
Supplementary Material

Anonymous Author(s)
Affiliation
Address
email

1 The appendix mainly consists of 5 parts.

- 2 • Appendix A: We present details about the key ideas and techniques used in the paper but
3 is not discussed in details due to page limits, including details of the model Neural DNF
4 and the learning algorithm BOAT.
- 5 • Appendix B: We discuss related works on interpretability research. we first answer why we
6 prefer inherently interpretable models, we only mention briefly in the main paper, here we
7 discuss more. We then talk about recent works on interpretable deep learning. We end with
8 a discussion of a comparison of interpretability of rule-based models and linear model.
- 9 • Appendix C: We discuss more about Disjunctive Normal Form (DNF). for classification,
10 previous works on learning DNF and more closely, some recent works to learn DNF by
11 gradient descent using a differentiable DNF function.
- 12 • Appendix D: We relate Neural DNF to the literature of neuro-symbolic integration.
- 13 • Appendix E: We present details of experiment setup and experimental results not covered
14 in the main paper.

15 Contents

16	A Algorithm Details	2
17	A.1 The modified <i>Bop</i>	2
18	A.2 BOAT	2
19	A.3 Improved SemHash	2
20	A.4 Initialization of \tilde{W}	3
21	A.5 Regularization of DNF: R_g and possible Alternatives	4
22	A.6 Possible direction of improvement	4
23	B Related Works on Interpretability	4
24	B.1 Inherent interpretability, not Post-hoc interpretation	4
25	B.2 Related works on Interpretable Deep Learning	5
26	B.3 why favouring the symbolic DNF for interpretability	5
27	C DNF	6
28	C.1 Differentiable replacement of the DNF function g	6
29	D Connections to Neural-symbolic Integration	7
30	E Experiment	8
31	E.1 Evaluation of the BOAT Optimizer	8
32	E.2 The 2d-XOR	10
33	E.3 MNIST	10
34	E.4 Other datasets in Section 4.2	11

35 A Algorithm Details

36 we introduce some key ideas and techniques and the details of algorithm

37 A.1 The modified Bop

38 Helwegen et al. [2019] proposes the *Bop* optimizer in the context of binarized neural networks of
39 value $\{-1, 1\}$. As suggested by Helwegen et al. [2019], the Bop can be viewed as a basic (gradient-
40 based) binary optimizer, in the same sense that SGD is a basic (gradient-based) continuous-valued
41 optimizer. The update rule for the original Bop is

$$w = \begin{cases} -w, & \text{if } |m| > \tau \text{ and } \text{sign}(w) = \text{sign}(m) \\ w, & \text{otherwise.} \end{cases}$$

42 And the modified bop for $\{0, 1\}$ used in this paper is

$$w = \begin{cases} 1 - w, & \text{if } |m| > \tau \text{ and } (w = 1 \text{ and } m > 0 \text{ or } w = 0 \text{ and } m < 0) \\ w, & \text{otherwise.} \end{cases}$$

43 The modification of suiting the case of $\{-1, 1\}$ to $\{0, 1\}$ is minor.

44 Also, in the implementation we add a bias correction procedure. Let \hat{m}_t be the non-corrected gradi-
45 ent momentum that is updated as $\hat{m}_t = \gamma\hat{m}_{t-1} + (1 - \gamma)\nabla$, The bias correction is given by

$$m_t = \hat{m}_t / (1 - \gamma^t)$$

46 We need this correction because in the original Bop paper Helwegen et al. [2019] uses random
47 initialization for m but in our implementation we make specific that m is zero-initialized. So we
48 will need this bias correction.

49 A.2 BOAT

50 Basically, BOAT consists of the modified bop and the proposed temperatured noise. We differ from
51 the Helwegen et al. [2019] as follows: (1) First, we fit into the case of $\{0, 1\}$ instead of $\{-1, 1\}$,
52 which is trivial; (2) we introduces the adaptively-temperatured noise controlled by the learnable
53 temperature parameter. (3) we add a bias-correction procedure.

54 We mentioned in the main paper that we suspect the noise smoothes the loss surfaces so to help the
55 optimization. But it is far from a rigorous statement and we do lack some theoretical understanding
56 of why and how this noise helps learning.

57 We refer interested readers to the original Bop paper Helwegen et al. [2019] or a more recent paper
58 by Meng et al. [2020] who discusses the connection of Bop and STE-Binary network and also
59 provides a bayesian perspective on binary network learning. We believe that the our method BOAT
60 can in principle be linked to approximate variational inference [Khan et al., 2018, Kingma et al.,
61 2015, Potapczynski et al., 2019]. In particular, the variational Adam [Khan et al., 2018] is very like
62 ours, except they works for continuous values and we add noise for binary parameters. We leave
63 further theoretical investigation of BOAT for future works.

64 A.3 Improved SemHash

65 Since the feature extractor ϕ processes the raw input x into a set of intermediate representations
66 $\{c_1, c_2, \dots, c_k\}$ called concept predicates. As ϕ is parameterized by a neural network θ , ϕ 's output
67 \tilde{c}_i is real-valued. We use a binary step function to discretize \tilde{c}_i into Boolean predicates:

$$c_i = \begin{cases} 1, & \text{if } \tilde{c}_i > 0 \\ 0, & \text{otherwise.} \end{cases}$$

68 However, since gradient through this step function is zero almost anywhere and thus prevents train-
69 ing, we utilize the *Improved SemHash*[Kaiser and Bengio, 2018] as one way to make the overall
70 model differentiable.

71 During training, Improved SemHash first draw Gaussian noise ϵ with mean 0 and standard deviation
72 1. The noise ϵ is added to \tilde{c} , two vectors c and c' are then computed.

$$c = \mathbf{1}(\tilde{c} + \epsilon)$$

Algorithm 1 The BOAT algorithm for learning Neural DNF

Hyperparameters: Accepting threshold $\tau > 0$; Exponential decay rate $\gamma \in [0, 1]$; Initial noise temperature $\sigma_0 \in [0, 0.5]$; Size of rule pool N . (Default: $\tau=10^{-6}$, $\gamma=1-10^{-5}$, $\sigma_0=0.2$, $N=64$.)

Input: Dataset \mathcal{D} ; **Output:** The DNF $g(\{\mathbf{W}, \mathbf{S}\})$; The neural network $\phi(\theta)$.

```
1: Initialize  $\theta$  randomly.
2: For every  $w$  in  $\{\mathbf{W}, \mathbf{S}\}$ : initialize  $w \in \{0, 1\}$  randomly,  $m_w \leftarrow 0, \sigma_w \leftarrow \sigma_0$ .
3: while stopping criterion not met do
4:   Sample mini-batch  $\{(x_i, y_i)\}^{\text{batch size}}$  from the training set  $\mathcal{D}$ .
5:   Compute the perturbed  $\{\tilde{\mathbf{W}}, \tilde{\mathbf{S}}\}$  where each entry  $\tilde{w}$  is perturbed according to  $\sigma_w$  (??).
6:   Use  $\theta$  and perturbed  $\tilde{\mathbf{W}}, \tilde{\mathbf{S}}$  to compute the objective function  $\sum_{x_i, y_i} L(g_{\tilde{\mathbf{W}}, \tilde{\mathbf{S}}}(\phi_\theta(x_i)), y_i)$ 
7:   for every binary parameter  $w$  in  $\{\mathbf{W}, \mathbf{S}\}$  do
8:     Compute gradient  $\nabla_w$  w.r.t the objective function computed at line 4.
9:     Update exponential moving average of gradient:  $m_w \leftarrow \gamma m_w + (1 - \gamma)\nabla_w$ .
10:    if  $|m_w| > \tau$  and ( $m_w < 0$  and  $w = 0$  or  $m_w > 0$  and  $w = 1$ ) then
11:       $w \leftarrow 1 - w$  ▷ Line 7-11: update binary parameters using Modified Bop
12:    Update  $\theta$  by Adam given the objective function computed at line 6.
13:    for every  $\sigma_w$  do
14:      Update  $\sigma_w$  by Adam given the objective function computed at line 6.
15:      Clip  $\sigma_w = \min(0.5, \max(\sigma_w, 0))$ 
```

