

A. Axiomatization of Semantic Loss: Details

This appendix provides further details on our axiomatization of semantic loss. We detail here a complete axiomatization of semantic loss, which will involve restating some axioms and propositions from the main paper.

The first axiom says that there is no loss when the logical constraint α is always true (it is a logical tautology), independent of the predicted probabilities p .

Axiom 5 (Truth). The semantic loss of a true sentence is zero: $\forall p, L^s(\text{true}, p) = 0$.

Next, when enforcing two constraints on disjoint sets of variables, we want the ability to compute semantic loss for the two constraints separately, and sum the results for their joint semantic loss.

Axiom 6 (Additive Independence). Let α be a sentence over \mathbf{X} with probabilities p . Let β be a sentence over \mathbf{Y} disjoint from \mathbf{X} with probabilities q . The semantic loss between sentence $\alpha \wedge \beta$ and the joint probability vector $[p \ q]$ decomposes additively: $L^s(\alpha \wedge \beta, [p \ q]) = L^s(\alpha, p) + L^s(\beta, q)$.

It directly follows from Axioms 5 and 6 that the probabilities of variables that are not used on the constraint do not affect the semantic loss.

Proposition 5 formalizes this intuition.

Proposition 5 (Locality). Let α be a sentence over \mathbf{X} with probabilities p . For any \mathbf{Y} disjoint from \mathbf{X} with probabilities q , the semantic loss $L^s(\alpha, [p \ q]) = L^s(\alpha, p)$.

Proof. Follows from the additive independence and truth axioms. Set $\beta = \text{true}$ in the additive independence axiom, and observe that this sets $L^s(\beta, q) = 0$ because of the truth axiom. \square

To maintain logical meaning, we postulate that semantic loss is monotone in the order of implication.

Axiom 7 (Monotonicity). If $\alpha \models \beta$, then the semantic loss $L^s(\alpha, p) \geq L^s(\beta, p)$ for any vector p .

Intuitively, as we add stricter requirements to the logical constraint, going from β to α and making it harder to satisfy, semantic loss cannot decrease. For example, when β enforces the output of a neural network to encode a subtree of a graph, and we tighten that requirement in α to be a path, semantic loss cannot decrease. Every path is also a tree and any solution to α is a solution to β .

A first consequence following the monotonicity axiom is that logically equivalent sentences must incur an identical semantic loss for the same probability vector p . Hence, the semantic loss is indeed a semantic property of the logical

sentence, and *does not depend on the syntax* of the sentence.

Proposition 6. If $\alpha \equiv \beta$, then the semantic loss $L^s(\alpha, p) = L^s(\beta, p)$ for any vector p .

A second consequence is that semantic loss must be non-negative.

Proposition 7 (Non-Negativity). Semantic loss is non-negative.

Proof. Because $\alpha \models \text{true}$ for all α , the monotonicity axiom implies that $\forall p, L^s(\alpha, p) \geq L^s(\text{true}, p)$. By the truth axiom, $L^s(\text{true}, p) = 0$, and therefore $L^s(\alpha, p) \geq 0$ for all choices of α and p . \square

A state \mathbf{x} is equivalently represented as a data vector, as well as a logical constraint that enforces a value for every variable in \mathbf{X} . When both the constraint and the predicted vector represent the same state (for example, $X_1 \wedge \neg X_2 \wedge X_3$ vs. $[1 \ 0 \ 1]$), there should be no semantic loss.

Axiom 8 (Identity). For any state \mathbf{x} , there is zero semantic loss between its representation as a sentence, and its representation as a deterministic vector: $\forall \mathbf{x}, L^s(\mathbf{x}, \mathbf{x}) = 0$.

The axioms above together imply that any vector satisfying the constraint must incur zero loss. For example, when our constraint α requires that the output vector encodes an arbitrary total ranking, and the vector \mathbf{x} correctly represents a single specific total ranking, there is no semantic loss.

Proposition 8 (Satisfaction). If $\mathbf{x} \models \alpha$, then the semantic loss $L^s(\alpha, \mathbf{x}) = 0$.

Proof of Proposition 8. The monotonicity axiom specializes to say that if $\mathbf{x} \models \alpha$, we have that $\forall p, L^s(\mathbf{x}, p) \geq L^s(\alpha, p)$. By choosing p to be \mathbf{x} , this implies $L^s(\mathbf{x}, \mathbf{x}) \geq L^s(\alpha, \mathbf{x})$. From the identity axiom, $L^s(\mathbf{x}, \mathbf{x}) = 0$, and therefore $0 \geq L^s(\alpha, \mathbf{x})$. Proposition 7 bounds the loss from below as $L^s(\alpha, \mathbf{x}) \geq 0$. \square

As a special case, logical literals (x or $\neg x$) constrain a single variable to take on a single value, and thus play a role similar to the labels used in supervised learning. Such constraints require an even tighter correspondence: semantic loss must act like a classical loss function (i.e., cross entropy).

