

---

# Causality Refined Diagnostic Prediction

---

Marcus Klasson\*, Kun Zhang†, Bo C. Bertilson‡, Cheng Zhang\*, Hedvig Kjellström\*

\*KTH Royal Institute of Technology    †Carnegie Mellon University    ‡Karolinska Institute  
{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.

## 1 Introduction

Using pain or discomfort drawings for diagnostic purposes has shown to be effective both in medical science [5, 4, 3] and machine learning [15]. Pain drawings have a long history in the health care system and the first research paper on the use of it was published in 1949 [8]. It is used in many clinics and especially in the assessment of spine-related pain [7, 14, 2, 13]. Figure 1 shows an example of a discomfort drawing (a modern version of the original pain drawing) and its assessed diagnostic labels. The drawing is filled out by the patient who shadows those parts of the body where he/she experiences discomfort. Health-care personnel then assess the drawing to decide on diagnosis or further examinations. The drawing has also been used as an input for diagnostic prediction using machine learning methods which is the focus of this paper [15].

Identifying causal relationship [12] among diagnostic labels is extremely useful. The diagnostic labels in Figure 1 includes symptom diagnoses, pattern diagnoses and psychophysiological diagnoses. The labels describe the symptoms, while pattern diagnosis and psychophysiological diagnosis explain the underlying cause of the symptoms. However, these labels are presented in an independent way since it is just the snap-shot of a single patient. With access to a large number of cases, the causal relationship among these labels could be identified, providing a medical knowledge-graph of the relationship between symptoms and causes. Firstly, this causal graph by itself would be valuable for the decision-making and education of health care personnel – e.g., a clinician can get support in interpreting a pain drawing by using the causal graph as a template to look for likely causes of the observed symptoms. Secondly, the causal graph can be employed to make better automatic diagnostic prediction from pain drawings – e.g., by using the identified relationships to refine the prediction results from other machine learning models [15].

Our work explores utilizing causal identification algorithms to refine the diagnostic prediction and interpretation using discomfort drawings. Our contribution is two-fold:



**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.

- We employ a causal identification algorithm, the Peter-Clark (PC) algorithm [11, 10], to determine the causal relationship among diagnostic labels.
- We add a post-processing step to a diagnostic prediction algorithm, the Inter-Battery Topic Model (IBTM) [16], using the identified causal graph to enhance both the detection accuracy and the interpretability of the result.

We present our method in Section 2, experiments in Section 3, and discuss future work in Section 4.

## 2 Method

Our method of refining diagnostic prediction is presented in this section. The method is summarized in Figure 2. The baseline model is IBTM (left panel) [16] as used in [15]. IBTM is a latent variable model that learns a joint embedding  $\theta$  of the discomfort drawing  $w$  and the diagnostic label  $a$ . Thus, it can be used for diagnostic prediction given an unseen discomfort drawing. Using training data, we identify a causal graph demonstrated in the middle panel. We present the causal identification method in Section 2.1 with results from the data. At test time, we predict the diagnostic labels for an unseen pain drawing with IBTM. We then refine the predicted labels in the causal graph and obtain the final structured output as in the right panel in Figure 2. We present this procedure in Section 2.2.

### 2.1 Causal Identification

The causal graph should be structured as a *directed acyclic graph* (DAG), where the relationships between the three types of diagnostic labels can be identified. The properties of a DAG is that all edges must be directed from one node to another and that the graph does not allow any cycles, meaning that there can not exist a path from one node that leads to itself [6]. With some other constraints on the DAG, e.g., that interventions on a cause variable does not change the generating process from this cause to its effect, the DAG has a causal interpretation and represents causal relations between the variables [9, 11].

To learn a DAG from data, a constraint-based search algorithm such as the PC (Peter-Clark) algorithm [11, 10] can be used. This algorithm starts from an undirected graph and performs a sequence of independence and conditional independence tests, and based on this builds an equivalence class, which may contain more than one DAG that shares the same independence and conditional independence

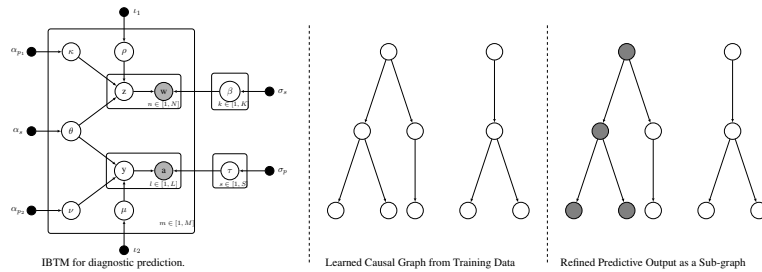


Figure 2: The method. The left panel shows the IBTM which is a multi-view latent space model that predict diagnostic labels given input. The middle panel demonstrates a causal graph which is learned from training data. The prediction is then used as the input for this causal graph to refine the result. The last panel demonstrates the final output, which is a subgraph instead of independent labels.




