



Project acronym	B4B
Project full name	Brains for Building's Energy Systems
Grant No	M00I32004
Project duration	May 1, 2021- December 31, 2025

Deliverable D1.04a

Impact of expert knowlegde on prior and conditional probabilities in DBN-based FDD: a fuzzy and LLM-based approach for robust diagnostic (FINAL)

Lead beneficiary	TU Delft
Lead authors	Ziao Wang, Chujie Lu, Arjen Meijer, Laure Itard
Publication date	31 December 2025
Deliverable Status	FINAL
File name	B4B WP1 D1.04a Expansion of the method impact of prior and conditionals probablities in DBN CONF.docx

Confidential until publication. For more information, please contact l.c.m.itard@tudelft.nl

The project was carried out with a Top Sector Energy subsidy from the Ministry of Economic Affairs and Climate, carried out by the Netherlands Enterprise Agency. The specific subsidy for this project concerns M00I subsidy round 2020

SUMMARY

This research focuses on knowledge-driven Diagnostic Bayesian Networks (DBN) for fault diagnosis in HVAC systems. In these DBN's network structure is derived from expert knowledge, physical principles, and literature. Current methodologies exhibit significant limitations regarding both model structure construction and parameter quantification. Diagnostic results are typically derived from a single, static DBN model, where fault-symptom connection may be biased toward a specific system configuration and specific expert knowledge, lacking the possibility for generalization and being prone to faults. Filling the Conditional Probability Tables (CPTs) presents a significant challenge. While some studies attempt to learn these parameters from existing historical data, this approach is frequently constrained by the scarcity of high-quality labelled fault datasets. As a result, researchers often revert to manual estimation, requiring experts to assign numerical probabilities based on experience. This reliance on direct numerical estimation from experts is time-consuming and prone to inconsistency, and is hindering the scalability of the approach.

To better capture expert knowledge, construct generalizable Bayesian Networks and assign realistic conditional probability while overcoming data scarcity, Fuzzy Set Theory (FST), can be used to model the ambiguity, vagueness, and subjectivity often present in human judgment. Despite its extensive application in other industries, the adaptation of FST remains largely underexplored in the domain of HVAC FDD. This report investigates the application of FST in Air Handling Unit (AHU) diagnosis.

Important steps in FST are expert elicitation, where domain experts estimate event probabilities using linguistic variables; fuzzification to convert them into a set or range of values; aggregation/fusion methods to combine the opinions of multiple experts; and finally, defuzzification and pruning to enable probability conversion and further calculation.

This study therefore addresses how to integrate divergent expert judgments into a DBN without diluting the diagnostically important fault-symptom structure. A fuzzy DBN-based workflow for HVAC fault diagnosis was developed, that integrates heterogeneous expert knowledge into a structured elicitation, fusion, and pruning process. Using 672 DBN configurations and 11 real AHU fault cases, the study examined how elicitation design choices, expert disagreement, and LLM-assisted knowledge input influence diagnostic performance and posterior behaviour.

The results show that the core challenge of multi-expert elicitation is not the presence of disagreement itself, but whether the fusion process can distinguish diagnostically useful structure from dispersed, incomplete, or misleading expert inputs. This study shows that expert knowledge should not be treated as uniform or directly interchangeable, because fusion outcomes depend on how fault-symptom structure is distributed across experts. When important diagnostic structure is shared, fusion can consolidate that common knowledge and produce a more stable posterior around the true fault. When it is fragmented across only a few experts, simple fusion risks diluting precisely the knowledge that matter most for diagnosis. The significance of this work therefore lies in showing how heterogeneous expert knowledge can be fused without simply averaging it away: robust fusion requires preserving critical minority knowledge, removing noisy knowledge, and retaining the fault-symptom relationships that carry real diagnostic value. In this study, AG-based weighting (Agreement Weight) with moderate pruning provided the most robust overall solution, while OD-based fusion (Objective-Distance Weight) remained especially valuable in low-consensus cases where preserving decisive but weakly shared links was essential. Overall, this study contributes a broader perspective on expert knowledge integration for engineering diagnosis. It shows that reliable multi-expert modeling is not achieved by maximizing consensus alone, but by designing elicitation and fusion procedures that can recover, preserve, and reorganize useful knowledge under disagreement.