

Logic-induced Diagnostic Reasoning for Semi-supervised Semantic Segmentation

Chen Liang, Wenguan Wang, Jiaxu Miao, Yi Yang*

ReLER, CCAI, Zhejiang University

<https://github.com/leonnop/LogicDiag>

Abstract

Recent advances in semi-supervised semantic segmentation have been heavily reliant on pseudo labeling to compensate for limited labeled data, disregarding the valuable relational knowledge among semantic concepts. To bridge this gap, we devise LOGICDIAG, a brand new neural-logic semi-supervised learning framework. Our key insight is that conflicts within pseudo labels, identified through symbolic knowledge, can serve as strong yet commonly ignored learning signals. LOGICDIAG resolves such conflicts via reasoning with logic-induced diagnoses, enabling the recovery of (potentially) erroneous pseudo labels, ultimately alleviating the notorious error accumulation problem. We showcase the practical application of LOGICDIAG in the data-hungry segmentation scenario, where we formalize the structured abstraction of semantic concepts as a set of logic rules. Extensive experiments on three standard semi-supervised semantic segmentation benchmarks demonstrate the effectiveness and generality of LOGICDIAG. Moreover, LOGICDIAG highlights the promising opportunities arising from the systematic integration of symbolic reasoning into the prevalent statistical, neural learning approaches.

1. Introduction

Deep learning has revolutionized computer vision tasks, yielding remarkable breakthroughs [1–5]. However, such advances are often only possible in the presence of large labeled training datasets, which are challenging to acquire. Semantic segmentation, in particular, poses difficulties due to the need for pixel-level manual labeling [6–13], which is time-consuming and labor-intensive, precluding the application of such methods, especially in domains like medical image analysis. To remedy this issue, there has been a growing interest in semi-supervised semantic segmentation, which aims to train segmentation models using a combination of limited labeled data and a large amount of unlabeled

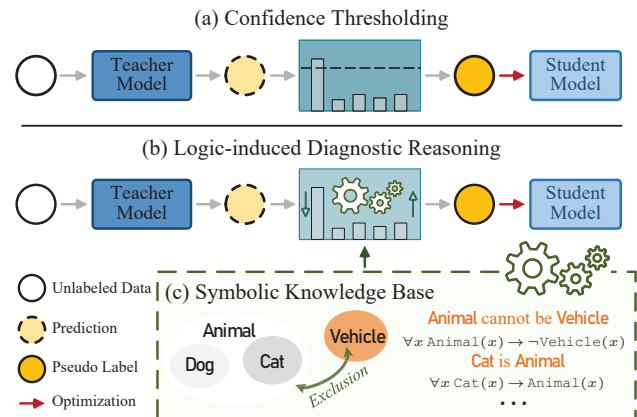


Figure 1: Prevalent data-driven pseudo-labeling methods typically rely on heuristic confidence thresholding (*i.e.*, (a)). We opt to leverage symbolic knowledge in the form of logic rules (*i.e.*, (c)), to diagnose and resolve potential errors within predictions (*i.e.*, (b)).

data [14–20]. Pseudo labeling [21, 22] constitutes one such technique, where the unlabeled data is assigned pseudo labels based on model predictions. The model is then iteratively trained using these pseudo labeled data as if they were labeled examples. Typically, this method is implemented within the teacher-student framework [17] (Fig. 1 (a)).

Despite its prevalence, the pseudo labeling paradigm, a statistical learning approach, is commonly perceived as unreliable, due to the accumulation of erroneous predictions [23]. Previous works [22, 24, 25] have attempted to mitigate this issue by rejecting pseudo labels with classification scores below a heuristic threshold, known as confidence thresholding [22]. However, relying solely on this strategy often proves unsatisfactory. **First**, the confirmation bias that may occur during the early stages of training can be difficult, if not impossible, to be rectified in subsequent learning [26]. **Second**, the purely data-driven nature of such threshold-based methods makes them challenging to interpret. **Third**, to achieve optimal performance, the threshold must be manually adjusted for each model and dataset, lim-

*Corresponding author: Yi Yang.

iting its practical application. Beyond the confines, symbolic reasoning, as a compelling alternative, offers appealing characteristics: it requires *little* or *no data* to generalize systematically. In light of this background, we suggest a promising direction towards an integration of statistical learning and symbolic reasoning to harness the strengths of both paradigms to achieve improved performance.

In this paper, we propose to explicitly compile the rich symbolic knowledge into the prevalent pseudo label based neural training regime (*cf.* Fig. 1 (b)). Our key insight is that symbolic knowledge can effectively resolve conflicts within the pseudo labels, which reveal potential model errors. For instance, leveraging the prior knowledge on compositionality, we can naturally identify the inconsistencies like classifying a pixel as both a `cat` and a `vehicle` (*cf.* Fig. 1 (c)). Correcting these errors progressively enhances the accuracy of pseudo labels, and thus mitigating the confirmation bias in iterative self-training. Building upon this insight, we introduce LOGICDIAG, a new SSL framework that employs logical reasoning to identify such conflicts based on the symbolic knowledge expressed in the form of first-order logic. The framework then suggests possible diagnoses to rectify the prediction. To further address the challenge of multiple diagnoses, we model the likelihood of each diagnosis being the actual fault based on a comprehensive measure of both predictive confidence and degree of conflicts according to the fuzzy logic. By doing so, our approach introduces a holistic neural-logic machine, that consolidates the benefits of powerful declarative languages, transparent internal functionality, and enhanced model performance.

LOGICDIAG is a principled framework that seamlessly integrates with mainstream semi-supervised learning methods. It requires only minor adjustments to the dense classification head. With LOGICDIAG, we can easily describe diverse symbolic knowledge using first-order logic and inject it into sub-symbolic pipelines. Taking dense segmentation as main battlefield, we capture and evaluate a central cognitive ability of human, *i.e.*, structured abstractions of visual concepts [27], by grounding three logic rules onto LOGICDIAG: *Composition*, *Decomposition*, and *Exclusion*.

Extensive experiments have validated the effectiveness of our LOGICDIAG, which exhibit solid performance gains (*i.e.*, 1.21%-4.15% mIoU) on three well-established benchmarks (*i.e.*, PASCAL VOC 2012 [28], Cityscapes [29], and COCO [30]). We particularly observe significant improvements in the label scarce settings, indicating LOGICDIAG’s superior utilization of unlabeled data. Besides, when employed onto the existing SSL frameworks, *e.g.*, AEL [15], MKD [14], the performance is consistently advanced. The results demonstrate the strong generality and promising performance of LOGICDIAG, that also evidence the great potential of the integrated neural-symbolic computing in the fundamental large-scale semi-supervised segmentation.

2. Related Work

Semi-Supervised Semantic Segmentation. Recent years have witnessed remarkable progress in image-level semi-supervised learning, driven primarily by two paradigms: self-training [21, 31, 32] and consistency regularization [33, 34]. Despite the impressive outcomes they produce, both paradigms heavily rely on the quality of pseudo-labels generated by the network itself, which makes them susceptible to confirmation bias [23] and leads to error accumulation throughout training. Considerable efforts have been devoted to addressing this issue through confidence thresholding [22, 35–37], curriculum scheduling [25, 38], sample ensembling across training iterations [17, 39], multiple augmented views [24], and/or neighboring examples [40–42]. Some others introduce perturbations in data [22, 36, 37], feature [43], and/or network [34, 44]. Similarly, in the context of semi-supervised semantic segmentation, which provides a pixel-level interpretation of visual semantics, prevalent approaches also follow the self-training paradigm [45–48] and consistency regularization paradigm [49–54]. Very recent endeavors have further pushed forward the frontier through incorporating dense perturbations [14, 47, 51, 55, 56], contrastive supervisions [57–60], or coping with imbalanced and long-tailed nature of pixel samples [15, 61].

Although significant advancements have been made, the problem of error accumulation remains far from being fully resolved. Existing methods heavily rely on an empirical threshold to justify pseudo labels, lacking explicit manipulation of rich symbolic knowledge. One common oversight in existing methods is the neglect of informative structures between semantic concepts. As a consequence, these approaches often yield barely satisfactory results in terms of both model performance and generality. One notable exception is [62]. While a label hierarchy is incorporated, it is only used for reducing labeling costs of fine-grained classes, resulting in limited improvement. This work pursues an integrated neural-logic framework that addresses these fundamental limitations, providing a refreshing viewpoint on semi-supervised semantic segmentation.

