

Image Classification with Consistent Supporting Evidence

Peiqi Wang

Ruizhi Liao

Daniel Moyer

Massachusetts Institute of Technology, Cambridge, MA, USA

Seth Berkowitz

Steven Horng

Beth Israel Deaconess Medical Center / Harvard Medical School, Boston, MA, USA

Polina Golland

Massachusetts Institute of Technology, Cambridge, MA, USA

WPQ@MIT.EDU

RUIZHI@CSAIL.MIT.EDU

DMOYER@CSAIL.MIT.EDU

SBERKOWI@BIDMC.HARVARD.EDU

SHORNG@BIDMC.HARVARD.EDU

POLINA@CSAIL.MIT.EDU

Abstract

Adoption of machine learning models in health-care requires end users’ trust in the system. Models that provide additional supportive evidence for their predictions promise to facilitate adoption. We define consistent evidence to be both compatible and sufficient with respect to model predictions. We propose measures of model inconsistency and regularizers that promote more consistent evidence. We demonstrate our ideas in the context of edema severity grading from chest radiographs. We demonstrate empirically that consistent models provide competitive performance while supporting interpretation.

Keywords: Interpretability, Medical Image Analysis

1. Introduction

Identifying radiological findings and inferring disease stages from medical images is common in clinical practice. Many models make predictions without explaining the conclusion. In contrast, human experts often provide specific explanation based on prior knowledge of human physiology to support their image-based diagnosis. We aim to build models that are transparent in the reasoning process, at an appropriate level of understanding consumable by end users, e.g., clinicians. What additional information should a machine learning model provide to gain the trust of its end users? We propose a solution motivated by an example of how radiologists themselves operate.



Input	Unregularized	Regularized
	Prediction: $\hat{y} = 0$ Evidence: • $\mathcal{I}(1)$: NULL • $\mathcal{I}(2)$: septal lines, interstitial abnormality • $\mathcal{I}(3)$: NULL Inconsistent	Prediction: $\hat{y} = 0$ Evidence: • $\mathcal{I}(1)$: NULL • $\mathcal{I}(2)$: NULL • $\mathcal{I}(3)$: NULL Consistent
	Prediction: $\hat{y} = 2$ Evidence: • $\mathcal{I}(1)$: hilar congestion • $\mathcal{I}(2)$: NULL • $\mathcal{I}(3)$: NULL Inconsistent	Prediction: $\hat{y} = 2$ Evidence: • $\mathcal{I}(1)$: hilar congestion • $\mathcal{I}(2)$: interstitial abnormality • $\mathcal{I}(3)$: NULL Consistent

Figure 1: Our model provides prediction of disease stage y and supporting evidence. We use $\mathcal{I}(c)$ to denote the set of evidence labels detected in the image that directly support disease stage c . We show examples of inconsistent evidence highlighted in red produced by the baseline model (left column). Our proposed regularizer corrects these mistakes so that predicted evidence label becomes compatible (top right) and sufficient (bottom right).

Radiological *findings* are concepts determined as useful by radiologists. Findings include image features, pathological states, and observations about the underlying physiology (Gluecker et al., 1999). The radiologists aggregate the findings to provide an overall interpretation of the image. They support the eventual diagnosis by providing an account of the identified findings based on prior knowledge of relationships

between findings and the patient’s physiological state. We propose and demonstrate an approach that recapitulates the reasoning process of domain experts. In addition to primary predictions, the model provides supporting evidence, i.e., findings, deemed useful by the end users.

It is critical that predictions and supporting evidence are consistent with each other. In practice, radiologists cannot draw their conclusions based on incompatible evidence, nor could they support their conclusions with insufficient evidence. Similarly, end users will question the credibility of a model when its predictions and accompanying evidence are incompatible or insufficient.

In this paper, we build explainable models that supplement their predictions with consistent supporting evidence, illustrated in Figure 1. We define measures of inconsistency between the model’s primary output and its supporting evidence and propose simple regularizers that encourage the classifier to be more consistent. We demonstrate that we can train consistent models without loss in performance in the context of pathology grading from a chest radiograph.

2. Related Work

2.1. Interpretable Machine Learning

While model interpretability is an important topic in machine learning, few methods take the end users’ needs into account. For example, some method localize image regions important for a prediction (Zhou et al., 2016; Selvaraju et al., 2016), but fail to express what properties of the image region are associated with the model output. Others aim to use a simpler model (Caruana et al., 2015) or to approximate the behavior of a complex model with a simpler one (Ribeiro et al., 2016). While effective for handling low dimensional tabular data where the covariates are physically meaningful attributes, such methods are less useful for extremely high dimensional imaging data. Our work provide clinically meaningful supporting evidence useful to end users of the system, rather than support the developers’ understanding of how the model reaches its decision.

Another approach is to train a classifier whose predictions rely on higher-level concepts. Unsupervised methods can make for a more interpretable model for general purpose tasks but cannot take advantage of strong domain knowledge ubiquitous in healthcare (Alvarez Melis and Jaakkola, 2018). Alternatively,

concept bottleneck models that learn concepts with supervision have been applied to arthritis grading (Koh et al., 2020), retinal disease classification (De Fauw et al., 2018), and other applications (Losch et al., 2019; Bucher et al., 2019). This strategy relies on the appropriate choice of the concepts to maintain good performance. In contrast, our model predictions do not solely rely on supporting evidence, thus avoiding the undue influence the choice of supporting evidence or concepts might have on the accuracy of the task label prediction.

