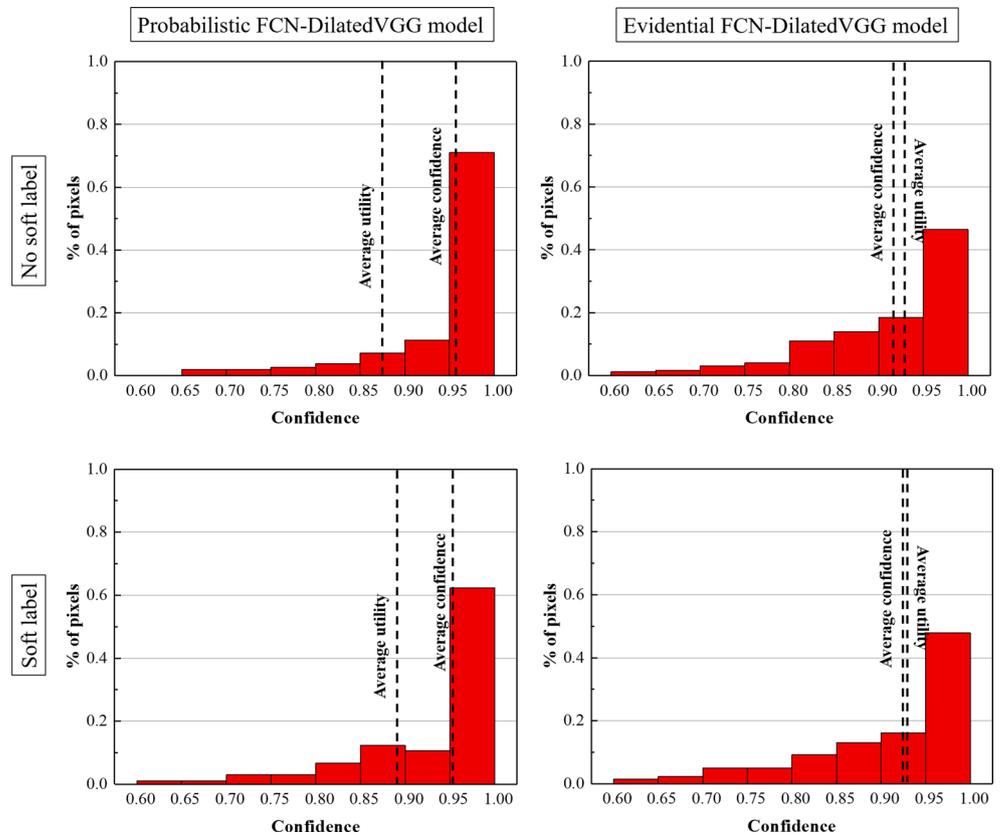


Fig. 12 Testing PU and UIoU vs. γ on the SIFT Flow database. The first and second columns are the models trained with/without soft labels, respectively

Fig. 13 Pixel rate histograms for the P-FCN-DilatedVGG (left) and E-FCN-DilatedVGG (right) models on the SIFT Flow database without (top)/with (bottom) soft labels



