Abstract

This thesis investigates Hidden Markov Models (HMMs) as a mathematical framework for modeling systems with hidden internal states and observable outputs. These models extend classical Markov chains by introducing a latent stochastic process that governs the system's behavior, while observable data are treated as emissions conditioned on these hidden states. Such a structure makes HMMs especially powerful in contexts where the true underlying dynamics of a system cannot be observed directly.

The work begins by establishing the theoretical foundation of finite-state Markov chains, including key concepts such as stationarity, recurrence, ergodicity, and long-run convergence. Building on this, the thesis introduces Hidden Markov Models in formal terms and discusses two fundamental algorithms: the Viterbi algorithm, which infers the most probable sequence of hidden states, and the Baum-Welch algorithm, which estimates model parameters from incomplete data using expectation-maximization.

A central application explored in this thesis is the modeling of credit rating dynamics—a problem of clear relevance in financial mathematics. Credit ratings, often issued in discrete categories by agencies such as Moody's, are observable but imperfect indicators of a firm's actual financial health. Since the true credit quality is influenced by numerous hidden factors and may evolve over time, it can be naturally modeled as a hidden process. HMMs provide a principled and flexible approach for capturing such latent transitions and interpreting observed rating behavior within a probabilistic framework.

Overall, the thesis contributes to both the theoretical understanding and practical utility of Hidden Markov Models. It highlights their value as a bridge between abstract probabilistic theory and applied modeling challenges in finance and beyond, particularly in contexts where uncertainty, noise, and hidden variables play a central role.