In our work, predictions and evidence are separate outputs of the model. Most closely related prior method focus on learning a mapping to the product space of the task label and supporting evidence (Hind et al., 2019; Codella et al., 2019), with application to text classification (Zaidan et al., 2007; Zhang et al., 2016). In contrast, we learn a structured output where known relationships between predictions and evidence are enforced. We inject domain specific knowledge and require our model to provide supporting evidence that is clinically feasible under a specific prediction.

A recently demonstrated unsupervised strategy that requires a forward model that relates supporting evidence to a subset of the input features is also relevant (Raghu et al., 2021). This method can be difficult to implement in our radiograph grading task as it assumes knowledge of an accurate forward model, from supporting evidence to a high dimensional image, which is infeasible for most medical imaging problems.

2.2. Logical Constraints

There are multiple ways to represent symbolic constraints. For example, trees have been used to express subsumption relationships between attributes, e.g., hierarchical annotation of medical images (Dimitrovski et al., 2011). Hierarchical multi-label learning aims to enforce such constraints (Bi and Kwok, 2011; Yan et al., 2015; Wehrmann et al., 2018; Giunchiglia and Lukasiewicz, 2020). Unfortunately, trees are overly restrictive and cannot express consistency constraints that are important in our application.

Alternatively, boolean statements can be quite expressive in representing logical constraints. Logical constraints on model outputs can be enforced by replacing logical operators with their subdifferentiable fuzzy t-norms (Diligenti et al., 2017; Li et al., 2019) or through the use of specialized loss functions (Xu

et al., 2018). Our approach to representing and enforcing logical constraints is both easier to interpret and simpler to implement. Moreover, we enforce logical constraints by making changes to how our model provides supporting evidence while keeping main task predictions unchanged.

3. Method

In this section, we define consistency of evidence and introduce measures of inconsistency. We provide an example application to ground our definitions. We construct novel loss functions that encourage consistency and discuss optimization that arises when training classifiers with consistent supporting evidence.

3.1. Problem Setup

Let x, y , and $z = (z_1, \dots, z_K)$ be random variables representing an image, a C -class task label, and K binary evidence labels. Let \mathcal{D}_t be a data set that includes pairs (x, y) and (x, z_k) for $k = 1, \dots, K$. In this work, we do not assume full tuples (x, y, z_1, \dots, z_K) are available. Learning joint predictors from available sub-tuples is an interesting direction that is outside the scope of this paper. Moreover, we allow the same image x to be included as part of several different pairs in the data set \mathcal{D}_t .

We use the training data set \mathcal{D}_t to build probabilistic classifiers $p(y = c | x; \theta)$ and $p(z_k = z | x; \theta)$ for $k = 1, \dots, K$. The maximum a posteriori (MAP) estimates of the task label y and evidence labels z is obtained via

$$\hat{y} = \arg \max_{c \in [C]} p(y = c | x; \theta), \quad (1)$$

$$\hat{z}_k = \arg \max_{z \in \{-1, +1\}} p(z_k = z | x; \theta), \quad k = 1, \dots, K. \quad (2)$$

We use $\hat{y}(x)$ to express explicit dependence of the predicted label \hat{y} on the input image x .

3.2. Consistent Evidence

We assume that domain experts provide domain specific knowledge in the form of logical constraints between the task label y and the evidence labels z . We identify two major logical constraints that are important in our application, specifically that supporting evidence should be compatible and sufficient with the task label.

Let $\mathcal{I}_1 : [C] \rightarrow \mathcal{P}([K])$ be the indexing function for evidence that is incompatible with a particular value of task label, where we use $\mathcal{P}(\cdot)$ to denote the power set. Specifically, if evidence labels $\{z_{i_1}, \dots, z_{i_M}\}$ are incompatible with task label $y = c$, then $\mathcal{I}_1(c) = \{i_1, \dots, i_M\}$. Let $\mathcal{I}_2 : [C] \rightarrow \mathcal{P}([K])$ be the indexing function for evidence that directly supports a particular value of task label. We assume that \mathcal{I}_1 and \mathcal{I}_2 are provided by domain experts.

Definition 1 (*Consistent Evidence*) *The task label $y \in [C]$ and the evidence label vector $z = (z_1, \dots, z_K) \in \{-1, +1\}^K$ are consistent if*

$$\forall k \in \mathcal{I}_1(y) : z_k = -1, \quad (3)$$

$$\exists k \in \mathcal{I}_2(y) : z_k = +1. \quad (4)$$

The first criterion specifies that no evidence is incompatible with the task label y . The second criterion specifies that there should be at least one direct evidence label present that supports the task label y .

In reality, perfectly consistent evidence may not be necessary or possible. For example, domain experts often specify constraints with a notion of uncertainty, e.g., $\mathcal{I}_1(y)$ is incompatible with y most of the time except for occasional corner cases. In addition, certain direct evidence might be so rare that it becomes impossible to include it in $\mathcal{I}_2(y)$. Therefore, it is perfectly sensible that there is no direct evidence present in some cases, if we have not included the corresponding evidence label in the construction. This motivates us to consider these constraints in probabilistic terms.