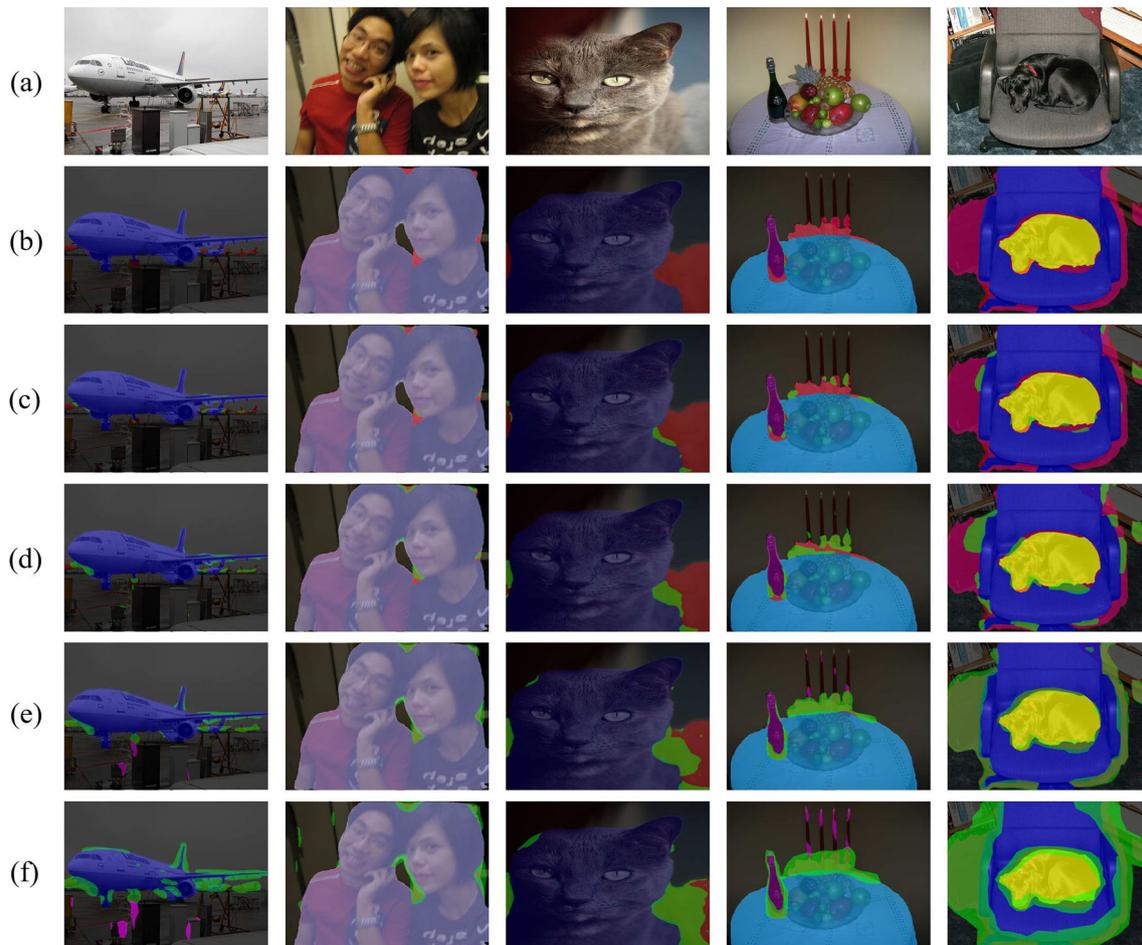


Fig. 14 Segmentation examples from the Pascal VOC 2011 database: **a** Original image, **b** Precise segmentation, **c** Imprecise segmentation with $\gamma = 0.6$, **d** Imprecise segmentation with $\gamma = 0.7$, **e** Imprecise segmentation with $\gamma = 0.8$, and **f** Imprecise segmentation with

$\gamma = 0.9$. Red masks are pixels incorrectly classified in the precise segmentation; green masks are pixels assigned to multi-class sets except set Ω ; pink masks are pixels assigned to set Ω ; other masks are pixels assigned to correct single-class sets

confident. Similar results are observed with the MIT-scene Parsing (Figs. 10 and 11) and SIFT Flow (Figs. 12 and 13) databases. We can thus conclude the DS and utility layers improve the performance of the FCN models in imprecise segmentation tasks by allowing the assignment of some ambiguous pixels to multi-class sets.

In Figs. 8, 10 and 12, we can see that the value of UIoU first increases and then decreases when γ increases from 0.5 to 1. To explain this behavior, Fig. 14 shows some segmentation examples generated by the E-FCN-8s model trained on the Pascal VOC database with soft labels. The first and second columns of Fig. 14 contain, respectively, the original images and their precise segmentation predicted masks, while the third to sixth columns show the imprecise segmentation results for values of γ ranging from 0.6 to 0.9. When γ increases from 0.5 to 0.8, the majority of the green masks (the areas whose pixels are assigned to multi-class sets) tends to cover the red masks (the areas whose pixels are incorrectly classified in the precise segmentation). This

observation can be explained by the fact that, in (19), the increase in the utility of the intersection between predicted and labeled areas is larger than the increase in the union between the two areas. As a result, UIoU increases when γ increases from 0.5 to 0.8. However, when γ increases from 0.8 to 1.0, the majority of the green masks cover the areas predicted correctly in the precise segmentation, which causes the increase in the utility of intersection to be smaller than the increase in the union areas. This phenomenon leads to the decrease of UIoU when γ is larger than 0.8.

The use of soft labels improves the performance of the FCN models for imprecision segmentation tasks. As shown in Fig. 8, the FCN models trained on the Pascal VOC database with soft labels have larger testing PU and UIoU than the ones without soft labels, which demonstrates the accuracy improvement using soft labels. Additionally, the use of soft labels can also improve the calibration of the FCN models. Figure 15 shows that the ECEs and bin gaps in the E-FCN and P-FCN models are

Fig. 15 Average utility histograms for P-FCN-8s (left) and E-FCN-8s (right) with $\gamma = 0.8$ on the Pascal VOC 2011 database without (top)/with (bottom) soft labels