Hierarchical Classification. Class-wise hierarchical dependencies have been studied in supervised tasks across several machine learning domains, *e.g.*, functional genomics [63–65], text categorization [66], object recognition [67–69], image classification [70, 71]. In the computer vision field, the class taxonomy is mainly explored through: i) semantic-aligned *label embedding* [72–74]; ii) hierarchy-coherent *loss constraint* [71, 75, 76]; or iii) structured *network architecture* [77–79]. Among the previous efforts, only a few attempts [80–84] towards label hierarchy-aware semantic segmentation have been made. Whereas, they all heavily rely on labeled hierarchical data and/or specialized neural architectural design, making them impractical for use when merely a handful of labeled data is available, as in SSL. In

contrast, our algorithm exploits class taxonomy as the symbolic knowledge for diagnosing conflicts in model predictions, enabling comprehensive exploration of hierarchical relations in both labeled and unlabeled data. The general design also makes our algorithm versatile and easily implementable to standard hierarchy-agnostic SSL architectures. **Neural-Symbolic Computing.** Building preferable computational methods for integrated statistical learning and symbolic reasoning is a long-standing challenge [85]. This active line of research, namely neural-symbolic computing (NSC), draws soaring attention in recent years [86–89]. NSC shows great potential to reconcile the robust learning capabilities of neural networks with the interpretability and reasoning abilities of symbolic representations [90,91], and thereby gains widespread recognition as a catalyst for the next generation of AI [92,93]. NSC has demonstrated its superiority across a wide range of domains, including mathematical reasoning [94–96], robotics control [97–99], as well as scientific discovery [100–103]. Its virtue of data efficiency has also attracted researchers from SSL fields [104–107], where the symbolic knowledge, usually expressed in logic, is mostly incorporated as a form of regularization through loss constraints applied to output space. Though being exciting, the advances have been primarily limited to “toy” tasks. The full potential and challenges of NSC for large-scale realistic problems remain largely unexplored. Our method is partly motivated by, but also distinct from, previous efforts that merely encourage valid output structures in a soft manner, thereby allowing errors to accumulate, especially during the initial stages of training.

To our best knowledge, this is the first work that promotes and implements an integrated neural-symbolic framework in large-scale vision-oriented SSL. Previous studies have primarily concentrated on the neural aspect, *i.e.*, sub-symbolic methods. In contrast, our method explicitly compiles symbolic knowledge into training regime of the neural network. This integration allows us to leverage the advantages of powerful declarative languages and transparent internal functionality. The encouraging results we have obtained provide compelling empirical evidence of the significant potential of NSC in the large-scale vision domain.

3. Methodology

We commence by formalizing modern semi-supervised semantic segmentation approaches, situating them within a sub-symbolic framework and highlighting their inherent limitations (§3.1). We then present LOGICDIAG, a general logic-induced diagnosis framework, that complements the current sub-symbolic approaches with a principled infusion of symbolic knowledge in the form of logic rules (§3.2). Finally, we showcase its practical application in the realm of visual semantic interpretation (§3.3).

Problem Statement. In the standard SSL setting, given an unknown distribution over visual space \mathcal{X} (*e.g.*, pixel space for segmentation) and category label space $\mathcal{Y} = \{1, \dots, C\}$ with C semantic categories, the goal is to find a predictor $h: \mathcal{X} \mapsto \mathcal{Y}$, such that the generalization error is minimized, based on the observed data \mathcal{D} , consisting of a labeled subset $\mathcal{D}^l = \{(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}\}_{i=1}^{N^l}$ and an unlabeled subset $\mathcal{D}^u = \{u_j \in \mathcal{X}\}_{j=1}^{N^u}$. Typically, N^l is rather small, *i.e.*, $N^u \gg N^l$, resulting in the issue of insufficient learning signal.

3.1. Sub-symbolic Semi-supervised Learning

Mainstream solutions predominantly revolve around a neural pipeline, known as sub-symbolic methods [108]. They rely on the consistency regularization paradigm [22, 109], that comprises four key components, as shown in Fig. 2 (a).

- *Data augmentors* $\mathcal{A}(\cdot)/\alpha(\cdot)$, that transform examples into strongly-/weakly-augmented views, respectively.
- *Base encoder* $f(\cdot)$, that maps augmented views into D -dimensional representations, *i.e.*, $\mathbf{x} = f \circ \alpha(x) \in \mathbb{R}^D$.
- *Prediction head* $g(\cdot)$, that gives the predictive probability distribution from the representation, *i.e.*, $\mathbf{o} = g(\mathbf{x}) \in \Delta^C$, and Δ^C is the C -way probability simplex.
- *Pseudo label processor* $\psi(\cdot)$, that converts the raw predictions into pseudo labels guided by heuristic priors or assumptions, *e.g.*, confidence thresholding [22,24,25,35].

The predictor h , consisting of the base encoder and prediction head (*i.e.*, $h = g \circ f$), is jointly optimized on the complete observed set \mathcal{D} , *i.e.*, $\{\mathcal{D}^l, \mathcal{D}^u\}$. For the labeled ones, the standard cross entropy, denoted by $\mathbb{H}(\cdot, \cdot)$, can be directly applied on the weakly augmented examples:

$$\mathcal{L}^l = \frac{1}{N^l} \sum_{i=1}^{N^l} \mathbb{H}((h \circ \alpha)(x_i), y_i). \quad (1)$$

For the unlabeled data, the pseudo labels are generated from the weakly augmented views of given examples. As a prevalent choice [109], the pseudo label processor $\psi(\cdot)$ filters out unreliable pseudo labels with a confidence threshold τ , *i.e.*, $\psi(\mathbf{o}) = \mathbb{I}(\max(\mathbf{o}) \geq \tau) \cdot \arg \max \mathbf{o}$. Here \mathbb{I} defines the identify function. Then, the predictor h is expected to yield consistent predictions on the strongly augmented views:

$$\mathcal{L}^u = \frac{1}{N^u} \sum_{j=1}^{N^u} \mathbb{H}((h \circ \mathcal{A})(u_j), (\psi \circ h \circ \alpha)(u_j)), \quad (2)$$

The entire network is supervised with both the two losses:

$$\mathcal{L} = \mathcal{L}^l + \lambda \mathcal{L}^u, \quad (3)$$

where the scalar hyperparameter λ trade off the two terms. Despite the prevalence it achieved, there remains three issues in mainstream SSL methods: **First**, they simply treat all categories equally in a flat view, overlooking the structured relationships between the visual concepts. The transferable knowledge residing in the label hierarchy is also discarded, which is of particular essential in the label scarce scenarios. **Second**, self-training with pseudo labels in turn

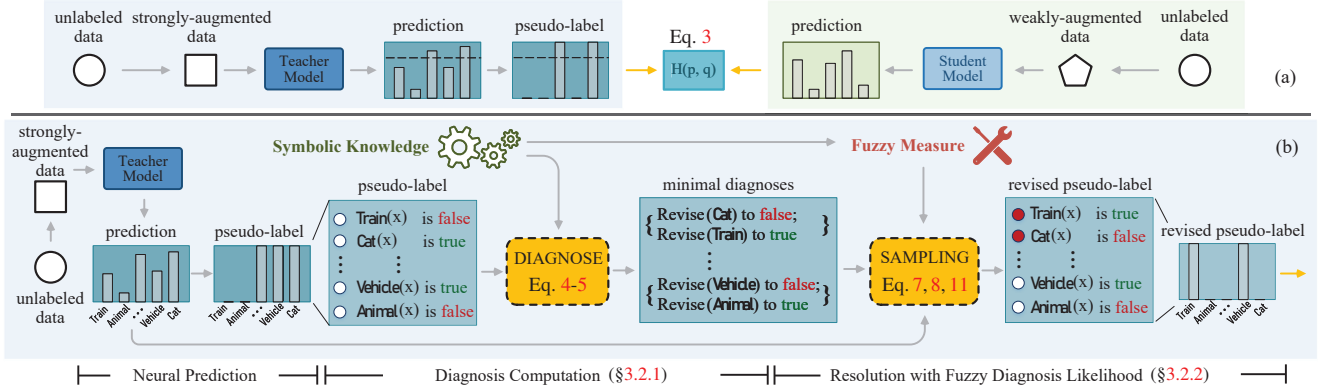


Figure 2: Illustrations of (a) sub-symbolic SSL pipeline (§3.1); (b) LOGICDIAG, logic-induced diagnostic reasoning framework (§3.2).

