

Chapter 2

Language models

2.1 Simple language models

2.1.1 Introduction

2.1.2 n-gram models

2.1.3 Training and evaluation

2.1.4 Smoothing and backoff

2.2 Large language models

2.2.1 Auto-regressive generation

2.2.2 Jurafsky & Martin's Introduction to LLMs

2.2.2.1 Introduction to LLMs

2.2.2.2 Conditional generation of texts: intuition

2.2.2.3 Prompting

2.2.2.4 Sampling for LLM generation

2.2.2.5 Pretraining LLMs

2.2.2.6 Finetuning

2.2.2.7 Evaluating LLMs

2.2.2.8 Ethical and safety issues

2.2.3 Final thoughts (from Yann LeCun)