73 c is the result after applying the binary step function and

$$c' = \max(0, \min(1, 1.2 * \text{sigmoid}(\tilde{c} + \epsilon) - 0.1))$$

74 c' is computed by the above function called saturating sigmoid function [Kaiser and Sutskever, 2015,
75 Kaiser and Bengio, 2016]. During training, c is used half of the time and c' is used for the other
76 half of the time in the forward pass; for the backward pass we define the gradient of c to \tilde{c} the same
77 as c' to \tilde{c} . During testing, the noise is disabled and c is used as output.

78 It is not clear according to the description of [Kaiser and Bengio, 2018] what ‘half of time’ for c
79 and c' means, that is, whether we use c for one mini-batch and c' for the next mini-batch, or we use
80 c and c' for half of the samples for each mini-batch. In our implementation, we choose the latter
81 option: we use c for half of the samples in the mini-batch, and use c' for the other half of samples in
82 the mini-batch, determined randomly.

83 The use of saturating sigmoid function, instead of just sigmoid function, is introduced first by Kaiser
84 and Sutskever [2015] who claims to have slight improvement. But we do not observe such improve-
85 ment in Neural DNF so in our implementation we simply computed

$$c' = \text{sigmoid}(\tilde{c} + \epsilon)$$

86 We use the simple sigmoid function instead of the saturating sigmoid function.

87 We name two reason of using Improved SemHash: (1) Improved SemHash does not need any manual
88 annealing of temperature [Bulat et al., 2019, Hansen et al., 2019] or additional loss functions. There
89 are some serveral alternative that need tuning of annealing, including the annealed sigmoid/tanh
90 [Bulat et al., 2019] and gumbel-softmax trick [Maddison et al., 2016], but we believe tuning this
91 annealing schedule is difficult. (2) Improved SemHash achieves good results compared with many
92 alternative solutions. Comparisons with other discretized latent representation learning can be found
93 at [Kaiser et al., 2018], semHash performs great despite being very simple. It also has been demon-
94 strated to be a robust discretization technique in various domains [Kaiser and Bengio, 2018, Chen
95 et al., 2019b, Kaiser et al., 2019]. But of course, we note that Improved SemHash is not the only
96 option for binarizing the extracted features.

97 A.4 Initialization of $\tilde{\mathbf{W}}$

98 We will initializes \mathbf{W} and \mathbf{S} randomly, each entry is drawn from a Bernoulli disctribution where
99 $p_{\text{Bernoulli}} < 0.1$. This is because we want \mathbf{W} and \mathbf{S} to be sparse.

100 **Remark:** note that if we are really careful about the initialization, there is a small issue for initializa-
101 tion of \mathbf{W} , because some w can the value of c or the negation value. In principle, a condition and the
102 negation of it cannot be set to 1 at the same time; but we do not consider it in our implementation.
103 It seems that it does not matter in our optimization using BOAT .

104 **A.5 Regularization of DNF: R_g and possible Alternatives**

105 The simplest choice of $\mathbb{R}_g(W)$ is to use a L2 regularization. However, this is not what we really
106 want. Recall that we wish to obtain a DNF, a set of rules where the number of total rules is small and
107 the length of each rule is small. So more technically we want a small number of columns of \mathbf{W} that
108 have non-zero elements and small number of non-zero elements for each column (that is indicated
109 by the membership vector \mathbf{S}). This can be realized by penalizing the number of rules as the L1-norm
110 of \mathbf{S} and the length of selected rules also by the L1-norm (note that as only rules selected by \mathbf{S} are
111 actually used), we use a grouped L1-norm by

$$\mathbb{R}_g(\mathbf{W}, \mathbf{S}) = \lambda_g \sum_j^N |\mathbf{S}_j|_1 |\mathbf{W}_{\cdot,j}|_1$$

112 which is very like group lasso.

113 The reason we use this the grouped L1-norm is that does variable selection at the group level and
114 in our case can effectively eliminate a group of weights by columns(rules). We support the use
115 of grouped L1-norm (like group lasso) instead of L2 norm by an empirical study on the cognitive
116 preference of rules Fürnkranz et al. which shows that there is no strong preference on shorter rules
117 but instead even a slight preference on longer rules. In other words, a small number of long rules is
118 preferred over many short rules.

119 We leave more sophisticated metrics like feature overlaps [Lakkaraju et al., 2016] as future work.
120 There are, indeed, many metrics but it remains hard to apply them for Neural DNF because some of
121 them is not differentiable. If we use a discrete optimization algorithm this is not a problem, but in
122 the case of our Neural DNF, we need the objective function to be end-to-end differentiable, so we
123 need the regularization terms to be differentiable as well.

124 **A.6 Possible direction of improvement**

125 In the main paper we only discuss the regularization for the second stage DNF g , so the overall
126 objective is given as

$$\mathbf{L} = \mathbb{L}_{loss} + \lambda_g \mathbb{R}_g(\mathbf{W}, \mathbf{S})$$

127 This is because we do not focus or design anything for the first stage neural network feature extractor.
128 If we consider this (for future works), we can extend the objective to

$$\mathbf{L} = \mathbb{L}_{loss} + \lambda_g \mathbb{R}_g(\mathbf{W}, \mathbf{S}) + \lambda_\phi \mathbb{R}_\phi(\theta)$$

129 by considering a regularization term for the feature extractor that somehow defines some inter-
130 pretability constraints.

131 **B Related Works on Interpretability**

132 The interest for interpretability is not new. It appears with the development of rule-based expert
133 systems in the mid-1980's Amel [2019]. Of course the current situation is different because recently
134 we have seen an increasing trend of interpretability research in machine learning [Lipton, 2018], in
135 particular interpretable deep learning[Xie et al., 2020, Fan et al., 2020].

136 There are many definitions on interpretability, we use Lipton [2018]’s Simulatability definition of
137 interpretability: for a prediction to be fully understood, the human should be able to re-calculate and
138 reach the same prediction given reasonable time.

139 **B.1 Inherent interpretability, not Post-hoc interpretation**

140 We summarize two main tracks of approaches for interpretability research, namely *post-hoc inter-*
141 *pretation* and *building inherently interpretable models*, and we favor the latter. Post-hoc interpreta-
142 tion builds a secondary model to explain the given pre-trained deep learning model. Representative
143 works include saliency maps [Simonyan et al., 2013], LIME [Ribeiro et al., 2016], Concept Activa-
144 tion Vectors [Kim et al., 2017]; some aim at giving counterfactual explanation [Dhurandhar et al.,
145 2018, Zhang et al., 2018, Grath et al., 2018]. However, the provided explanation is in fact provided
146 by a secondary model, not the original one, so it might not correspond faithfully to how the original
147 blackbox model actually computes its prediction. Recent works further suggest that post-hoc inter-
148 pretations are not robust [Adebayo et al., 2018, Fen et al., 2019, Alvarez-Melis and Jaakkola, 2018]

149 and can even be misleading [Lakkaraju and Bastani, 2019, Rudin, 2019], and more specifically, the
150 counterfactual explanations have the ‘unjustify’ issue [Laugel et al., 2019].