Axiom 9 (Label-Literal Correspondence). The semantic loss of a single literal is proportionate to the cross-entropy loss for the equivalent data label: $L^s(x, p) \propto -\log(p)$ and $L^s(\neg x, p) \propto -\log(1 - p)$.

Next, we have the symmetry axioms.

Axiom 10 (Value Symmetry). For all p and α , we have that $L^s(\alpha, p) = L^s(\bar{\alpha}, 1 - p)$ where $\bar{\alpha}$ replaces every variable in α by its negation.

Axiom 11 (Variable Symmetry). Let α be a sentence over \mathbf{X} with probabilities \mathbf{p} . Let π be a permutation of the variables \mathbf{X} , let $\pi(\alpha)$ be the sentence obtained by replacing variables x by $\pi(x)$, and let $\pi(\mathbf{p})$ be the corresponding permuted vector of probabilities. Then, $L^s(\alpha, \mathbf{p}) = L^s(\pi(\alpha), \pi(\mathbf{p}))$.

The value and variable symmetry axioms together imply the equality of the multiplicative constants in the label-literal duality axiom for all literals.

Lemma 9. *There exists a single constant K such that $L^s(X, p) = -K \log(p)$ and $L^s(\neg X, p) = -K \log(1 - p)$ for any literal x .*

Proof. Value symmetry implies that $L^s(X_i, \mathbf{p}) = L^s(\neg X_i, 1 - \mathbf{p})$. Using label-literal correspondence, this implies $K_1 \log(p_i) = K_2 \log(1 - (1 - p_i))$ for the multiplicative constants K_1 and K_2 that are left unspecified by that axiom. This implies that the constants are identical. A similar argument based on variable symmetry proves equality between the multiplicative constants for different i . \square

Finally, this allows us to prove the following form of semantic loss for a state \mathbf{x} .

Lemma 10. *For state \mathbf{x} and vector \mathbf{p} , we have $L^s(\mathbf{x}, \mathbf{p}) \propto -\sum_{i:\mathbf{x} \models X_i} \log p_i - \sum_{i:\mathbf{x} \models \neg X_i} \log(1 - p_i)$.*

Proof of Lemma 10. A state \mathbf{x} is a conjunction of independent literals, and therefore subject to the additive independence axiom. Each literal’s loss in this sum is defined by Lemma 9. \square

The following and final axiom requires that semantic loss is proportionate to the logarithm of a function that is additive for mutually exclusive sentences.

Axiom 12 (Exponential Additivity). Let α and β be mutually exclusive sentences (i.e., $\alpha \wedge \beta$ is unsatisfiable), and let $f^s(K, \alpha, \mathbf{p}) = K^{-L^s(\alpha, \mathbf{p})}$. Then, there exists a positive constant K such that $f^s(K, \alpha \vee \beta, \mathbf{p}) = f^s(K, \alpha, \mathbf{p}) + f^s(K, \beta, \mathbf{p})$.

We are now able to state and prove the main uniqueness theorem.

Theorem 11 (Uniqueness). *The semantic loss function in Definition 1 satisfies all axioms in Appendix A and is the only function that does so, up to a multiplicative constant.*

Proof of Theorem 11. The truth axiom states that $\forall \mathbf{p}, f^s(K, \text{true}, \mathbf{p}) = 1$ for all positive constants K . This is the first Kolmogorov axiom of probability. The second Kolmogorov axiom for $f^s(K, \cdot, \mathbf{p})$ follows from the additive independence axiom of semantic loss. The

third Kolmogorov axiom (for the finite discrete case) is given by the exponential additivity axiom of semantic loss. Hence, $f^s(K, \cdot, \mathbf{p})$ is a probability distribution for some choice of K , which implies the definition up to a multiplicative constant. \square

B. Specification of the Convolutional Neural Network Model

Table 6 shows the slight architectural difference between the CNN used in ladder nets and ours. The major difference lies in the choice of ReLU. Note we add standard padded cropping to preprocess images and an additional fully connected layer at the end of the model, neither is used in ladder nets. We only make those slight modification so that the baseline performance reported by Rasmus et al. (2015) can be reproduced.

