

Supplemental Materials

498 A Details on Neurally-Guided Symbolic Abstraction

499 We here provide details on the neurally-guided symbolic abstraction algorithm.

500 A.1 Algorithm of Neurally-Guided Symbolic Abstraction

501 We show the algorithm of neurally-guided symbolic abstraction in Algorithm 1.

Algorithm 1 Neurally-Guided Symbolic Abstraction

Input: C_0, π_θ , hyperparameters (N_{beam}, T_{beam})

```

1:  $C_{to\_open} \leftarrow C_0$ 
2:  $\mathcal{C} \leftarrow \emptyset$ 
3:  $t = 0$ 
4: while  $t < T_{beam}$  do
5:    $C_{beam} \leftarrow \emptyset$ 
6:   for  $C_i \in C_{to\_open}$  do
7:      $\mathcal{C} = \mathcal{C} \cup \{C_i\}$ 
8:     for  $R \in \rho(C_i)$  do
9:       # Evaluate each clause
10:       $score = eval(R, \pi_\theta)$ 
11:      # select top-k rules
12:       $C_{beam} = top\_k(C_{beam}, R, score, N_{beam})$ 
13:      # selected rules are refined next
14:    $C_{to\_open} = C_{beam}$ 
15:    $t = t + 1$ 
16: return  $\mathcal{C}$ 

```

502 A.2 Rule Generation

503 At line 8 in Algorithm 1, given action rule C , we generate new action rules using the following
504 refinement operation:

$$\rho(C) = \{X_A \leftarrow X_S^{(1)}, \dots, X_S^{(n)}, Y_S \mid Y_S \in \mathcal{G}_S^* \wedge Y_S \neq X_S^{(i)}\}, \quad (7)$$

505 where \mathcal{G}_S^* is a non-ground state atoms. This operation is a specification of (*downward*) *refinement*
506 *operator*, which a fundamental technique for rule learning in ILP [Nienhuys-Cheng and de Wolf,
507 1997], for action rules to solve RL tasks.

508 We use mode declarations [Muggleton, 1995, Cropper et al., 2022] to define the search space,
509 *i.e.* \mathcal{G}_S^* in Eq. 7 which are defined as follows. A mode declaration is either a head declaration
510 $modeh(r, p(mdt_1, \dots, mdt_n))$ or a body declaration $modeb(r, p(mdt_1, \dots, mdt_n))$, where $r \in \mathbb{N}$ is
511 an integer, p is a predicate, and mdt_i is a mode datatype. A mode datatype is a tuple (pm, dt) , where
512 pm is a place-marker and dt is a datatype. A place-marker is either $\#$, which represents constants, or
513 $+$ (resp. $-$), which represents input (resp. output) variables. r represents the number of the usages of
514 the predicate to compose a solution. Given a set of mode declarations, we can determine a finite set
515 of rules to be generated by the rule refinement.

516 Now we describe mode declarations we used in our experiments. For Getout, we used the following
517 mode declarations:

```

modeb(2, type(-object, +type))
modeb(1, closeby(+object, +object))
modeb(1, on_left(+object, +object))
modeb(1, on_right(+object, +object))
modeb(1, have_key(+object))
modeb(1, not_have_key(+object))

```

518 For 3Fishes, we used the following mode declarations:

```
modeb(2, type(-object, +type))
modeb(1, closeby(+object, +object))
modeb(1, on_top(+object, +object))
modeb(1, at_bottom(+object, +object))
modeb(1, on_left(+object, +object))
modeb(1, on_right(+object, +object))
modeb(1, bigger_than(+object, +object))
modeb(1, high_level(+object, +object))
modeb(1, low_level(+object, +object))
```

519 For Loot, we used the following mode declarations:

```
modeb(2, type(-object, +type))
modeb(2, color(+object, #color))
modeb(1, closeby(+object, +object))
modeb(1, on_top(+object, +object))
modeb(1, at_bottom(+object, +object))
modeb(1, on_left(+object, +object))
modeb(1, on_right(+object, +object))
modeb(1, have_key(+object))
```