151 The second track, on the contrary, is to build an inherently interpretable deep learning model, so
152 that the explanation it provided is exactly how it calculates the prediction. Inherently interpretable
153 models are more preferred to examine and use when deployed in real world applications for vari-
154 ous advantages (See [Khandani et al., 2010, Florez-Lopez and Ramon-Jeronimo, 2015]): it derives
155 explanations to justify decision (legal issue), so people are less likely to refuse to use, easier to be
156 combined with experts. An inherently interpretable classifier uses a interpretable prediction pro-
157 cess to compute the prediction, and provides such computation process itself as the explanation for
158 prediction.

159 **B.2 Related works on Interpretable Deep Learning**

160 For tabular datasets where each feature is already-meaningful, linear model or rule-based model are
161 well-established choices of interpretable models [Amel, 2019], but for other data types such as image
162 where each feature dimension itself is not meaningful, interpretable model is harder to construct.

163 A reasonable approach, which we call the two-stage paradigm, is to first construct intermediate
164 representations $\phi(x)$ that is interpretable, and on top of that a simple interpretable classifier g is
165 applied as the second stage such that the prediction is given by $\hat{y} = f(x) = g(\phi(x))$.

166 Most works on interpretable deep learning [Melis and Jaakkola, 2018, Chen et al., 2019a, Vaughan
167 et al., 2018] choose linear model as the second-stage model g . Indeed, any neural network can be
168 intuitively viewed as a two-stage model where the second-stage is a linear model, as long we treat the
169 network up to the last hidden layer as a generic feature extractor ϕ . But the notion of interpretability
170 means we wish to make certain heuristics to make $\phi(x)$ interpretable.

171 Vaughan et al. [2018] formularize g as a linear model that $g(\phi(x)) = \sum_i w_i \phi_i(x)$ where $\phi(x)$ is
172 a set of ridge functions, each of which are produced by a neural network. It claims to be more
173 interpretable than general networks, because such a function has simpler partial derivatives that can
174 simplify saliency analysis, statistical analysis and so on.

175 Chen et al. [2019a] also uses a linear model g and proposes a prototype-based design for $\phi(x)$.
176 It learns interpretable $\phi(x)$ in the sense that each dimension of $\phi(x)$ is the similarity score to an
177 image patch. It provides extra interpretability tailed for vision tasks as the similarity scores $\phi(x)$
178 can be visualized with the corresponding ‘prototype’ image patches. Melis and Jaakkola [2018]
179 takes a step further that they use neural networks to produce not only $\phi(x)$ but also the coefficient of
180 the linear model, such that prediction function is given by $g(\phi(x)) = \sum_i w_i(x) \phi_i(x)$ where $\phi(x)$
181 (called ‘concepts’) are regularized by auto-encoding reconstruction loss and $w_i(x)$ can be intuitively
182 understood as an input-dependent relevance score for a concept $\phi_i(x)$. Here we can see that the
183 implementation of $\phi(x)$ is domain-specific and customizable.

184 Theses works extend standard deep neural networks to interpretable ones by improving the design
185 of the feature extractor ϕ and use a linear model for g . This is because a linear model, or its general
186 form of general additive model, can be integrated into the automatic differentiation very easily.

187 Taking an opposite direction, we differ with previous works by considering to a rule-based classifier
188 for g and do not put our focus on ϕ . However, because of the discrete struture, integrating a rule-
189 based model g remains an hard and unexplored direction. But in terms of interpretability we favor
190 rule-based models over linear models, not only because the discrete structure is more intuitive for
191 human to follow, but also that it can provide only the satisfied rules for explanation, unlike that for
192 linear model for which we need to present the full model coefficients.

193 **B.3 why favouring the symbolic DNF for interpretability**

194 we argue that there are two major reasons for preferring rules over linear models as the second-stage
195 classifier: (1) **Rules as combinations of conditions are more interpretable than feature impor-**
196 **tance** (at least for classification). It is widely acknowledged that the rule-based models are inter-
197 pretable [Freitas, 2014, Huysmans et al., 2011, Wachter et al., 2018] because the rules give explana-
198 tions by explicitly describing the decision boundary as logical combinations of conditions. There-
199 fore, rules can more naturally provide counterfactual explanations in the form of ‘*Would changing a*
200 *certain factor have changed the decision?*’ which is often considered to be the most important prop-
201 erty of explanation [Doshi-Velez et al., 2017, Wachter et al., 2018, Grath et al., 2018, Miller, 2019].

202 On the other hand, coefficients as feature importances are arguably ‘harder to use and to understand
203 for a non-expert user’ [Wachter et al., 2018]. We conjecture that rules are probably closer to human’s
204 mental model for classification, which can be corroborated by the fact that rules have been the most
205 intuitive choice of model for human categorization learning [Bruner et al., 1956]. (2) **Using rules as
206 a second-stage classifier g requires the feature extractor ϕ to produce Boolean output** ($\{0, 1\}$),
207 which is less complex than continuous ones. Compared with continuous values, Boolean output is
208 simpler as it has only two states, making it easier for human to probe its meaning or to potentially
209 align it with human knowledge. as future works

210 **Rules can derive counterfactual explanations while linear models cannot** (at least, not easy for
211 linear models). The explicit decision boundary of rules not only can give factual explanations,
212 namely giving the conjunctive conditions of the satisfied rule, but also give counterfactual expla-
213 nations, by simply presenting that could have changed the prediction. Unlike weighted feature im-
214 portance, rules explicitly presents all the sufficient conditions for prediction and thus can naturally
215 handles counterfactual explanations in the form of ‘*Would changing a certain factor have changed
216 the decision?*’. Indeed, the counterfactual explanation is often considered to be the most impor-
217 tant property of explanation, confirmed not only from a more practical perspective [Doshi-Velez
218 et al., 2017, Wachter et al., 2018, Grath et al., 2018] and also from cognitive/psychological research
219 [Miller, 2019].

220 C DNF

221 Here we first explain the reason of choosing DNF as the rule module.

222 We choose DNF, one of the most powerful and historically significant symbolic methods, for its
223 interpretability and its generality. It has a simple and transparent ‘*OR-of-ANDs*’ prediction process:
224 if at least one *AND* clause (a conjunction of Boolean predicates) is satisfied, it predicts positive;
225 otherwise it predicts negative. DNF is interpretable not only that the discrete structure is intuitive
226 to follow, each conjunctive clauses (AND) of DNF can be viewed as subtype for explanation, i.e. the
227 DNF can be decomposed into individual local patterns. We also appreciate the generality of DNF,
228 as any propositional logic formula has a equivalent DNF and thus any rule-based binary classifier
229 including decision set/list/tree can be expressed as a DNF.