(Eq. 2) inevitably accumulates errors and causes confirmation bias [23], which severely affects the model performance. Although low-confidence filtering (*i.e.*, ψ) can alleviate the problem, errors accumulated in the early training stage are difficult to be corrected in subsequent training, especially for the poorly-behaved categories [26]. **Third**, the label filtering process relies on the empirical adjustment of a threshold τ , which requires specialization for each model and dataset to promote the performance. This significantly limits the flexibility and generality of current methods.

Accordingly, we suggest that it is now imperative to rethink prevailing sub-symbolic pipeline, in which the filtering process $\psi(\cdot)$ has been more detrimental than beneficial.

3.2. Logic-induced Diagnostic Reasoning

A well-performed SSL model not only generates accurate sample labels but also aligns with our background knowledge of the world, *e.g.*, *bird is an animal*, *sky is above the grass*. By resolving conflicting predictions through reasoning based on the knowledge, we can potentially correct erroneous pseudo-labels, which leads to a more effective learning process of the neural aspect (*i.e.*, $h(\cdot)$) and ultimately reduces the accumulation of errors. To this end, our approach, LOGICDIAG, incorporates the symbolic aspect (*i.e.*, $\psi(\cdot)$), which reasons about malfunctions, dubbed *conflict*, in the model’s output (*i.e.*, pseudo label) according to the symbolic knowledge, and offers possible solutions, dubbed *diagnosis* [110], to resolve the conflicts within predictions.

3.2.1 Diagnosis Computation

We formalize our target with a triple $\mathcal{S} = \langle \mathcal{D}_x, h, \mathcal{K} \rangle$:

- $\mathcal{D}_x = \{x_i\}_{i=1}^{N^l + N^u}$ defines the collection of pixel samples from both labeled and unlabeled datasets;
- h is the neural predictor, that generates a set of binary pseudo labels, denoted as $\mathcal{O} = \{o_j\}_{j=1}^{|\mathcal{O}|}$, with one label for each semantic concept such as *bird*, *animal*, *etc.* Each o determines whether a given pixel x belongs to a specific

concept or not, which can be achieved by applying a binarizing operation b , to the output of h , *i.e.*, $\mathcal{O} = b \circ h(x)$.¹

- \mathcal{K} is a finite set of rules expressed in *first-order logic* (FOL) that captures the world knowledge on normality.

We provide a gentle introduction on FOL, which comprises four parts: **i)** *constants* representing specific pixel instances x_i or truth (*i.e.*, \top : true, \perp : false); **ii)** *variables* ranging over these constants, denoted by x ; **iii)** *predicates* evaluating the semantics of variables to be true or false (*e.g.*, *bird*(x) is true states the fact that pixel x belongs to a category *bird*); **iv)** *connectives* (*e.g.*, \wedge : and, \vee : or, \neg : not, \rightarrow : imply) and *quantifiers* (*i.e.*, \forall : for all, \exists : exist) over finite predicates. To simplify subsequent formulations without altering the meaning, we may omit the explicit mention of the pixel variable x in certain predicates (*e.g.*, $o(x)$ to o).

A *conflict* arises when the assumption that all outputs are normal is inconsistent with our symbolic knowledge \mathcal{K} :

$$\mathcal{K} \wedge \mathcal{O} \wedge \bigwedge_{o \in \mathcal{O}} t(o) \vdash \perp. \quad (4)$$

Here \vdash is logical entailment. We define the unary predicate $t(\cdot)$ over the output o such that $t(o)$ is true when o is normal in terms of consistency, while $\neg t(o)$ is true when o is faulty. To explain the inconsistency above, a *diagnosis* is proposed by assuming that some outputs in a set $\omega \subseteq \mathcal{O}$ are faulty:

$$\mathcal{K} \wedge \mathcal{O} \wedge \bigwedge_{o \in \omega} \neg t(o) \wedge \bigwedge_{o \in \mathcal{O} \setminus \omega} t(o) \not\vdash \perp. \quad (5)$$

For all possible diagnoses, we restrict our focus solely to the ones that contain only false elements, in the world being modeled, known as minimal diagnosis [110]. Formally, a diagnosis ω is minimal *iff.* any strict subset $\omega' \subset \omega$ is not a diagnosis. Thus far, the problem has been reduced to a Boolean satisfiability problem. We present to solve it with greedy algorithms. In practice, considering the rather small search space (*i.e.*, $|\mathcal{K}|$ and $|\mathcal{O}|$), we keep our pipeline straightforward without employing sequential approximation (*e.g.*, MCMC). At this point, a diagnosis has already

¹Here we extend h to generate binary outputs by attaching a sigmoid function, without modifying the main architecture of the network (*cf.* §3.3).

been computed and the pseudo label \mathcal{O} can be revised to the logically consistent \mathcal{O}' according to the following equation:

$$\mathcal{O}' = \{\neg o\}_{o \in \omega} \cup \{o\}_{o \in \mathcal{O} \setminus \omega}. \quad (6)$$

However, there might be multiple diagnoses that are consistent with our target, it is not immediately clear which one is the ‘correct’ revision to make. A naïve approach would be to uniformly sample from all possible minimal diagnoses and revise them one at a time, which is inefficient and ineffective, as demonstrated by the empirical results (c.f. §4.3). To address this issue, we take one step further to model the likelihood of the diagnosis being the actual faulty.

3.2.2 Resolution with Fuzzy Diagnosis Likelihood

We can derive the likelihood of a correct diagnosis given the triple system $\mathcal{S} = \langle \mathcal{D}_x, h, \mathcal{K} \rangle$ under the common assumption of independent failure [111], using the multiplication rule:

$$P(\omega|\mathcal{S}) = \prod_{o \in \omega^+} P(t(o)|\mathcal{S}) \prod_{o \in \omega^-} (1 - P(t(o)|\mathcal{S})), \quad (7)$$

where ω is the diagnosis; ω^+ and ω^- are the sets of outputs assigned normal and abnormal behavior modes, respectively. To comprehensively estimate the probability of an output o being normal, denoted as $P(t(o)|\mathcal{S})$, we consider both the degree of conflict and predictive confidence:

$$P(t(o)|\mathcal{S}) := \begin{cases} P(o|\mathcal{S}) \cdot (1 - c(o; \mathcal{S})), & \text{if } o \vdash \top \\ (1 - P(o|\mathcal{S})) \cdot (1 - c(o; \mathcal{S})), & \text{if } o \vdash \perp \end{cases} \quad (8)$$

Here, function c maps each output $o \in \mathcal{O}$ to its conflict degree $c(o; \mathcal{S}) \in [0, 1]$; and $P(o|\mathcal{S}) = P(o|x) = h(x) \in [0, 1]$ represents the predictive probability. Intuitively, if an output violates the established knowledge but has a high predictive confidence, it should be penalized more than the others (i.e., a small valued $c(o; \mathcal{S})$), increasing the likelihood of correct diagnosis. Conversely, high confidence indicates low probability in abnormality, thereby balancing the likelihood.