Definition 1 is a specification over the values that random variables can take. The same definition applies to the true data distribution $\{y, z\}$ and to the predicted distribution $\{\hat{y}, \hat{z}\}$. In practice, we construct training data $\{y, z\}$ to be perfectly consistent and demonstrate a training method that encourages the model outputs $\{\hat{y}, \hat{z}\}$ to be consistent as well.

Moreover, we are not restricted to predicting the findings in $\cup_{c \in [C]} \mathcal{I}_2(c)$. If some findings provide useful information but are not directly supportive, they can still be included in the set of evidence labels.

3.3. Example Application

This section illustrates the construction of indexing functions \mathcal{I}_1 and \mathcal{I}_2 for the pulmonary edema grading task that motivated our work.

Pulmonary edema is defined as an abnormal accumulation of fluid in the lungs. Higher hydrostatic

pressure in the vasculature causes more severe symptoms. Typically, radiologists grade the severity of edema based on findings that are typical of the most severe stage of pulmonary edema (Gluecker et al., 1999).

We use a categorization that identifies four edema severity levels, in order of increasing severity: *no edema* (0), *mild edema* (1), *moderate edema* (2), and *severe edema* (3) (Liao et al., 2019; Horng et al., 2021). The edema severity grading task involves assigning a severity level $y \in \{0, 1, 2, 3\}$ to a test image. In this task, there are 4 classes, i.e., $C = 4$.

Table 1: Findings that directly support a particular severity level.

Severity y	Findings $\mathcal{I}_2(y)$
0 (none)	-
1 (mild)	vascular congestion hilar congestion peribronchial cuffing
2 (moderate)	septal lines interstitial abnormality
3 (severe)	air bronchograms parenchymal opacity

Table 2: Examples of consistent evidence.

Severity y	Evidence z
1	peribronchial cuffing
2	vascular congestion septal lines interstitial abnormality

In our work, we identify $K = 7$ supporting evidence labels deemed useful by clinicians, as shown in Table 1. They are canonical radiological manifestation of the underlying pathology. End users expect presence of these findings to be indicative of a specific edema severity level.

As an example, radiologists grade an image as *moderate edema* if they observe *septal lines* (short parallel lines at the periphery of the lung) or *intersti-*

Table 3: Examples of inconsistent evidence. First two examples are incompatible. The latter two examples are insufficient.

Severity y	Evidence z
1	hilar congestion septal lines
2	vascular congestion interstitial abnormality air bronchograms
1	-
3	septal lines

tial abnormality (excess fluids in the supporting tissue within the lung). Note that presence of evidence from a lower value of edema severity is not inconsistent. For example, radiologists may at the same time observe presence of *vascular congestion* (enlargement of pulmonary veins) and *septal lines* in a moderate edema case.

In the severity grading task, we consider an evidence label as incompatible if its presence directly supports a higher level severity level. Thus define

$$\mathcal{I}_1(c) = \bigcup_{c' > c} \mathcal{I}_2(c'). \quad (5)$$

As an example, a model that grades an image as *moderate edema* should not use *air bronchograms* (opacification of alveoli) as supporting evidence.

We consider evidence as insufficient when no direct evidence for edema severity grading is present. As an example, a model which grades an image as *severe edema* cannot rely on *septal lines* only to support its prediction.

Tables 2 and 3 illustrate further examples of consistent and inconsistent evidence, respectively.

3.4. Measuring Inconsistency

We quantify the inconsistency probabilistically based on Definition 1. First, we define a measure of incompatibility as the probability that there is an incompatible evidence label

$$\mathbb{P} \left[\bigcup_{k \in \mathcal{I}_1(y)} \{z_k = +1\} \right]. \quad (6)$$

To facilitate computation, we upper bound this probability using union bound by

$$\mathcal{R}_1(y, z) = \sum_{k \in \mathcal{I}_1(y)} \mathbb{P}[z_k = +1]. \quad (7)$$

We provide an estimate of incompatibility over data set \mathcal{D} by taking expectation over its empirical distribution

$$\mathcal{R}_1(\mathcal{D}) = \mathbb{E}_{(y,z) \sim \mathcal{D}} \left[\sum_{k \in \mathcal{I}_1(y)} \mathbb{1}[z_k = +1] \right], \quad (8)$$

where we have replaced $\mathbb{P}[z_k = +1]$ with $\mathbb{1}[z_k = +1]$ since z_k is binary valued. Intuitively, $\mathcal{R}_1(\mathcal{D})$ is the average count of evidence labels incompatible with the task label.

In this work, we use Equation 8 as a measure of incompatible evidence. We also note that one could eliminate the dependency on the size of $\mathcal{I}_1(y)$ by defining

$$\mathbb{E}_{(y,z) \sim \mathcal{D}} \left[\frac{1}{|\mathcal{I}_1(y)|} \sum_{k \in \mathcal{I}_1(y)} \mathbb{1}[z_k = +1] \right]. \quad (9)$$

We prefer the measure of incompatibility defined in Equation 8 for measuring incompatibility as it arises naturally from bounding the probability of incompatible evidence.