Figure 3: Subgraph of diagnostic labels from Tetrad displaying the causal structure between DLI L4-L5, L5 radiculopathy, and patellofemoral pain syndrome (PFS) and right shinbone discomfort. We do not show the full graph due to space limit.

relations in the data. As long as the conditional independence facts are known for the data, the PC algorithm can efficiently output such an equivalence class for both discrete and continuous data [1, 10].

We make the following additional constraints in the causal search (in addition to the independence and conditional independence constraints discovered from data): the diagnostic labels are separated into three categories, *causes*  $C$  (the psychophysiological diagnostic labels), *reasons*  $R$  (the pattern diagnostic labels) and *symptoms*  $S$  (the symptom diagnostic labels), where  $C$ ,  $R$  and  $S$  are sets of these diagnoses. We also assume that these categories are structured as  $C \rightarrow R \rightarrow S$ , so that  $C$  has a direct cause on  $R$  and that  $R$  directly causes  $S$ . Edges are forbidden within each individual set. We found that with those additional constraints, the output of PC was a DAG, with all edges oriented. Figure 3 shows a subgraph from a DAG built in Tetrad [1] with samples of the available diagnostic labels. The results tell that a disco-ligament injury at segment L4-L5 has a direct effect on the L5 nerve radiculopathy on both sides of the body. An injury on the left side of the nerve causes patellofemoral pain syndrome (PFS) on both sides of the body, while an injury on the right side causes discomforts on the right shinbone. Such a result is consistent with the distribution and connectivity of the nerves in the human body, indicating that the algorithm correctly identified the causal relationship among diagnostic labels in the dataset.

## 2.2 Refine Predictive Results using Causal Graph

The probability of each diagnostic label is predicted using IBTM. The conditional probability from cause to effect is obtained from the causal identification step. Given the learned DAG, we can then straightforwardly estimate the conditional distribution of each variable conditioning on its parents (or direct causes) from the given data.

Suppose we have obtained the prediction produced by IBTM. IBTM predicts each of the outputs separately, ignoring the dependence relationships among them. As a consequence, it does not make full use of the dependence information in the data and the result is to be improved. Given the good performance of IBTM in single-output prediction, however, to achieve better predictions, it is sensible to leverage 1) the predicted marginal distributions for all outputs by IBTM and 2) the relationships among the variables implied by the qualitative DAG structure and the quantitative conditional distributions. In other words, we can further update the prediction from IBTM,  $p_{IBTM}(x)$ , with the causal relationships given by graph  $G$ . The probability for node  $x$  is denoted by  $p(x)$  and these are initialized with the  $p_{IBTM}(x)$ . For each time step  $t$ , we iterate through all of the nodes and update its neighboring nodes  $x'$  with learning rate  $\epsilon$ . The procedure of post-refining the predictive result is summarized in Algorithm 1.

