
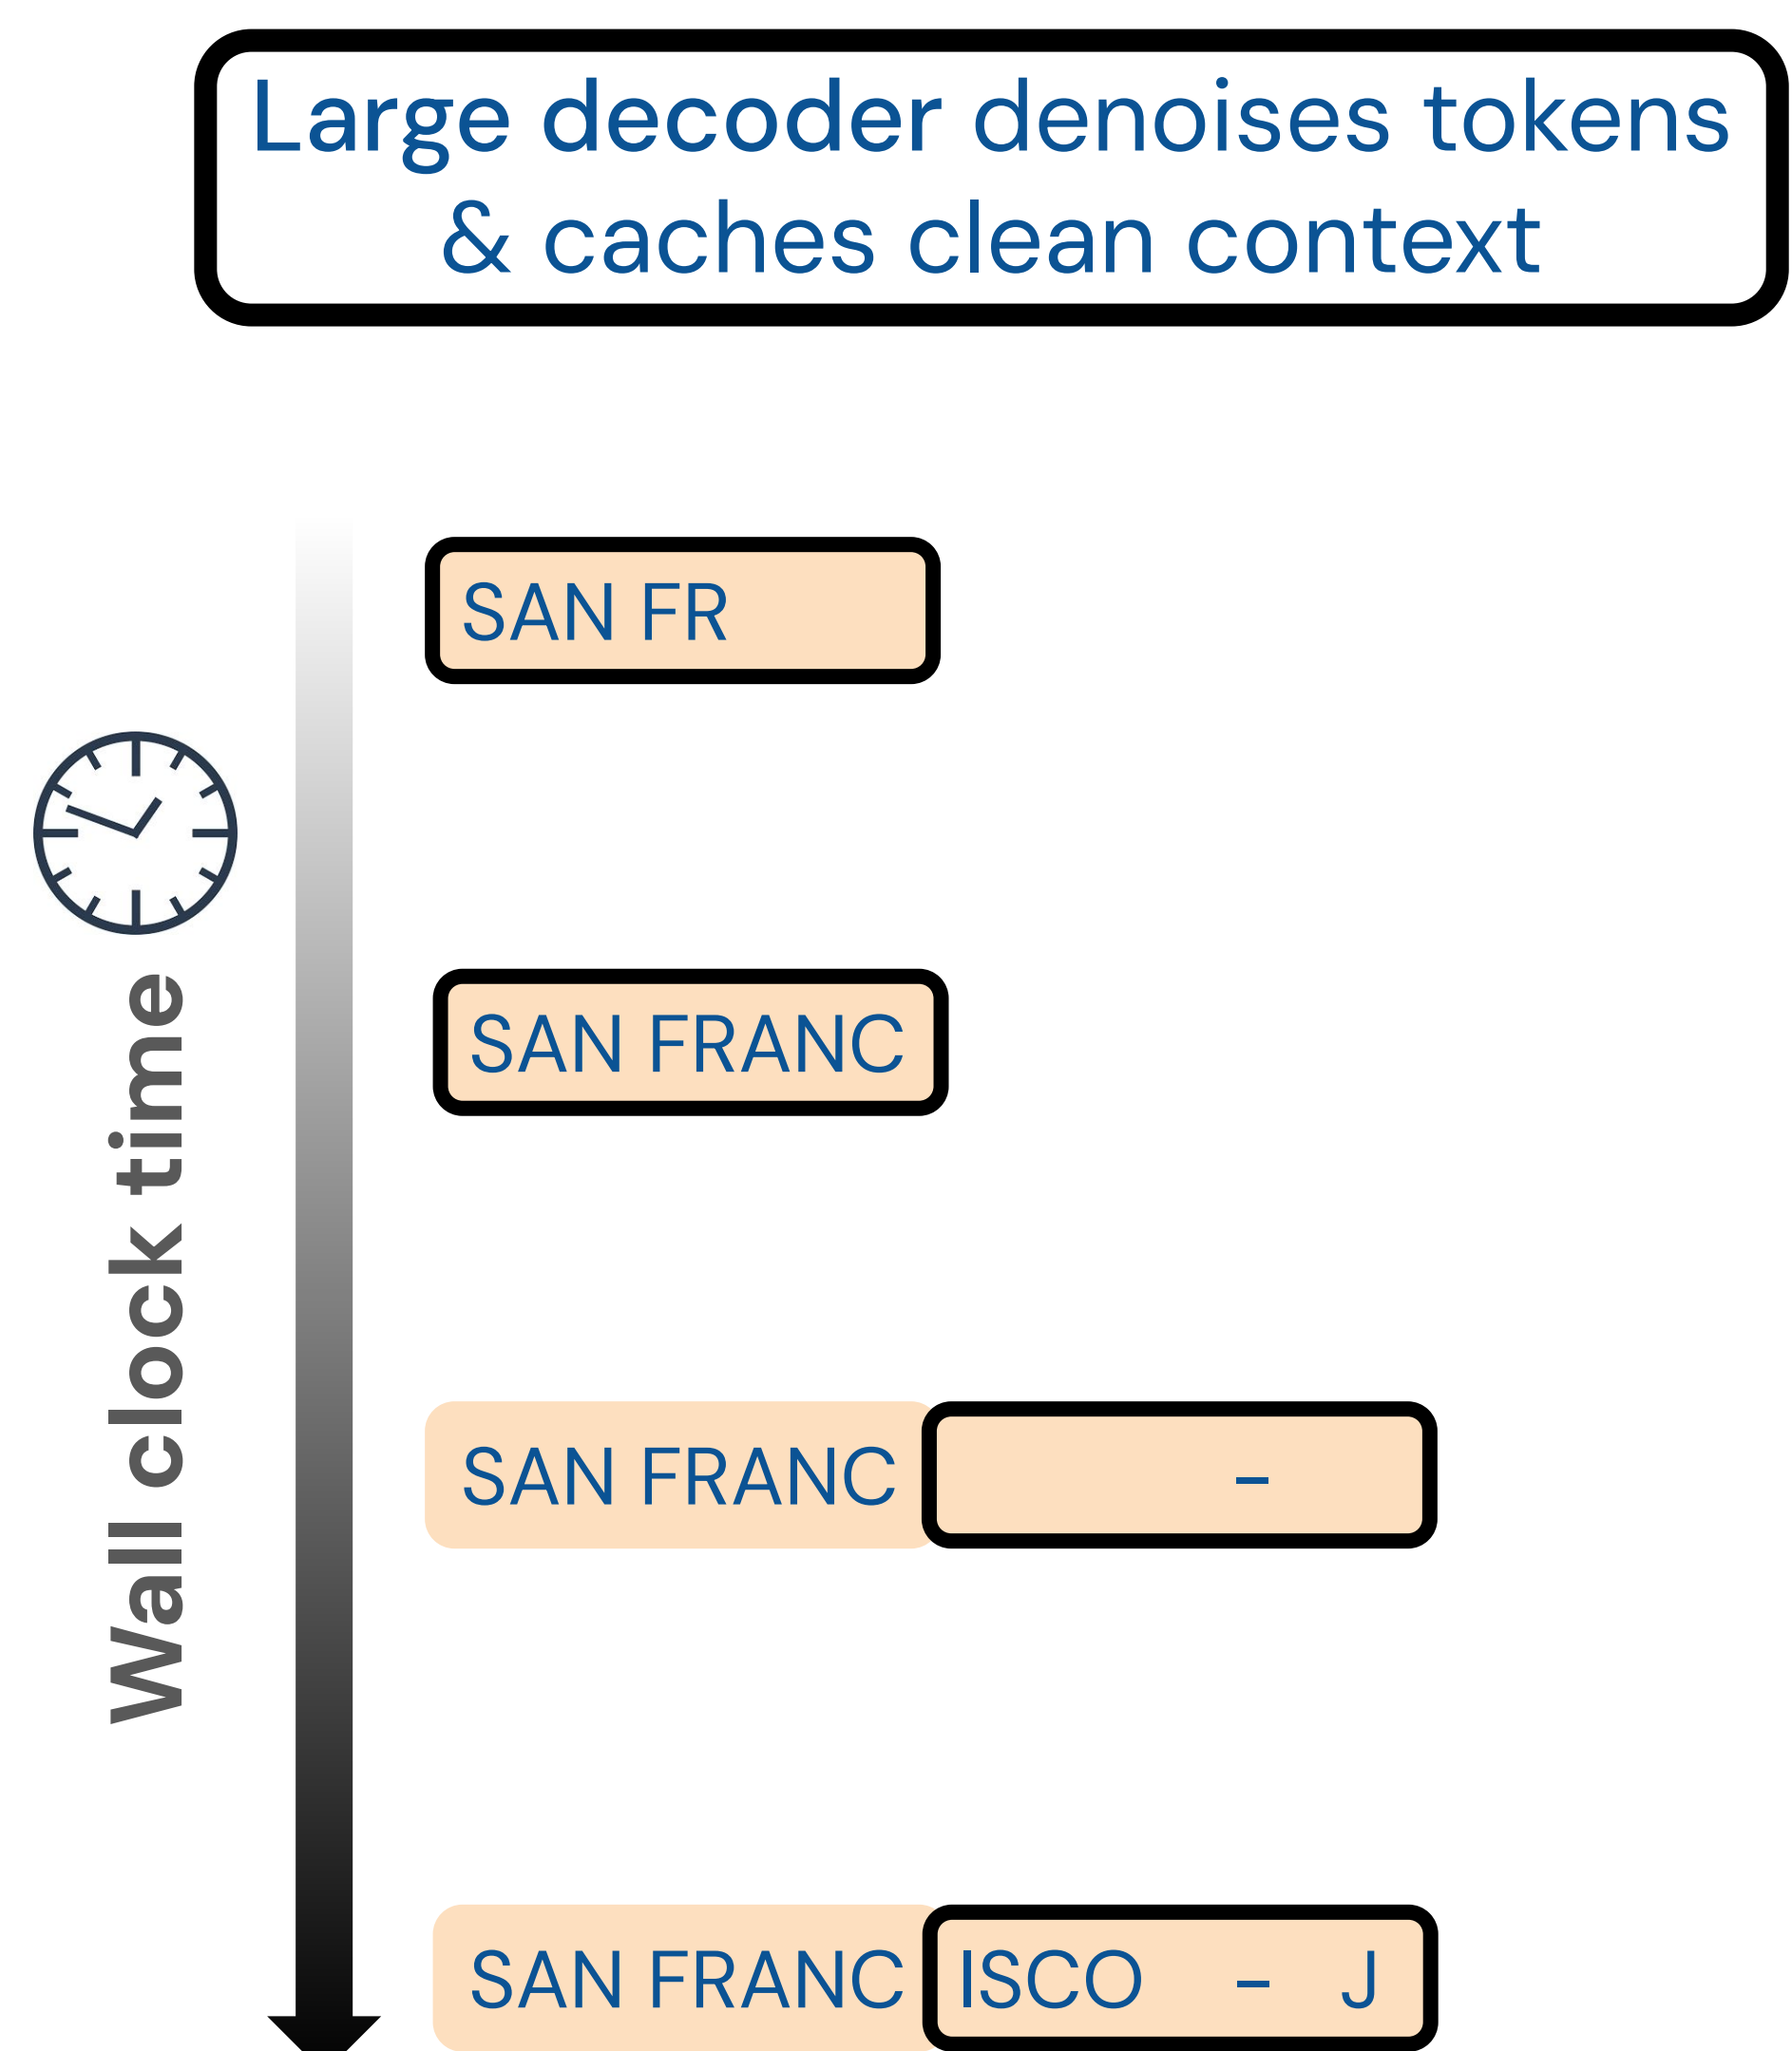