520 B Additional Results

521 B.1 Weights learning

522 Fig. 6 shows the NUDGE agent $\pi_{(\mathcal{C}, \mathbf{W})}$ parameterized by rules \mathcal{C} and weights \mathbf{W} before training
523 (top) and after training (bottom) on the GetOut environment. Each element on the x-axis of the plots
524 corresponds to an action rule. In this examples, we have 10 action rules $\mathcal{C} = \{C_0, C_1, \dots, C_9\}$, and
525 we assign $M = 5$ weights *i.e.* $\mathbf{W} = [\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_4]$. The distributions of rule weighs with *softmax*
526 are getting peaked by learning to maximize the return. The right 4 rules are redundant rules, and
527 theses rules get low weights after learning.

528 B.2 Deduction Pipeline

529 Fig. 7 provides the deduction pipeline of a NUDGE agent on 3 different states. Facts can be deduced
530 from an object detection method, or directly given by the object centric environment. For state #1,
531 the agent chooses to jump as the jump action is prioritized over the other ones and all atoms that
532 compose this rules' body have high valuation (including closeby). In state #2, the agent chose to go
533 **left** as the rule left_key is selected. In state #3, the agent selects **right** as the rule right_door has
534 the highest forward chaining evaluation.

535 B.3 Policies of every logic environment.

536 We show the logic policies obtained by NUDGE in GetOut, 3Fishes, and Loot in Fig. 8, *e.g.* the first
537 line of GetOut, “0.574 : jump(X) :- closeby(01, 02), type(01, agent), type(02, enemy).”, repre-
538 sents that the action rule is chosen by the weight vector \mathbf{w}_1 with a value 0.574. NUDGE agents
539 have several weight vectors $\mathbf{w}_1, \dots, \mathbf{w}_M$ and thus several chosen action rules are shown for each
540 environment.

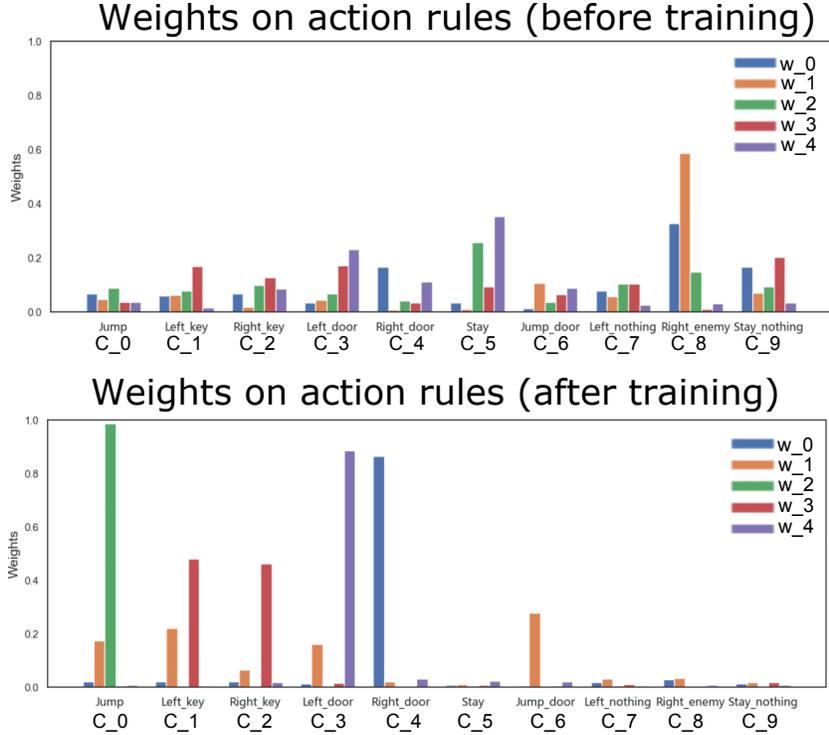


Figure 6: Weights on action rules via softmax before training (top) and after training (bottom) on NUDGE in GetOut. Each element on the x-axis of the plots corresponds to an action rule. NUDGE learns to get high returns while identifying useful action rules to solve the RL task. The right 5 rules are redundant rules, and these rules get low weights after learning.