230 DNF have a long history but learning DNF is still a very hard problem, not to mention that we
231 need to add interpretability constraints. Theoretical results on learning DNF in the PAC (probably
232 approximately correct) setting are often unrealistic in practice and are hard to incorporate extra in-
233 terpretability constraints. On the other hand, the practical use of rule learning of a DNF form attracts
234 more attention from data mining community, namely, descriptive pattern discovery [Novak et al.,
235 2009]. Seminal algorithms include CN2 and RIPPER (constrained for binary classification). More
236 recent state of the art machine learning algorithms for learning DNF can work quite well on small
237 tabular datasets, but not very scalable on high dimensional datasets, we name few representative
238 recent works using greedy heuristics [Obermann and Waack, 2015, 2016]. Bayesian approxima-
239 tion [Wang et al., 2015], linear programming [Su et al., 2015]. We name the interpretable decision
240 set by Lakkaraju et al. [2016] and Bayesian rule set by Wang et al. [2017] as two recent work as
241 representative.

242 However, we note that these approach is not compatible here if we wish to jointly optimize g with
243 the neural network feature extractor ϕ in an end-to-end way. It is because that first, constructing W
244 by rule mining becomes non-sense if the neural module is currently being trained; second, discrete
245 optimization methods for learning S is not compatible with gradient-based optimization.

246 C.1 Differentiable replacement of the DNF function g

247 The differentiable replacement of the DNF function we use in the paper is adopted from Payani and
248 Fekri [2019], Wang et al. [2019b]. However, as we mentioned in the footnote in the main paper, we
249 believe that this relaxation is not new but rather re-invented. Simialr formulations on differentiable
250 operations see [Sajjadi et al., 2016, Arabshahi et al., Nomura et al., 1992]. It is also likely that we
251 miss critical references on neuro-fuzzy system research in the 90s.

252 Basically, we construct differentiable g following the approaches of [Payani and Fekri, 2019, Wang
253 et al., 2019b], where neural networks that execute differentiable logical operations are learned for in-
254 ductive logic programming [Payani and Fekri, 2019] and multiple DNF layers are stacked to classify
255 on tabular datasets [Wang et al., 2019b].

256 We here introduce some works that also uses a similar differentiable replacement of the DNF func-
257 tion below. Sajjadi et al. [2016] learns a DNF in the context of tackling the *moving target* problem
258 (herd-effect) [Fahlman and Lebiere, 1990]. For each rule there is no selection of conditions, but
259 rather, all conditions are taking as conjunction. One thing worth mentioning is that Sajjadi et al.
260 [2016] do not actually train the parameters of the disjunctive function but only gives a good initial-
261 ization.

262 Payani and Fekri [2019] propose the *neural logical networks* to solve inductive logic programming,
263 optimized in a gradient-based way. It does not add any extra regularizations to push the parameters
264 to binary values, in the other hand it carefully discusses the problem of initialization to solve the
265 problem.

266 Wang et al. [2019b] proposes to stack many DNF layers to make accurate and interpretable pre-
267 dictions (at least the authors claim that the resulted networks is interpretable), although we doubt
268 it as the stacked DNF may in fact corresponding to a rule set that is so complicated for human too
269 understand. Wang et al. [2019b] also propose the randomized binarization to make the parameters
270 of the matrix \hat{W} to be close to binary values. We view this as a binarization version of dropout, at
271 each round, some parameters in \hat{W} are thresholded to 0 or 1 and get fixed, and other parameters
272 are optimized as usual by backprop. However, it seems in the experiment of [Wang et al., 2019b]
273 that the rate of binarization plays an important role, often the rate need to set in a very high value
274 (for example, 80 90%) We also think that, it is not very efficient in terms of learning, by fixing such
275 a high percentage of parameters to thresholded values. Because there binarized parameters are not
276 getting updated at all (no gradient).

277 **Optimization issues.** The presense of both discrete and continuous parameters in Neural DNF
278 introduces a hard optimization problem. A straightforward solution is to use a EM-like approach, by
279 applying well-studied optimization techniques for neural and rule-based module separately. But we
280 argue such separate training scheme might not be as efficient as joint end-to-end training. Another
281 solution is to end-to-endly optimizing the overall model, which is challenging. We discuss in the
282 main paper of two straightforward alternatives for joint learning: DNF-Real and DNF-STE and we
283 also propose BOAT , the key algorithm in our paper.

284 **Possible future directions.** We consider possible directions about moving from DNF (proposi-
285 tional logic) to more flexible and powerful rule languages, And we are hoping that the proposed
286 BOAT can possibly help.

287 For example, there is an inductive program synthesis project called ‘TerpreT’ [Gaunt et al., 2016],
288 where the authors design a special language TerpreT for inductive program synthesis. It is designed
289 to separate the program specifaion and the inference algorithm so that it can be optimized by
290 different backends such as SGD, relaxed linear programming and non-machine-learning solvers so
291 that compare all these optimization backends. Its key surprising finding is that in terms of empirical
292 performance, SGD is dominated by more traditional constraint solvers. It also gives a very simple
293 case where SGD can fail consistently with a exponential number of local minima, yet constraint
294 solvers can work it out very easily. Note that in [Gaunt et al., 2016], the binary parameters is not in
295 fact binary but only transformed by sigmoid/tanh/etc, so it is like the DNF-Real [Payani and Fekri,
296 2019] we mentioned in the main paper.

297 **D Connections to Neural-symbolic Integration**

298 It turns out that interpretable deep learning is very close to neuro-symbolic integration, just slightly
299 different focus. Neural DNF can also be seen as achieving an interpretable model by neuro-symbolic
300 integration. In fact, this ‘explainability or interpretability through neuro-symbolic integration’ ap-
301 proach is recently gaining attention (for example, see David Cox’s AAAI/IAAI 2020 invited talk).

302 Neuro-symbolic integration [Besold et al., 2017], which aims at combining neural methods with
303 symbolic-logic methods, learning and reasoning, seems to be very related to our Neural DNF .
304 However, the literature on this topic is vast and offers a multitude of approaches covering different
305 settings, making it difficult to discuss related works completely. We follow a most recent survey
306 [Garcez et al., 2019], which divides neuro-symbolic integarion into the two categories: Horizontal
307 and Vertical.

308 Horizontal integration contains most of the research, aims at integrating neural and symbolic tech-
309 niques into one inseparable model: either by making neural model behaves more like a symbolic

310 model (also known as deep learning with logical constraints), which using logical knowledge to im-
311 prove neural network learning, or the making symbolic method neural, using a neural network as
312 interpreter to execute symbolic programs. Some approaches for Horizontal integration are

- 313 • Put logic as constraints on a DNN: a deep net is extended with an extra regularization term
314 derived from some logical property or from extra heavy annotations. But it is not guaran-
315 teed that such a DNN can make consistent symbolic prediction, as suggested by Xu et al.
316 [2017] that deep nets trained with additional logical regularizations cannot consistently
317 make predictions that is from the logic they were trained on.
- 318 • Enable a DNN to execute logical programs Cohen et al. [2017]. We can make a logic
319 program differentiable by using differentiable functions and let a neural network to execute
320 the logical program, such as querying a database. These approaches predicts a transparent
321 structure as output, but its prediction process is not transparent or following logical property
322 at all.

323 Vertical integration, the category to which our Neural DNF belong, assembles a deep learning model
324 and a symbolic model in a sequential manner. It intends to use deep learning for pattern recognition
325 (perception), and the symbolic model for high-level reasoning. However vertical integration but
326 does not receive much attention even it has a strong neuroscience inspiration.

