Diagnosing Multiple Interacting Defects with Combination Descriptions

Nancy E. Reed
Computer Science Department, University of Minnesota
4-192 EE/CS Bldg., 200 Union St. S.E., Minneapolis, MN 55455
reed@cs.umn.edu

Cases with multiple defects can be difficult to diagnose because the defects can interact, meaning that the observable cues are not a sum of the cues for the component defects. Diagnostic methods that use cue-to-defect relationships fail when interactions between defects change the observable cues. The primary alternative, model-based methods, are limited to domains with accurate and complete models, along with initialization data. Using these traditional methods, when defects interact and models aren't available, each possible defect combination must be included in the knowledge base. This results in an explosion of possible alternatives, greatly increased knowledge acquisition effort, slower processing, and increased maintenance effort.

This research develops a computational diagnostic model that can diagnose multiple defects, even when cues are altered or missing, by using descriptions of cue combinations. We develop a description and classification of the ways cues change when defects interact. Each type of cue may combine in a different way, so each type has a separate description. The combination methods use the expectations of component defects to diagnose multiple interacting defects, instead of requiring a description of each possible defect combination.

In a medical domain (diagnosis of congenital heart defects), we found that cues combine with one another in a small number of ways: all cues may appear *(union)*, the values of the cues may be *added* (quantitatively or qualitatively), or *dominant* cues may mask other cues present. Cues of each type combine in one of these basic ways, or use a combination of a few of the basic ways, depending on characteristics of the cues, case, or domain.

Figure 1 shows two types of cues, murmurs and heart sounds, for an example case. The case has two defects, atrial septal defect (ASD) and aortic stenosis (AS).

The cues are different than expected because the defects interact. Two murmurs are expected, a loud systolic ejection murmur in the pulmonary area for ASD, and a moderate systolic ejection murmur in the aortic area for AS. Only one murmur is observed. Loud murmurs mask softer ones occurring at the same time, so the absence of the (softer) expected murmur for AS is explained. The observed murmur supports ASD alone or ASD+AS. The observed heart sound S2 is normal, while both defect expectations are abnormal. ASD produces a wide, fixed split S2 while AS produces a narrow, variably split S2. The wide and narrow widths combine additively to explain a normal width, while the variable and fixed expectations combine additively to a variable split, which is a normal S2. Neither ASD nor AS is supported alone, but ASD+AS is supported by the normal S2 cue.

If we had used cue-to-defect relationships and matching on this case, we would have explained only the observed murmur (with one of the component defects), leaving three missing abnormal expectations unexplained.

The diagnostic model is tested by constructing a program with a knowledge base in pediatric cardiology (Fallot) and testing it on cases of single and multiple defects from hospital files. Fallot uses a combination of recognition-based reasoning (Thompson et al. 1983) and the cue combination descriptions. This program correctly diagnoses cases with multiple interacting defects for which conventional methods fail.

References

W. B. Thompson, P. E. Johnson, and J. B. Moen. 1983. Recognition-based diagnostic reasoning. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 236–238.
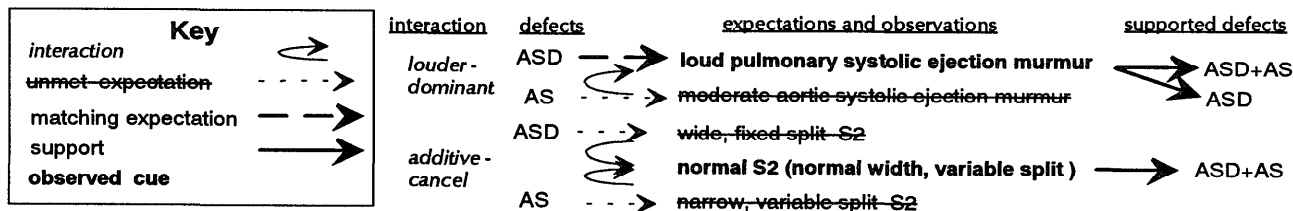
Figure 1: Explanation for murmur and heart sound cues for a case