Causality Refined Diagnostic Prediction

Marcus Klasson*, Kun Zhang†, Bo C. Bertilson‡, Cheng Zhang*, Hedvig Kjellström*
*{mklas, chengz, hedvig}@kth.se, †kunz1@cmu.edu, ‡bo.bertilson@ki.se

Abstract

Applying machine learning in the health care domain has shown promising results in recent years. Interpretable outputs from learning algorithms are desirable for decision making by health care personnel. In this work, we explore the possibility of utilizing causal relationships to refine diagnostic prediction. We focus on the task of diagnostic prediction using discomfort drawings, and explore two ways to employ causal identification to improve the diagnostic results. Firstly, we use causal identification to infer the causal relationships among diagnostic labels which, by itself, provides interpretable results to aid the decision making and training of health-care personnel. Secondly, we suggest a post-processing approach where the inferred causal relationships are used to refine the prediction accuracy of a multi-view probabilistic model. Experimental results show firstly that causal identification is capable of detecting the causal relationships among diagnostic labels correctly, and secondly that there is potential for improving pain diagnostics prediction accuracy using the causal relationships.

Discomfort Drawings

- A patient shades the areas of the body where the patient experiences discomfort in a discomfort drawing [1]
- Figure 1 shows an example of a pain drawing and its assessed diagnostic labels

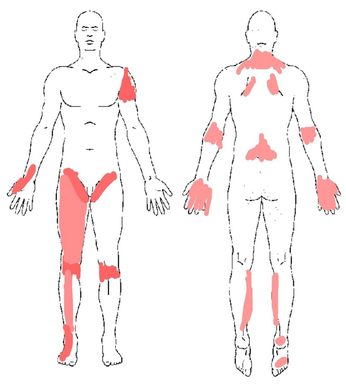

Symptom diagnoses: Neck discomfort, B Scapula discomfort, R Shoulder impingement, Interscapular discomfort, B Medial elbow discomfort, Lumbago, R Thumb discomfort, B Adductor tendonitis, R Front thigh discomfort, B PFS, B Calf discomfort, R Shin discomfort, R Calcaneal pain, R Arch discomfort

Pattern diagnoses: B C7 Radiculopathy, B L1 Radiculopathy, B L5 Radiculopathy, R S1 Radiculopathy, R C6 Radiculopathy

Psychophysiological diagnoses: DLI C6-C7, DLI L4-L5, DLI L5-S1, DLI S1-S2, DLI S2-S3

Figure 1: A discomfort drawing (left) and assessment from a medical expert (right). R stands for right side, L for left side and B for bilateral. PFS refers to patellofemoral pain syndrome and DLI to disco-ligament injury.

Causal Identification

The diagnostic labels are separated into three sets of diagnoses:

- *Causes* C (psychophysiological diagnoses)
- *Reasons* R (pattern diagnoses)
- *Symptoms* S (symptom diagnoses)

We assume that these sets are structured as $C \rightarrow R \rightarrow S$. Figure 2 shows a subgraph of a *directed acyclic graph* (DAG) learned with the PC (Peter-Clark) algorithm [2] from samples of the available diagnostic labels.

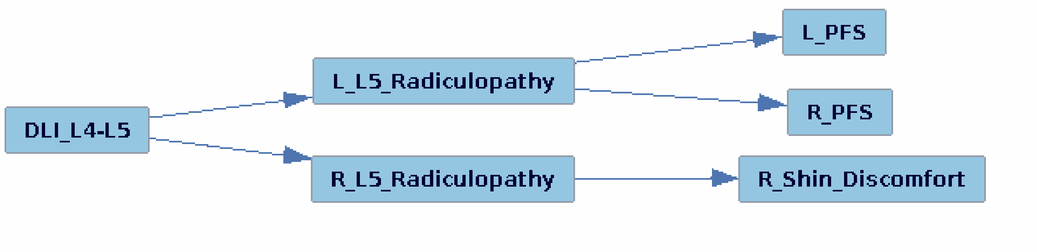

Figure 2: Subgraph of diagnostic labels displaying the causal structure between DLI L4-L5, L5 radiculopathy, and patellofemoral pain syndrome (PFS) and right shinbone discomfort.

References

[1] H. Palmer. Pain charts; a description of a technique whereby functional pain may be diagnosed from organic pain. *The New Zealand medical journal*, 48(264):187–213, 1949.

[2] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search, Second Edition (Adaptive Computation and Machine Learning)*. The MIT Press, 2001.

[3] C. Zhang, H. Kjellström, C. H. Ek, and B. C. Bertilson. Diagnostic prediction using discomfort drawings with IBTM. In *Machine Learning in Health Care Conference*, 2016.

Method

The proposed method is summarized in Figure 3.