To measure this degree of conflict, we resort to the fuzzy logic [112], a form of soft probabilistic logic, specifically, the *Goguen fuzzy logic* [113] and *Gödel fuzzy logic* [114]. Fuzzy logic generalizes FOL to uncertain inputs, where the variables have truth values in $[0, 1]$, e.g., the predictive probability with respect to each visual concept in our case. The logical connectives (e.g., \wedge, \vee, \neg) are approximated with fuzzy operators (i.e., *t-norm, t-conorm, fuzzy negation*):

$$\phi \wedge \varphi := \phi \cdot \varphi, \quad \phi \vee \varphi := \max(\phi, \varphi), \quad \neg \phi := 1 - \phi. \quad (9)$$

Besides, the existential and universal quantifier (i.e., \forall, \exists) are approximated in a form of generalized mean [115, 116]:

$$\begin{aligned} \exists \phi(x) &:= \left(\frac{1}{|\mathcal{D}_x|} \sum_{x \in \mathcal{D}_x} \phi(x)^q \right)^{\frac{1}{q}}, \\ \forall \phi(x) &:= 1 - \left(\frac{1}{|\mathcal{D}_x|} \sum_{x \in \mathcal{D}_x} (1 - \phi(x)^q) \right)^{\frac{1}{q}}, \end{aligned} \quad (10)$$

where $q \in \mathbb{Z}$. Instead of being strictly true or false, fuzzy logic offers a soft measure on how much a logic rule is fired, i.e., truth degree, which can be interpreted as the complement to the desired conflict degree:

$$c(o; \mathcal{S}) = 1 - \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} g(o, k; \mathcal{S}), \quad (11)$$

where we define $g(\cdot)$ as the fuzzy truth measurement.

Following the derivation, we are now ready to assess the extent to which a series of outputs violate the rules in our knowledge \mathcal{K} (c.f. Eq. 11), so as to estimate the likelihood of actual diagnoses (c.f. Eq. 7-8). Then, the naïve utilization is to select the diagnosis with the highest probability concerning pixel x_i , i.e., $\arg \max_{\omega_i} P(\omega_i|x_i, h, \mathcal{K})$. However, such local estimation is prone to get stuck in spurious corrections, as evidenced by our experiments (c.f. §4.3). Instead, we resort to Monte Carlo estimation to calculate the posterior distribution. Our empirical findings suggest that optimizing with just one sample per datapoint is sufficient.

Overall, as shown in Fig. 2 (b), the reasoning and learning aspects of LOGICDIAG work iteratively. The neural model $h(\cdot)$ first poses assumptions about current observations, then the symbolic model $\psi(\cdot)$ reasons over the symbolic knowledge \mathcal{K} to determine the diagnoses, which in turn facilitates the learning process of the neural model.

3.3. Diagnosis with Visual Semantics: An Example

LOGICDIAG is model-agnostic to the neural aspect of SSL pipelines, and is also compatible with general symbolic knowledge described in FOL. As a result, LOGICDIAG can further be advanced by embracing the development of new architectures or incorporating more knowledge. To showcase the practical deployment of LOGICDIAG, we examine its application in segmentation scenario, with particular interest in structured visual semantics [80] that are commonly overlooked in the mainstream. Specifically, the semantic concepts and their relations are formed as a tree-shaped label hierarchy $\mathcal{T} = \langle \mathcal{O}, \mathcal{E} \rangle$ (c.f. Fig. 3 (a)). The node set \mathcal{O} is the union of nodes from L levels of abstraction, denoted as $\mathcal{O} = \cup_{l=1}^L \mathcal{O}_l$. The leaf nodes, \mathcal{O}_1 , represent the most specific concepts, namely category labels, such as *bird*, *cat*, where $\mathcal{O}_1 = \mathcal{Y}$, $|\mathcal{O}_1| = C$. The internal nodes represent higher-level concepts such as *vehicle*, *animal*, and the root nodes \mathcal{O}_L represent the most abstract concepts, such as *object*. Besides, the edge set \mathcal{E} encodes relational knowledge among all these concepts, with directed edges $u \rightarrow v \in \mathcal{E}$ denoting a *part-of* relation between two concepts $u, v \in \mathcal{O}$ in adjacent levels, e.g., *animal* \rightarrow *bird*.

Recall that we define the $\langle \text{data, model, knowledge} \rangle$ triple $\langle \mathcal{D}_x, h, \mathcal{K} \rangle$ for LOGICDIAG (c.f. §3.2.1). In the context of structured semantic concept exploration, the model h yields $|\mathcal{O}|$ binary pseudo labels, denoting nodes within the label hierarchy. Besides, $\mathcal{K}_{\mathcal{T}}$ contains FOL rules, describing the




Figure 3: Illustrations of (a) label hierarchy \mathcal{T} ; (b) Composition (Eq. 12); (c) Decomposition (Eq. 13); (d) Exclusion (Eq. 14) rules (§3.3).

structured symbolic knowledge according to the label hierarchy \mathcal{T} . Inspired by previous efforts [80, 117, 118] in hierarchical classification, we define $\mathcal{K}_{\mathcal{T}}$ with three types of rules, i.e., *composition*, *decomposition*, and *exclusion*.

• **Composition Rule** (\mathcal{K}_C). *If one class is labeled true, its parent (i.e., superclass) is labeled true* (Fig. 3 (b)):

$$\forall x(o(x) \rightarrow p_o(x)), \quad (12)$$

where p_o is the parent node of o in \mathcal{T} , i.e., $p_o \rightarrow o \in \mathcal{E}$. For example, “bird is (a subclass of) animal” shall be interpreted as: $\forall x(\text{bird}(x) \rightarrow \text{animal}(x))$.

• **Decomposition Rule** (\mathcal{K}_D). *If one class is labeled true, at least one of its children (i.e., subclasses) is true* (Fig. 3 (c)):

$$\forall x(o(x) \rightarrow \bigvee_{r_o \in R_o} r_o(x)), \quad (13)$$

where R_o is the set of children node(s) of o in \mathcal{T} , i.e., $o \rightarrow r_o \in \mathcal{E}$. For example, “animal subsumes (is the superclass of) bird, dog, ..., cat” shall be interpreted as:

$$\forall x(\text{animal}(x) \rightarrow \text{bird}(x) \vee \text{dog}(x) \vee \dots \vee \text{cat}(x)).$$

• **Exclusion Rule** (\mathcal{K}_E). *If one class is labeled true, all its sibling classes are labeled false* (Fig. 3 (d)):

$$\forall x(o(x) \rightarrow \bigwedge_{s_o \in S_o} \neg s_o(x)), \quad (14)$$

where S_o is the set of sibling node(s) of o in \mathcal{T} . For example, “bird cannot be train, ..., nor cat” shall be interpreted as: $\forall x(\text{bird}(x) \rightarrow \neg \text{train}(x) \wedge \dots \wedge \neg \text{cat}(x))$.

According to the fuzzy measure defined in Eq. 9-10, we derive the truth degrees as follows (Note, $\phi \rightarrow \varphi \Leftrightarrow \neg \phi \vee \varphi$):

• **Composition Rule** (Eq. 12): $g(o, \mathcal{K}_C | \mathcal{S}) =$

$$1 - \left[\frac{1}{|\mathcal{D}_x|} \sum_{x \in \mathcal{D}_x} (P(o|x) - P(o|x) \cdot P(p_o|x))^q \right]^{\frac{1}{q}}. \quad (15)$$

• **Decomposition Rule** (Eq. 13): $g(o, \mathcal{K}_D | \mathcal{S}) =$

$$1 - \left[\frac{1}{|\mathcal{D}_x|} \sum_{x \in \mathcal{D}_x} (P(o|x) - P(o|x) \cdot \max_{r_o \in R_o} P(r_o|x))^q \right]^{\frac{1}{q}}. \quad (16)$$

• **Exclusion Rule** (Eq. 14): $g(o, \mathcal{K}_E | \mathcal{S}) =$

$$1 - \frac{1}{|S_o|} \sum_{s_o \in S_o} \left[\frac{1}{|\mathcal{D}_x|} \sum_{x \in \mathcal{D}_x} (P(o|x) \cdot P(s_o|x))^q \right]^{\frac{1}{q}}, \quad (17)$$

where we translate *one-vs-multiple* exclusion (Eq. 14) to the equivalent *one-vs-one* form to avoid sorites paradox [119].

We are now prepared to proceed with the calculations (c.f. Eq. 7, 8, 11) and sample from the diagnosis likelihood

$P(\omega | \mathcal{S})$, where the pseudo labels can be revised accordingly (c.f. Eq. 6). Overall LOGICDIAG is supervised with Eq. 3.

Implementation Detail. In practice, computing the full semantics of the universal quantification \forall is infeasible due to the large learning corpora, where we use batch-training as sampling based approximation [115]. To facilitate efficient distributed training, we implement diagnostic reasoning steps using matrix multiplications. Besides, we employ point-wise supervision [121] to further optimize efficiency while maintaining high performance standards.