# Encoder-Decoder Diffusion Language Models for Efficient Training and Inference

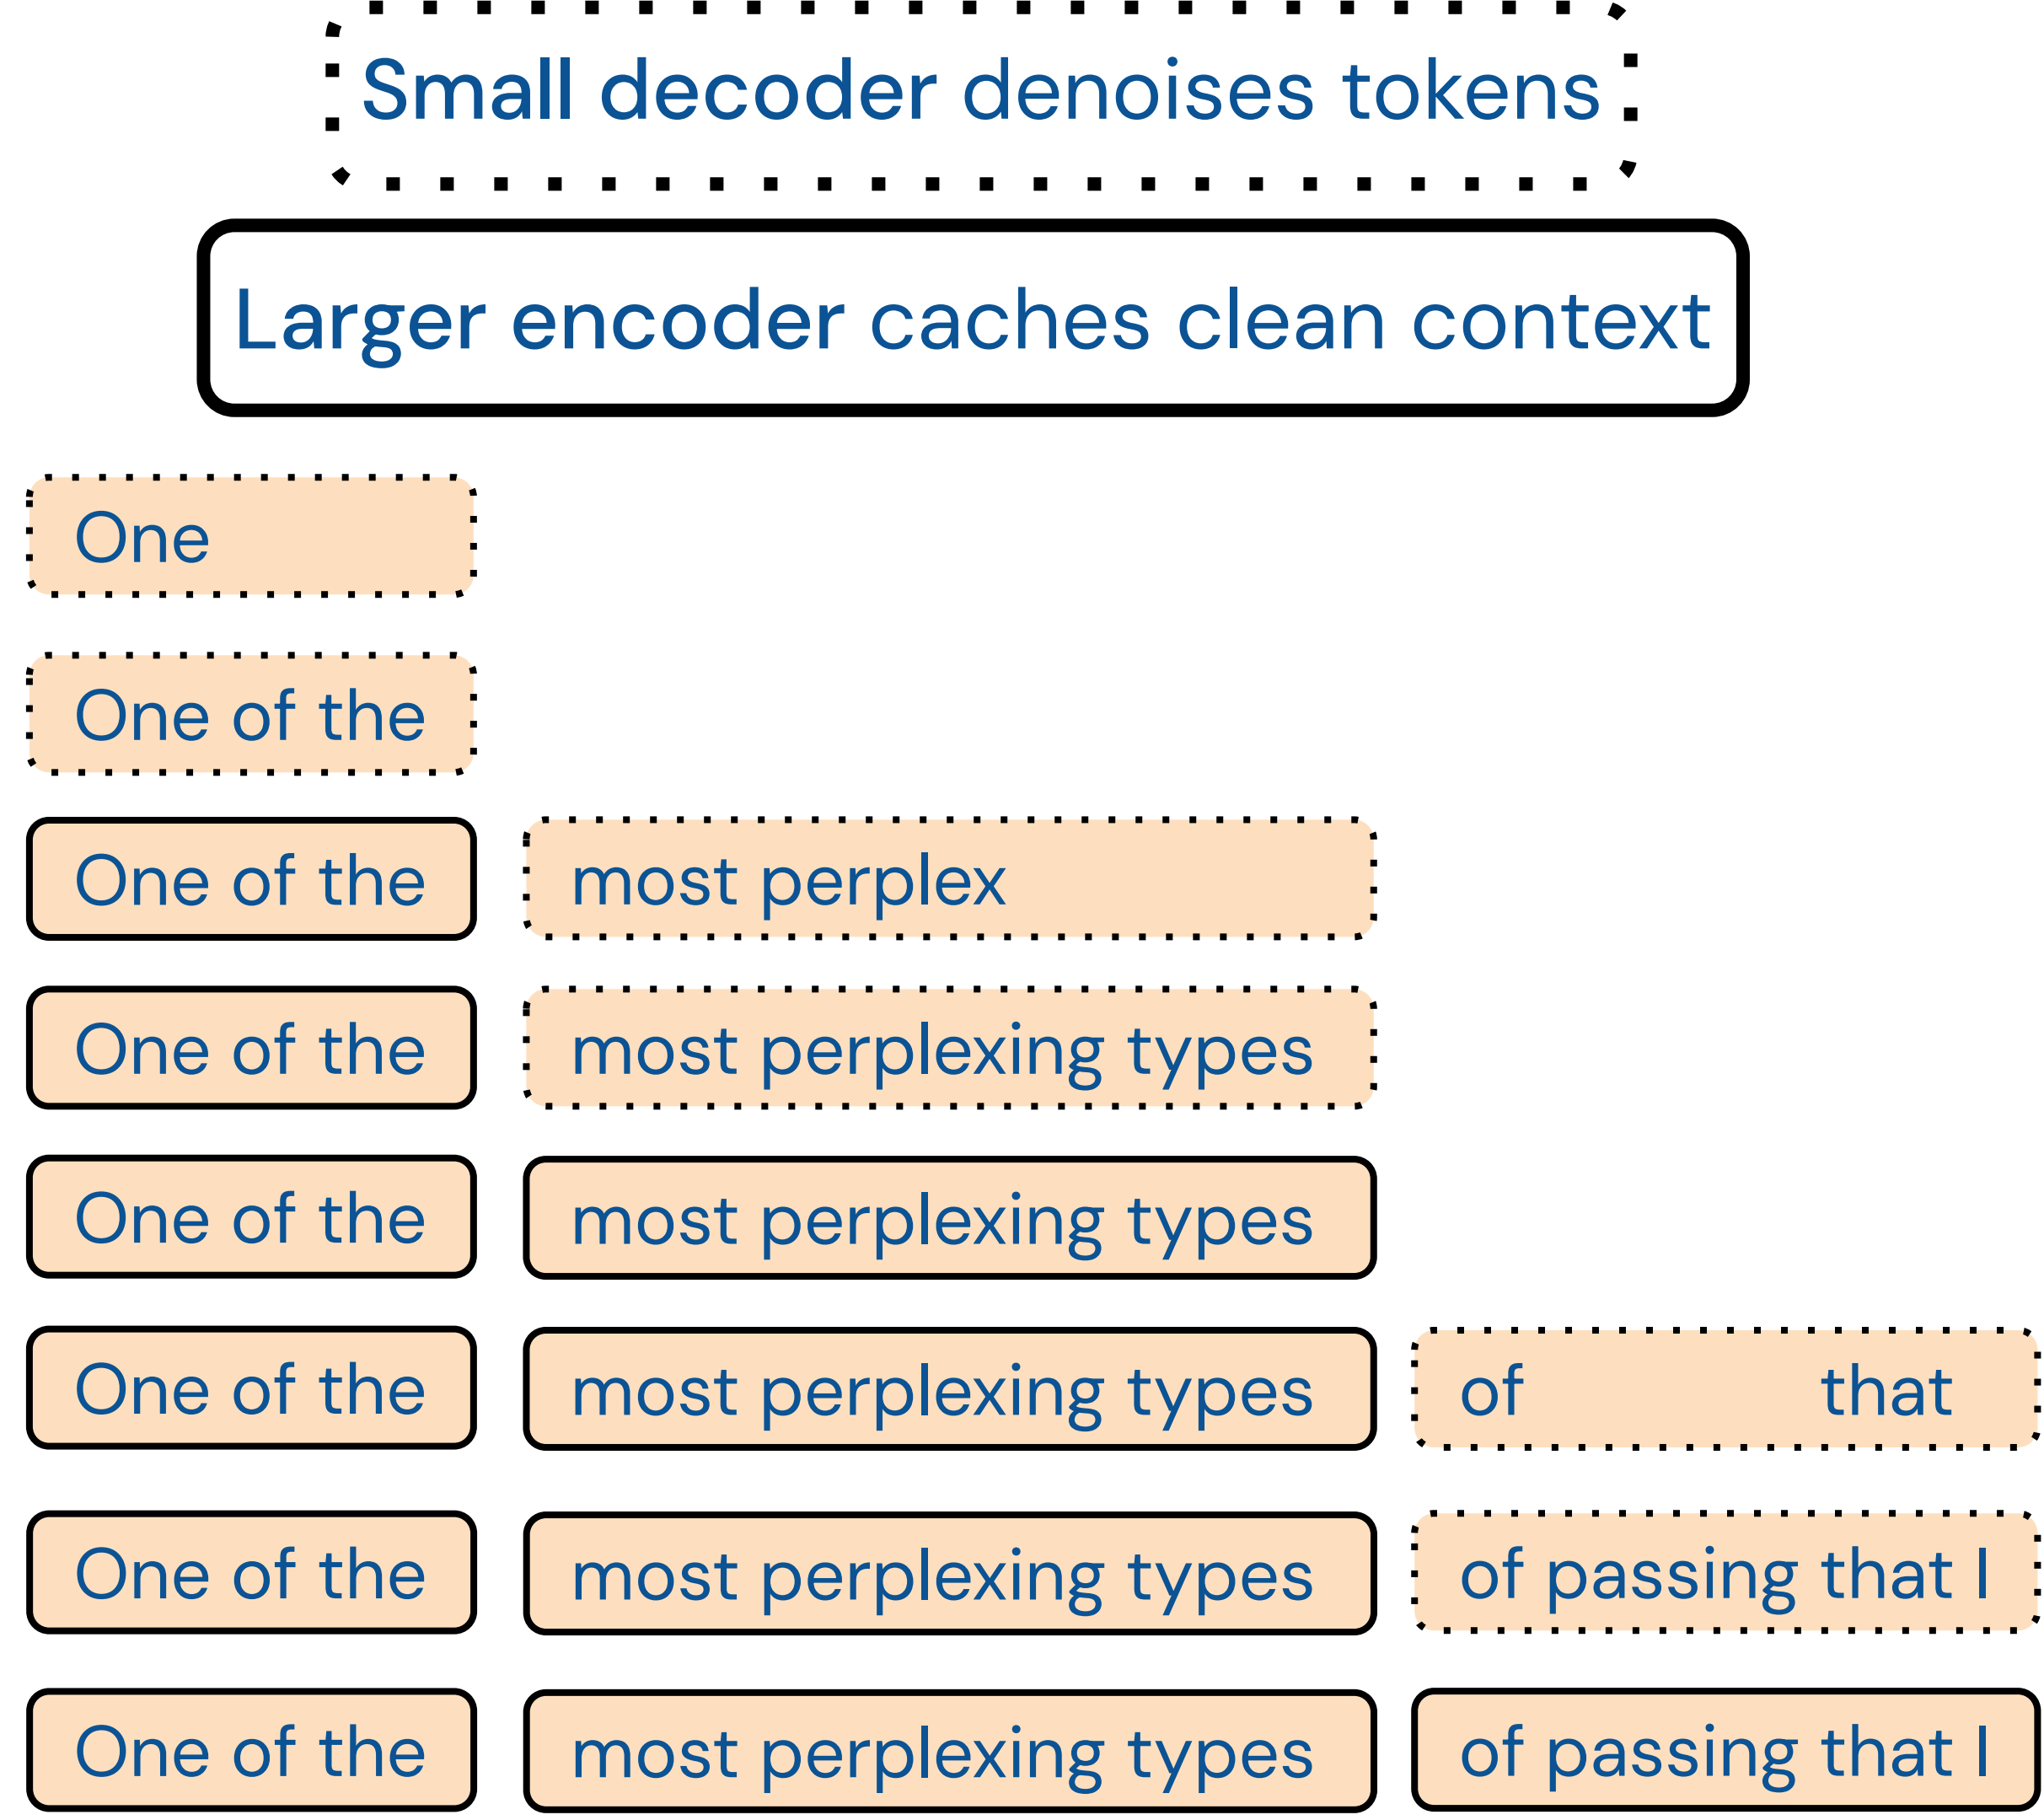
Marianne Arriola\*, Yair