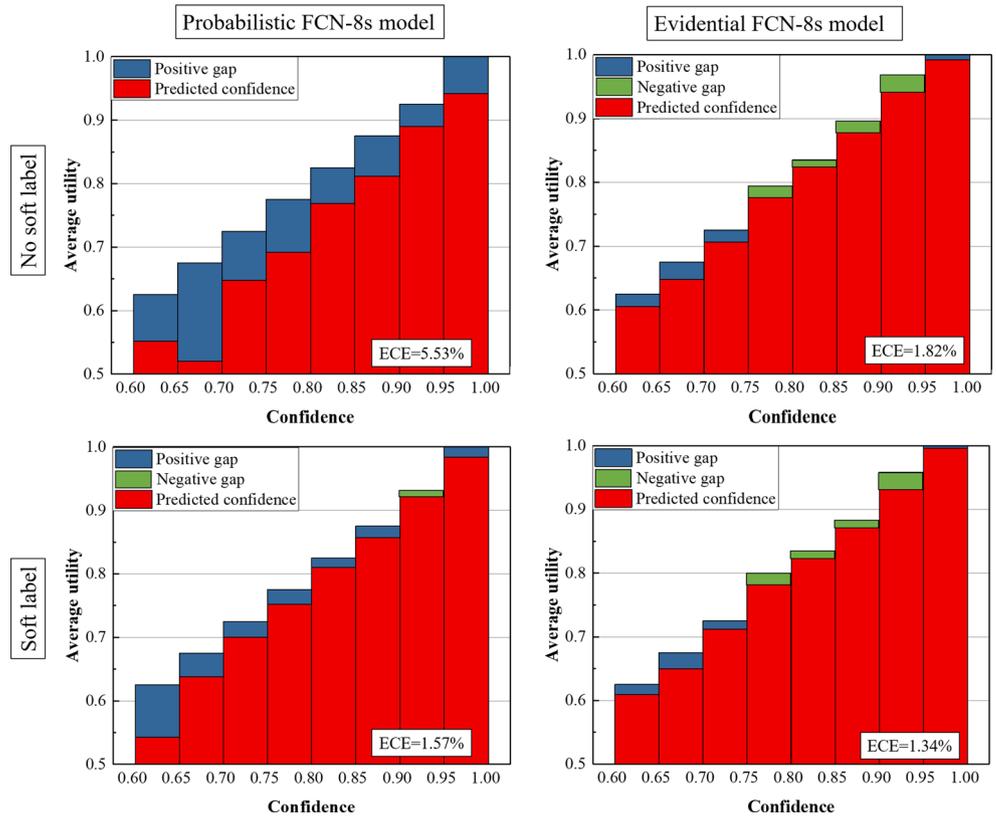


Fig. 16 Average utility histograms for P-FCN-SegNet (left) and E-FCN-SegNet (right) with $\gamma = 0.8$ on the MIT-scene Parsing database without (top)/with (bottom) soft labels

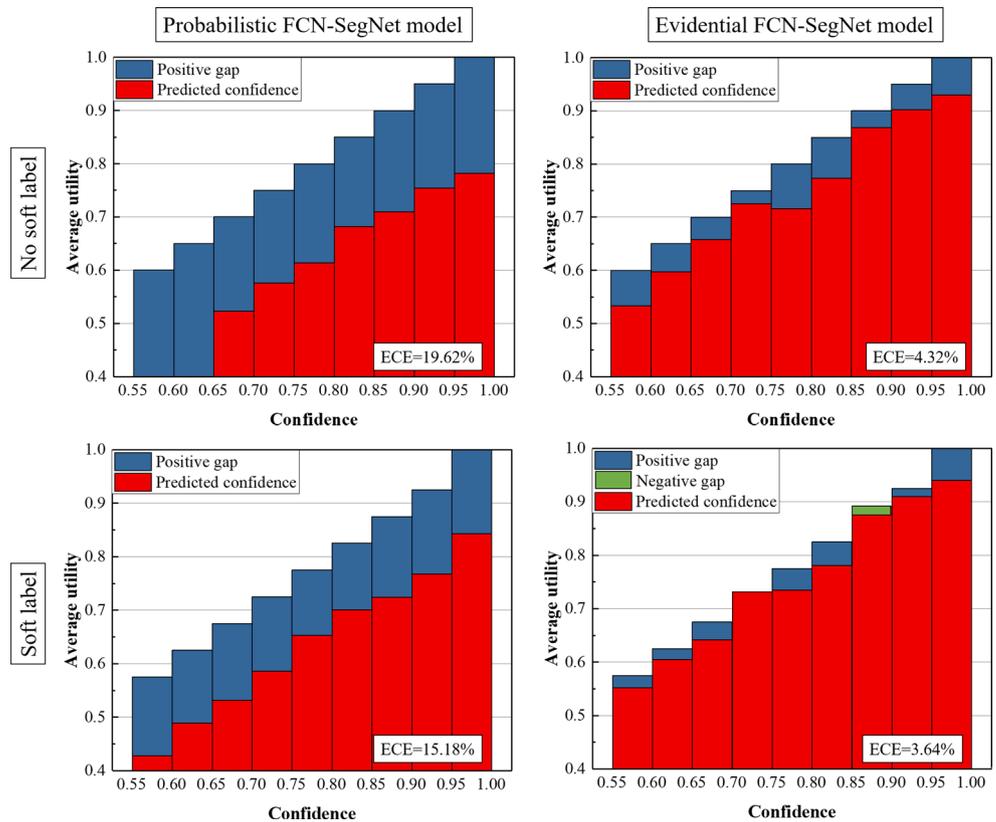


Fig. 17 Average utility histograms for P-FCN-DilatedVGG (left) and E-FCN-DilatedVGG (right) with $\gamma = 0.8$ on the SIFT Flow database without (top)/with (bottom) soft labels