Similarly, we define a measure of insufficiency as the probability that there is no sufficient evidence

$$\mathbb{P} \left[\bigcap_{k \in \mathcal{I}_2(y)} \{z_k = -1\} \right], \quad (10)$$

which leads to an upper bound

$$\mathcal{R}_2(y, z) = \min_{k \in \mathcal{I}_2(y)} \mathbb{P}[z_k = -1] \quad (11)$$

and its empirical estimate

$$\mathcal{R}_2(\mathcal{D}) = \mathbb{E}_{(y,z) \sim \mathcal{D}} \left[\min_{k \in \mathcal{I}_2(y)} [1 - \mathbb{1}[z_k = +1]] \right] \quad (12)$$

$$= 1 - \mathbb{E}_{(y,z) \sim \mathcal{D}} \left[\max_{k \in \mathcal{I}_2(y)} \mathbb{1}[z_k = +1] \right]. \quad (13)$$

Note that $\mathcal{R}_2(\mathcal{D})$ is the average count of absence of direct evidence.

Now we can provide an upper bound on probability of inconsistent evidence $\mathcal{R}(y, z) = \mathcal{R}_1(y, z) + \mathcal{R}_2(y, z)$ and its empirical estimate $\mathcal{R}(\mathcal{D}) = \mathcal{R}_1(\mathcal{D}) + \mathcal{R}_2(\mathcal{D})$.

3.5. Consistency Regularization

Models trained naively to predict labels y and z jointly are not guaranteed to be consistent. Here, we provide regularizers that encourage supporting evidence to be more consistent.

Observe that Equations 7 and 11 are upper bounds on the true probability of model being inconsistent. We can simply use these upper bounds, or modification thereof, as regularizers. We opt to use cross entropy to avoid inconsistent evidence.

To penalize incompatibility, we define

$$\mathcal{R}_1(\theta) = -\mathbb{E}_{x \sim \mathcal{D}} \left[\sum_{k \in \mathcal{I}_1(\hat{y}(x))} \ln p(z_k = -1 | x; \theta) \right]. \quad (14)$$

Intuitively, $\mathcal{R}_1(\theta)$ penalizes evidence probability that is incompatible with the predicted task label. Including $\mathcal{R}_1(\theta)$ in the loss function is equivalent to supplying pseudo negative samples for evidence obtained from the predicted task label \hat{y} .

Instead of penalizing incompatibility with respect to MAP estimate of the task label $\hat{y}(x)$, we can penalize incompatibility for each value of task label weighted by the posterior probability, i.e.,

$$\tilde{\mathcal{R}}_1(\theta) = -\mathbb{E}_{x \sim \mathcal{D}} \left[\sum_{c \in [C]} \sum_{k \in \mathcal{I}_1(c)} p(y = c | x; \theta) \cdot \ln p(z_k = -1 | x; \theta) \right]. \quad (15)$$

In contrast to Equation 14 where gradients cannot flow through $\hat{y}(x)$ due to the arg max operator, Equation 15 provides a softer regularizer that affects the predictions of both the task and the evidence labels.

Similarly, we define

$$\mathcal{R}_2(\theta) = -\mathbb{E}_{x \sim \mathcal{D}} \left[\ln \max_{k \in \mathcal{I}_2(\hat{y}(x))} p(z_k = +1 | x; \theta) \right]. \quad (16)$$

Intuitively, $\mathcal{R}_2(\theta)$ encourages presence of *some* evidence to support predicted task label. Including $\mathcal{R}_2(\theta)$ in the loss function is equivalent to supplying pseudo positive samples obtained from the predicted task label \hat{y} .

Similar to Equation 15, we can penalize insufficiency using posterior probability as weights,

$$\tilde{\mathcal{R}}_2(\theta) = -\mathbb{E}_{x \sim \mathcal{D}} \left[\sum_{c \in [C]} p(y = c \mid x; \theta) \cdot \ln \max_{k \in \mathcal{L}_2(c)} p(z_k = +1 \mid x; \theta) \right]. \quad (17)$$

In our work, we focus on regularizers $\mathcal{R}_1(\theta)$ and $\mathcal{R}_2(\theta)$. In Appendix B, we provide preliminary comparison of regularizers $\mathcal{R}_1(\theta)$ and $\mathcal{R}_2(\theta)$ with the *soft* regularizers $\tilde{\mathcal{R}}_1(\theta)$ and $\tilde{\mathcal{R}}_2(\theta)$. We leave further investigation of soft regularizers to future work.

3.6. Optimization

We apply deep multitask learning for joint predictions of y, z_1, \dots, z_K . In particular, we parameterize $p(y \mid x; \theta)$ and $p(z_k \mid x; \theta)$ for $k = 1, \dots, K$ with neural network $f(x; \theta)$ and assume function f outputs logits over $K + 1$ marginals.