C. Hyper-parameter Tuning Details

Validation sets are used for tuning the weight associated with semantic loss, the only hyper-parameter that causes noticeable difference in performance for our method. For our semi-supervised classification experiments, we perform a grid search over $\{0.001, 0.005, 0.01, 0.05, 0.1\}$ to find the optimal value. Empirically, 0.005 always gives the best or nearly the best results and we report its results on all experiments.

For the FASHION dataset specifically, because MNIST and FASHION share the same image size and structure, methods developed in MNIST should be able to directly perform on FASHION without heavy modifications. Because of this, we use the same hyper-parameters when evaluating our method. However, for the sake of fairness, we subject ladder nets to a small-scale parameter tuning in case its performance is more volatile.


For the grids experiment, the only hyper parameter that needed to be tuned was again the weight given to semantic loss, which after trying $\{0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1\}$ was selected to be 0.5 based on validation results. For the preference learning experiment, we initially chose the semantic loss weight from $\{0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1\}$ to be 0.1 based on validation, and then further tuned the weight to 0.25.

D. Specification of Complex Constraint Models

Grids To compile our grid constraint, we first use Nishino et al. (2017) to generate a constraint for each source destination pair. Then, we conjoin each of these with indicators specifying which source and destination




(a) Confidently Correct



(b) Unconfidently Correct



(c) Unconfidently Incorrect



(d) Confidently Incorrect

Figure 5: FASHION pictures grouped by how confidently the supervised base model classifies them correctly. With semantic loss, the final semi-supervised model predicts all correctly and confidently.

Table 6: Specifications of CNNs in Ladder Net and our proposed method.

CNN in Ladder Net	CNN in this paper
Input 32×32 RGB image	
	Resizing to 36×36 with padding Cropping Back
Whitening Contrast Normalization Gaussian Noise with std. of 0.3	
3×3 conv. 96 BN LeakyReLU	3×3 conv. 96 BN ReLU
3×3 conv. 96 BN LeakyReLU	3×3 conv. 96 BN ReLU
3×3 conv. 96 BN LeakyReLU	3×3 conv. 96 BN ReLU
2×2 max-pooling stride 2 BN	
3×3 conv. 192 BN LeakyReLU	3×3 conv. 192 BN ReLU
3×3 conv. 192 BN LeakyReLU	3×3 conv. 192 BN ReLU
3×3 conv. 192 BN LeakyReLU	3×3 conv. 192 BN ReLU
2×2 max-pooling stride 2 BN	
3×3 conv. 192 BN LeakyReLU	3×3 conv. 192 BN ReLU
1×1 conv. 192 BN LeakyReLU	3×3 conv. 192 BN ReLU
1×1 conv. 10 BN LeakyReLU	1×1 conv. 10 BN ReLU
Global meanpool BN	
	Fully connected BN
10-way softmax	

pair must be used, and finally we disjoin all of these together to form our constraint.

To generate the data, we begin by randomly removing one third of edges. We then filter out connected components with fewer than 5 nodes to reduce degenerate cases, and proceed with randomly selecting pairs of nodes to create data points.

The predictive model we employ as our baseline is a 5 layer MLP with 50 hidden sigmoid units per layer. It is trained using Adam Optimizer, with full data batches (Kingma & Ba, 2015). Early stopping with respect to validation loss is used as a regularizer.

Preference Learning We split each user’s ordering into their ordering over sushis 1,2,3,5,7,8, which we use as the features, and their ordering over 4,6,9,10 which are the labels we predict. The constraint is compiled directly from

logic, as this can be done in a straightforward manner for an n-item ordering.

The predictive model we use here is a 3 layer MLP with 25 hidden sigmoid units per layer. It is trained using Adam Optimizer with full data batches (Kingma & Ba, 2015). Early stopping with respect to validation loss is used as a regularizer.

E. Probabilistic Soft Logic Encodings

We here give both encodings on the exactly-one constraint over three x_1, x_2, x_3 . The first encoding is:

$$(\neg x_1 \wedge x_2 \wedge x_3) \vee (x_1 \wedge \neg x_2 \wedge x_3) \vee (x_1 \wedge x_2 \wedge \neg x_3)$$

The second encoding is:

$$(x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2) \wedge (\neg x_1 \vee \neg x_3) \wedge (\neg x_2 \vee \neg x_3)$$

Both encodings extend to cases whether the number of variables is arbitrary.

The norm functions used for these experiments are as described in Kimmig et al. (2012), with the loss for an interpretation I being defined as follows:

$$\begin{aligned}x_1 \wedge x_2 &= \max\{0, I(x_1) + I(x_2) - 1\} \\x_1 \vee x_2 &= \min\{I(x_1) + I(x_2), 1\} \\ \neg x_1 &= 1 - I(x_1)\end{aligned}$$