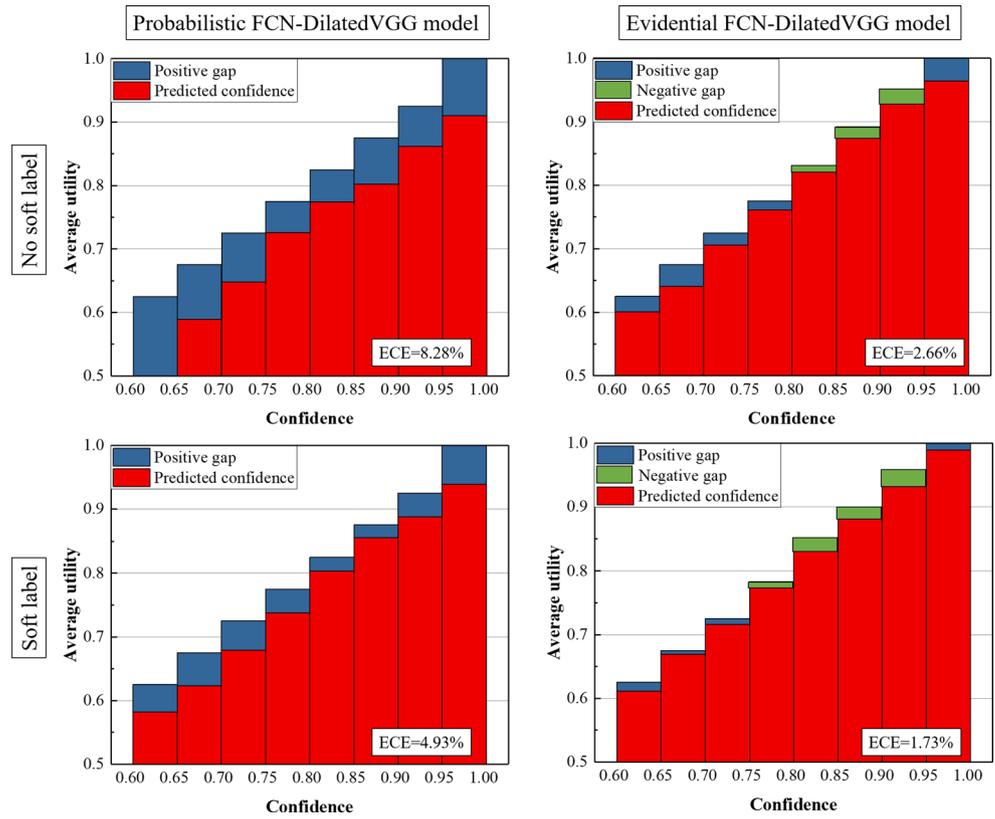


Fig. 18 Proportion of pixels assigned to Ω as a function of γ for novelty detection on the combination of MIT-scene Parsing and SIFT Flow databases (top) and the testing set from the Pascal VOC 2011 database (bottom) when the learning set is from the Pascal VOC database without (left)/with (right) soft labels

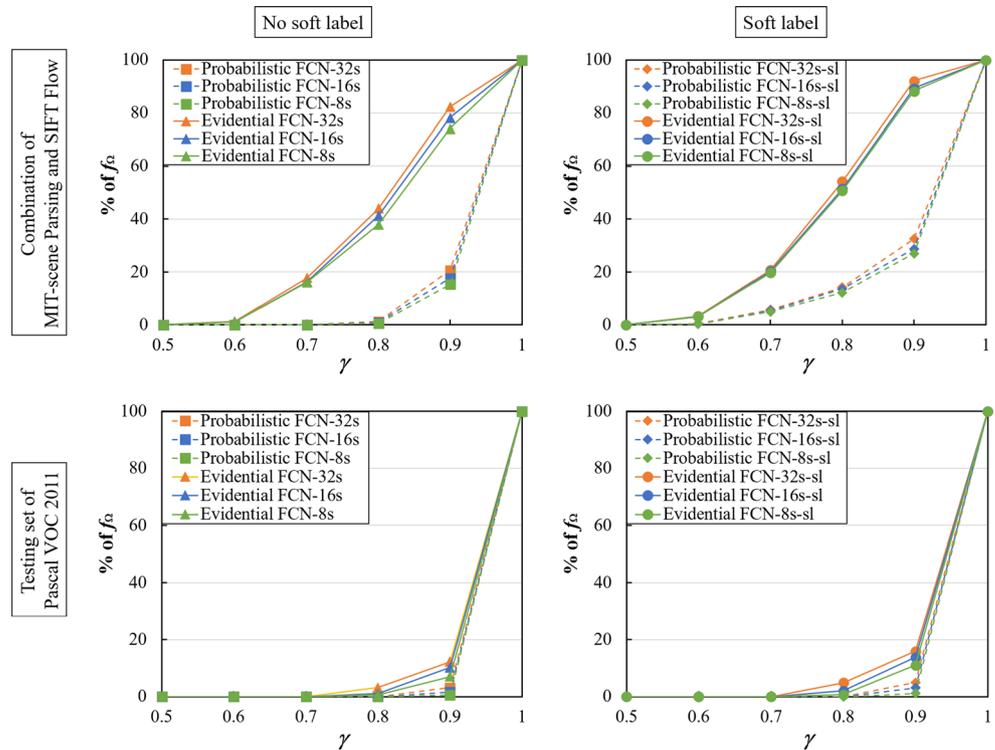
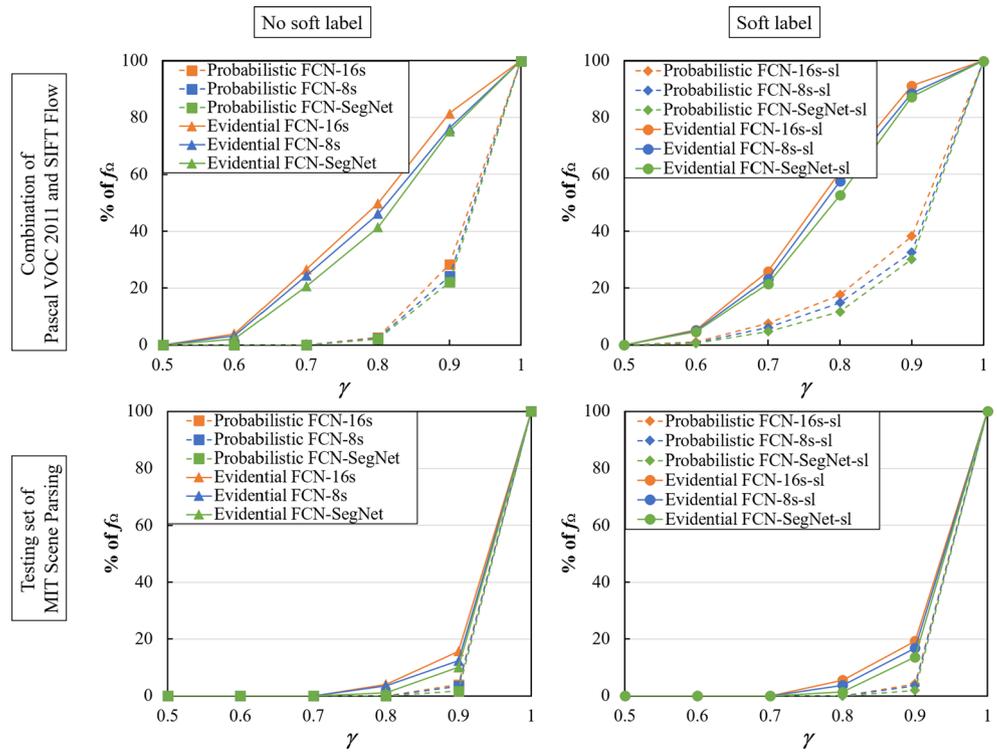


