

455 **A Broader Impact**

456 Our proposed INVERT method significantly contributes to enhancing the transparency and safety
457 of deep neural networks. By providing human understandable and interpretable explanations for
458 neurons in black-box models, our approach offers valuable insights into their internal operations,
459 improving understanding. Moreover, our method is able to identify potentially harmful representations
460 through concept augmentation, which can be easily extended to cover a wider range of concepts.
461 An important advantage of our method is its notable reduction in computational cost compared to
462 previous approaches. This reduction not only improves efficiency but also minimizes the harmful
463 environmental impact associated with excessive GPU usage.

464 It is important to note that we cannot make definitive claims regarding specific groups of people
465 benefiting from or being disadvantaged by our method. The general applicability and potential
466 implications of our approach should be explored further and with caution.

467 **B Prior work**

468 Let’s consider a function, $g : \mathbb{D} \rightarrow \mathbb{R}^{k \times k}$, that signifies a convolutional neuron within a model that
469 produces activation maps of dimensions $k \times k$, along with a concept $c \in \mathcal{C}$. Both Network Dissection
470 [20] and Compositional Explanations of Neurons [22] methodologies make use of the Intersection
471 over Union (IoU) similarity metric to measure the degree of correlation between a function and a
472 concept. A prerequisite for these methodologies are segmentation masks pertaining to the concepts,
473 meaning for every concept $c \in \mathcal{C}$, there exists a corresponding function $M : \mathbb{D} \rightarrow \{0, 1\}^{h \times w}$, which
474 generates a binary mask for the specific concept, of the same size as the original input.

475 To evaluate the correlation between function g and concept c , the multi-dimensional outputs from g are
476 subjected to thresholding based on neuron-specific percentiles (i.e., values above chosen percentiles
477 are converted to 1 and the remaining to 0), and upsampled to match the dimensions of the original
478 image. We can define the resulting function that produces binary masks of the same size as the input
479 as $G : \mathbb{D} \rightarrow \{0, 1\}^{h \times w}$. The final similarity (IoU) score between g and c can be computed as the
480 Intersection over Union score between concept masks M and function G :

$$d_{IoU}(g, c) = \frac{\sum_{x \in \mathcal{D}} \mathbf{1}(M(x) \cap G(x))}{\sum_{x \in \mathcal{D}} \mathbf{1}(M(x) \cup G(x))}. \tag{4}$$

481 Given that previous methodologies were able to procure explanations solely from convolutional
482 neurons, we carried out a comparison with INVERT by computing INVERT explanations through
483 the mean of the activations across the activation maps of convolutional neurons. Specifically, in
484 section 4.2, the method of Compositional Explanations of neurons was applied using a 7x7 input map
485 for each feature. Conversely, the INVERT approach uses a strategy that computes a scalar value by
486 calculating the average of the input map.

487 **C INVERT algorithm**

488 Given a neural representation $f : \mathbb{D} \rightarrow \mathbb{R}$, a dataset $\mathcal{D} \subset \mathbb{D}$, and a set of concepts $\mathcal{C} \in \mathbb{C}$,
489 the INVERT approach seeks to identify a compositional concept φ^* , which is formed as a logical
490 operation on the concepts, to optimize AUC similarity $d(f, \varphi^*(\mathcal{C}))$. For this purpose, we utilized an
491 optimization process similar to that of the CompExpl methodology [22], employing Beam search to
492 find the optimal compositional concept.

493 This method requires the configuration of certain parameters, namely the predetermined formula
494 length $L \in \mathbb{N}$, the beam size $B \in \mathbb{N}$, and additionally, the threshold $T \in (0, 1/2)$. Beam search
495 intends to iteratively combine concepts, starting with the atomic concepts (primitives) from \mathcal{C} . At
496 every iteration of the process, the top B best-performing compositional concepts are selected, and
497 all feasible formulas are computed with primitives. Subsequently, only the top B best-performing
498 concepts are selected, and the process continues until the formula reaches the predetermined length.

499 In detail, firstly, we define a set of primitives $\bar{\Phi}$ — a set of compositional concepts that correspond to
500 the set of concepts \mathcal{C} and their negation. The set $\bar{\Phi}$ comprises $2k$ compositional concepts, with each

501 concept corresponding to either the base concept or its negation. Next, all $2k$ concepts are evaluated in
502 terms of AUC similarity with a given function, and the top B best performing compositional concepts,
503 that satisfy $p(\varphi(\mathcal{C})) \geq T$ are selected, leading to the formation of the set Φ^* where $|\Phi^*| = B$,
504 referred to as a Beam. These are the top B best-performing compositional concepts with a length
505 of 1, satisfying the requisite condition on their positive fraction in the dataset. Subsequently, the
506 following operations are iteratively performed until the predetermined formula length L is met:

- 507 1. Each of the B compositional concepts in the beam Φ^* is combined with all primitives using
508 either the AND or OR operation, thereby augmenting the formula length by 1, resulting in
509 a total of $4Bk$ new formulas.
- 510 2. All newly generated formulas are evaluated based on their similarity to the representation,
511 and the beam Φ^* is updated to include the top B performing formulas, which satisfy the
512 condition $p(\varphi(\mathcal{C})) \geq T$.

513 Upon reaching the predetermined formula length L , the Beam-Search procedure concludes by
514 identifying the compositional concept φ^* with the highest observed AUC.