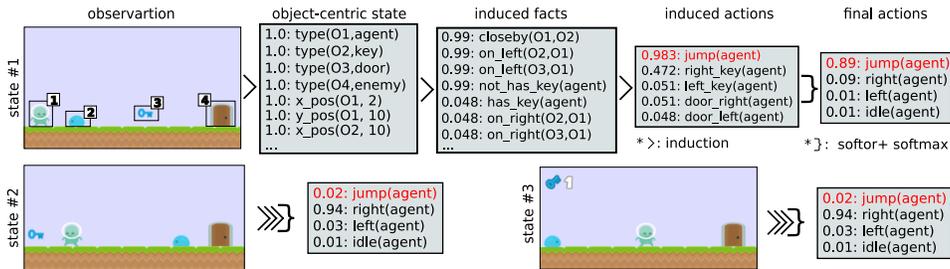


Figure 7: The logic reasoning of NUDGE agents makes them interpretable. The detailed logic pipeline for the input state #1 of the *Getout* environment and the condensed action selection for state #2 and state #3.

541 **C Illustrations of our environments**

542 We showcase in Fig. 9 one state of the 3 object-centric environments and their variations. In **GetOut**
 543 (blue humanoid agent), the goal is to obtain a key, then go to a door, while avoiding a moving enemy.
 544 **GetOut-2En** is a variation with 2 enemies. In **3Fishes**, the agent controls a green fish and is
 545 confronted with 2 other fishes, one smaller (that the agent need to “eat”, *i.e.* go to) and one bigger,
 546 that the agent needs to dodge. A variation is **3Fishes-C**, where the agent can eat green fishes and
 547 dodge red ones, all fishes have the same size. Finally, in **Loot**, the (orange) agent is exposed with 1
 548 or 2 chests and their corresponding (*i.e.* same color) keys. In **Loot-C**, the chests have different colors.
 549 All 3 environment are *stationary* in the sense of Delfosse et al. [2021].

```

# GetOut
0.574: jump(X) :- closeby(01,02), type(01,agent), type(02,enemy).
0.315: right_go_to_door(X) :- have_key(X), on_left(01,02), type(01,agent), type(02,door).
0.296: right_go_to_door(X) :- have_key(X), on_left(01,02), type(01,agent), type(02,door).
0.291: right_go_get_key(X) :- not_have_key(X), on_left(01,02), type(01,agent),
    type(02,key).
0.562: right_go_to_door(X) :- have_key(X), on_left(01,02), type(01,agent), type(02,door).

#3Fishes
0.779: right_to_eat(X) :- is_bigger_than(01,02), on_left(02,01), type(01,agent),
    type(02,fish).
0.445: down_to_dodge(X) :- is_bigger_than(02,01), on_left(02,01), type(01,agent),
    type(02,fish).
0.579: down_to_eat(X) :- high_level(01,02), is_smaller_than(02,01), type(01,agent),
    type(02,fish).
0.699: up_to_dodge(X) :- closeby(02,01), is_smaller_than(01,02), low_level(02,01),
    type(01,agent), type(02,fish).
0.601: up_to_eat(X) :- is_bigger_than(02,01), on_left(02,01), type(01,agent),
    type(02,fish).
0.581: left_to_eat(X) :- closeby(01,02), on_right(01,02), type(01,agent), type(02,fish).

# Loot
0.844: up_to_door(X) :- close(01,02), have_key(02), on_top(02,01), type(01,agent),
    type(02,door).
0.268: right_to_key(X) :- close(01,02), on_right(02,01), type(01,agent), type(02,key).
0.732: right_to_door(X) :- close(01,02), have_key(02), on_left(01,02), type(01,agent),
    type(02,door).
0.508: up_to_key(X) :- close(01,02), on_top(02,01), type(01,agent), type(02,key).
0.995: left_to_door(X) :- close(01,02), have_key(02), on_left(02,01), type(01,agent),
    type(02,door).
0.414: down_to_key(X) :- close(01,02), on_top(01,02), type(01,agent), type(02,key).
0.992: down_to_door(X) :- close(01,02), have_key(02), on_top(01,02), type(01,agent),
    type(02,door).
0.447: left_to_key(X) :- close(01,02), on_left(02,01), type(01,agent), type(02,key).