- Use *Inter-Battery Topic Model* (IBTM) for diagnostic prediction given an unseen discomfort drawing (left panel) [3]
- Identify a causal graph from training data (middle panel)
- Refine the predicted labels in the causal graph and obtain the final structured output (right panel)

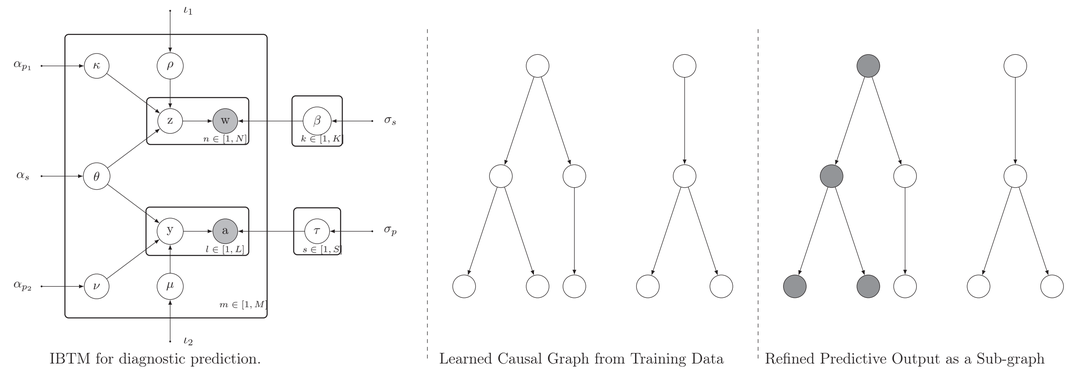

Figure 3: The left panel shows the IBTM, a multi-view latent space model, that predict diagnostic labels given input. The middle panel demonstrates a causal graph which is learned from training data. The IBTM predictions are then used as input for the causal graph to refine the result. The last panel demonstrates the final output, which is a subgraph instead of independent labels.

Algorithm for Updating Causal Graph

Algorithm 1: Update MPDs $p(x)$ in causal graph G.

for all nodes $x \in G$ **do**
  $p^{(0)}(x) \leftarrow p_{IBTM}(x)$
end for
for $t = 1, \ldots, \tau$ **do**
  for all nodes $x \in G$ **do**
    for all neighbours$(x) = x'$ **do**
      Compute $p^{(t)*}(x')$ using x and conditional probabilities between them.
      $p^{(t)}(x') \leftarrow \epsilon\, p^{(t)*}(x') + (1 - \epsilon)\, p^{(t-1)}(x')$
    end for
  end for
end for

Experimental Results

- Randomly split the data in two halves for training and testing
- The number of diagnostic labels to predict is determined by mean shift clustering for every test drawing
- Average the F1-measure over all test examples on each iteration for different update rates ϵ and state sizes K (Figure 4)
- Table 1 shows an example where the post-processing at iteration 5 with $\epsilon = 0.003$ and $K = 30$ enhances the IBTM results

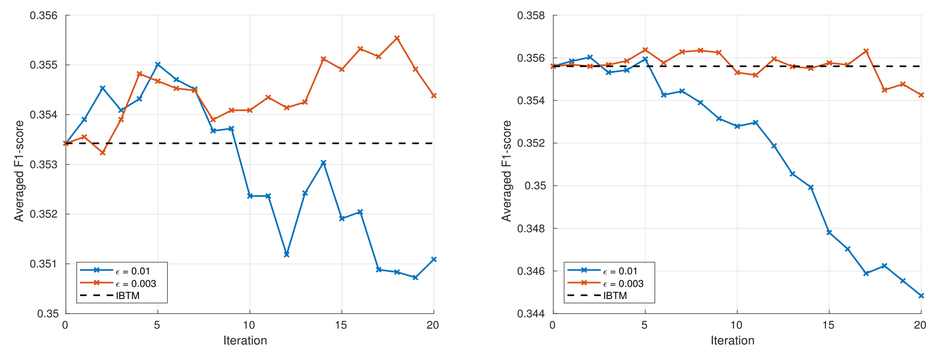

Figure 4: F1-measures averaged over five random data splits w.r.t iterations of updating the causal graph, where IBTM was trained with $K = 30$ (left) and $K = 50$ (right). Two different update rates ϵ were used, and the black dashed line is the averaged results from IBTM.

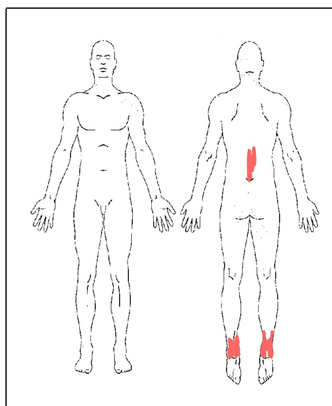

5 Prd IBTM: L Calcaneal pain; DLI L5-S1; L S1 Rdc; R S1 Rdc; L C2 Rdc;

5 Prd DAG: L Calcaneal pain; R Calcaneal pain*; DLI L5-S1; R S1 Rdc; L S1 Rdc;

6 GT: Lumbago; R Calcaneal pain*; L Calcaneal pain; DLI L5-S1; R S1 Rdc; L S1 Rdc;

F1 IBTM 72.73%

F1 DAG 90.91%

Diff. 18.18%

Table 1: Example of unseen discomfort drawing (left), predicted diagnostic labels (middle) and resulting F-measures (right). Predictions made by IBTM are followed after **Prd IBTM**, the refined predictions are after **Prd DAG** and ground truth labels after **GT**. Successfully predicted labels are marked in blue, otherwise in red.