4. Experiment

4.1. Experimental Setup

Datasets. We evaluate LOGICDIAG on standard datasets:

- **PASCAL VOC 2012** [28] is a famous semantic segmentation dataset, consisting $\sim 4k$ samples in *original* dataset, that are split into 1,464/1,449/1,456 images for `train/val/test`, respectively. It provides annotations for 21 categories including the background, which are grouped into 4 superclasses. Following the conventions [57, 109], 9,118 coarsely labeled images in SBD [122] are adopted to complement `train` data, namely the *augmented* set.
- **Cityscapes** [29] has 5k fine-annotated urban scene images, with 2,975/500/1,524 for `train/val/test`, that defines 19 semantic categories organized within 6 superclasses. Please refer to the supplementary for details.
- **COCO** [30] features dense annotations for 80 object categories (e.g., animals, furniture, etc.), within diverse indoor and outdoor scenes. COCO stands out as the largest among the three benchmarks, comprising 118k/5k images for `train/val`. The dataset officially provides a three-level semantic hierarchy, covering 2/12/80 concepts.

Partition Protocol. We conduct evaluation under standard partition protocols [57, 59]. For PASCAL VOC/Cityscapes, we sample 1/2, 1/4, 1/8, and 1/16 of the whole training set as labeled data. For COCO, we use smaller ratios, i.e., 1/32, 1/64, 1/128, 1/256, 1/512, considering the larger size of the dataset. We adopt the same sampled data to the state-of-the-arts [51, 57] to enable meaningful comparisons.

Base Network Architecture. We take DeepLabV3+ [120] as base segmentation architecture (i.e., $h(\cdot)$, cf. §3), where ResNet101 [1] and Xception65 [123] pretrained on Ima-

| Method | PASCAL VOC 2012 <i>original</i> | | | | PASCAL VOC 2012 <i>augmented</i> | | | |
|--------------------------|---------------------------------------|---------------------------------------|---------------------------------------|---------------------------------------|---------------------------------------|---------------------------------------|---------------------------------------|---------------------------------------|
| | 1/2 (732) | 1/4 (366) | 1/8 (183) | 1/16 (92) | 1/2 (5291) | 1/4 (2646) | 1/8 (1323) | 1/16 (662) |
| MT [17] [NeurIPS17] | 69.16 | 63.01 | 55.81 | 48.70 | 77.61 | 76.62 | 73.20 | 70.59 |
| GCT [50] [ECCV20] | 70.67 | 64.71 | 54.98 | 46.04 | 77.14 | 75.25 | 73.30 | 69.77 |
| CutMixSeg [49] [BMVC20] | 69.84 | 68.36 | 63.20 | 55.58 | 75.89 | 74.25 | 72.69 | 72.56 |
| PseudoSeg [51] [ICLR21] | 72.41 | 69.14 | 65.50 | 57.60 | - | - | - | - |
| AEL [15] [NeurIPS21] | - | - | - | - | 80.29 | 78.06 | 77.57 | 77.20 |
| CPS [109] [CVPR21] | 75.88 | 71.71 | 67.42 | 64.07 | 78.64 | 77.55 | 75.83 | 72.18 |
| PC2Seg [59] [ICCV21] | 73.05 | 69.78 | 66.28 | 57.00 | - | - | - | - |
| PS-MT [56] [CVPR22] | 78.42 | 76.57 | 69.58 | 65.80 | 79.76 | 78.72 | 78.20 | 75.50 |
| ST++ [47] [CVPR22] | 77.30 | 74.60 | 71.00 | 65.20 | - | 77.90 | 77.90 | 74.70 |
| U2PL [57] [CVPR22] | 76.16 | 73.66 | 69.15 | 67.98 | 80.50 | 79.30 | 79.01 | 77.21 |
| GTA-Seg [16] [NeurIPS22] | 78.37 | 75.57 | 73.16 | 70.02 | 81.01 | 80.57 | 80.47 | 77.82 |
| MKD [14] [NeurIPS22] | 78.66 | 76.76 | 74.63 | 69.10 | 80.60 | 79.55 | 79.74 | 78.44 |
| Ours | 79.39\uparrow0.73 | 77.93\uparrow1.17 | 76.66\uparrow2.03 | 73.25\uparrow4.15 | 81.00\uparrow0.40 | 80.62\uparrow1.07 | 80.24\uparrow0.50 | 79.65\uparrow1.21 |
| Ours + AEL [15] | 79.56 \uparrow 0.17 | 78.16 \uparrow 0.23 | 76.86 \uparrow 0.20 | 73.65 \uparrow 0.41 | 81.11 \uparrow 0.11 | 80.78 \uparrow 0.16 | 80.47 \uparrow 0.23 | 79.77 \uparrow 0.12 |
| Ours + MKD [14] | 80.06 \uparrow 0.67 | 78.43 \uparrow 0.50 | 77.18 \uparrow 0.52 | 74.70 \uparrow 1.45 | 81.21 \uparrow 0.21 | 80.95 \uparrow 0.33 | 80.53 \uparrow 0.29 | 80.08 \uparrow 0.43 |

Table 1: **Quantitative results** (§4.2) on PASCAL VOC 2012 [28] val. All methods are built upon DeepLabV3+ [120]-ResNet101 [1].

| Method | 1/2 (1488) | 1/4 (744) | 1/8 (372) | 1/16 (186) |
|--------------------------|--------------|--------------|--------------|--------------|
| MT [17] [NeurIPS17] | 78.59 | 76.53 | 73.71 | 68.08 |
| CCT [52] [CVPR20] | 78.29 | 76.35 | 74.48 | 69.64 |
| GCT [50] [ECCV20] | 78.58 | 76.45 | 72.96 | 66.90 |
| CutMixSeg [49] [BMVC20] | 78.95 | 77.24 | 75.83 | 72.13 |
| AEL [15] [NeurIPS21] | 80.28 | 79.01 | 77.90 | 75.83 |
| CPS [109] [CVPR21] | 76.81 | 74.58 | 74.31 | 69.78 |
| S-Baseline [46] [ICCV21] | 78.70 | 77.80 | 74.10 | - |
| U2PL [57] [CVPR22] | 79.05 | 76.47 | 74.37 | 70.30 |
| GTA-Seg [16] [NeurIPS22] | 76.08 | 72.02 | 69.38 | 62.95 |
| MKD [14] [NeurIPS22] | 80.74 | 78.28 | 75.98 | 75.31 |
| Ours | 80.95 | 80.21 | 78.90 | 76.83 |

Table 2: **Quantitative results** (§4.2) on Cityscapes [29] val.

geNet [124] are adopted as the backbone networks.

Training. We implement LOGICDIAG on MMSegmentation [125] following standardized training settings and common weak/strong data augmentations [14, 47, 57] (*i.e.*, \mathcal{A}/α , *cf.* §3). For PASCAL VOC/Cityscapes/COCO, images are cropped to $513 \times 513/769 \times 769/513 \times 513$ and models are trained for 40k/40k/40k iterations with 16/16/16 batch size. The SGD optimizer is adopted with a weight decay of $5e-4/5e-4/5e-4$. The learning rate is set to 0.0025/0.01/0.01, and is scheduled following the polynomial annealing policy [126], with a power of 0.9/0.9/0.9. We set $\lambda = 5$ to balance the unsupervised loss (*cf.* Eq. 3), and opt to assign $q=5$ for logic quantifier approximation (*cf.* Eq. 10).

Inference. Following [14, 59], for PASCAL VOC/COCO, we keep the aspect ratio of test images and rescale the short side to 512. For Cityscapes, we use sliding window inference with 769×769 window size. All results are reported without any test-time augmentation for the sake of fairness.

Evaluation Metric. Following the conventions [57, 109], we use mean intersection-over-union (mIoU) for evaluation. In diagnostic experiments, we further report mIoU at each hierarchy level l (denoted as mIoU^l) for thorough analysis.

| Method | 1/32 (3697) | 1/64 (1849) | 1/128 (925) | 1/256 (463) | 1/512 (232) |
|-------------------------|--------------|--------------|--------------|--------------|--------------|
| Supervised | 42.24 | 37.80 | 33.60 | 27.96 | 22.94 |
| PseudoSeg [51] [ICLR21] | 43.64 | 41.75 | 39.11 | 37.11 | 29.78 |
| PC2Seg [59] [ICCV21] | 46.05 | 43.67 | 40.12 | 37.53 | 29.94 |
| MKD [14] [NeurIPS22] | 47.25 | 45.50 | 42.32 | 38.04 | 30.24 |
| Ours | 50.51 | 48.83 | 45.35 | 40.28 | 33.07 |

Table 3: **Quantitative results** (§4.2) on COCO [30] val, based on DeepLabV3+ [120]-Xception65 [123] architecture.