Given a classification loss function $L(\cdot, \cdot)$, the objective is simply the empirical risk,

$$\mathcal{L}(\theta) = \mathbb{E}_{(x,y) \sim \mathcal{D}_t} [L(y, f(x; \theta))] \quad (18)$$

$$+ \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{(x,z_k) \sim \mathcal{D}_t} [L(z_k, f(x; \theta))]. \quad (19)$$

We add consistency regularization to multitask classification loss, which yields a regularized empirical risk minimization problem

$$\min_{\theta} \mathcal{L}(\theta) + \omega_1 \mathcal{R}_1(\theta) + \omega_2 \mathcal{R}_2(\theta), \quad (20)$$

where $\omega_1, \omega_2 \in \mathbb{R}^+$ are coefficients that control the degree of regularization.

4. Implementation Details

We use residual networks to parameterize our probabilistic classifiers (He et al., 2016). The network is modified to output a $(C + K)$ -dimensional vector representing the posterior marginal probabilities for y, z_1, \dots, z_K .

We use weighted cross entropy loss as $L(\cdot, \cdot)$ to handle class imbalances. We employ the Adam optimizer with a constant learning rate of $2 \cdot 10^{-4}$ with mini-batch size of 32 for stochastic optimization of network parameters (Kingma and Ba, 2015). Each

gradient update involves random sampling of a label (task or evidence), assembling a mini-batch of data corresponding to the sampled label, computing the objective function, and updating parameters with backpropagated gradients. This approach enables us to learn even if some labels are missing for some images.

We normalize images to zero mean and unit variance, and resize them to 224x224 pixels. We apply random image augmentations to images, e.g., crop, horizontal flip, brightness and contrast variations, to alleviate model overfitting.

We implement Equation 16 by substituting the max operator with a soft maximum operator, i.e., $\text{LSE}_{i \in [n]}(x_i) = \log \sum_i \exp(x_i)$. This way, we enforce sufficiency of evidence by upscaling probabilities of direct evidence that are larger to begin with.

We use exactly the same network architecture, data augmentation, and optimization parameters to isolate the impact of the proposed regularization on consistency and performance.

We compute mean and standard deviation statistics for inconsistency and test prediction from 3 runs with different random seed.




Figure 2: Model inconsistency across different values of \hat{y} evaluated on the test set $\hat{\mathcal{D}}$ for a naively trained model, i.e., $\omega_1 = \omega_2 = 0$. The majority of inconsistent evidence comes from incompatible evidence associated with small values of the predicted task label \hat{y} .




Figure 3: The effect of varying strength of regularizations during training on model consistency. Here $\mathcal{R}_1, \mathcal{R}_2$ is short hand for $\mathcal{R}_1(\hat{\mathcal{D}}), \mathcal{R}_2(\hat{\mathcal{D}})$, respectively. The proposed regularizers $\mathcal{R}_1(\theta), \mathcal{R}_2(\theta)$ encourage the model to provide more compatible evidence (left) or more sufficient evidence (middle), respectively. The application of regularizers at the same time encourages the model to provide more consistent evidence (right).

5. Experiments

5.1. Data

We use a subset of 238,086 frontal-view chest X-ray from the MIMIC-CXR data set (Johnson et al., 2019). We split the data set into training (217,016), validation (10,445), and test (10,625) sets randomly. The performance of predicted evidence is computed over this test set. There is no patient overlap between training, validation and test sets.

Edema severity labels are extracted from associated reports by searching for keywords that are indicative of a specific disease stage. The 7,802 labeled image/report pairs are split into training (6,656), validation (648), and test (498) set. The test set was corrected for keyword matching errors by an expert radiologist, as detailed in prior work (Chauhan et al., 2020). We use $\hat{\mathcal{D}}$ to denote this test set that includes images and predicted labels (x, \hat{y}, \hat{z}) . All subsequent evaluations of model consistency and performance is computed using $\hat{\mathcal{D}}$.

5.2. Model Inconsistency

We examine model inconsistency overall and over partitions of data with respect to values of predicted label \hat{y} . The sum of model inconsistency over the partitions gives the quantities $\mathcal{R}_1(\hat{\mathcal{D}})$ in Equation 8 and $\mathcal{R}_2(\hat{\mathcal{D}})$ in Equation 13.

Figure 2 reports model inconsistency over partitions of \hat{y} for a model that is trained without consistency regularization, i.e., $\omega_1 = \omega_2 = 0$. We observe that $\mathcal{R}_1(\hat{\mathcal{D}})$ is typically larger than $\mathcal{R}_2(\hat{\mathcal{D}})$ due to the fact that $\mathcal{R}_1(\hat{\mathcal{D}})$ is essentially an average count of po-

tentially many incorrect evidence labels, while $\mathcal{R}_2(\hat{\mathcal{D}})$ is an average count of missing evidence, and therefore is upper bounded by 1. We also observe a downward trend in values of $\mathcal{R}_1(\hat{\mathcal{D}})$ with increasing values for \hat{y} . This is reasonable as there are many ways to make mistake with a small \hat{y} , while no way to provide conflicting evidence when $\hat{y} = 3$ in our framework.

5.3. Consistency Regularization

To demonstrate that proposed regularization promotes model consistency, we vary values of ω_1, ω_2 in the objective function and train multiple models. We select the most accurate model on the validation set and compute inconsistency on the test set $\hat{\mathcal{D}}$.

