Hypergraph Models for Predicting *Fusarium* Head Blight Risk

Sajjad Heydari[1] Sean Walkowiak[2] Christopher Henry[1]

[1]University of Manitoba, Department of Computer Science [2]Canadian Grain Commission

Background and Motivation

Fusarium Head Blight (FHB) is a major disease of wheat in the Canadian Prairies, causing yield loss, grain downgrading, and mycotoxin contamination. Its development depends on complex interactions among weather, host genetics, and pathogen biology. Warm and humid conditions during anthesis strongly favour infection, and multiple *Fusarium* species with different pathogenic profiles contribute to disease risk.

Producers rely on timely and accurate FHB risk assessments to support decisions such as wheat variety selection and fungicide application. Reliable prediction tools can improve both yield outcomes and grain quality, yet current operational models are limited in how they capture the factors that influence disease pressure.

Problem Statement

The objective is to predict the percentage of FHB-damaged kernels for each wheat variety at each collection centre and year. The target used in this work is the proportion of infected grain by weight:

$$\text{FHB\%} = \frac{\text{weight of infected kernels}}{\text{total sample weight}} \times 100.$$

The data required for this prediction has a naturally irregular and heterogeneous structure. Each collection centre is linked to a variable number of weather stations that change from year to year, and stations often have missing or partial monthly records. Historical FHB outcomes differ across varieties and regions, and varieties enter and leave production over time, creating inconsistent history lengths. Spatial proximity also influences disease outcome, but this relationship is not captured by fixed-size feature vectors.

These characteristics make the problem difficult to model with traditional approaches that expect uniform inputs. A representation is needed that can incorporate varying numbers of sensors, handle missing or uneven historical information, and reflect the spatial relationships among regions and stations.

Data and Sensory Network

This study uses data from the Canadian Grain Commission's Harvest Sample Program, covering the years 1995 to 2024. Each observation represents a specific combination of collection centre, wheat variety, and year, along with the measured proportion of infected grain by weight.

Monthly weather information is incorporated from nearby weather stations, using variables such as temperature, precipitation, wind, and snow cover. For each collection centre and year, we consider a set of surrounding stations, and the composition of this set changes over time as stations are added, removed, or report incomplete records. Historical FHB information is drawn from previous seasons at the same collection centre and its neighbours, and includes outcomes for multiple varieties that have been grown in that area over time.

Together, these components naturally form a sensory network. Weather stations provide environmental signals, historical records across varieties and years provide temporal signals, and collection centres (along with their neighbours) anchor these signals geographically. This network structure captures how environmental conditions, local disease history, and spatial context jointly contribute to observed FHB levels.

Collection Centres and Weather Stations

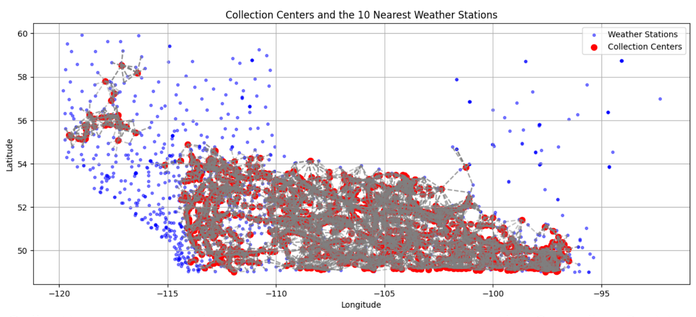

Collection centres and nearby weather stations across the Canadian Prairies.

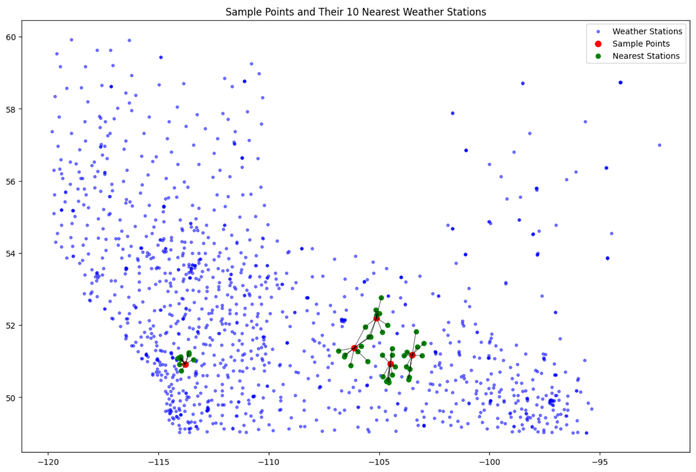

Sample subset of collection centres and their ten closest weather stations for visual clarity.