Reproducibility. Our models are implemented in PyTorch. All experiments are conducted on four Tesla A100 GPUs.

4.2. Quantitative Comparison Result

PASCAL VOC 2012 *original*. Table 1 (left) summarizes the quantitative comparisons under varying label amounts, from which we take three major observations: **First**, Our method indeed surpasses the SOTAs across all settings and establishes the new state-of-the-arts of **79.39%/77.93%/76.66%/73.25%** under 1/2-1/16 partitions, indicating the effectiveness of our neural-logic framework. **Second**, The superiority of our method is more significant on fewer labeled data, with the largest margin achieved when only 1/16 labels are available, *i.e.*, **4.15%** over MKD, demonstrating its potential in extremely label-scarce scenarios. **Third**, Our framework is fully compatible with mainstream SSL methods [14, 57]. The performance can be consistently lifted when equipped with additional perturbation techniques.

PASCAL VOC 2012 *augmented*. Table 1 (right) demonstrates the comparison results on the PASCAL VOC 2012 val using *augmented* set for training, where our method again sets new state-of-the-arts across all partition protocols, yielding an average gain of **0.80%** upon the previous SOTA [14]. When only provided with scarce labeled data, *i.e.*, 92 labeled images, our method still performs impressive owing to the compactly incorporated symbolic logic.

Cityscapes. Table 2 quantitatively compares our method

| Method | mIoU ² | mIoU ¹ | Speed <i>min.</i> | \mathcal{K}_C Eq. 12 | \mathcal{K}_D Eq. 13 | \mathcal{K}_E Eq. 14 | mIoU ² | mIoU ¹ | Resolution Strategy | mIoU ² | mIoU ¹ |
|---------------------------------------|-------------------|-------------------|-----------------------|---------------------------|---------------------------|---------------------------|-------------------|-------------------|------------------------|-------------------|-------------------|
| Baseline | - | 68.02 | 85.3 | | | | 84.34 | 69.15 | Uniform | 86.31 | 71.83 |
| + Hierarchical Prediction (§3.3) | 84.34 | 69.15 | 86.5 ^{+1.1%} | ✓ | | | 86.92 | 72.25 | Predictive | 86.65 | 72.05 |
| + Diagnosis Computation (§3.2.1) | 86.31 | 71.83 | 88.4 ^{+3.7%} | | ✓ | | 87.10 | 72.47 | Greedy | 87.29 | 72.61 |
| + Fuzzy Diagnosis Likelihood (§3.2.2) | 87.91 | 73.25 | 89.3 ^{+4.7%} | ✓ | ✓ | ✓ | 85.00 | 71.31 | Sampling | 87.91 | 73.25 |
| | | | | ✓ | ✓ | ✓ | 87.91 | 73.25 | | | |

(a) Key Components

(b) Semantic Logic Rules

(c) Resolution Strategy

Table 4: **Ablative experiments** on PASCAL VOC 2012 [28] val with 1/16 *augmented* training set. Please refer to §4.3 for more details.

against the competitors on Cityscapes val. In spite of the presence of complex street scenes, our method still delivers a solid overtaking trend across different partitions, with an averaged advancement of **1.65%** in terms of mIoU.

COCO. Table 3 presents the model performance on COCO val. As observed, by incorporating the large semantic hierarchy in COCO, LOGICDIAG provides even greater performance gains against the leading method (*i.e.*, MKD [14]) across all partitions by **2.94%** mIoU on average. The experimental results confirm again the efficacy of LOGICDIAG.

4.3. Diagnostic Experiment

For in-depth analysis, we perform a set of ablative studies on PASCAL VOC 2012 [28] val with 1/16 *augmented* set. Please refer to the supplementary for more experiments.

Key Component Analysis. In Table 4a, we first validate the importance of our proposed components by attaching them one at a time. The 1st row reports the result of a bare baseline model - DeepLabV3+ with plain consistency regularization [22]. Next, in the 2nd row, we convert the prediction mode from flat to hierarchical, which already boosts performance and supports our claim that hierarchical semantics can provide additional training signals implicitly. Moreover, the 3rd row gives the score when the minimal diagnosis set is further computed, and we uniformly sample one conflict from it to resolve. As seen, this leads to moderate improvement caused by the explicit introduction of symbolic knowledge that potentially resolves conflicts. Finally, as shown in the 4th row, through sampling from the diagnosis likelihood, the biggest improvement is achieved, demonstrating the necessity of the fuzzy measurement that guides the resolution.

Semantic Logic Rules. Then, we investigate the effectiveness of logic-induced hierarchy rules (§3.3) in Table 4b. Starting from the baseline (1st row), we individually add Composition (*cf.* Eq. 12), Decomposition (*cf.* Eq. 13), and Exclusion (*cf.* Eq. 14) rules, denoted as \mathcal{K}_C , \mathcal{K}_D , \mathcal{K}_E , *resp.*, into the proposed framework, resulting in the scores listed in the 2nd to 4th rows. The last row exhibits the outcome achieved with our full training regime. Upon examining the table, three observations can be made. **First**, incorporating each of the logic rule results in consistent performance gains, indicating that the set of rules captures diverse facets of visual semantics, and indeed benefits SSL models within our symbolic resolution framework. **Second**, the

best performance is attained by combining all three logic rules, highlighting the significance of comprehensive interpretation of structured semantic concepts. **Third**, this also implies that integrating extra symbolic knowledge has great potential to further enhance current sub-symbolic SSL.

Resolution Strategy. Table 4c reveals the impact of conflict resolution strategies (§3.2.2). The default strategy, ‘Sampling’, utilizes the Monte Carlo method to estimate the actual posterior distribution $P(\omega|\mathcal{S})$ (*cf.* Eq. 7). Here we investigate three alternatives: ‘Uniform’, which assigns equal weight to all valid diagnoses; ‘Predictive’, which samples solely based on predictive probability $P(o|\mathcal{S})$ without considering the degree of conflicts; ‘Greedy’, which resolves conflicts by prioritizing the highest probability according to $P(\omega|\mathcal{S})$. The results show that ‘Predictive’, ‘Greedy’, and ‘Sampling’ all outperform ‘Uniform’, indicating the necessity of likelihood modeling. Of these strategies, ‘Sampling’ stands out as the most effective, providing strong evidence for its capacity to escape from the spurious corrections.

Training Speed. For completeness, we include the training time of 10K iterations in the last column of Table 4a. Our training regime (the last row) incurs a trivial delay of $\sim 4.7\%$.

Inference Speed. During inference, the ancestral classes in the hierarchical classification head can be safely disregarded without introducing additional computational burden.

5. Conclusion

This paper introduces LOGICDIAG, a neural-logic SSL framework that consolidates the benefits from both symbolic reasoning and sub-symbolic learning. Through resolving conflicts within pseudo labels using logic-induced diagnoses, LOGICDIAG systematically compiles rich symbolic knowledge into the neural training pipeline. Experimental findings illustrate that LOGICDIAG outperforms existing SSL frameworks, especially in label scarce settings. The enhanced flexibility and generality of LOGICDIAG showcase the immense potential of this holistic neural-logic paradigm in pixel-wise semi-supervised learning. We believe this paper opens a new avenue for future exploration in the field.

Acknowledgement. This work was partly supported by the Fundamental Research Funds for the Central Universities (No. 226-2022-00051), and the National Natural Science Foundation of China (No. 62293554).