Figure 3 demonstrates the effects of regularization on model consistency. We observe that the regularizers $\mathcal{R}_1(\theta)$ and $\mathcal{R}_2(\theta)$ are effective in reducing the respective intended model inconsistency, indicated by a reduction of $\mathcal{R}_1(\hat{\mathcal{D}})$ in Figure 3(a) and $\mathcal{R}_2(\hat{\mathcal{D}})$ in Figure 3(b) respectively. Additionally, we observe that penalizing $\mathcal{R}_1(\hat{\mathcal{D}})$ inadvertently makes $\mathcal{R}_2(\hat{\mathcal{D}})$ larger and vice versa. This makes intuitive sense, since a model that is more likely to predicts absence of evidence will (i) less likely to provide incompatible evidence and (ii) less likely to provide some direct evidence. We observe that we can reduce both types of inconsistency by regularizing with both loss terms, as shown in Figure 3(c).

It important to note that even though $\mathcal{R}_2(\hat{\mathcal{D}})$ is relatively small in models trained with $\omega_1 = 0$, regularizing with $\mathcal{R}_2(\hat{\mathcal{D}})$ is necessary as we want to avoid situations in Figure 3(a) where $\mathcal{R}_2(\hat{\mathcal{D}})$ becomes intolerably large.




Figure 4: Correctly and incorrectly classified test images with supporting evidence given by a consistent model ($\omega_1 = \omega_2 = 10$). We use $\mathcal{I}(c)$ to denote the set of evidence labels detected in the image that directly support disease stage c .

5.4. Interpretability

Figure 4 illustrates how a consistent model (trained with $\omega_1 = \omega_2 = 10$) provides supporting evidence for randomly sampled test images. We provide correctly and incorrectly classified test images for each severity level. We observe that the regularized model provides consistent evidence in all 8 examples, even in cases where model prediction of task label is not correct.

How does providing consistent supporting evidence build trust in the model? We note that supporting findings are already described in radiological reports and can be easily mined for training and verified in an image. When the supporting evidence is clearly correct, it builds additional trust in the predicted task label. When the consistent but wrong evidence is presented, it is easy to see in the image and helps the end users understand why the main task label is

wrong. Our method avoid confusions that arise from model providing inconsistent evidence.

Crucially, evidence labels should not only be consistent, but also correct. To this point, we reported performance of evidence detection in Section 5.5. Our proposed regularizers offers a complementary tool to help end users understand why the model erred. Our method can be integrated with technologies, e.g., RCNN, GradCAM (Girshick et al., 2014; Selvaraju et al., 2016), that provide localization, i.e., confirmation that the model is focusing on the correct regions in the image.

5.5. Performance-Consistency Tradeoff

Next, we show that we can achieve good model consistency without compromising predictive performance. We vary ω_1, ω_2 together in the objective function and




Figure 5: The effect of regularization on model inconsistency (left 2) and performance of predicting task label y (middle 2) and evidence labels z (right 2). Here $\mathcal{R}_1, \mathcal{R}_2$ is short hand for $\mathcal{R}_1(\hat{\mathcal{D}}), \mathcal{R}_2(\hat{\mathcal{D}})$, respectively. When regularizing for both $\mathcal{R}_1, \mathcal{R}_2$, e.g., down the diagonals of the matrix, we notice dramatic decrease in model inconsistency and competitive performance for the regularized model.

train multiple models. We select for the most accurate model on the validation set for subsequent evaluations.

Figure 5 demonstrates that we can ensure satisfactory model consistency. At the same time, the regularized model achieves similar performance on the severity grading task. The improvement in performance can be attributed empirically to fact that heavily regularized models over-fit less. However, we do observe a drop in the average performance of the model for predicting evidence. The drop in predicted evidence performance is tolerable if we consider that the model rarely provides inconsistent supporting evidence.

Figure 5 reinforces previous observation that penalizing $\mathcal{R}_1(\hat{\mathcal{D}})$ makes $\mathcal{R}_2(\hat{\mathcal{D}})$ higher and vice versa, when ω_i for $i = 1, 2$ being held constant is different from 0. Figure 3 shows the first column, the first row, and the diagonal slices of the grid in left 2 sub-figures in Figure 5. We refer the reader to Table 4 in Appendix A for detailed numbers of inconsistency and performance along the diagonal slice of the grid.

6. Conclusion

We argue for supplementing model predictions with supporting evidence that is deemed useful by end users. We defined a notion of consistent evidence via incorporating domain specific constraints. Then, we proposed ways to measure and enforce such constraints during model training. We evaluated our method on the pulmonary edema severity grading task, which provides a grounding for our consistent evidence framework. We demonstrated that consistent models remain competitive on the main task.

Acknowledgments

This work was supported in part by NIH NIBIB NAC P41EB015902 grant, MIT Lincoln Laboratory, MIT JClinic, MIT Deshpande Center, and Philips.