Fig. 19 Proportion of pixels assigned to Ω as a function of γ for novelty detection on the combination of Pascal VOC 2011 and SIFT Flow (top) and the testing set of the MIT-scene Parsing database (bottom) when the learning set is from the Pascal VOC database without (left)/with (right) soft labels



smaller when using the learning set with soft labels. These results demonstrate the feasibility of processing pixels with ambiguous information using soft labels when training FCN

models. An improvement of accuracy and calibration due to learning from soft labels is also observed with the MIT-scene Parsing and SIFT Flow databases, as shown,

Fig. 20 Proportion of pixels assigned to Ω as a function of γ for novelty detection on the combination of Pascal VOC 2011 and MIT-scene Parsing (top) and the testing set of the SIFT Flow database (bottom) when the learning set is from the Pascal VOC database without (left)/with (right) soft labels

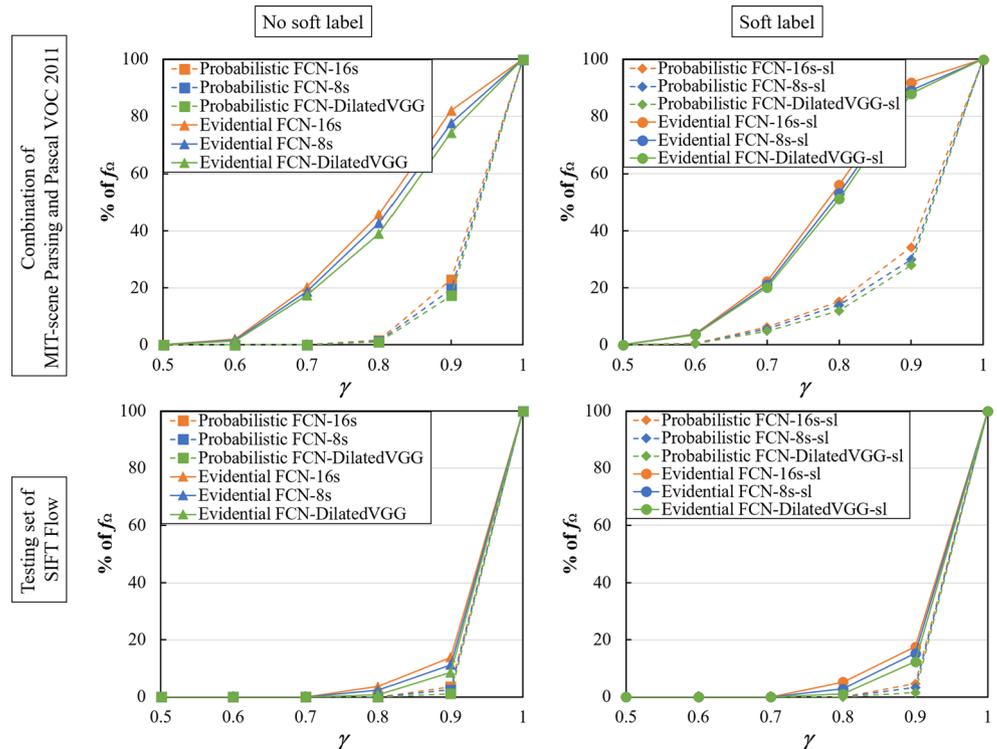
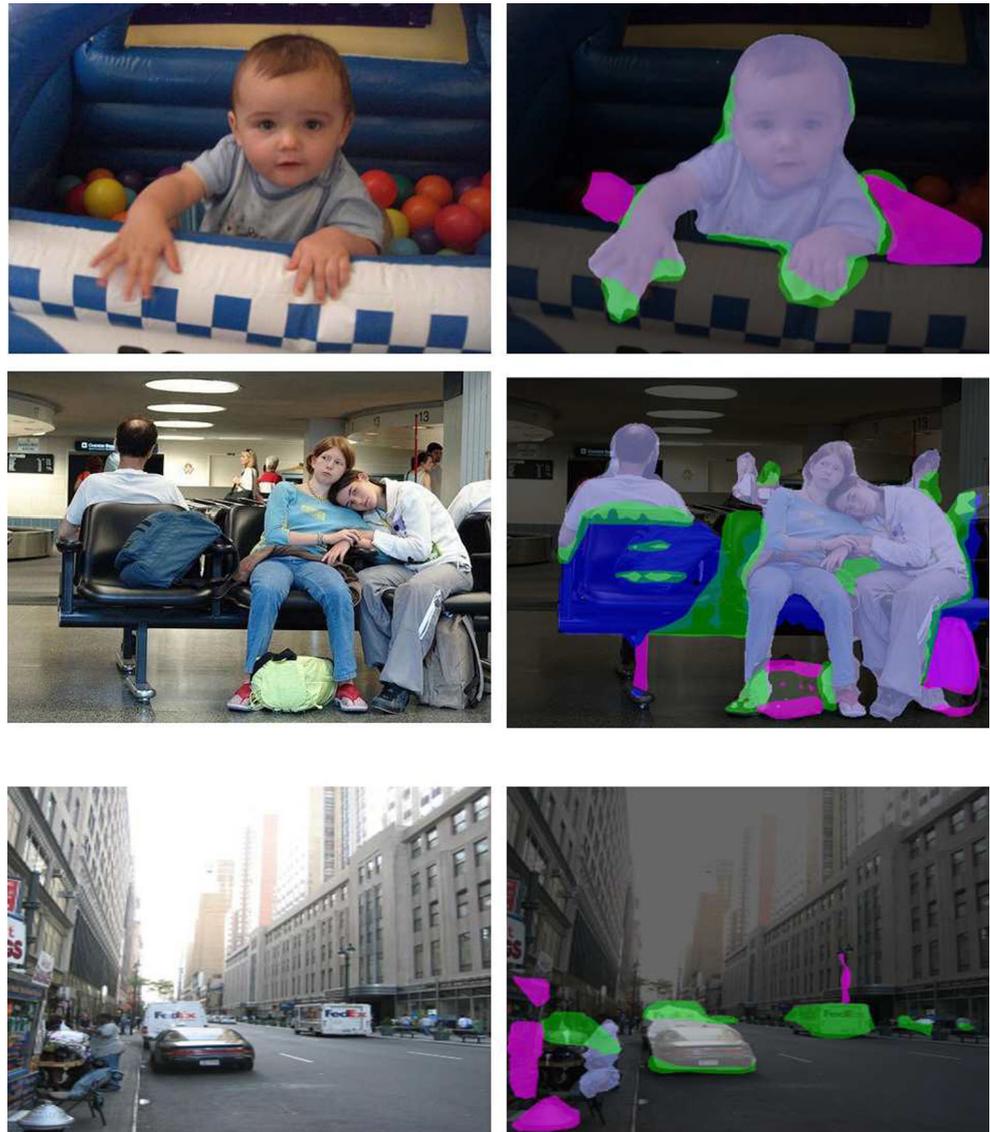


