Quantile Regression & Pinball Loss

Usually, I think in terms of mean and standard deviation, but that assumes that the uncertainty is symmetric around the mean and that might not be the case, for example in case of noise with skew-normal distribution in the data. I had never encountered the term "Pinball loss". So, I asked ChatGPT to explain it to me: "Pinball loss" helps as follows: Unlike typical regression which predicts the mean of the target variable, quantile regression predicts a specific quantile and Pinball loss can penalize over- and under-predictions asymmetrically, based on the quantile. If you're predicting the median (q(uantile)=0.5), the pinball loss penalizes over- and under-estimates equally, but if you're predicting a higher quantile (q=0.9), under-estimates are penalized more heavily than over-estimates and for q=0.4 it is the other way around. Interesting