References

- David Alvarez Melis and Tommi Jaakkola. Towards Robust Interpretability with Self-Explaining Neural Networks. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Wei Bi and James T. Kwok. Multi-label classification on tree- and DAG-structured hierarchies. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, pages 17–24, Madison, WI, USA, June 2011. Omnipress. ISBN 978-1-4503-0619-5.
- Maxime Bucher, Stéphane Herbin, and Frédéric Jurie. Semantic Bottleneck for Computer Vision Tasks. In C. V. Jawahar, Hongdong Li, Greg Mori, and Konrad Schindler, editors, *Computer Vision – ACCV 2018*, Lecture Notes in Computer Science, pages 695–712, Cham, 2019. Springer International Publishing. ISBN 978-3-030-20890-5. doi: 10.1007/978-3-030-20890-5_44.
- Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*, pages 1721–1730, Sydney, NSW, Australia, 2015. ACM Press. ISBN 978-1-4503-3664-2. doi: 10.1145/2783258.2788613.
- Geeticka Chauhan, Ruizhi Liao, William Wells, Jacob Andreas, Xin Wang, Seth Berkowitz, Steven Horng, Peter Szolovits, and Polina Golland. Joint modeling of chest radiographs and radiology reports for pulmonary edema assessment. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 529–539. Springer, 2020.
- Noel C. F. Codella, Michael Hind, Karthikeyan Natesan Ramamurthy, Murray Campbell, Amit Dhurandhar, Kush R. Varshney, Dennis Wei, and Aleksandra Mojsilović. Teaching AI to Explain its Decisions Using Embeddings and Multi-Task Learning. *ICML Workshop on Human in the Loop Learning*, June 2019.
- Jeffrey De Fauw, Joseph R. Ledsam, Bernardino Romera-Paredes, Stanislav Nikolov, Nenad Tomasev, Sam Blackwell, Harry Askham, Xavier Glorot, Brendan O’Donoghue, Daniel Visentin, George van den Driessche, Balaji Lakshminarayanan, Clemens Meyer, Faith Mackinder, Simon Bouton, Kareem Ayoub, Reena Chopra, Dominic King, Alan Karthikesalingam, Cían O. Hughes, Rosalind Raine, Julian Hughes, Dawn A. Sim, Catherine Egan, Adnan Tufail, Hugh Montgomery, Demis Hassabis, Geraint Rees, Trevor Back, Peng T. Khaw, Mustafa Suleyman, Julien Cornebise, Pearse A. Keane, and Olaf Ronneberger. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine*, 24(9):1342–1350, September 2018. ISSN 1078-8956, 1546-170X. doi: 10.1038/s41591-018-0107-6.
- Michelangelo Diligenti, Soumali Roychowdhury, and Marco Gori. Integrating Prior Knowledge into Deep Learning. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 920–923, December 2017. doi: 10.1109/ICMLA.2017.00-37.
- Ivica Dimitrovski, Dragi Kocev, Suzana Loskovska, and Sašo Džeroski. Hierarchical annotation of medical images. *Pattern Recognition*, 44(10-11): 2436–2449, October 2011. ISSN 00313203. doi: 10.1016/j.patcog.2011.03.026.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 580–587, June 2014. doi: 10.1109/CVPR.2014.81.
- Eleonora Giunchiglia and Thomas Lukasiewicz. Coherent hierarchical multi-label classification networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9662–9673. Curran Associates, Inc., 2020.
- Thomas Gluecker, Patrizio Capasso, Pierre Schnyder, François Gudinchet, Marie-Denise Schaller, Jean-Pierre Revelly, René Chiolero, Peter Vock, and Stéphan Wicky. Clinical and Radiologic Features of Pulmonary Edema. *RadioGraphics*, 19(6): 1507–1531, November 1999. ISSN 0271-5333. doi: 10.1148/radiographics.19.6.g99no211507.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recogni-

- tion. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, June 2016. doi: 10.1109/CVPR.2016.90.
- Michael Hind, Dennis Wei, Murray Campbell, Noel C. F. Codella, Amit Dhurandhar, Aleksandra Majsilović, Karthikeyan Natesan Ramamurthy, and Kush R. Varshney. TED: Teaching AI to Explain its Decisions. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 123–129, Honolulu HI USA, January 2019. ACM. ISBN 978-1-4503-6324-2. doi: 10.1145/3306618.3314273.
- Steven Horng, Ruizhi Liao, Xin Wang, Sandeep Dalal, Polina Golland, and Seth J. Berkowitz. Deep Learning to Quantify Pulmonary Edema in Chest Radiographs. *Radiology: Artificial Intelligence*, 3(2):e190228, March 2021. ISSN 2638-6100. doi: 10.1148/ryai.2021190228.
- Alistair E. W. Johnson, Tom J. Pollard, Seth J. Berkowitz, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-ying Deng, Roger G. Mark, and Steven Horng. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data*, 6(1):317, December 2019. ISSN 2052-4463. doi: 10.1038/s41597-019-0322-0.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *ICLR*, 2015.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept Bottleneck Models. In *ICML*, January 2020.
- Tao Li, Vivek Gupta, Maitrey Mehta, and Vivek Srikumar. A Logic-Driven Framework for Consistency of Neural Models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3924–3935, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1405.
- Ruizhi Liao, Jonathan E. Rubin, Grace Lam, Seth J. Berkowitz, Sandeep Dalal, William Wells, Steven Horng, and Polina Golland. Semi-supervised Learning for Quantification of Pulmonary Edema in Chest X-Ray Images. *ArXiv*, abs/1902.10785, 2019.
- Max Losch, Mario Fritz, and Bernt Schiele. Interpretability Beyond Classification Output: Semantic Bottleneck Networks. *arXiv:1907.10882 [cs]*, July 2019.
- Aniruddh Raghu, John Gutttag, Katherine Young, Eugene Pomerantsev, Adrian V. Dalca, and Collin M. Stultz. Learning to predict with supporting evidence: Applications to clinical risk prediction. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 95–104, Virtual Event USA, April 2021. ACM. ISBN 978-1-4503-8359-2. doi: 10.1145/3450439.3451869.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ”Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’16*, pages 1135–1144, New York, NY, USA, August 2016. Association for Computing Machinery. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939778.
- Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. *International Journal of Computer Vision*, 2016. ISSN 0920-5691, 1573-1405. doi: 10.1007/s11263-019-01228-7.
- Jonatas Wehrmann, Ricardo Cerri, and Rodrigo Barros. Hierarchical Multi-Label Classification Networks. In *International Conference on Machine Learning*, pages 5075–5084. PMLR, July 2018.
- Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Broeck. A Semantic Loss Function for Deep Learning with Symbolic Knowledge. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5502–5511. PMLR, July 2018.
- Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu. HD-CNN: Hierarchical Deep Convolutional Neural Networks for Large Scale Visual Recognition. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2740–2748, Santiago, Chile, December 2015. IEEE. ISBN 978-1-4673-8391-2. doi: 10.1109/ICCV.2015.314.
- Omar Zaidan, Jason Eisner, and Christine Piatko. Using ”Annotator Rationales” to Improve Ma-