Schiff\*, Hao Phung,  
Aaron Gokaslan, Volodymyr Kuleshov  
Cornell Tech



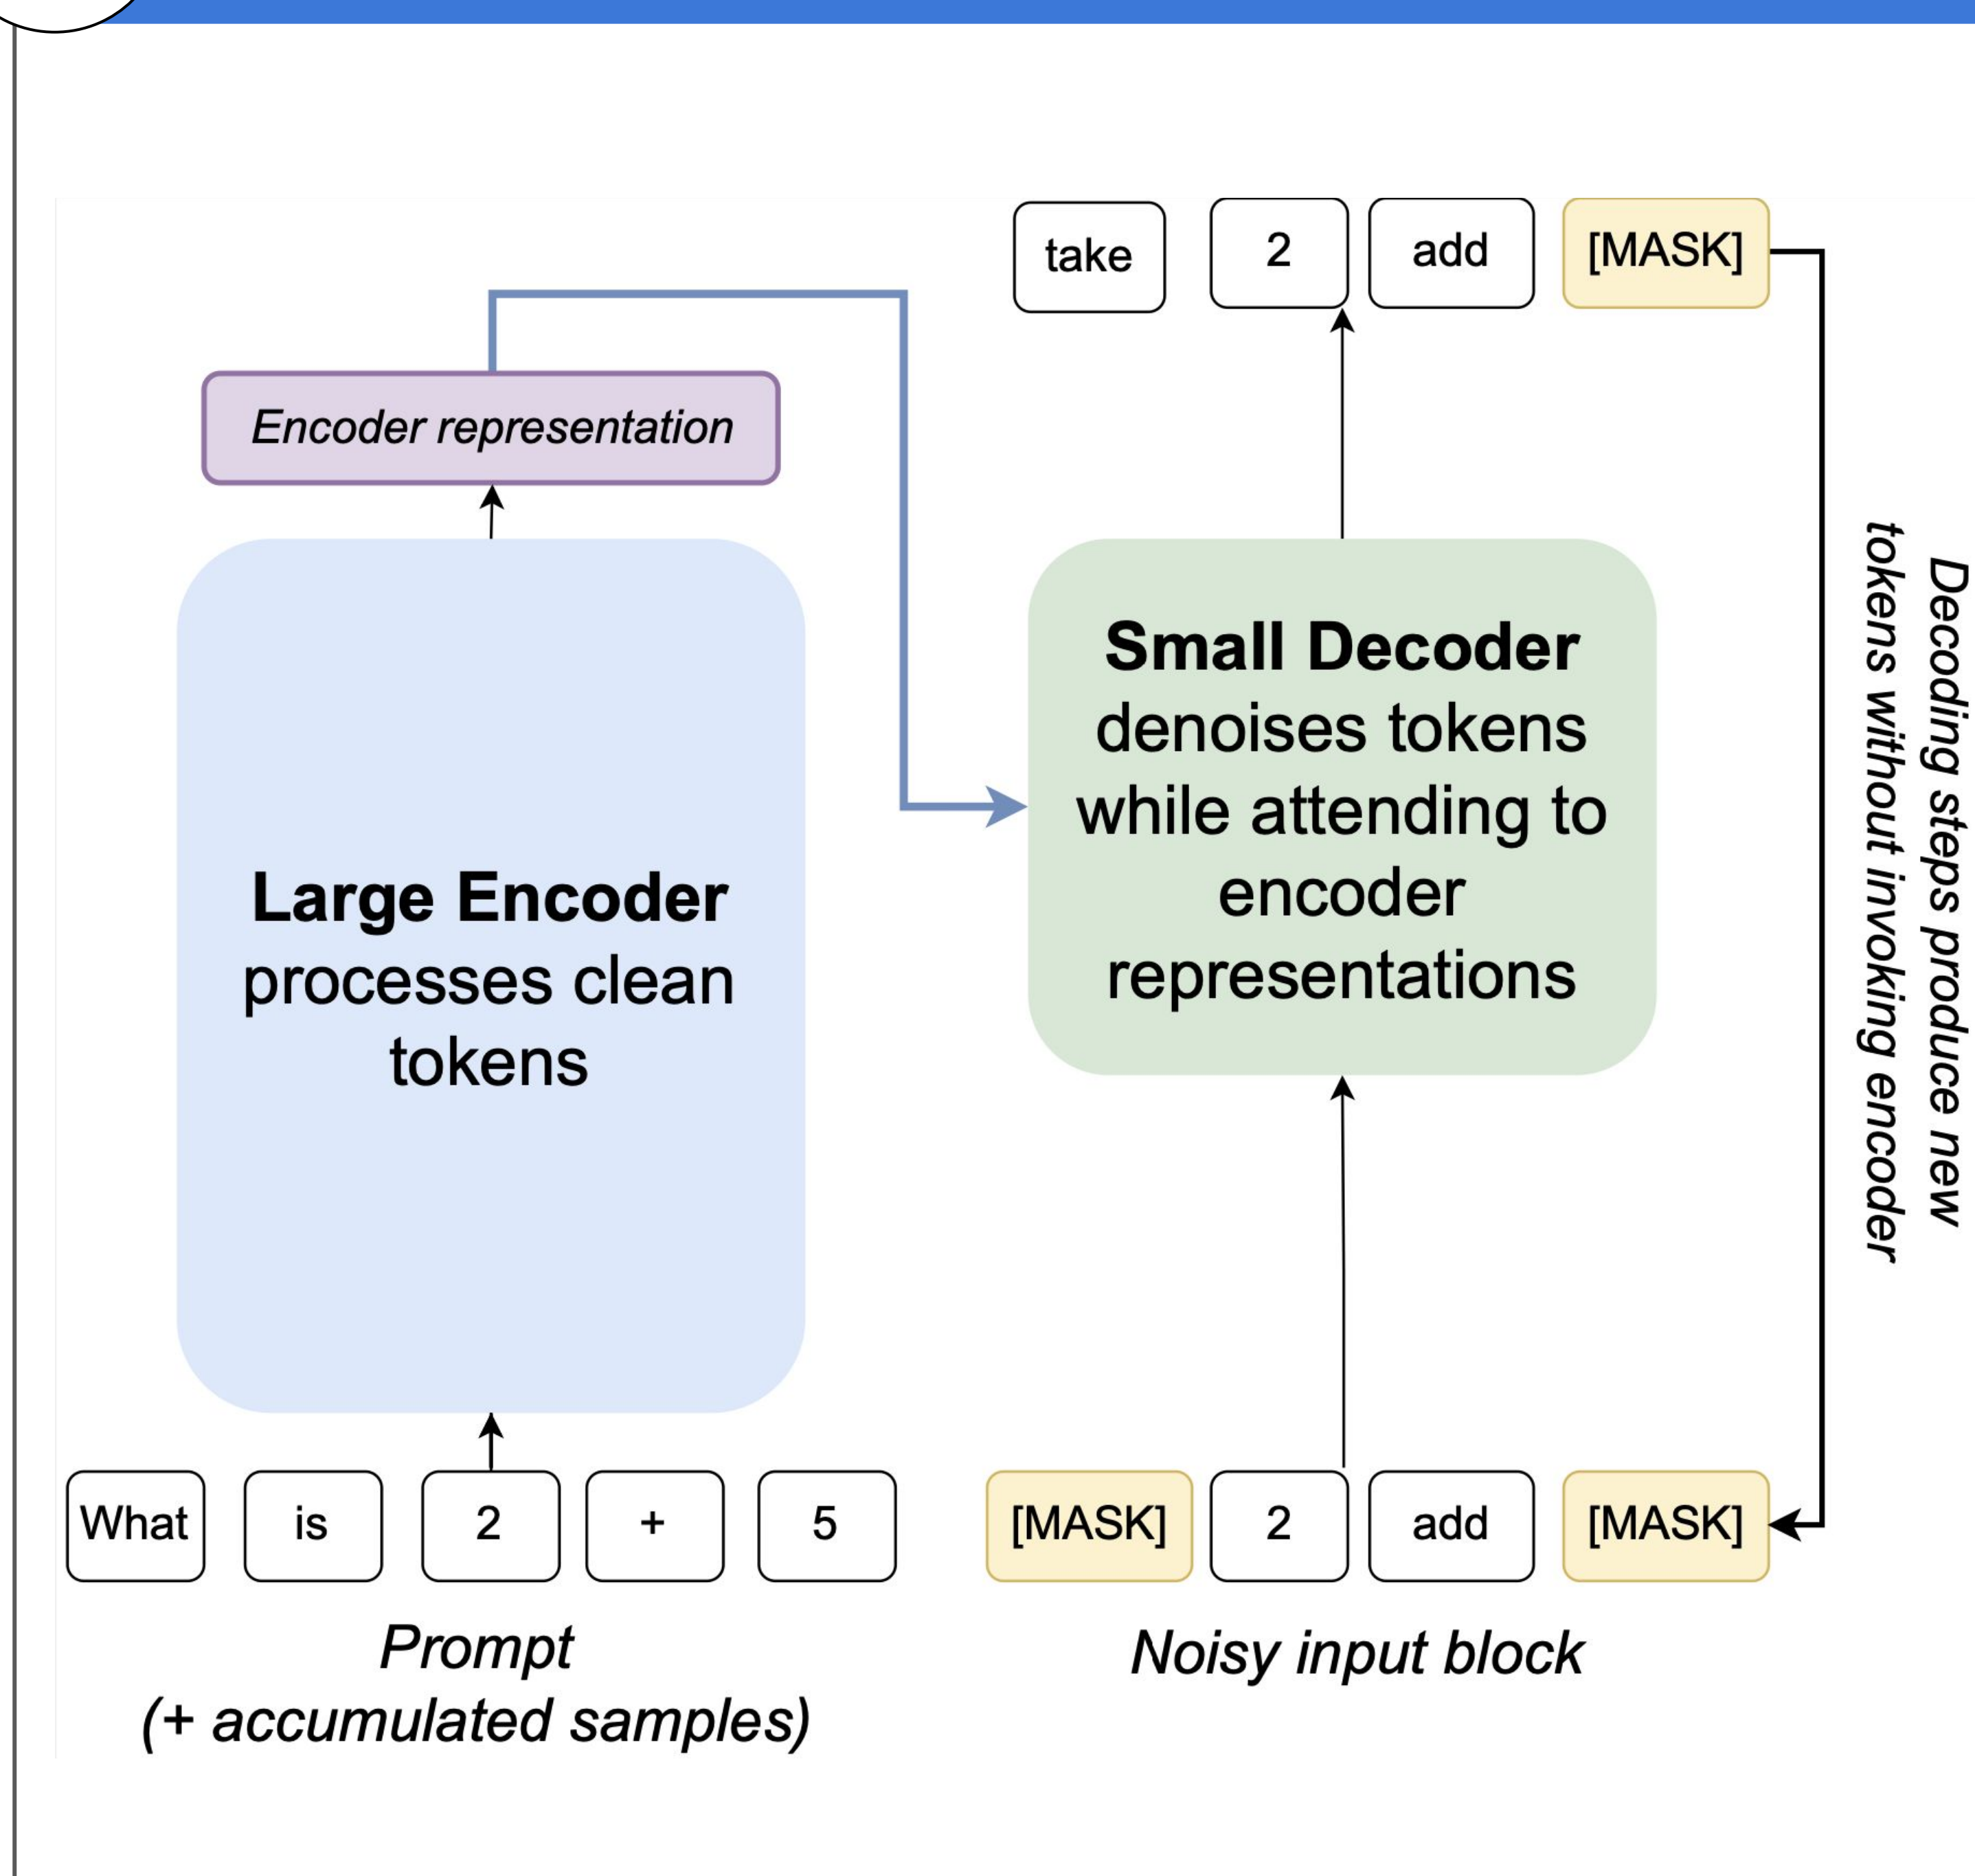