Why Hypergraphs?

Each FHB observation brings together many pieces of information: a collection centre, several nearby weather stations, multiple years of disease outcomes from that area, and the varieties grown over time. These pieces do not form simple one-to-one relationships. Instead, each setting links a whole group of related factors at once.

A hypergraph is a tool for representing these kinds of groups. Rather than connecting two points at a time, a hypergraph can link an entire set of items together. This makes it possible to treat an observation as a single unit that includes the collection centre, its surrounding weather stations, the recent FHB history in that region, and the variety being tested.

This approach keeps the full structure of the data. It naturally handles situations where some centres have more weather stations than others, where history lengths differ among varieties, and where neighbouring regions influence disease pressure. Hypergraphs allow the model to use all of these connections directly, reflecting how FHB risk arises from combined environmental, historical, and regional factors.

In practice, the connections in this work form a weighted hypergraph, where the strength of each link's member reflects factors such as distance between weather stations and collection centres.

Example of a Hypergraph

A hypergraph can link any number of items at once. The figure below shows hyperedges of different sizes: pairs of nodes connected by lines, groups of three connected by triangles, and groups of four connected by cloud-shaped links. This illustrates how hypergraphs naturally represent relationships that involve multiple elements at the same time.

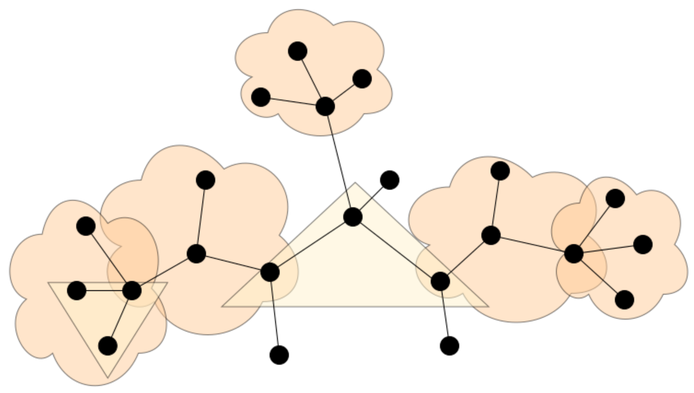

A simple hypergraph with hyperedges of size 2, 3, and 4.

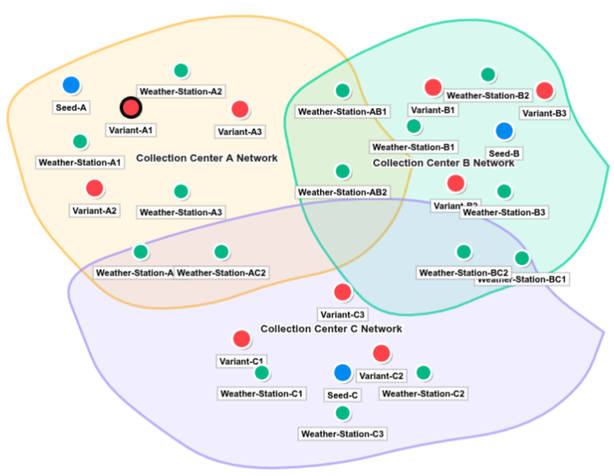

Hypergraph view of simplified data: each region connects to its weather stations, varieties, and neighbouring centres. Overlapping areas represent shared influence across locations and seasons.

Hypergraph Message Passing Neural Network

Once the data is organized as a hypergraph, the model applies a message passing neural network to combine information across the connected elements. Each hyperedge represents one observation, linking the collection centre, the weather stations used for that season, and the relevant historical FHB records. During message passing, each node in the hyperedge sends information into the hyperedge, which then returns an updated signal back to its nodes. Repeating these exchanges allows the model to integrate signals from local weather, multi-year disease history, and neighbouring regions.

Through this process, the model learns how strongly each component should influence FHB risk. Relevant or nearby weather stations naturally contribute more to the prediction, and historical outcomes shape how varieties have performed in the region over time. Because the model operates directly on the hypergraph, it can work with different numbers of stations, varying history lengths, and spatial relationships without forcing the data into a fixed format.

The model outputs a continuous prediction of the percentage of FHB-damaged kernels for that collection centre, variety, and year.