```

Figure 8: **NUDGE produces an interpretable policy as set of weighted rules.** Weighted action rules discovered by NUDGE in the each logic environment.

550 D Hyperparameters and rules sets

551 D.1 Hyperparameters

552 We here provide the hyperparameters used in our experiments. We set the clip parameter $\epsilon_{clip} = 0.2$,
553 the discount factor $\gamma = 0.99$. We use the Adam optimizer, with $1e - 3$ as actor learning rate, $3e - 4$
554 as critic learning rate. The episode length is 500 timesteps. The policy is updated every 1000 steps
555 We train every algorithm for $800k$ steps on each environment, apart from neural PPO, that needed
556 $5M$ steps on Loot. We use an epsilon greedy strategy with $\epsilon = \max(e^{-\frac{episode}{500}}, 0.02)$.

557 D.2 Rules set

558 All the rules set \mathcal{C} of the different NUDGE and logic agents are available at <https://anonymous.4open.science/r/LogicRL-C43B> in the folder nsfr/nsfr/data/lang.

560 E Details of Differentiable Forward Reasoning

561 We provide details of differentiable forward reasoning used in NUDGE. We denote a valuation vector
562 at time step t as $\mathbf{v}^{(t)} \in [0, 1]^G$. We also denote the i -th element of vector \mathbf{x} by $\mathbf{x}[i]$, and the (i, j) -th
563 element of matrix \mathbf{X} by $\mathbf{X}[i, j]$. The same applies to higher dimensional tensors.

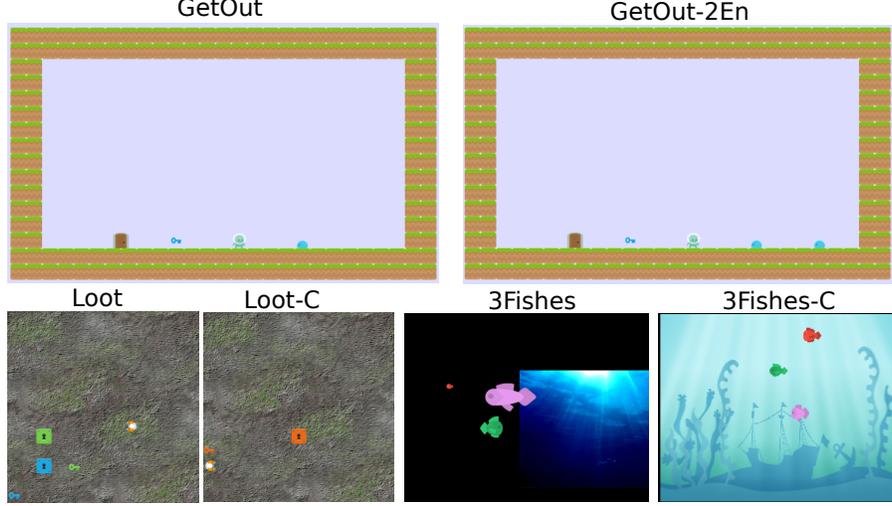


Figure 9: Pictures of our environments (**GetOut**, **Loot** and **3Fishes**) and their variations (**GetOut-2En**, **Loot-C** and **3Fishes-C**). All these environments can provide object-centric state descriptions (instead of pixel-based states).

564 E.1 Differentiable Forward Reasoning

565 We compose the reasoning function $f_{(\mathcal{C}, \mathbf{W})}^{reason} : [0, 1]^G \rightarrow [0, 1]^{G_A}$, which takes the initial valuation
 566 vector and returns valuation vector for induced action atoms. We describe each step in detail.

567 **(Step 1) Encode Logic Programs to Tensors.** To achieve differentiable forward reasoning, each
 568 action rule is encoded to a tensor representation. Let S be the number of the maximum number of
 569 substitutions for existentially quantified variables in \mathcal{C} , and L be the maximum length of the body of
 570 rules in \mathcal{C} . Each action rule $C_i \in \mathcal{C}$ is encoded to a tensor $\mathbf{I}_i \in \mathbb{N}^{G \times S \times L}$, which contain the indices
 571 of body atoms. Intuitively, $\mathbf{I}_i[j, k, l]$ is the index of the l -th fact (subgoal) in the body of the i -th rule
 572 to derive the j -th fact with the k -th substitution for existentially quantified variables.