- 327 • SATNet by Wang et al. [2019a] opens a framework to design flexible combinatorial func-
328 tion in a smoothed differentiable way and the combinatorial relationship is learned instead
329 of hard-coded grammar. But still, we do not have a control or a sense of understanding to
330 interpret the prediction.
- 331 • Manhaeve et al. [2018] use DNN for perception and symbolic reasoning given a defined
332 grammar . They use a DNN to handle perception and outputs a classification as a neural
333 predicate, and then use this predicates to do reasoning. It learn a Probabilistic ProLog
334 program by gradient descent enabled by compiling a Sentential Decision Diagram. But the
335 relationship of these predicates are hard-coded by the used grammar, not learned; and the
336 symbols (that is, the extracted feature $\phi(x)$) is not learned from scratch.

337 Learning neural predicates that represent the abstract features/symbols ('cat's mouth', etc) and learn-
338 ing the relationship among these variables are a novel thing! And this is also very important if we
339 wish to do symbolic prediction on raw unstructured data. It is generally true that most of the methods
340 will need pre-training or heavy annotations of the symbols extracted by a DNN.

341 But note that for our Neural DNF both the features $\phi(x)$ and the rules g are learned from scratch.
342 This is in fact quite non-trivial if we consider Neural DNF in the context of neuro-symbolic integra-
343 tion.

344 **E Experiment**

345 **E.1 Evaluation of the BOAT Optimizer**

346 Here we evaluate BOAT on datasets with Boolean features so that we learn only the DNF g . We use
347 a synthetic dataset adopted from Payani and Fekri [2019] which consists of 10-bit Boolean strings
348 and a randomly-generated DNF as ground-truth.

349 First, to show the proposed noise is indeed necessary and helps the optimization, we learn the DNF
350 with/without the noise on multiple datasets generated with different random seeds (so the ground-
351 truth DNF is different) and plot the loss curves. From fig. 1a we observe that multiple runs of BOAT
352 give very similar and stable convergence, while in fig. 1b, surprisingly, runs consistently fail to
353 converge without the noise. We view this as strong evidence for the necessity of the proposed noise.
354 On the other hand, if noise temperature is initialized larger, the convergence is slower (fig. 1b).

355 There are some other issues with a large noise rate: (1) The slower learning of a larger noise rate
356 can also be explained from the multiplicity nature of the DNF function we use (eq 2 and 3 in the
357 main paper). This becomes more severe in high dimensional inputs. And finding a small noise rate
358 initialization value will help. (2) when the noise rate initialized too high (and the hyperparameter of
359 adaptivity rate not set right), it is also possible that the binary parameters cannot get huge gradient
360 enough to flip, while it is only the continuous parameters including the noise rate are getting updated.
361 This is an undesired result. In this case, find a smaller initialization rate.

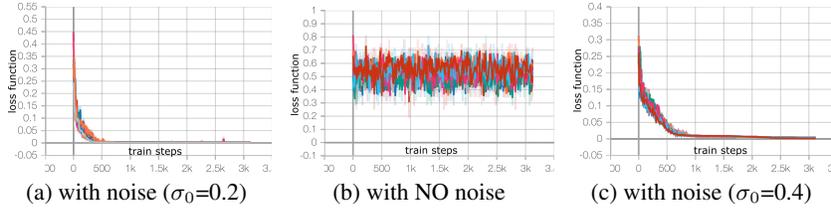


Figure 1: Loss curve with 10 differently-generated synthetic dataset

362 Next, we apply BOAT and compared the convergence speed with the following baselines: (1)
 363 DNF-Real [Payani and Fekri, 2019, Wang et al., 2019b] and (2) DNF-STE which are discussed
 364 in ???. (3) A Linear Model. (4) A multi-layer perceptron. Note that we also use the adap-
 365 tive noise for DNF-STE since otherwise optimization will fail. As shown in fig. 2, BOAT
 366 gives the fastest convergence, while MLP, DNF-STE and DNF-Real converge much more slowly.
 367 The linear model does not converge. As for the reason of DNF-STE’s slower and less stable
 368 convergence, we believe it is because that unlike the modified *Bop*, DNF-STE does not
 369 have the mechanism to prevent rapid flipping and thus optimization becomes more unstable.
 370 As for performance, all above methods except linear
 371 model achieve 100% F1 score on test set and BOAT
 372 can always find the ground truth DNF on different
 373 synthetic datasets generated with different random
 374 seeds. **Remark:** Note that since we use random initial-
 375 ization for W, S , a natural suspicion is that the
 376 ground truth DNF happens to be discovered by the
 377 random initialization, not learned. To refute it,
 378 we can use zero initialization and find that BOAT
 379 converges similarly to random initialization. We report
 380 it in the appendix together with the performance re-
 381 sults on some UCI datasets.

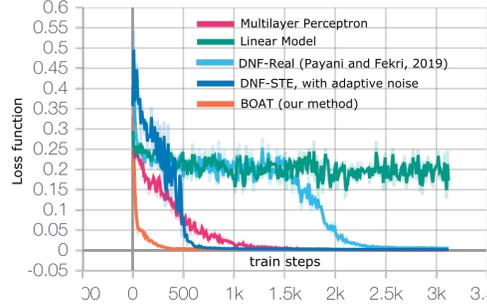


Figure 2: Loss Curve on the synthetic dataset

382 Here we evaluate only the second stage of our model,
 383 i.e., we use some datasets that the input features are
 384 already Boolean attributes, so we can use an identity
 385 function for the first stage neural module ϕ .

386 By removing ϕ , we focus on the evaluation of the learning algorithm, i.e., the proposed BOAT . We
 387 decide to use the synthetic Boolean bit-string dataset introduced by [Payani and Fekri, 2019] where
 388 this dataset is used for comparing the convergence of learning of the introduced DNF layer.

389 We randomly draw 5000 bit-strings where each bit is 0 or 1 (uniformly drawn). We then also ran-
 390 domly generate a ground-truth DNF and use this ground-truth DNF to assign labels to the bit-strings.
 391 Hopefully this ground-truth DNF should be recovered and indeed it is recovered. The ground-truth
 392 DNF consists of 5 clauses (rules) and each clause (rule) has 3 conditions. The conditions include
 393 negations. We random choose 4000 strings as the training set.

394 The generation of this synthetic Boolean bit-string dataset can be found at `dataloaders.py`. Note
 395 that in [Payani and Fekri, 2019], the bit is drawn with prob 0.75 to be zero and 0.25 to be one, and
 396 in our case we just use prob 0.5 to be zero and 0.5 to be one. This is because we also use negation
 397 in the generation of ground-truth DNF.

398 The DNF structure is $2K \rightarrow N \rightarrow 1$. We set $N = 64$ for the DNF as the default value for our method
 399 as well as DNF-Real and DNF-STE. For MLP and Linear model, we concatenate the input with
 400 its negation, and use a three-layer architecture ($2K \rightarrow N \rightarrow 1$) for MLP, and the linear model is in
 401 essential a two-layer perceptron ($2K \rightarrow 1$). As the DNF layer uses negations, we will also do the
 402 same for the linear model and multi-layer perceptron(MLP). We simple concatenate the input feature
 403 with its negations so that the linear model is of structure ($2K \rightarrow 1$) and the MLP ($2K \rightarrow N \rightarrow 1$) For
 404 our method and all baselines we compare, we use Adam with initial learning rate 0.001. We use
 405 $\lambda_g = 0.0001$ as the default value.

406 Note that for compared baselines, if we use a larger learning rate such as 0.01 and since the synthetic
 407 Boolean bit-string dataset is not very difficult, the learning of DNF will be quicker (thus the cover-
 408 gence figure will look different). However for our method, if we also set the initial learning rate for
 409 the noise from 0.001 to 0.01, the learning of BOAT will also be quicker (and a similar convergence
 410 figure should be able to be reproduced).