References

- [1] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 6, 7
- [2] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1
- [3] Ross Girshick. Fast r-cnn. In *ICCV*, 2015. 1
- [4] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 1
- [5] Yi Yang, Yueting Zhuang, and Yunhe Pan. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, pages 1551–1558, 2021. 1
- [6] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 1
- [7] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. In *ICCV*, 2021. 1
- [8] Chen Liang, Wenguan Wang, Jiayu Miao, and Yi Yang. Gmmseg: Gaussian mixture based generative semantic segmentation models. In *NeurIPS*, 2022. 1
- [9] James Liang, Tianfei Zhou, Dongfang Liu, and Wenguan Wang. Clustseg: Clustering for universal segmentation. In *ICML*, 2023. 1
- [10] Tianfei Zhou, Wenguan Wang, Ender Konukoglu, and Luc Van Gool. Rethinking semantic segmentation: A prototype view. In *CVPR*, 2022. 1
- [11] Jiayu Miao, Yunchao Wei, Yu Wu, Chen Liang, Guangrui Li, and Yi Yang. Vspw: A large-scale dataset for video scene parsing in the wild. In *CVPR*, 2021. 1
- [12] Wenguan Wang, Guolei Sun, and Luc Van Gool. Looking beyond single images for weakly supervised semantic segmentation learning. *IEEE TPAMI*, 2022. 1
- [13] Chen Liang, Wenguan Wang, Tianfei Zhou, Jiayu Miao, Yawei Luo, and Yi Yang. Local-global context aware transformer for language-guided video segmentation. *IEEE TPAMI*, 45(8):10055–10069, 2023. 1
- [14] Jianlong Yuan, Jinchao Ge, Qi Qian, Zhibin Wang, Fan Wang, and Yifan Liu. Semi-supervised semantic segmentation with mutual knowledge distillation. In *NeurIPS*, 2022. 1, 2, 7, 8
- [15] Hanzhe Hu, Fangyun Wei, Han Hu, Qiwei Ye, Jinshi Cui, and Liwei Wang. Semi-supervised semantic segmentation via adaptive equalization learning. In *NeurIPS*, 2021. 1, 2, 7
- [16] Ying Jin, Jiaqi Wang, and Dahua Lin. Semi-supervised semantic segmentation via gentle teaching assistant. In *NeurIPS*, 2022. 1, 7
- [17] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 1, 2, 7
- [18] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by multi-scale foreground-background integration. *IEEE TPAMI*, 44(9):4701–4712, 2021. 1
- [19] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. In *NeurIPS*, 2021. 1
- [20] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. In *NeurIPS*, 2022. 1
- [21] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshops*, 2013. 1, 2
- [22] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020. 1, 2, 3, 8
- [23] Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *IEEE IJCNN*, 2020. 1, 2, 4
- [24] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 1, 2, 3
- [25] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In *NeurIPS*, 2021. 1, 2, 3
- [26] Baixu Chen, Junguang Jiang, Ximei Wang, Jianmin Wang, and Mingsheng Long. Debaised self-training for semi-supervised learning. In *NeurIPS*, 2022. 1, 4
- [27] Johannes Bill, Hrag Pailian, Samuel J Gershman, and Jan Drugowitsch. Hierarchical structure is employed by humans during visual motion perception. *Proceedings of the National Academy of Sciences*, 117(39):24581–24589, 2020. 2
- [28] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 2015. 2, 6, 7, 8
- [29] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 2, 6, 7
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 6, 7

- [31] Kamal Nigam and Rayid Ghani. Analyzing the effectiveness and applicability of co-training. In *CIKM*, 2000. 2
- [32] Yves Grandvalet, Yoshua Bengio, et al. Semi-supervised learning by entropy minimization. In *NeurIPS*, 2004. 2
- [33] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *NeurIPS*, 2014. 2
- [34] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NeurIPS*, 2016. 2
- [35] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *ICLR*, 2019. 2, 3
- [36] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020. 2
- [37] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020. 2
- [38] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-supervised learning with dynamic thresholding. In *ICML*, 2021. 2
- [39] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2
- [40] Ahmet Iscen, Giorgos Toliias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *CVPR*, 2019. 2
- [41] Weiwei Shi, Yihong Gong, Chris Ding, Zhiheng MaXiaoYu Tao, and Nanning Zheng. Transductive semi-supervised deep learning using min-max features. In *ECCV*, 2018. 2
- [42] Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch: Semi-supervised learning with contrastive graph regularization. In *CVPR*, 2021. 2
- [43] Chia-Wen Kuo, Chih-Yao Ma, Jia-Bin Huang, and Zsolt Kira. Featmatch: Feature-based augmentation for semi-supervised learning. In *ECCV*, 2020. 2
- [44] Antti Rasmus, Mathias Berglund, Mikko Honkela, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *NeurIPS*, 2015. 2
- [45] Zhengyang Feng, Qianyu Zhou, Guangliang Cheng, Xin Tan, Jianping Shi, and Lizhuang Ma. Dmt: Dynamic mutual training for semi-supervised learning. *PR*, 2022. 2
- [46] Jianlong Yuan, Yifan Liu, Chunhua Shen, Zhibin Wang, and Hao Li. A simple baseline for semi-supervised semantic segmentation with strong data augmentation. In *ICCV*, 2021. 2, 7
- [47] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. St++: Make self-training work better for semi-supervised semantic segmentation. In *CVPR*, 2022. 2, 7
- [48] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Bayer. S4l: Self-supervised semi-supervised learning. In *ICCV*, 2019. 2
- [49] Geoff French, Timo Aila, Samuli Laine, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, high-dimensional perturbations. In *BMVC*, 2020. 2, 7
- [50] Zhanhan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *ECCV*, 2020. 2, 7
- [51] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, Xiao Bian, Jia-Bin Huang, and Tomas Pfister. Pseudoseg: Designing pseudo labels for semantic segmentation. In *ICLR*, 2021. 2, 6, 7
- [52] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *CVPR*, 2020. 2, 7
- [53] Jongmok Kim, Jooyoung Jang, and Hyunwoo Park. Structured consistency loss for semi-supervised semantic segmentation. In *CoRR*, 2020. 2
- [54] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic segmentation with directional context-aware consistency. In *CVPR*, 2021. 2
- [55] Peng Tu, Yawen Huang, Feng Zheng, Zhenyu He, Liujuan Cao, and Ling Shao. Guidedmix-net: Semi-supervised semantic segmentation by using labeled images as reference. In *AAAI*, 2022. 2
- [56] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *CVPR*, 2022. 2, 7
- [57] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo-labels. In *CVPR*, 2022. 2, 6, 7
- [58] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J Davison. Bootstrapping semantic segmentation with regional contrast. In *ICLR*, 2022. 2
- [59] Yuanyi Zhong, Bodi Yuan, Hong Wu, Zhiqiang Yuan, Jian Peng, and Yu-Xiong Wang. Pixel contrastive-consistent semi-supervised semantic segmentation. In *ICCV*, 2021. 2, 6, 7
- [60] Yanning Zhou, Hang Xu, Wei Zhang, Bin Gao, and Pheng-Ann Heng. C3-semiseg: Contrastive semi-supervised segmentation via cross-set learning and dynamic class-balancing. In *ICCV*, 2021. 2
- [61] Jiashuo Fan, Bin Gao, Huan Jin, and Lihui Jiang. Ucc: Uncertainty guided cross-head co-training for semi-supervised semantic segmentation. In *CVPR*, 2022. 2
- [62] Ashima Garg, Shaurya Bagga, Yashvardhan Singh, and Saket Anand. Hiermatch: Leveraging label hierarchies for improving semi-supervised learning. In *WACV*, 2022. 2
- [63] Zafer Barutcuoglu, Robert E Schapire, and Olga G Troyanskaya. Hierarchical multi-label prediction of gene function. *Bioinformatics*, 2006. 2