573 For example, let $R_0 = \text{jump}(\text{agent}) : \text{-type}(01, \text{agent}), \text{type}(02, \text{enemy}), \text{closeby}(01, 02) \in \mathcal{C}$
 574 and $F_2 = \text{jump}(\text{agent}) \in \mathcal{G}$, and we assume that constants for objects are $\{\text{obj1}, \text{obj2}\}$. R_0 has
 575 existentially quantified variables 01 and 02 on the body, so we obtain ground rules by substituting
 576 constants. By considering the possible substitutions for 01 and 02, namely $\{01/\text{obj1}, 02/\text{obj2}\}$
 577 and $\{01/\text{obj2}, 02/\text{obj1}\}$, we have *two* ground rules, as shown in top of Table 3. Bottom rows
 578 of Table 3 shows elements of tensor $\mathbf{I}_{0, :, 0, :}$ and $\mathbf{I}_{0, :, 1, :}$. Facts \mathcal{G} and the indices are represented on
 579 the upper rows in the table. For example, $\mathbf{I}_{0, 2, 0, :} = [3, 6, 7]$ because R_0 entails $\text{jump}(\text{agent})$
 580 with the first ($k = 0$) substitution $\tau = \{01/\text{obj1}, 02/\text{obj2}\}$. Then the subgoal atoms are
 581 $\{\text{type}(\text{obj1}, \text{agent}), \text{type}(\text{obj2}, \text{enemy}), \text{closeby}(\text{obj1}, \text{obj2})\}$, which have indices $[3, 6, 7]$, re-
 582 spectively. The atoms which have a different predicate, e.g., $\text{closeby}(\text{obj1}, \text{obj2})$, will never be
 583 entailed by clause R_0 . Therefore, the corresponding values are filled with 0, which represents the
 584 index of the *false* atom.

585 **(Step 2) Assign Rule Weights.** We assign weights to compose the policy with several action rules
 586 as follows: (i) We fix the target programs' size as M , *i.e.*, where we try to find a policy with M
 587 action rules. (ii) We introduce C -dim weights $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_M]$. (iii) We take the *softmax* of each
 588 weight vector $\mathbf{w}_j \in \mathbf{W}$ and softly choose M action rules out of C action rules to compose the policy.

589 **(Step 3) Perform Differentiable Inference.** We compute 1-step forward reasoning using weighted
 590 action rules, then we recursively perform reasoning to compute T -step reasoning.

591 **[(i) Reasoning using an action rule]** First, for each action rule $C_i \in \mathcal{C}$, we evaluate body atoms for
 592 different grounding of C_i by computing $b_{i,j,k}^{(t)} \in [0, 1]$:

$$b_{i,j,k}^{(t)} = \prod_{1 \leq l \leq L} \text{gather}(\mathbf{v}^{(t)}, \mathbf{I}_i)[j, k, l] \quad (8)$$

($k = 0$) $\text{jump}(\text{agent}) : -\text{type}(\text{obj1}, \text{agent}), \text{type}(\text{obj2}, \text{enemy}), \text{closeby}(\text{obj1}, \text{obj2}).$
($k = 1$) $\text{jump}(\text{agent}) : -\text{type}(\text{obj2}, \text{agent}), \text{type}(\text{obj1}, \text{enemy}), \text{closeby}(\text{obj2}, \text{obj1}).$

j	0	1	2	3	4	5
\mathcal{G}	\perp	\top	$\text{jump}(\text{agent})$	$\text{type}(\text{obj1}, \text{agent})$	$\text{type}(\text{obj2}, \text{agent})$	$\text{type}(\text{obj1}, \text{enemy})$
$\mathbf{I}_{0,j,0,:}$	[0, 0, 0]	[1, 1, 1]	[3, 6, 7]	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]
$\mathbf{I}_{0,j,1,:}$	[0, 0, 0]	[1, 1, 1]	[4, 5, 8]	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]

j	6	7	8	...
\mathcal{G}	$\text{type}(\text{obj2}, \text{enemy})$	$\text{closeby}(\text{obj1}, \text{obj2})$	$\text{closeby}(\text{obj2}, \text{obj1})$...
$\mathbf{I}_{0,j,0,:}$	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]	...
$\mathbf{I}_{0,j,1,:}$	[0, 0, 0]	[0, 0, 0]	[0, 0, 0]	...