Fig. 21 Examples of novelty detection from the MIT-scene Parsing database and their segmentation masks given by the E-FCN-8s model trained using the Pascal VOC database with soft labels when γ equals 0.8. Red masks are pixels incorrectly assigned in the precise segmentation; green masks are pixels assigned to multi-class sets except set Ω ; pink masks are pixels assigned to set Ω ; other masks are pixels assigned to correct single-class sets



respectively, in Figs. 16 and 17. Therefore, we can conclude that the use of soft labels improves the accuracy and calibration of FCN models.

4.4 Novelty detection results

For novelty detection, a pixel is considered as an outlier or an ambiguous sample if it is assigned to set Ω . Figures 18, 19 and 20 show the results of novelty detection using the E-FCN and P-FCN models when the learning set is extracted, respectively, from the Pascal VOC, MIT-scene Parsing and SIFT Flow databases, and the test set is composed of images from the other two databases. In each testing set composed of two databases, only the pixels whose class is not represented in the corresponding learning set are reported in Figs. 18-20. The E-FCN models assign outliers

and some known-class pixels to set Ω for values of γ between 0.7 and 0.9, while the P-FCN models do not. This observation shows that the E-FCN models are more efficient than the probabilistic ones for rejecting outliers together with ambiguous samples. The proposed architecture thus has the potential to perform novelty detection once given a reasonable value of tolerance to imprecision. However, none of the FCN models performs well when γ is less than 0.7 since these models favor precise decisions.

The E-FCN models tend to reject unknown objects whose features are very different from those of the known objects in the learning set. For example, Fig. 21 shows images from the MIT-scene Parsing database in which pixels representing ‘bag’, ‘street light’ and ‘ball’ objects are rejected by an E-FCN-8s model trained using the Pascal VOC database, which does not contain these objects. As

Table 5 Percentage of pixels from some unknown classes in the MIT-scene Parsing and SIFT Flow databases classified by an E-FCN-8s model trained on the Pascal VOC database into some sets of classes

	bag	True class						
		street light	ball	seat	bench	bed	wall	
Assigned set	Ω	68.4	77.3	75.2	7.8	4.7	15.9	4.9
	{bottle, ...}	21.8	16.3	16.1	48.5	39.7	30.3	0.2
	{chair, ...}	11.3	9.2	8.5	84.7	81.7	58.6	0.3
	{background, ...}	15.2	13.7	11.5	58.7	48.6	46.9	88.0
	Others	4.2	2.4	1.5	2.7	3.5	5.2	3.7

The model was trained with soft labels and $\gamma = 0.8$. For instance, 68.4% of the pixels representing a bag were rejected (i.e., assigned to Ω), and 84.7% of pixels representing a seat were assigned to a set of classes containing the class “chair”

shown in Table 5, 75.2% of the pixels representing a ball in the MIT-scene Parsing and SIFT Flow databases are assigned to Ω , while 16.1% are assigned to a set of classes containing “bottle”. For the “bag” and “street light” classes, these numbers are, respectively, 68.4%/21.8% and 77.3%/16.3%. Some unknown objects are not so easily rejected because of their similarity with known objects. For instance, 84.7% of the pixels representing a seat and 81.7% of pixels representing a bench are assigned to a set of classes containing “chair”, and 88% of “wall” pixels are assigned to a set of classes containing “background”.

We can also observe that the FCN models trained using a leaning set with soft labels reject more outliers than those trained without soft labels, as shown in Figs. 18, 19 and 20. This is because the use of soft labels makes the FCN models more cautious and better calibrated, as discussed in Section 4.3. More precisely, for ambiguous pixels or outliers, the output mass functions of the FCN models trained with soft labels are more uniform than those computed by FCN models trained without soft labels. As a result, ambiguous pixels and outliers are more easily assigned to set Ω . We can thus conclude that soft labels have the potential to enhance novelty detection performance.