- [64] Ricardo Cerri, Rodrigo C Barros, and Andre CPLF de Carvalho. A genetic algorithm for hierarchical multi-label classification. In *SAC*, 2012. 2
- [65] Yu Li, Sheng Wang, Ramzan Umarov, Bingqing Xie, Ming Fan, Lihua Li, and Xin Gao. Deepre: sequence-based enzyme ec number prediction by deep learning. *Bioinformatics*, 2018. 2
- [66] Juho Rousu, Craig Saunders, Sandor Szedmak, and John Shawe-Taylor. Kernel-based learning of hierarchical multi-label classification models. *JMLR*, 2006. 2
- [67] Marcin Marszałek and Cordelia Schmid. Semantic hierarchies for visual object recognition. In *CVPR*, 2007. 2
- [68] Kristen Grauman, Fei Sha, and Sung Hwang. Learning a tree of metrics with disjoint visual features. In *NeurIPS*, 2011. 2
- [69] Sung Hwang, Kristen Grauman, and Fei Sha. Semantic kernel forests from multiple taxonomies. In *NeurIPS*, 2012. 2
- [70] Samy Bengio, Jason Weston, and David Grangier. Label embedding trees for large multi-class tasks. In *NeurIPS*, 2010. 2
- [71] Jia Deng, Alexander C Berg, Kai Li, and Li Fei-Fei. What does classifying more than 10,000 image categories tell us? In *ECCV*, 2010. 2
- [72] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *CVPR*, 2016. 2
- [73] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *NeurIPS*, 2013. 2
- [74] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, 2015. 2
- [75] Nakul Verma, Dhruv Mahajan, Sundararajan Sellamankam, and Vinod Nair. Learning hierarchical similarity metrics. In *CVPR*, 2012. 2
- [76] Bin Zhao, Fei Li, and Eric Xing. Large-scale category structure aware image categorization. In *NeurIPS*, 2011. 2
- [77] Karim Ahmed, Mohammad Haris Baig, and Lorenzo Torresani. Network of experts for large-scale image categorization. In *ECCV*, 2016. 2
- [78] Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu. Hd-cnn: hierarchical deep convolutional neural networks for large scale visual recognition. In *ICCV*, 2015. 2
- [79] Alon Zweig and Daphna Weinshall. Exploiting object hierarchy: Combining models from different category levels. In *ICCV*, 2007. 2
- [80] Liulei Li, Tianfei Zhou, Wenguan Wang, Jianwu Li, and Yi Yang. Deep hierarchical semantic segmentation. In *CVPR*, 2022. 2, 5, 6
- [81] Wenguan Wang, Hailong Zhu, Jifeng Dai, Yanwei Pang, Jianbing Shen, and Ling Shao. Hierarchical human parsing with typed part-relation reasoning. In *CVPR*, 2020. 2
- [82] Wenguan Wang, Tianfei Zhou, Siyuan Qi, Jianbing Shen, and Song-Chun Zhu. Hierarchical human semantic parsing with comprehensive part-relation modeling. *IEEE TPAMI*, 44(7):3508–3522, 2021. 2
- [83] Wenguan Wang, Zhijie Zhang, Siyuan Qi, Jianbing Shen, Yanwei Pang, and Ling Shao. Learning compositional neural information fusion for human parsing. In *ICCV*, 2019. 2
- [84] Oindrila Saha, Zezhou Cheng, and Subhransu Maji. Improving few-shot part segmentation using coarse supervision. In *ECCV*, 2022. 2
- [85] Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 1943. 3
- [86] Eleonora Giunchiglia and Thomas Lukasiewicz. Multi-label classification neural networks with hard logical constraints. *JAIR*, 72, 2021. 3
- [87] Komal Teru, Etienne Denis, and Will Hamilton. Inductive relation prediction by subgraph reasoning. In *ICML*, 2020. 3
- [88] Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. Kagnet: Knowledge-aware graph networks for commonsense reasoning. In *EMNLP*, 2019. 3
- [89] Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. In *NeurIPS*, 2017. 3
- [90] Artur d’Avila Garcez, Sebastian Bader, Howard Bowman, Luis C Lamb, Leo de Penning, BV Illuminoo, Hoifung Poon, and COPPE Gerson Zaverucha. Neural-symbolic learning and reasoning: A survey and interpretation. *Neuro-Symbolic Artificial Intelligence: The State of the Art*, 342, 2022. 3
- [91] Wenguan Wang and Yi Yang. Towards data-and knowledge-driven artificial intelligence: A survey on neuro-symbolic computing. *arXiv preprint arXiv:2210.15889*, 2022. 3
- [92] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017. 3
- [93] Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018. 3
- [94] Guillaume Lample and François Charton. Deep learning for symbolic mathematics. In *ICLR*, 2019. 3
- [95] Qing Li, Siyuan Huang, Yining Hong, Yixin Chen, Ying Nian Wu, and Song-Chun Zhu. Closed loop neural-symbolic learning via integrating neural perception, grammar parsing, and symbolic reasoning. In *ICML*, 2020. 3
- [96] Forough Arabshahi, Sameer Singh, and Animashree Anandkumar. Combining symbolic expressions and black-box function evaluations in neural programs. In *ICLR*, 2018. 3

- [97] Jiankai Sun, Hao Sun, Tian Han, and Bolei Zhou. Neuro-symbolic program search for autonomous driving decision module design. In *CoRL*, 2021. 3
- [98] Tom Silver, Ashay Athalye, Joshua B Tenenbaum, Tomas Lozano-Perez, and Leslie Pack Kaelbling. Learning neuro-symbolic skills for bilevel planning. In *CoRL*, 2022. 3
- [99] Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, and Yuke Zhu. Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs. In *ICRA*, 2021. 3
- [100] Miles Cranmer, Alvaro Sanchez Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David Spergel, and Shirley Ho. Discovering symbolic models from deep learning with inductive biases. In *NeurIPS*, 2020. 3
- [101] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873), 2021. 3
- [102] Ameesh Shah, Eric Zhan, Jennifer Sun, Abhinav Verma, Yisong Yue, and Swarat Chaudhuri. Learning differentiable programs with admissible neural heuristics. In *NeurIPS*, 2020. 3
- [103] Albert Tseng, Jennifer J Sun, and Yisong Yue. Automatic synthesis of diverse weak supervision sources for behavior analysis. In *CVPR*, 2022. 3
- [104] Giuseppe Marra, Francesco Giannini, Michelangelo Diligenti, and Marco Gori. Integrating learning and reasoning with deep logic models. In *ECML PKDD*, 2019. 3
- [105] Emile van Krieken, E Acar, and Frank van Harmelen. Semi-supervised learning using differentiable reasoning. *Journal of Applied Logic*, 6(4), 2019. 3
- [106] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Broeck. A semantic loss function for deep learning with symbolic knowledge. In *ICML*, 2018. 3
- [107] Efthymia Tsamoura, Timothy Hospedales, and Loizos Michael. Neural-symbolic integration: A compositional perspective. In *AAAI*, 2021. 3
- [108] A Garcez, M Gori, LC Lamb, L Serafini, M Spranger, and SN Tran. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *Journal of Applied Logics*, 6(4), 2019. 3
- [109] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *CVPR*, 2021. 3, 6, 7
- [110] Raymond Reiter. A theory of diagnosis from first principles. *Artificial intelligence*, 32(1), 1987. 4
- [111] Patrick Rodler. Random vs. best-first: Impact of sampling strategies on decision making in model-based diagnosis. In *AAAI*, 2022. 5
- [112] Bart Kosko and Satoru Isaka. Fuzzy logic. *Scientific American*, 269(1), 1993. 5
- [113] Petr Hájek. *Metamathematics of fuzzy logic*, volume 4. Springer Science & Business Media, 2013. 5
- [114] Kurt Gödel. *Kurt Gödel: collected works: volume I: publications 1929-1936*, volume 1. 1986. 5
- [115] Emile van Krieken, Erman Acar, and Frank van Harmelen. Analyzing differentiable fuzzy logic operators. *Artificial Intelligence*, 2022. 5, 6
- [116] Samy Badreddine, Artur d’Avila Garcez, Luciano Serafini, and Michael Spranger. Logic tensor networks. *Artificial Intelligence*, 303, 2022. 5
- [117] Jia Deng, Nan Ding, Yangqing Jia, Andrea Frome, Kevin Murphy, Samy Bengio, Yuan Li, Hartmut Neven, and Hartwig Adam. Large-scale object classification using label relation graphs. In *ECCV*, 2014. 6
- [118] Wei Bi and James T Kwok. Multi-label classification on tree-and dag-structured hierarchies. In *ICML*, 2011. 6
- [119] Joseph A Goguen. The logic of inexact concepts. *Synthese*, pages 325–373, 1969. 6
- [120] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6, 7
- [121] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *CVPR*, 2020. 6
- [122] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 6
- [123] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 2017. 6, 7
- [124] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 7
- [125] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 7
- [126] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 7