Table 3: Example of ground rules (top) and elements in the index tensor (bottom). Each fact has its index, and the index tensor contains the indices of the facts to compute forward inferences.

593 where $\text{gather} : [0, 1]^G \times \mathbb{N}^{G \times S \times L} \rightarrow [0, 1]^{G \times S \times L}$ is:

$$\text{gather}(\mathbf{x}, \mathbf{Y})[j, k, l] = \mathbf{x}[\mathbf{Y}[j, k, l]]. \quad (9)$$

594 The gather function replaces the indices of the body state atoms by the current valuation values
595 in $\mathbf{v}^{(t)}$. To take logical *and* across the subgoals in the body, we take the product across valuations.
596 $b_{i,j,k}^{(t)}$ represents the valuation of body atoms for i -th rule using k -th substitution for the existentially
597 quantified variables to deduce j -th fact at time t .

598 Now we take logical *or* softly to combine all of the different grounding for C_i by computing
599 $c_{i,j}^{(t)} \in [0, 1]$:

$$c_{i,j}^{(t)} = \text{softor}^\gamma(b_{i,j,1}^{(t)}, \dots, b_{i,j,S}^{(t)}) \quad (10)$$

600 where softor^γ is a smooth logical *or* function:

$$\text{softor}^\gamma(x_1, \dots, x_n) = \gamma \log \sum_{1 \leq i \leq n} \exp(x_i/\gamma), \quad (11)$$

601 where $\gamma > 0$ is a smooth parameter. Eq. 11 is an approximation of the *max* function over probabilistic
602 values based on the *log-sum-exp* approach [Cuturi and Blondel, 2017].

603 **[(ii) Combine results from different action rules]** Now we apply different action rules using the
604 assigned weights by computing $h_{j,m}^{(t)} \in [0, 1]$:

$$h_{j,m}^{(t)} = \sum_{1 \leq i \leq C} w_{m,i}^* \cdot c_{i,j}^{(t)}, \quad (12)$$

605 where $w_{m,i}^* = \exp(w_{m,i}) / \sum_{i'} \exp(w_{m,i'})$, and $w_{m,i} = \mathbf{w}_m[i]$. Note that $w_{m,i}^*$ is interpreted as a
606 probability that action rule $C_i \in \mathcal{C}$ is the m -th component of the policy. Now we complete the 1-step
607 forward reasoning by combining the results from different weights:

$$r_j^{(t)} = \text{softor}^\gamma(h_{j,1}^{(t)}, \dots, h_{j,M}^{(t)}). \quad (13)$$

608 Taking softor^γ means that we compose the policy using M softly chosen action rules out of C
609 candidates of rules.

610 **[(iii) Multi-step reasoning]** We perform T -step forward reasoning by computing $r_j^{(t)}$ recursively for
611 T times: $v_j^{(t+1)} = \text{softor}^\gamma(r_j^{(t)}, v_j^{(t)})$. Finally, we compute $\mathbf{v}^{(T)} \in [0, 1]^G$ and returns $\mathbf{v}_A \in [0, 1]^{G_A}$
612 by extracting only output for action atoms from $\mathbf{v}^{(T)}$. The whole reasoning computation Eq. 8-13
613 can be implemented using only efficient tensor operations. See App. E.2 for a detailed description.

614 E.2 Implementation Details

615 Here we provide implementational details of the differentiable forward reasoning. The whole
616 reasoning computations in NUDGE can be implemented as a neural network that performs forward

617 reasoning and can efficiently process a batch of examples in parallel on GPUs, which is a non-trivial
 618 function of logical reasoners.