5 Conclusions

In this paper, we have presented a new approach based on the combination of DS theory and FCN for image semantic segmentation. In the proposed model, called evidential fully convolutional network (E-FCN), an encoder-decoder architecture first extracts pixel-wise feature maps from an input image. A Dempster-Shafer layer then computes mass functions at each pixel location based on distances to prototypes. Finally, a utility layer performs semantic segmentation based on pixel-wise mass functions. The proposed model can be trained using a learning set with soft labels in an end-to-end way.

The main finding of this study is that the proposed combination of FCNs and ENNs makes it possible to improve accuracy and calibration of FCN models by assigning ambiguous pixels to multi-class sets, while maintaining the good performance of FCNs in precise segmentation tasks. The E-FCN model is able to select a set of classes when the object representation does not allow us to select a single class unambiguously, which easily leads to incorrect decision-making in probabilistic FCNs. This result provides a new direction to improve the performance of FCN models for semantic segmentation. The learning strategy using soft labels further improves the accuracy and calibration of the FCN models. Additionally, the proposed approach makes it possible to reject outliers together with ambiguous pixels when the tolerance to imprecision is between 0.7 and 0.9.

Future work will focus on two main aspects. First, we will investigate multi-FCN-model information fusion for semantic segmentation based on the definition of soft labels, using an approach similar to that introduced in [36]. Other advanced evidential classifiers, such as the contextual-discounting evidential K -nearest neighbor [14] will also be considered to improve the performance of the proposed neural network architecture.

Acknowledgements This research was supported by a scholarship from the China Scholarship Council and by the Labex MS2T (reference ANR-11-IDEX-0004-02).

Declarations

Competing interests The authors declare that they have no conflict of interest.