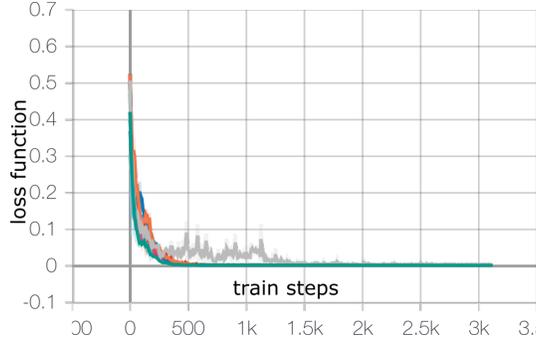


Figure 3: Loss curve with 10 differently-generated synthetic dataset using zero-initialization

411 **random initialization and zero initialization.** As we’ve mentioned in the main paper, one might
 412 suspect that the ground-truth DNF is not learned by our algorithm BOAT, but is only happened to be
 413 discovered by the random initialization. So here we also run our method with 10 different random
 414 seeds (differently-generated dataset) but with all-zero-initialization for \mathbf{W} and \mathbf{S} . We can observe
 415 that the learning is still successful.

416 **On real-world tabular datasets.** In order to further investigate the performance, we also applied
 417 our method (learning the DNF g only) on several real-world datasets (table 1). We discretize the
 418 input features as preprocessing. Since these datasets are more complicated than the above synthetic
 419 dataset, we set N to 128; we set λ_g to $1e-4$ as the default value. We find that our method perform
 420 competitively well and does not have performance decrease when the noise is removed. DNF-
 421 Real’s performance slightly decreases after parameters are thresholded for Banknote dataset. One
 422 point worth noting is that learned DNF can achieve 100% F1-score on the tic-tac-toe dataset while
 423 weighted-sum-style models (linear Model/MLP/SENN¹) can only approach to 100%. We think it
 424 is because the tic-tac-toe data is gathered from the status of a combinatorial game and thus may be
 425 more easily learned by rule-based models.

	DNF-BOAT	DNF-REAL	DNF-STE	Linear	MLP	SENN
	With / Without noise	Before / After thresholding	With / Without noise	Model		
Banknote	92.11%/92.11%	92.56%/91.05%	91.49%/91.49%	92.56%	92.11%	92.11%
tic-tac-toe	100%/100%	100%/100%	93.59%/93.59%	98.79%	99.59%	99.59%
Blogger	81.81%/81.81%	81.81%/81.81%	78.78%/78.78%	71.42%	86.95%	86.95%

Table 1: Test F1 score of learned DNF

426 Note that learning rule-based models using neural networks is already a challenging task, which we
 427 will not further investigate here.

428 E.2 The 2d-XOR

429 Four gaussians are drawn using scikit’s ‘make_blob’ function with cluster mean (0,5), (5,0), (10,5)
 430 and (5,10).

431 E.3 MNIST

432 We use the standard train-test split for MNIST (in total 10000 test samples.)

433 We use a randomly-initialized LeNet-like convolutional network as the feature extractor ϕ to produce
 434 5 concept predicates. Since MNIST has 10 class, we use a separate \mathbf{W} , \mathbf{S} for each class. After

¹Note that SENN here is the second stage module of [Melis and Jaakkola, 2018] which is a MLP that takes input and generate the coefficients of a linear model.

435 training Neural DNF finds one or two rules for each class and can achieve over 99% test accuracy.
 436 In order to give an example of Neural DNF ’s explanation, we first apply Neural DNF on MNIST as
 437 a direct comparison to the explanations provided by linear models (e.g., the MNIST example from
 438 [Melis and Jaakkola, 2018]).

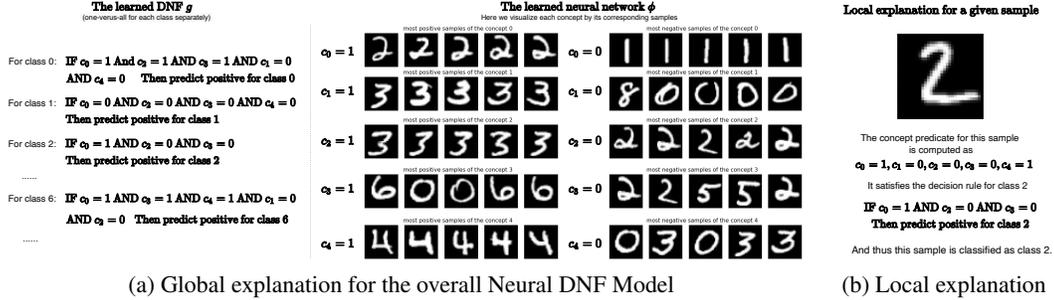


Figure 4: Explanations provided by Neural DNF on MNIST

439 Here we provide the explanations that Neural DNF derives: (1) In fig. 4a, we provide a *global*
 440 explanation for Neural DNF that explains the working of the overall model, by presenting the rules
 441 for each class as well as the image samples that triggers each concept predicate ($c = 1$) or not
 442 ($c = 0$). As it is necessary to understand what each concept c means, we provide representative
 443 samples for $c = 1$ and $c = 0$ by retrieving samples that maximize or minimize the pre-binarization
 444 value \tilde{c} . (2) In fig. 4b, Neural DNF is also able to derive *local* explanations which explain the
 445 classification for a particular sample. Given a test sample, we present the value of its concept
 446 predicate (computed by ϕ) as well as the satisfied rule by which prediction is computed. Note in
 447 particular that the satisfied rules are not only the explanation but also how the predication is exactly
 448 computed, thus *inherently* interpretable.

449 Unlike explaining by feature importance that can only indicate some of $\phi(x)$ contributes more signif-
 450 icantly to the final prediction, Neural DNF’s explanation describes the decision boundary explicitly
 451 as combinations of conditions. For example, in fig. 4b as the decision rule for class 2 is ‘If $c_0=1$ and
 452 $c_2=0$ and $c_3=0$ Then class 2’, it becomes clear that only c_0, c_2, c_3 are essential for predicting class 2
 453 while c_1, c_4 are not. In other words, for this test sample, we know precisely that if we change any of
 454 c_1 or c_4 , the prediction will remain but changing of any of c_0, c_2, c_3 will give a different prediction.

455 The detailed decision rule for each class is:

456 Decision rule for class 0 is IF $c_0, c_2, c_3 = 1$ AND $c_1, c_4 = 0$ THEN predict class 0

457 Decision rule for class 1 is IF $c_1 = 1$ AND $c_0, c_2, c_3, c_4 = 0$ THEN predict class 1

458 Decision rule for class 2 is IF $c_0 = 1$ AND $c_2, c_3 = 0$ THEN predict class 2

459 Decision rule for class 3 is IF $c_0, c_1, c_2 = 1$ AND $c_3, c_4 = 0$ THEN predict class 3

460 Decision rule for class 4 is IF $c_3, c_4 = 1$ AND $c_0, c_2 = 0$ THEN predict class 4

461 Decision rule for class 5 is IF $c_1, c_2 = 1$ AND $c_0, c_3 = 0$ OR $c_0, c_1, c_2 = 1$ AND $c_0, c_1, c_2,$
 462 $c_4 = 0$ THEN predict class 5

463 Decision rule for class 6 is IF $c_0, c_3, c_4 = 1$ AND $c_1, c_2 = 0$ THEN predict class 6

464 Decision rule for class 7 is IF $c_0, c_1, c_3 = 1$ AND $c_2 = 0$ THEN predict class 7

465 Decision rule for class 8 is IF $c_0, c_2, c_4 = 1$ AND $c_3 = 0$ OR $c_2, c_4 = 1$ AND $c_1, c_3 = 0$ THEN
 466 predict class 8

467 Decision rule for class 9 is IF $c_1, c_2, c_3 = 1$ AND $c_0 = 0$ THEN predict class 9

468 E.4 Other datasets in Section 4.2

469 We evaluate our method on datasets as follows: MNIST, KMNIST, SVHN, CIFAR10. These are not
 470 new but very standard image datasets and we just use standard train-test split comes with the pytorch
 471 backend. Note that the four datasets are image datasets so use an convolutional neural network.