619 Each clause $C_i \in \mathcal{C}$ is compiled into a differentiable function that performs forward reasoning using
 620 the tensor. The clause function is computed as:

$$\mathbf{C}_i^{(t)} = \text{softor}_3^\gamma \left(\text{prod}_2 \left(\text{gather}_1 \left(\tilde{\mathbf{V}}^{(t)}, \tilde{\mathbf{I}} \right) \right) \right), \quad (14)$$

621 where $\text{gather}_1(\mathbf{X}, \mathbf{Y})_{i,j,k,l} = \mathbf{X}_{i,\mathbf{Y}_{i,j,k,l,k,l}}$ ² obtains valuations for body atoms of the clause C_i
 622 from the valuation tensor and the index tensor. prod_2 returns the product along dimension 2, *i.e.* the
 623 product of valuations of body atoms for each grounding of C_i . The softor^γ function is applied along
 624 dimension 3, on all the grounding (or possible substitutions) of C_i .

625 softor_d^γ is a function for taking logical *or* softly along dimension d :

$$\text{softor}_d^\gamma(\mathbf{X}) = \gamma \log \left(\text{sum}_d \exp(\mathbf{X}/\gamma) \right), \quad (15)$$

626 where $\gamma > 0$ is a smoothing parameter, sum_d is the sum function along dimension d . The results
 627 from each clause $\mathbf{C}_i^t \in \mathbb{R}^{B \times G}$ is stacked into tensor $\mathbf{C}^{(t)} \in \mathbb{R}^{C \times B \times G}$.

628 Finally, the T -time step inference is computed by amalgamating the inference results recursively. We
 629 take the softmax of the clause weights, $\mathbf{W} \in \mathbb{R}^{M \times C}$, and softly choose M clauses out of C clauses
 630 to compose the logic program:

$$\mathbf{W}^* = \text{softmax}_1(\mathbf{W}). \quad (16)$$

631 where softmax_1 is a softmax function over dimension 1. The clause weights $\mathbf{W}^* \in \mathbb{R}^{M \times C}$ and
 632 the output of the clause function $\mathbf{C}^{(t)} \in \mathbb{R}^{C \times B \times G}$ are expanded (via copy) to the same shape
 633 $\tilde{\mathbf{W}}^*, \tilde{\mathbf{C}}^{(t)} \in \mathbb{R}^{M \times C \times B \times G}$. The tensor $\mathbf{H}^{(t)} \in \mathbb{R}^{M \times B \times G}$ is computed as

$$\mathbf{H}^{(t)} = \text{sum}_1(\tilde{\mathbf{W}}^* \odot \tilde{\mathbf{C}}), \quad (17)$$

634 where \odot is element-wise multiplication. Each value $\mathbf{H}_{i,j,k}^{(t)}$ represents the weight of k -th ground atom
 635 using i -th clause weights for the j -th example in the batch. Finally, we compute tensor $\mathbf{R}^{(t)} \in \mathbb{R}^{B \times G}$
 636 corresponding to the fact that logic program is a set of clauses:

$$\mathbf{R}^{(t)} = \text{softor}_0^\gamma(\mathbf{H}^{(t)}). \quad (18)$$

637 With r the 1-step forward-chaining reasoning function:

$$r(\mathbf{V}^{(t)}; \mathbf{I}, \mathbf{W}) = \mathbf{R}^{(t)}, \quad (19)$$

638 we compute the T -step reasoning using:

$$\mathbf{V}^{(t+1)} = \text{softor}_1^\gamma \left(\text{stack}_1 \left(\mathbf{V}^{(t)}, r(\mathbf{V}^{(t)}; \mathbf{I}, \mathbf{W}) \right) \right), \quad (20)$$

639 where $\mathbf{I} \in \mathbb{N}^{C \times G \times S \times L}$ is a precomputed index tensor, and $\mathbf{W} \in \mathbb{R}^{M \times C}$ is clause weights. After T -
 640 step reasoning, the probabilities over action atoms \mathcal{G}_A are extracted from $\mathbf{V}^{(T)}$ as $\mathbf{V}_A \in [0, 1]^{B \times G_A}$.

²done with pytorch.org/docs/torch.gather