Fatigue is a major contributor to workplace accidents and productivity loss. Current methods for managing fatigue in the workplace typically incorporate fatigue modelling, but such models are often inaccurate for individual predictions and input demands are too burdensome. Consumer sleep technology may alleviate these burdens, but device performance within such models must be properly evaluated. If accurate models of fatigue can be utilized using sleep tracking devices, they may be more routinely implemented into high-risk workplaces. This can ultimately improve the management of fatigue in shift work and will result in improved workplace safety, productivity, and worker health and wellbeing.

Chapter 2 outlines the development and initial validation of a custom machine-learning model of fatigue. This model was capable of estimating vigilance in a simulated shift work context based on data solely from the preceding sleep as estimated by an under-mattress sensor. Model performance was comparable to existing models of fatigue that typically require manual or long-term input. The results highlighted the potential for predicting fatigue based on device-derived sleep, alone.

Chapter 3 comparatively evaluated multiple models of fatigue. Existing biomathematical models and a machine learning model were tested in a simulated shift work context and in a regular sleep context. The two existing models only weakly explained fatigue in the shift work context. All three models accurately predicted a lack of significant fatigue in the regular sleep context but showed poor accuracy for predicting continuous fatigue. Thus, current models of fatigue show poor accuracy with device-estimated sleep, and new signals or modelling methods may be necessary.

Chapter 4 evaluated the performance of an under-mattress device, the Withings Sleep Analyzer (WSA), for estimating sleep and wake compared to polysomnography. This study showed that the WSA had comparable accuracy to existing sleep tracking devices during regular nighttime sleep opportunities but was less accurate during daytime sleep opportunities. Device performance also had low variability across good sleeps but was more variable in daytime sleep opportunities and in individuals with a sleep disorder. These considerations are warranted when using the WSA or other such devices in different populations.

Chapter 5 examined and characterized estimates of sleep from the WSA in relation to several cognitive tasks spanning several cognitive domains. Results showed that few devices-based sleep metrics were directly related to cognitive outcomes. However, non-parametric methods showed reasonable fitness of models for vigilance and working memory tasks, indicating that these domains should be of primary interest in future models of fatigue.

Together, this thesis indicates great promise for implementing and extending models of fatigue using sleep tracking devices. This work has demonstrated the development and testing of a fatigue model in a specific context and suggests that existing models are likely insufficiently accurate for automated, long-term fatigue management. Development of new fatigue models must consider the strengths and limitations of the sleep tracking devices being used. This thesis is part of a long-term goal to simplify and automate fatigue management in high-risk workplaces and constitutes a significant step towards this end.