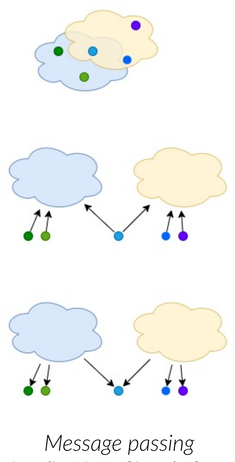

Message passing visualization. Signals from weather stations, historical records, and neighbouring regions are combined through the hyperedge to produce the final FHB risk estimate.

Predictive Performance

Model performance was evaluated at the collection-centre level using October–June weather data, spatial context, and a five-year history of FHB outcomes.

Traditional models. Support Vector Machines reached an RMSE of **18.7%**, with feed-forward neural networks performing similarly.

Transformer models. Performance depended heavily on how missing station data was handled: zero-filling produced **32.1%** RMSE, while removing partial station records improved accuracy to **16.4%**.

HMPNN. HMPNN uses the hypergraph structure to incorporate spatial relationships, multi-year history, and the stations active in each season. It handles missing or variable station sets naturally and achieved an RMSE of **5.34%**, substantially outperforming both traditional and transformer baselines.

Key Findings and Insights

Structure matters. FHB risk depends on spatial patterns, multi-year disease history, and variety behaviour. These relationships are lost when data is flattened into simple weather-only predictors.

Hypergraphs capture real agronomic relationships. By grouping each observation's weather stations, regional history, and variety information, HMPNN reflects how FHB pressure develops in practice rather than treating each record as an independent row of numbers.

History length strongly influences accuracy. Prediction quality improves markedly as more years of FHB history are included. Accuracy rises quickly when moving from one to about five years of history, after which gains become smaller. This suggests that a five-year window offers a practical balance between predictive performance and the cost and availability of historical data.

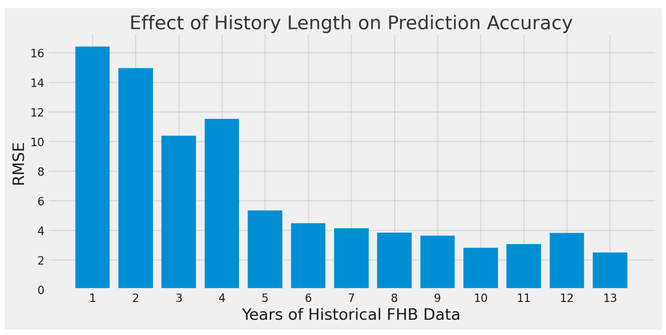

Effect of history length on prediction accuracy (HMPNN). Large gains occur as history increases up to about five years, with diminishing returns thereafter.

HMPNN handles missing and variable station data. Unlike transformer models, which degrade sharply when station records are incomplete, HMPNN operates directly on the hypergraph and can use whatever stations are active in a given season without discarding partial records.

Spatial context improves risk estimation. Neighbouring collection centres share weather patterns and disease pressure. Using spatial relationships strengthens prediction and captures regional FHB dynamics that single-location models miss.

Structure-aware learning outperforms baselines. Under the same October–June and five-year history setup, HMPNN achieved an RMSE of **5.34%**, compared with **18–19%** for baseline models and **16–32%** for transformer models, demonstrating a substantial improvement in FHB risk prediction.

Limitations

Climate and production changes. The model assumes that past weather and disease patterns are representative of future conditions. Shifts in temperature, precipitation, or drought frequency may reduce the reliability of predictions under changing climate trends.

Variety turnover. Wheat varieties enter and leave production over time. Newly released varieties may lack sufficient history for strong predictive performance, while older varieties provide historical signals that may not reflect current genetics or management practices.

Collection centre mismatch. Producers are not required to submit samples to their nearest collection centre. As a result, the weather associated with the centre may not fully match conditions in the field where the crop was grown.

Unobserved agronomic practices. Crop rotation, fungicide timing, tillage, and other management decisions influence FHB risk but are not captured in the dataset. These missing factors introduce variability that the model cannot account for directly.

Weather station coverage. Some regions have many nearby stations while others have few. Although HMPNN can handle missing or partial station data, sparse coverage still limits the precision of environmental signals in certain areas.

Acknowledgements

We thank the Canadian Grain Commission for providing access to the Harvest Sample Program data used in this study. We also acknowledge the guidance and support of Dr. Sean Walkowiak (CGC) and Dr. Christopher Henry (University of Manitoba). This work was further supported by the TerraByte research group, whose expertise and feedback contributed to the development of the modelling framework.

Contact: Sajjad Heydari — University of Manitoba, Department of Computer Science heydaris@myumanitoba.ca