515 D Comparing Fuzzy Logic operators

516 Fuzzy logic operators [38] serve as essential instruments within the domain of fuzzy logic, a mathe-
517 matical construct designed for modeling and handling data that is imprecise or vague. This contrasts
518 with conventional logic where an element strictly either belongs to a set or not; fuzzy logic allows for
519 the degree of membership to vary from 0 to 1, thereby allowing for partial membership.

520 In the present study, our objective was to contrast different categories of fuzzy logic operators and
521 examine their behavior concerning the proposed AUC metric. To fulfill this aim, we employed four
522 distinct deep learning image classification models: AlexNet [33], DenseNet161 [34], EfficientNet
523 B4 [35], and ViT 16 L [36]. Each of these models was pre-trained on the ImageNet dataset. We
524 focused on 1000 neural representations in the output logit (pre-SoftMax) layer for each model, for
525 which we recognized the 'ground-truth' concept — the corresponding ImageNet class. For fuzzy
526 logic operators' testing, we mapped the output of each individual representation to the set $[0, 1]$ by
527 normalizing each representation's output using their corresponding mean and standard deviation
528 across ImageNet dataset and applied a Sigmoid transformation. We tested four different Fuzzy logic
529 operators, specifically Gödel, Product, Łukasiewicz, and Yager with parameter $p = 2$, as illustrated
530 in Table 2.

531 For performance evaluation, we generated random compositional concepts of a given length and
532 computed the AUC similarity between fuzzy logic norms applied to functions corresponding to
533 these concepts. For instance, given the random compositional concept $\varphi = c_i \text{ OR } c_j$, we derive
534 compositional representations as per each of the four examined methods (e.g., the Gödel operator
535 produces function $h_G = \max(f_i, f_j)$). These compositional representations are then evaluated in
536 terms of AUC similarity with the compositional concept.

537 We conducted the evaluation in two modes, that is, assessing the performance of the OR (T-conorm)
538 operator and the performance of the AND (T-norm) operator. For each mode, we assembled 1000
539 random compositional concepts by sampling L concepts without replacement and calculated the
540 AUC between compositional concepts. Note that for the second mode, AND (T-norm), random
541 compositional concepts were assembled using the AND NOT operation, given the mutual exclusivity
542 of ImageNet labels.

543 Figures 8 and 9 depict the mean AUC similarity between random compositional concepts of varying
544 lengths and the corresponding compositional representations, which were assembled using four
545 distinct fuzzy logic operators. From these figures, it becomes evident that Gödel fuzzy logic operators
546 demonstrate the most significant robustness to the length of the formula, consistently attaining
547 superior AUC in contrast to other operators. Consequently, we can infer that Gödel's operator
548 emerges as the optimal choice for implementing fuzzy logic operations on these representations.

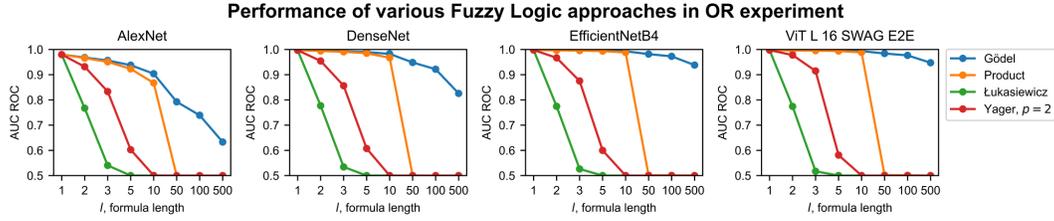


Figure 8: Average AUC similarity between random compositional OR concepts and corresponding compositional representations employing various Fuzzy logic operators (Higher is better) evaluated across four distinct models.

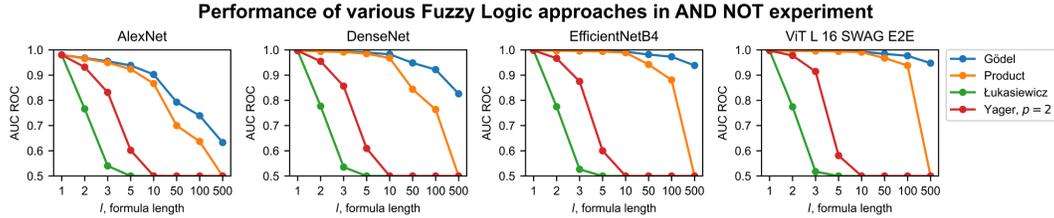


Figure 9: Average AUC similarity between random compositional AND NOT concepts and corresponding compositional representations employing various Fuzzy logic operators (Higher is better) evaluated across four distinct models.

Table 2: List of different fuzzy operators

	NOT(a)	AND(a, b) (T -norm)	OR(a, b) (T -conorm)
<i>Gödel</i>	$1 - a$	$\min(a, b)$	$\max(a, b)$
<i>Product</i>		$a \cdot b$	$a + b - a \cdot b$
<i>Łukasiewicz</i>		$\max(a + b - 1, 0)$	$\min(a + b, 1)$
<i>Yager, $p = 2$</i>		$\max(1 - ((1 - a)^2 + (1 - b)^2)^{\frac{1}{2}}, 0)$	$\min((a^2 + b^2)^{\frac{1}{2}}, 1)$