472 We use a fixed adam learning rate and fixed γ for the modified *Bop* optimizer. So there is no learning
 473 rate decay or other scheduling. For MNIST and KMNIST we use 5 concept predicate and run for
 474 100 epoch. For SVHN and CIFAR10 we use 32 concept predicate and run for 200 epoch as these
 475 two datasets are more challenging.

476 Note that we present results of Neural DNF using lazy tie-breaking: by selecting the first encoun-
 477 tered positive class in ascending order (e.g., when class 1, 4, 7 are all predicted as positive, we select
 478 class 1.). We can also use a random tie-breaking: when multiple classes are predicted as positive,
 479 we randomly pick one. We run 10 times on test set and report mean and standard deviation

Table 2: Test Accuracy of Neural DNF on some image datasets with two different tie-breaking

	MNIST	KMNIST	SVHN	CIFAR10
lazy tie-breaking	99.08%	95.43%	90.13%	67.91%
random tie-breaking	99.08% \pm 0.02%	95.43% \pm 0.03%	90.29% \pm 0.04%	68.26% \pm 0.12%

480 References

- 481 Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim.
 482 Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, pages
 483 9505–9515, 2018.
- 484 David Alvarez-Melis and Tommi S Jaakkola. On the robustness of interpretability methods. *arXiv*
 485 *preprint arXiv:1806.08049*, 2018.
- 486 Kay R Amel. From shallow to deep interactions between knowledge representation, reasoning and
 487 machine learning. In *Proceedings 13th International Conference Scala Uncertainty Mgmt (SUM*
 488 *2019)*, Compiègne, LNCS, pages 16–18, 2019.
- 489 P Arabshahi, TP Caudell, JJ Choi, and BJ Song. Steepest descent adaptation of min-max fuzzy
 490 if-then rules.
- 491 Tarek R Besold, Artur d’Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pas-
 492 cal Hitzler, Kai-Uwe Kühnberger, Luis C Lamb, Daniel Lowd, Priscila Machado Vieira Lima,
 493 et al. Neural-symbolic learning and reasoning: A survey and interpretation. *arXiv preprint*
 494 *arXiv:1711.03902*, 2017.
- 495 Jerome S Bruner, Jacqueline J Goodnow, and George A Austin. A study of thinking. 1956.
- 496 Adrian Bulat, Georgios Tzimiropoulos, Jean Kossaifi, and Maja Pantic. Improved training of binary
 497 networks for human pose estimation and image recognition. *arXiv preprint arXiv:1904.05868*,
 498 2019.
- 499 Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks
 500 like that: deep learning for interpretable image recognition. In *Advances in Neural Information*
 501 *Processing Systems*, pages 8928–8939, 2019a.
- 502 Zhoung Chen, Yang Li, Samy Bengio, and Si Si. You look twice: Gaternet for dynamic fil-
 503 ter selection in cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern*
 504 *Recognition*, pages 9172–9180, 2019b.
- 505 William W Cohen, Fan Yang, and Kathryn Rivard Mazaitis. Tensorlog: Deep learning meets prob-
 506 abilistic dbs. *arXiv preprint arXiv:1707.05390*, 2017.
- 507 Amit Dhurandhar, Pin-Yu Chen, Ronny Luss, Chun-Chen Tu, Paishun Ting, Karthikeyan Shan-
 508 mugam, and Payel Das. Explanations based on the missing: Towards contrastive explanations
 509 with pertinent negatives. In *Advances in neural information processing systems*, pages 592–603,
 510 2018.
- 511 Finale Doshi-Velez, Mason Kortz, Ryan Budish, Chris Bavitz, Sam Gershman, David O’Brien,
 512 Stuart Schieber, James Waldo, David Weinberger, and Alexandra Wood. Accountability of ai
 513 under the law: The role of explanation. *arXiv preprint arXiv:1711.01134*, 2017.

- 514 Scott E Fahlman and Christian Lebiere. The cascade-correlation learning architecture. In *Advances*
515 *in neural information processing systems*, pages 524–532, 1990.
- 516 Fenglei Fan, Jinjun Xiong, and Ge Wang. On interpretability of artificial neural networks. *arXiv*
517 *preprint arXiv:2001.02522*, 2020.
- 518 Hui Fen, Kuangyan Song, Madeilene Udell, Yiming Sun, Yujia Zhang, et al. Why should
519 you trust my interpretation? understanding uncertainty in lime predictions. *arXiv preprint*
520 *arXiv:1904.12991*, 2019.
- 521 Raquel Florez-Lopez and Juan Manuel Ramon-Jeronimo. Enhancing accuracy and interpretability
522 of ensemble strategies in credit risk assessment. a correlated-adjusted decision forest proposal.
523 *Expert Systems with Applications*, 42(13):5737–5753, 2015.
- 524 Alex A Freitas. Comprehensible classification models: a position paper. *ACM SIGKDD explorations*
525 *newsletter*, 15(1):1–10, 2014.
- 526 Johannes Fürnkranz, Tomáš Kliegr, and Heiko Paulheim. On cognitive preferences and the plausi-
527 bility of rule-based models. *Machine Learning*, pages 1–46.
- 528 A Garcez, M Gori, LC Lamb, L Serafini, M Spranger, and SN Tran. Neural-symbolic computing:
529 An effective methodology for principled integration of machine learning and reasoning. *Journal*
530 *of Applied Logics*, 6(4), 2019.
- 531 Alexander L. Gaunt, Marc Brockschmidt, Rishabh Singh, Nate Kushman, Pushmeet Kohli, Jonathan
532 Taylor, and Daniel Tarlow. Terpret: A probabilistic programming language for program induction.
533 *CoRR*, abs/1608.04428, 2016. URL <http://arxiv.org/abs/1608.04428>.
- 534 Rory Mc Grath, Luca Costabello, Chan Le Van, Paul Sweeney, Farbod Kamiab, Zhao Shen, and
535 Freddy Lécué. Interpretable credit application predictions with counterfactual explanations.
536 *CoRR*, abs/1811.05245, 2018. URL <http://arxiv.org/abs/1811.05245>.
- 537 Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma.
538 Unsupervised neural generative semantic hashing. In *Proceedings of the 42nd International ACM*
539 *SIGIR Conference on Research and Development in Information Retrieval*, pages 735–744, 2019.
- 540 Koen Helwegen, James Widdicombe, Lukas Geiger, Zechun Liu, Kwang-Ting Cheng, and Roeland
541 Nusselder. Latent weights do not exist: Rethinking binarized neural network optimization. In
542 *Advances in neural information processing systems*, pages 7531–7542, 2019.
- 543 Johan Huysmans, Karel Dejaeger, Christophe Mues, Jan Vanthienen, and Bart Baesens. An empir-
544 ical evaluation of the comprehensibility of decision table, tree and rule based predictive models.
545 *Decision Support Systems*, 51(1):141–154, 2011.
- 546 Łukasz Kaiser and Samy Bengio. Can active memory replace attention? In *Advances in Neural*
547 *Information Processing Systems*, pages 3781–3789, 2016.
- 548 Łukasz Kaiser and Samy Bengio. Discrete autoencoders for sequence models. *arXiv preprint*
549 *arXiv:1801.09797*, 2018.
- 550 Łukasz Kaiser and Ilya Sutskever. Neural gpus learn algorithms. *arXiv preprint arXiv:1511.08228*,
551 2015.
- 552 Łukasz Kaiser, Samy Bengio, Aurko Roy, Ashish Vaswani, Niki Parmar, Jakob Uszkoreit, and Noam
553 Shazeer. Fast decoding in sequence models using discrete latent variables. In *International Con-*
554 *ference on Machine Learning*, pages 2390–2399, 2018.
- 555 Łukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad
556 Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based
557 reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- 558 Mohammad Emtiyaz Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivas-
559 tava. Fast and scalable bayesian deep learning by weight-perturbation in adam. *arXiv preprint*
560 *arXiv:1806.04854*, 2018.