References

1. Badrinarayanan V, Kendall A, Cipolla R (2017) Segnet: a deep convolutional encoder-decoder architecture for image

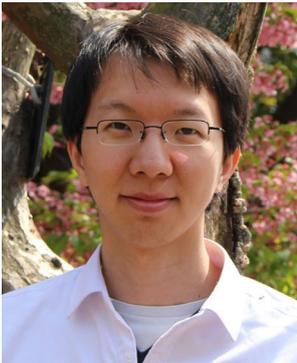
- segmentation. *IEEE Trans Pattern Anal Mach Intell* 39(12):2481–2495
2. Biggio B, Nelson B, Laskov P (2011) Support vector machines under adversarial label noise. In: Asian conference on machine learning, pp 97–112
 3. Chen LC, Papandreou G, Kokkinos I, Murphy K, Yuille AL (2017) Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans Pattern Anal Mach Intell* 40(4):834–848
 4. Chen XI, Wang Ph, Hao Ys, Zhao M (2018) Evidential KNN-based condition monitoring and early warning method with applications in power plant. *Neurocomputing* 315:18–32
 5. Côme E., Oukhellou L, Dencœur T, Aknin P (2009) Learning from partially supervised data using mixture models and belief functions. *Pattern Recogn* 42(3):334–348
 6. Cordts M, Omran M, Ramos S, Rehfeld T, Enzweiler M, Benenson R, Franke U, Roth S, Schiele B (2016) The cityscapes dataset for semantic urban scene understanding In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 3213–3223
 7. Dempster AP (1967) Upper and lower probabilities induced by a multivalued mapping. *Ann Math Stat* 38:325–339
 8. Dencœur T (1997) Analysis of evidence-theoretic decision rules for pattern classification. *Pattern Recogn* 30(7):1095–1107
 9. Dencœur T (2000) A neural network classifier based on Dempster-Shafer theory. *IEEE Trans Syst Man Cybern-Part A: Syst Hum* 30(2):131–150
 10. Denoeux T (2016) 40 years of Dempster-Shafer theory. *Int J Approx Reason* 79:1–6
 11. Denoeux T (2019) Decision-making with belief functions: a review. *Int J Approx Reason* 109:87–110
 12. Dencœur T (2019) Logistic regression, neural networks and Dempster-Shafer theory: A new perspective. *Knowl-Based Syst* 176:54–67
 13. Dencœur T, Dubois D, Prade H (2020) Representations of uncertainty in artificial intelligence: Beyond probability and possibility. In: A guided tour of artificial intelligence research, vol 1, chap. 4, pp 119–150. Springer, Berlin
 14. Dencœur T, Kanjanatarakul O, Sriboonchitta S (2019) A new evidential k-nearest neighbor rule based on contextual discounting with partially supervised learning. *Int J Approx Reason* 113:287–302
 15. Dencœur T, Smets P (2006) Classification using belief functions: relationship between case-based and model-based approaches. *IEEE Trans Syst Man Cybern Part B (Cybern)* 36(6):1395–1406
 16. Ess A, Müller T, Grabner H, Van Gool L (2009) Segmentation-based urban traffic scene understanding. In: *BMVC*, vol 1. Citeseer, pp 2
 17. Everingham M, Eslami SA, Van Gool L, Williams CK, Winn J, Zisserman A (2015) The pascal visual object classes challenge: a retrospective. *Int J Comput Vis* 111(1):98–136
 18. Forouzanfar M, Forghani N, Teshnehlab M (2010) Parameter optimization of improved fuzzy c-means clustering algorithm for brain mr image segmentation. *Eng Appl Artif Intell* 23(2):160–168
 19. Guettari N, Capelle-Laizé AS, Carré P (2016) Blind image steganalysis based on evidential K-Nearest Neighbors In: Proceedings of the 2016 IEEE International Conference on Image Processing, Phoenix, pp 2742–2746
 20. Guo C, Pleiss G, Sun Y, Weinberger KQ (2017) On calibration of modern neural networks. [arXiv:1706.04599](https://arxiv.org/abs/1706.04599)
 21. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 770–778
 22. Krähenbühl P, Koltun V (2011) Efficient inference in fully connected crfs with gaussian edge potentials. In: Advances in neural information processing systems, pp 109–117
 23. Kumar N, Berg AC, Belhumeur PN, Nayar SK (2009) Attribute and simile classifiers for face verification In: Proceedings of the 12th International Conference on Computer Vision. IEEE, Kyoto, pp 365–372
 24. Lian C, Ruan S, Dencœur T (2015) An evidential classifier based on feature selection and two-step classification strategy. *Pattern Recogn* 48:2318–2327
 25. Long J, Shelhamer E, Darrell T (2015) Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, pp 3431–3440
 26. Ma L, Dencœur T (2021) Partial classification in the belief function framework. *Knowl-Based Syst* 214:106742
 27. Natarajan N, Dhillon IS, Ravikumar PK, Tewari A (2013) Learning with noisy labels. In: Advances in neural information processing systems, pp 1196–1204
 28. Noh H, Hong S, Han B (2015) Learning deconvolution network for semantic segmentation. In: Proceedings of the IEEE International Conference on Computer Vision, pp 1520–1528
 29. O’Hagan M (1988) Aggregating template or rule antecedents in real-time expert systems with fuzzy set logic. In: Twenty-second asilomar conference on signals, systems and computers, vol 2, pp 681–689
 30. Shafer G (1976) A mathematical theory of evidence. Princeton University Press Princeton
 31. Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)
 32. Smets P (1993) Belief functions: the disjunctive rule of combination and the generalized Bayesian theorem. *Int J Approx Reason* 9(1):1–35
 33. Su ZG, Denoeux T, Hao YS, Zhao M (2018) Evidential k-NN classification with enhanced performance via optimizing a class of parametric conjunctive t-rules. *Knowl-Based Syst* 142:7–16
 34. Tighe J, Lazebnik S (2010) Superparsing: scalable nonparametric image parsing with superpixels. In: European conference on computer vision. Springer, pp 352–365
 35. Tong Z, Xu P, Dencœur T (2019) ConvNet and Dempster-Shafer theory for object recognition. In: Processing of the 13th international conference on scalable uncertainty management. Springer International Publishing, Cham, pp 368–381
 36. Xu P, Davoine F, Bordes JB, Zhao H, Dencœur T (2016) Multimodal information fusion for urban scene understanding. *Mach Vis Appl* 27(3):331–349
 37. Yager RR (1988) On ordered weighted averaging aggregation operators in multicriteria decision-making. *IEEE Trans Syst Man Cybern* 18(1):183–190
 38. Yager RR, Liu L (2008) Classic works of the Dempster-Shafer theory of belief functions, vol 219. Springer, Berlin
 39. Yoon Y, Jeon HG, Yoo D, Lee JY, So Kweon I (2015) Learning a deep convolutional network for light-field image super-resolution. In: Proceedings of the IEEE international conference on computer vision workshops, pp 24–32
 40. Yuan B, Yue X, Lv Y, Denoeux T (2020) Evidential deep neural networks for uncertain data classification. In: International conference on knowledge science, engineering and management. Springer, pp 427–437

41. Zeiler MD, Fergus R (2014) Visualizing and understanding convolutional networks. In: European conference on computer vision. Springer, pp 818–833
42. Zeiler MD, Taylor GW, Fergus R (2011) Adaptive deconvolutional networks for mid and high level feature learning. In: 2011 International conference on computer vision. IEEE, pp 2018–2025
43. Zhou B, Zhao H, Puig X, Fidler S, Barriuso A, Torralba A (2016) Semantic understanding of scenes through the ade20k dataset. arXiv:1608.05442

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Zheng Tong received his B.S. degree in civil engineering from Northeast Forestry University, China in 2015 and his M.S. degree in transportation engineering from Chang'an University, China in 2018. He is currently a Ph.D. student at the Université de Technologie de Compiègne, France, under the supervision of Thierry Dencœux. His research interests concern pattern recognition, deep learning, and the application of belief functions.



Philippe Xu received the B.S. and M.S. degrees in computer science from the École Normale Supérieure de Cachan, France, in 2009 and 2011, respectively, and the Ph.D. degree from the Université de Technologie de Compiègne (UTC), France, in 2014. Since 2015, he has been an Associate Professor with the Computer Science Department, UTC, where he is carrying out his research at Heudiasyc UMR 7253 (UTC/CNRS). Since 2017, he

has also been a member of the SIVALab, a joint laboratory between Renault, UTC, and CNRS. His research interests include information fusion, computer vision, and machine learning applied to localization and perception for autonomous vehicles.



Thierry Dencœux is a Full Professor (Exceptional Class) with the Department of Information Processing Engineering at the Université de Technologie de Compiègne, France, and a senior member of the French Academic Institute (*Institut Universitaire de France*). His research interests concern reasoning and decision-making under uncertainty and, more generally, the management of uncertainty in intelligent systems. His main contributions are in the

theory of belief functions with applications to statistical inference, pattern recognition, machine learning and information fusion. He has published more than 300 papers in this area. He is the Editor-in-Chief of the *International Journal of Approximate Reasoning*.