---

**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  $t = 1, \dots, \tau$  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
  
```

---




Figure 4: F-measures averaged over five random data splits with respect to 20 iterations of updating the causal graph, where IBTM was trained with  $K = 30$  (left) and  $K = 50$  (right). Two different update rates  $\epsilon$  are used, and the black dashed line is the averaged results from IBTM.

### 3 Experiments

We present our preliminary experimental results on refining diagnostic prediction. We use the same pain drawing dataset and parameter settings as in [15]. The dataset consists of 174 pain drawings and 280 diagnostic labels, where 150 of them are symptom diagnoses ( $S$ ), 61 are pattern diagnoses ( $R$ ) and 69 are psychophysiological diagnoses ( $C$ ). We split the dataset randomly in two halves for training and testing. As in [15], the number of diagnostic labels to predict is determined by mean shift clustering for every test drawing.

We run the updating procedure for 20 time steps and average the F-measure over all test examples on each iteration. We demonstrate the results with update rate  $\epsilon = 0.01$  and  $\epsilon = 0.003$  in Figure 4. IBTM was trained with the number of topics  $K = 30$  (left) and  $K = 50$  (right). The results show an increasing trend in the beginning iterations when  $\epsilon = 0.003$  and  $\epsilon = 0.01$  when  $K = 30$ .

Table 1 shows an example where the causal graph post-processing enhances the IBTM results. We see that R Calcaneal pain was not reported by IBTM because the probability was not high enough. On the other hand, with causal refining, the probability of Right Calcaneal pain increases, because it is highly probable when S1 Rdc is diagnosed. With our method, the result can be visualized to the clinician as a subgraph in the causal graph (omitted here due to space limitation).

### 4 Discussion

In this work, we employ a causal identification method, PC, to determine the causal relationship among various diagnostic labels. The identified causal graph is highly interpretable and is used to refine results from a diagnostic prediction model, IBTM.

Our future work extends in two directions. Firstly, we will explore ways to visualize such causal graphs to medical students as an education tool, and to visualize subgraphs to clinicians, to enhance interpretability of automatic diagnostic results. Secondly, while our current automatic prediction method employs causal reasoning in a post-processing manner where IBTM and a causal graph are learned separately, we will in future work focus on structured prediction where casual reasoning is integrated into the machine learning model in an end-to-end training manner.

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.

	<p><b>5 Prd IBTM:</b> L Calcaneal pain; DLI L5-S1; L S1 Rdc; R S1 Rdc; L C2 Rdc;</p> <p><b>5 Prd DAG:</b> L Calcaneal pain; R Calcaneal pain*; DLI L5-S1; R S1 Rdc; L S1 Rdc;</p> <p><b>6 GT:</b> Lumbago; R Calcaneal pain*; L Calcaneal pain; DLI L5-S1; R S1 Rdc; L S1 Rdc;</p>	<p><b>F1 IBTM</b> 72.73%</p> <p><b>F1 DAG</b> 90.91%</p> <p><b>Diff.</b> 18.18%</p>
--	--	---

## References

- [1] Tetrad manual. <http://www.phil.cmu.edu/tetrad/>, June 2017. Online; accessed November 30 2017.
- [2] M. J. Albeck. A critical assessment of clinical diagnosis of disc herniation in patients with monoradicular sciatica. *Acta neurochirurgica*, 138(1):40–44, 1996.
- [3] B. C. Bertilson, E. Brosjö, H. Billing, and L. E. Strender. Assessment of nerve involvement in the lumbar spine: agreement between magnetic resonance imaging, physical examination and pain drawing findings. *BMC musculoskeletal disorders*, 11(1):1, 2010.
- [4] B. C. Bertilson, M. Grunnesjö, S.E. Johansson, and L.E. Strende. Pain drawing in the assessment of neurogenic pain and dysfunction in the neck/shoulder region: Inter-examiner reliability and concordance with clinical examination. *Pain medicine*, 8(2):134–146, 2007.
- [5] B. C. Bertilson, M. Grunnesjö, and L. E. Strender. Reliability of clinical tests in the assessment of patients with neck/shoulder problems? impact of history. *Spine*, 28(19):2222–2231, 2003.
- [6] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.
- [7] D. Ohnmeiss, H. Vanharanta, and J. Ekholm. Relation between pain location and disc pathology: a study of pain drawings and ct/discography. *The Clinical journal of pain*, 15(3):210–217, 1999.
- [8] 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.
- [9] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.
- [10] P. Spirtes. Introduction to causal inference. *Journal of Machine Learning Research 11*, pages 1643–1662, 2010.
- [11] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search, Second Edition (Adaptive Computation and Machine Learning)*. The MIT Press, 2001.
- [12] P. Spirtes and K. Zhang. Causal discovery and inference: concepts and recent methodological advances. In *Applied informatics*, volume 3, page 3. Springer, 2016.
- [13] Y. Tanaka, S. Kokubun, T. Sato, and H. Ozawa. Cervical roots as origin of pain in the neck or scapular regions. *Spine*, 31(17):E568–E573, 2006.
- [14] N. Vucetic, H Määttänen, and O Svensson. Pain and pathology in lumbar disc hernia. *Clinical orthopaedics and related research*, 320:65–72, 1995.
- [15] 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.
- [16] C. Zhang, H. Kjellström, and C.H. Ek. Inter-battery topic representation learning. In *European Conference on Computer Vision*, 2016.