- 561 Amir E Khandani, Adlar J Kim, and Andrew W Lo. Consumer credit-risk models via machine-
562 learning algorithms. *Journal of Banking & Finance*, 34(11):2767–2787, 2010.
- 563 Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory
564 Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation
565 vectors (tcav). *arXiv preprint arXiv:1711.11279*, 2017.
- 566 Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameteri-
567 zation trick. In *Advances in neural information processing systems*, pages 2575–2583, 2015.
- 568 Himabindu Lakkaraju and Osbert Bastani. ” how do i fool you? ”: Manipulating user trust via
569 misleading black box explanations. *arXiv preprint arXiv:1911.06473*, 2019.
- 570 Himabindu Lakkaraju, Stephen H Bach, and Jure Leskovec. Interpretable decision sets: A joint
571 framework for description and prediction. In *Proceedings of the 22nd ACM SIGKDD interna-
572 tional conference on knowledge discovery and data mining*, pages 1675–1684, 2016.
- 573 Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, and Marcin Detyniecki.
574 The dangers of post-hoc interpretability: Unjustified counterfactual explanations. *arXiv preprint
575 arXiv:1907.09294*, 2019.
- 576 Zachary C Lipton. The mythos of model interpretability. *Queue*, 16(3):31–57, 2018.
- 577 Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous
578 relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- 579 Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt.
580 Deepproblog: Neural probabilistic logic programming. *CoRR*, abs/1805.10872, 2018. URL
581 <http://arxiv.org/abs/1805.10872>.
- 582 David Alvarez Melis and Tommi Jaakkola. Towards robust interpretability with self-explaining
583 neural networks. In *Advances in Neural Information Processing Systems*, pages 7775–7784, 2018.
- 584 Xiangming Meng, Roman Bachmann, and Mohammad Emtiyaz Khan. Training binary neural net-
585 works using the bayesian learning rule. *arXiv preprint arXiv:2002.10778*, 2020.
- 586 Tim Miller. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intel-
587 ligence*, 267:1–38, 2019.
- 588 Hiroyoshi Nomura, Isao Hayashi, and Noboru Wakami. A learning method of fuzzy inference rules
589 by descent method. In *[1992 Proceedings] IEEE International Conference on Fuzzy Systems*,
590 pages 203–210. IEEE, 1992.
- 591 Petra Kralj Novak, Nada Lavrač, and Geoffrey I Webb. Supervised descriptive rule discovery: A
592 unifying survey of contrast set, emerging pattern and subgroup mining. *Journal of Machine
593 Learning Research*, 10(Feb):377–403, 2009.
- 594 Lennart Obermann and Stephan Waack. Demonstrating non-inferiority of easy interpretable meth-
595 ods for insolvency prediction. *Expert Systems with Applications*, 42(23):9117–9128, 2015.
- 596 Lennart Obermann and Stephan Waack. Interpretable multiclass models for corporate credit rating
597 capable of expressing doubt. *Frontiers in Applied Mathematics and Statistics*, 2:16, 2016.
- 598 Ali Payani and Faramarz Fekri. Learning algorithms via neural logic networks. *arXiv preprint
599 arXiv:1904.01554*, 2019.
- 600 Andres Potapczynski, Gabriel Loaiza-Ganem, and John P Cunningham. Invertible gaussian repara-
601 rameterization: Revisiting the gumbel-softmax. *arXiv preprint arXiv:1912.09588*, 2019.
- 602 Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ” why should i trust you? ” explaining the
603 predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference
604 on knowledge discovery and data mining*, pages 1135–1144, 2016.
- 605 Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and
606 use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.

- 607 Mehdi Sajjadi, Mojtaba Seyedhosseini, and Tolga Tasdizen. Disjunctive normal networks. *Neuro-*
608 *computing*, 218:276–285, 2016.
- 609 Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Vi-
610 sualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- 611 Guolong Su, Dennis Wei, Kush R Varshney, and Dmitry M Malioutov. Interpretable two-level
612 boolean rule learning for classification. *arXiv preprint arXiv:1511.07361*, 2015.
- 613 Joel Vaughan, Agus Sudjianto, Erind Brahimi, Jie Chen, and Vijayan N Nair. Explainable neural
614 networks based on additive index models. *arXiv preprint arXiv:1806.01933*, 2018.
- 615 Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening
616 the black box: Automated decisions and the gdpr. *Harvard Journal of Law & Technology*, 31(2),
617 2018.
- 618 Po-Wei Wang, Priya L Donti, Bryan Wilder, and Zico Kolter. Satnet: Bridging deep learning and
619 logical reasoning using a differentiable satisfiability solver. *arXiv preprint arXiv:1905.12149*,
620 2019a.
- 621 Tong Wang, Cynthia Rudin, Finale Doshi-Velez, Yimin Liu, Erica Klampfl, and Perry MacNeille.
622 Or’s of and’s for interpretable classification, with application to context-aware recommender sys-
623 tems. *CoRR*, abs/1504.07614, 2015. URL <http://arxiv.org/abs/1504.07614>.
- 624 Tong Wang, Cynthia Rudin, Finale Doshi-Velez, Yimin Liu, Erica Klampfl, and Perry MacNeille. A
625 bayesian framework for learning rule sets for interpretable classification. *The Journal of Machine*
626 *Learning Research*, 18(1):2357–2393, 2017.
- 627 Zhuo Wang, Wei Zhang, Nannan Liu, and Jiayong Wang. Transparent classification with multilayer
628 logical perceptrons and random binarization. *ArXiv*, abs/1912.04695, 2019b.
- 629 Ning Xie, Gabrielle Ras, Marcel van Gerven, and Derek Doran. Explainable deep learning: A field
630 guide for the uninitiated. *arXiv preprint arXiv:2004.14545*, 2020.
- 631 Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic loss
632 function for deep learning with symbolic knowledge. *arXiv preprint arXiv:1711.11157*, 2017.
- 633 Xin Zhang, Armando Solar-Lezama, and Rishabh Singh. Interpreting neural network judgments
634 via minimal, stable, and symbolic corrections. In *Advances in Neural Information Processing*
635 *Systems*, pages 4874–4885, 2018.