chine Learning for Text Categorization. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 260–267, Rochester, New York, April 2007. Association for Computational Linguistics.

Ye Zhang, Iain Marshall, and Byron C. Wallace. Rationale-Augmented Convolutional Neural Networks for Text Classification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 795–804, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1076.

Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning Deep Features for Discriminative Localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2921–2929, Las Vegas, NV, USA, June 2016. IEEE. ISBN 978-1-4673-8851-1. doi: 10.1109/CVPR.2016.319.

Appendix A. Table for Figure 5.

Table 4 provides detailed numbers on inconsistency and performance.

Appendix B. Training with Soft Regularizer

We compare behavior of model trained using hard regularizers $\mathcal{R}_1(\theta)$, $\mathcal{R}_2(\theta)$ versus the soft regularizers $\tilde{\mathcal{R}}_1(\theta)$, $\tilde{\mathcal{R}}_2(\theta)$.

In Figure 6, we note that the soft regularizers reduce model inconsistency while maintain model consistency in a similar manner to hard regularizers. Unlike those trained with hard regularizers, models trained with soft regularizers seem to avoid the decrease in performance for evidence detection.

Table 4: The effect of regularization on model inconsistency and performance of predicting task label y and evidence label z . * stands for baseline model that is trained to predict task label y . Standard deviation is shown after \pm . We notice dramatic decrease in model inconsistency and competitive performance for the regularized model.

ω_1, ω_2	*	0.0,0.0	3.0,3.0	10.0,6.0	30.0,10.0
$\mathcal{R}_1(\hat{\mathcal{D}})$	-	0.507 ± 0.084	0.277 ± 0.103	0.218 ± 0.205	0.112 ± 0.049
$\mathcal{R}_2(\hat{\mathcal{D}})$	-	0.048 ± 0.011	0.023 ± 0.019	0.030 ± 0.023	0.019 ± 0.007
acc (y)	0.524 ± 0.017	0.524 ± 0.005	0.523 ± 0.008	0.512 ± 0.014	0.499 ± 0.022
auc (y)	0.836 ± 0.008	0.852 ± 0.004	0.865 ± 0.004	0.862 ± 0.002	0.846 ± 0.006
acc (vascular congestion)	-	0.795 ± 0.018	0.793 ± 0.026	0.777 ± 0.005	0.763 ± 0.015
acc (hilar congestion)	-	0.771 ± 0.047	0.801 ± 0.031	0.759 ± 0.016	0.679 ± 0.040
acc (peribronchial cuffing)	-	0.804 ± 0.041	0.802 ± 0.032	0.795 ± 0.026	0.789 ± 0.040
acc (septal lines)	-	0.869 ± 0.010	0.859 ± 0.043	0.796 ± 0.037	0.845 ± 0.023
acc (interstitial abnormality)	-	0.649 ± 0.011	0.647 ± 0.002	0.630 ± 0.014	0.606 ± 0.006
acc (air bronchograms)	-	0.863 ± 0.019	0.881 ± 0.030	0.888 ± 0.015	0.860 ± 0.018
acc (parenchymal opacity)	-	0.706 ± 0.018	0.764 ± 0.020	0.756 ± 0.009	0.738 ± 0.009




Figure 6: The effect of regularization on model inconsistency (left 2) and performance of predicting task label y (middle 2) and evidence labels z (right 2), when model is trained using the hard regularizers $\mathcal{R}_1(\theta), \mathcal{R}_2(\theta)$ (top) versus the soft regularizers $\tilde{\mathcal{R}}_1(\theta), \tilde{\mathcal{R}}_2(\theta)$ (bottom). Here $\mathcal{R}_1, \mathcal{R}_2$ is short hand for $\mathcal{R}_1(\hat{\mathcal{D}}), \mathcal{R}_2(\hat{\mathcal{D}})$, respectively.