## 1 Slow, decoder-only diffusion



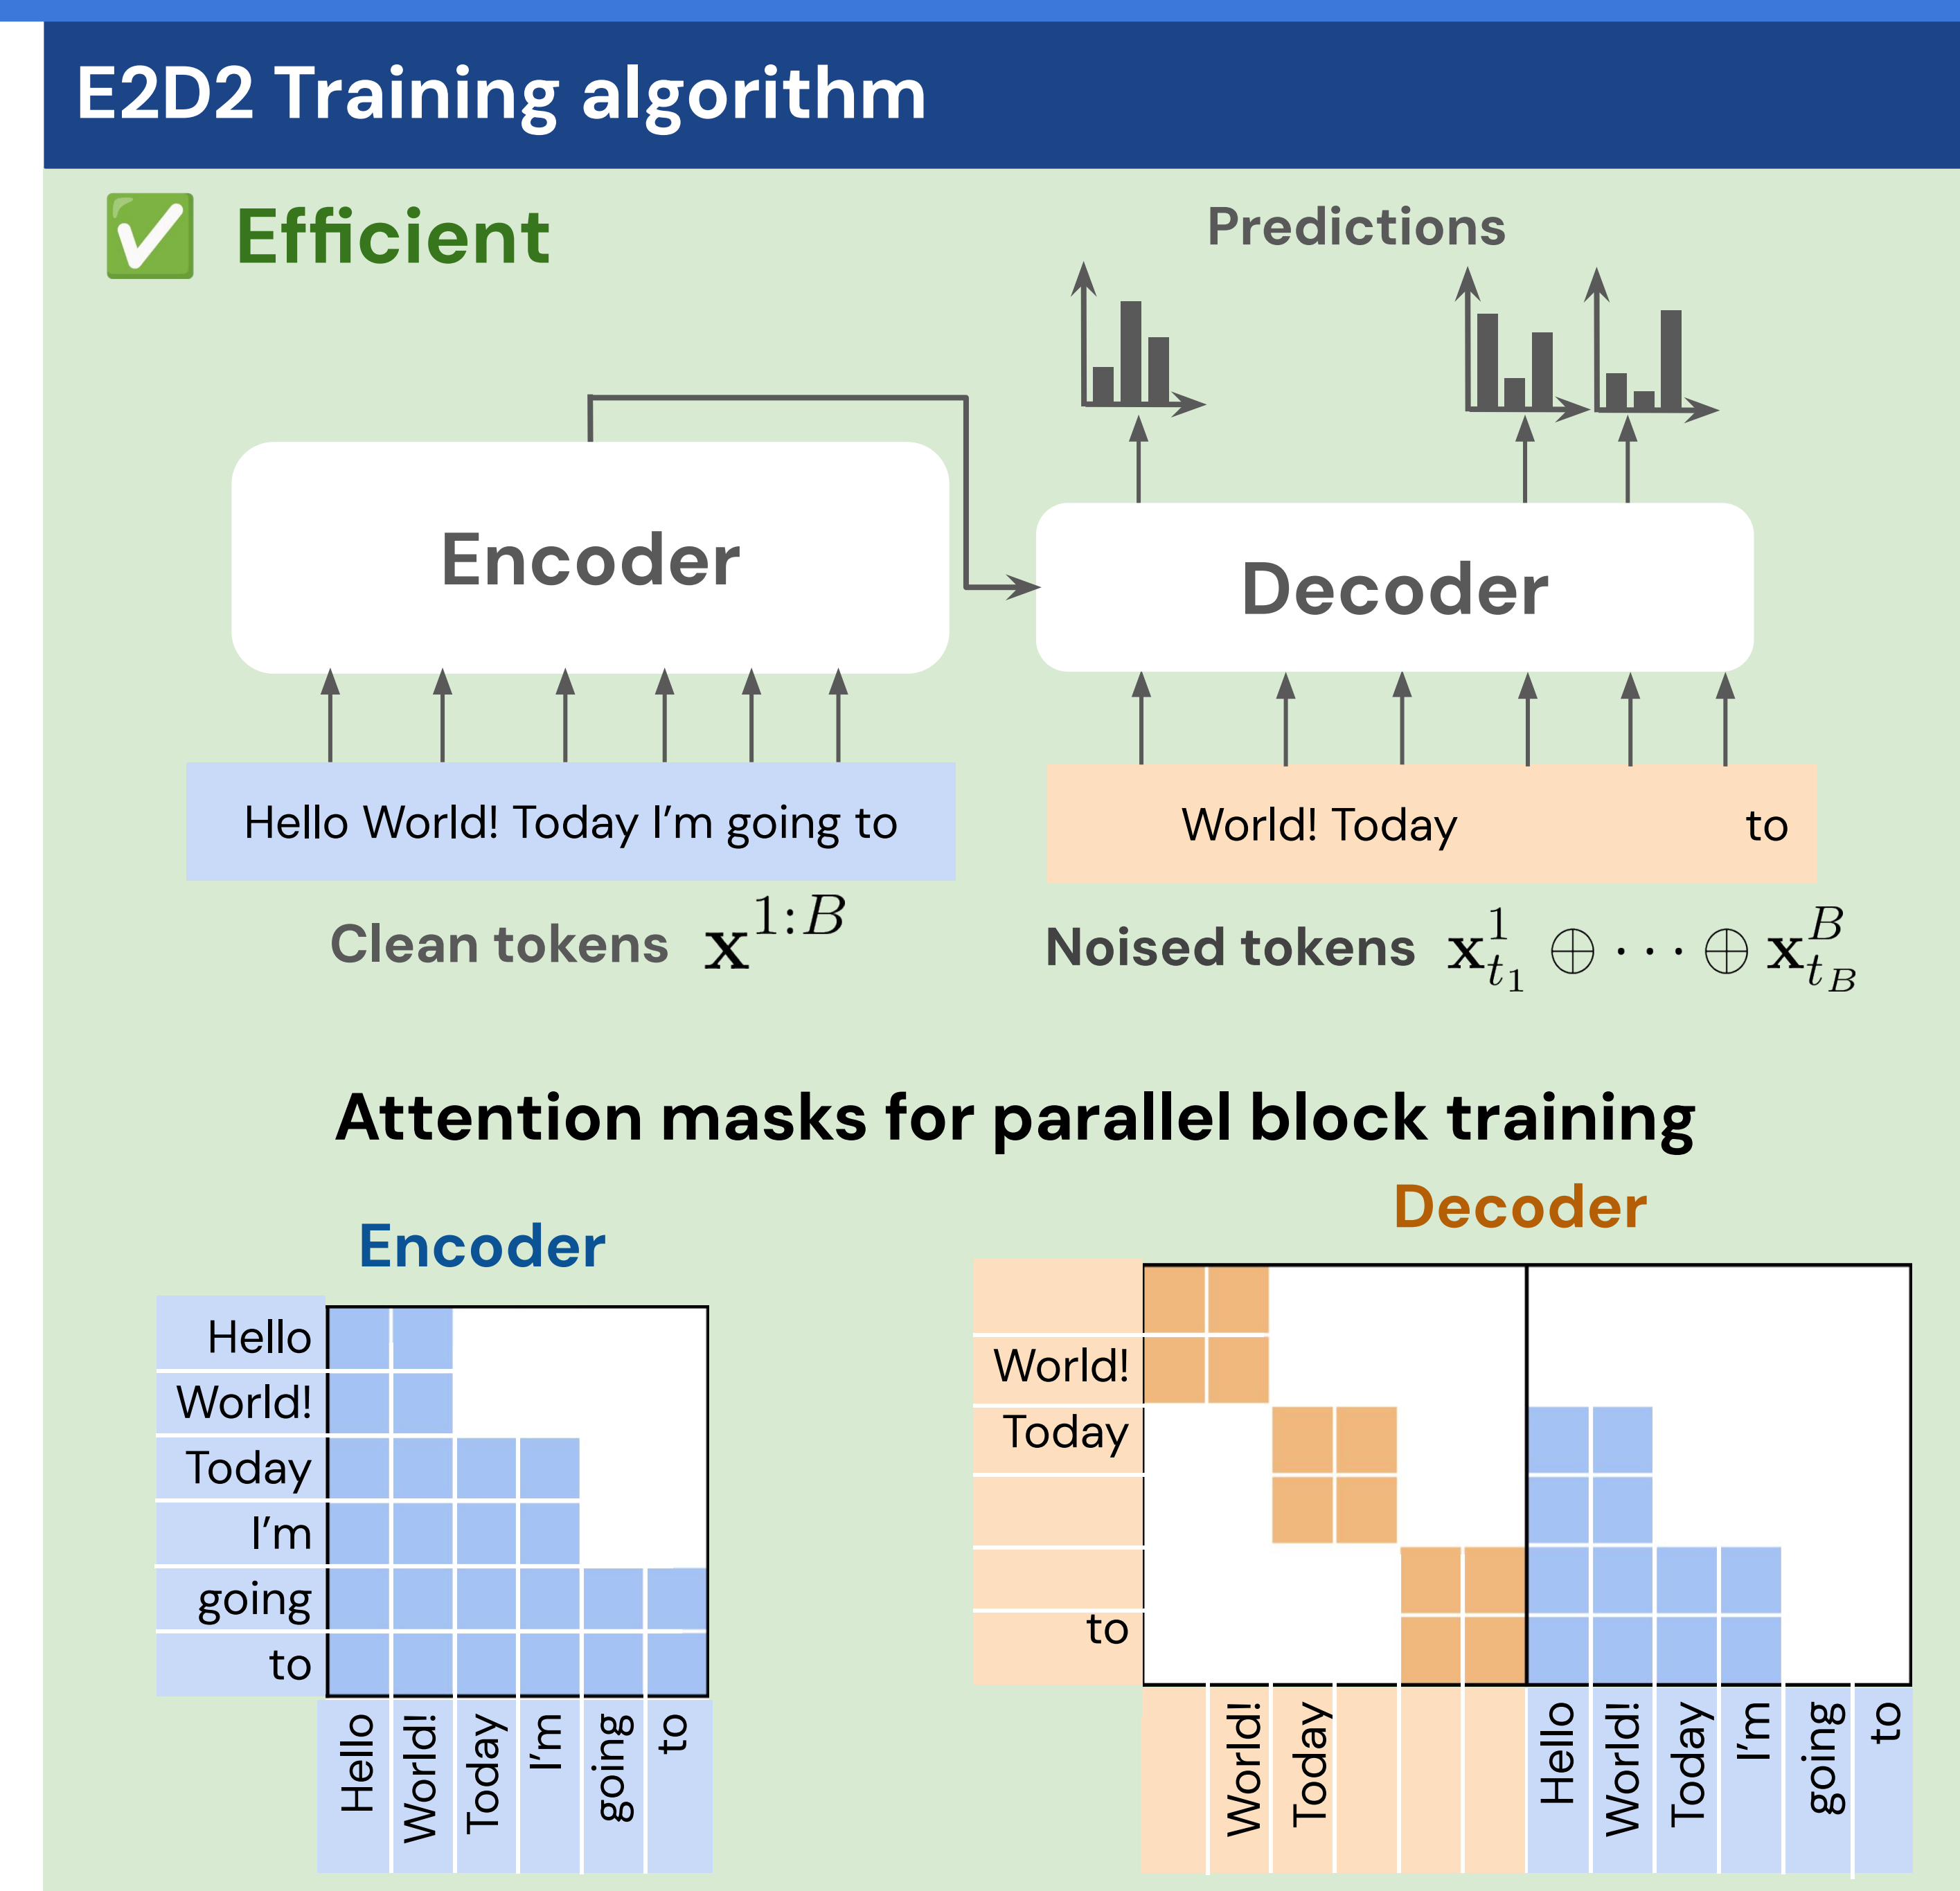
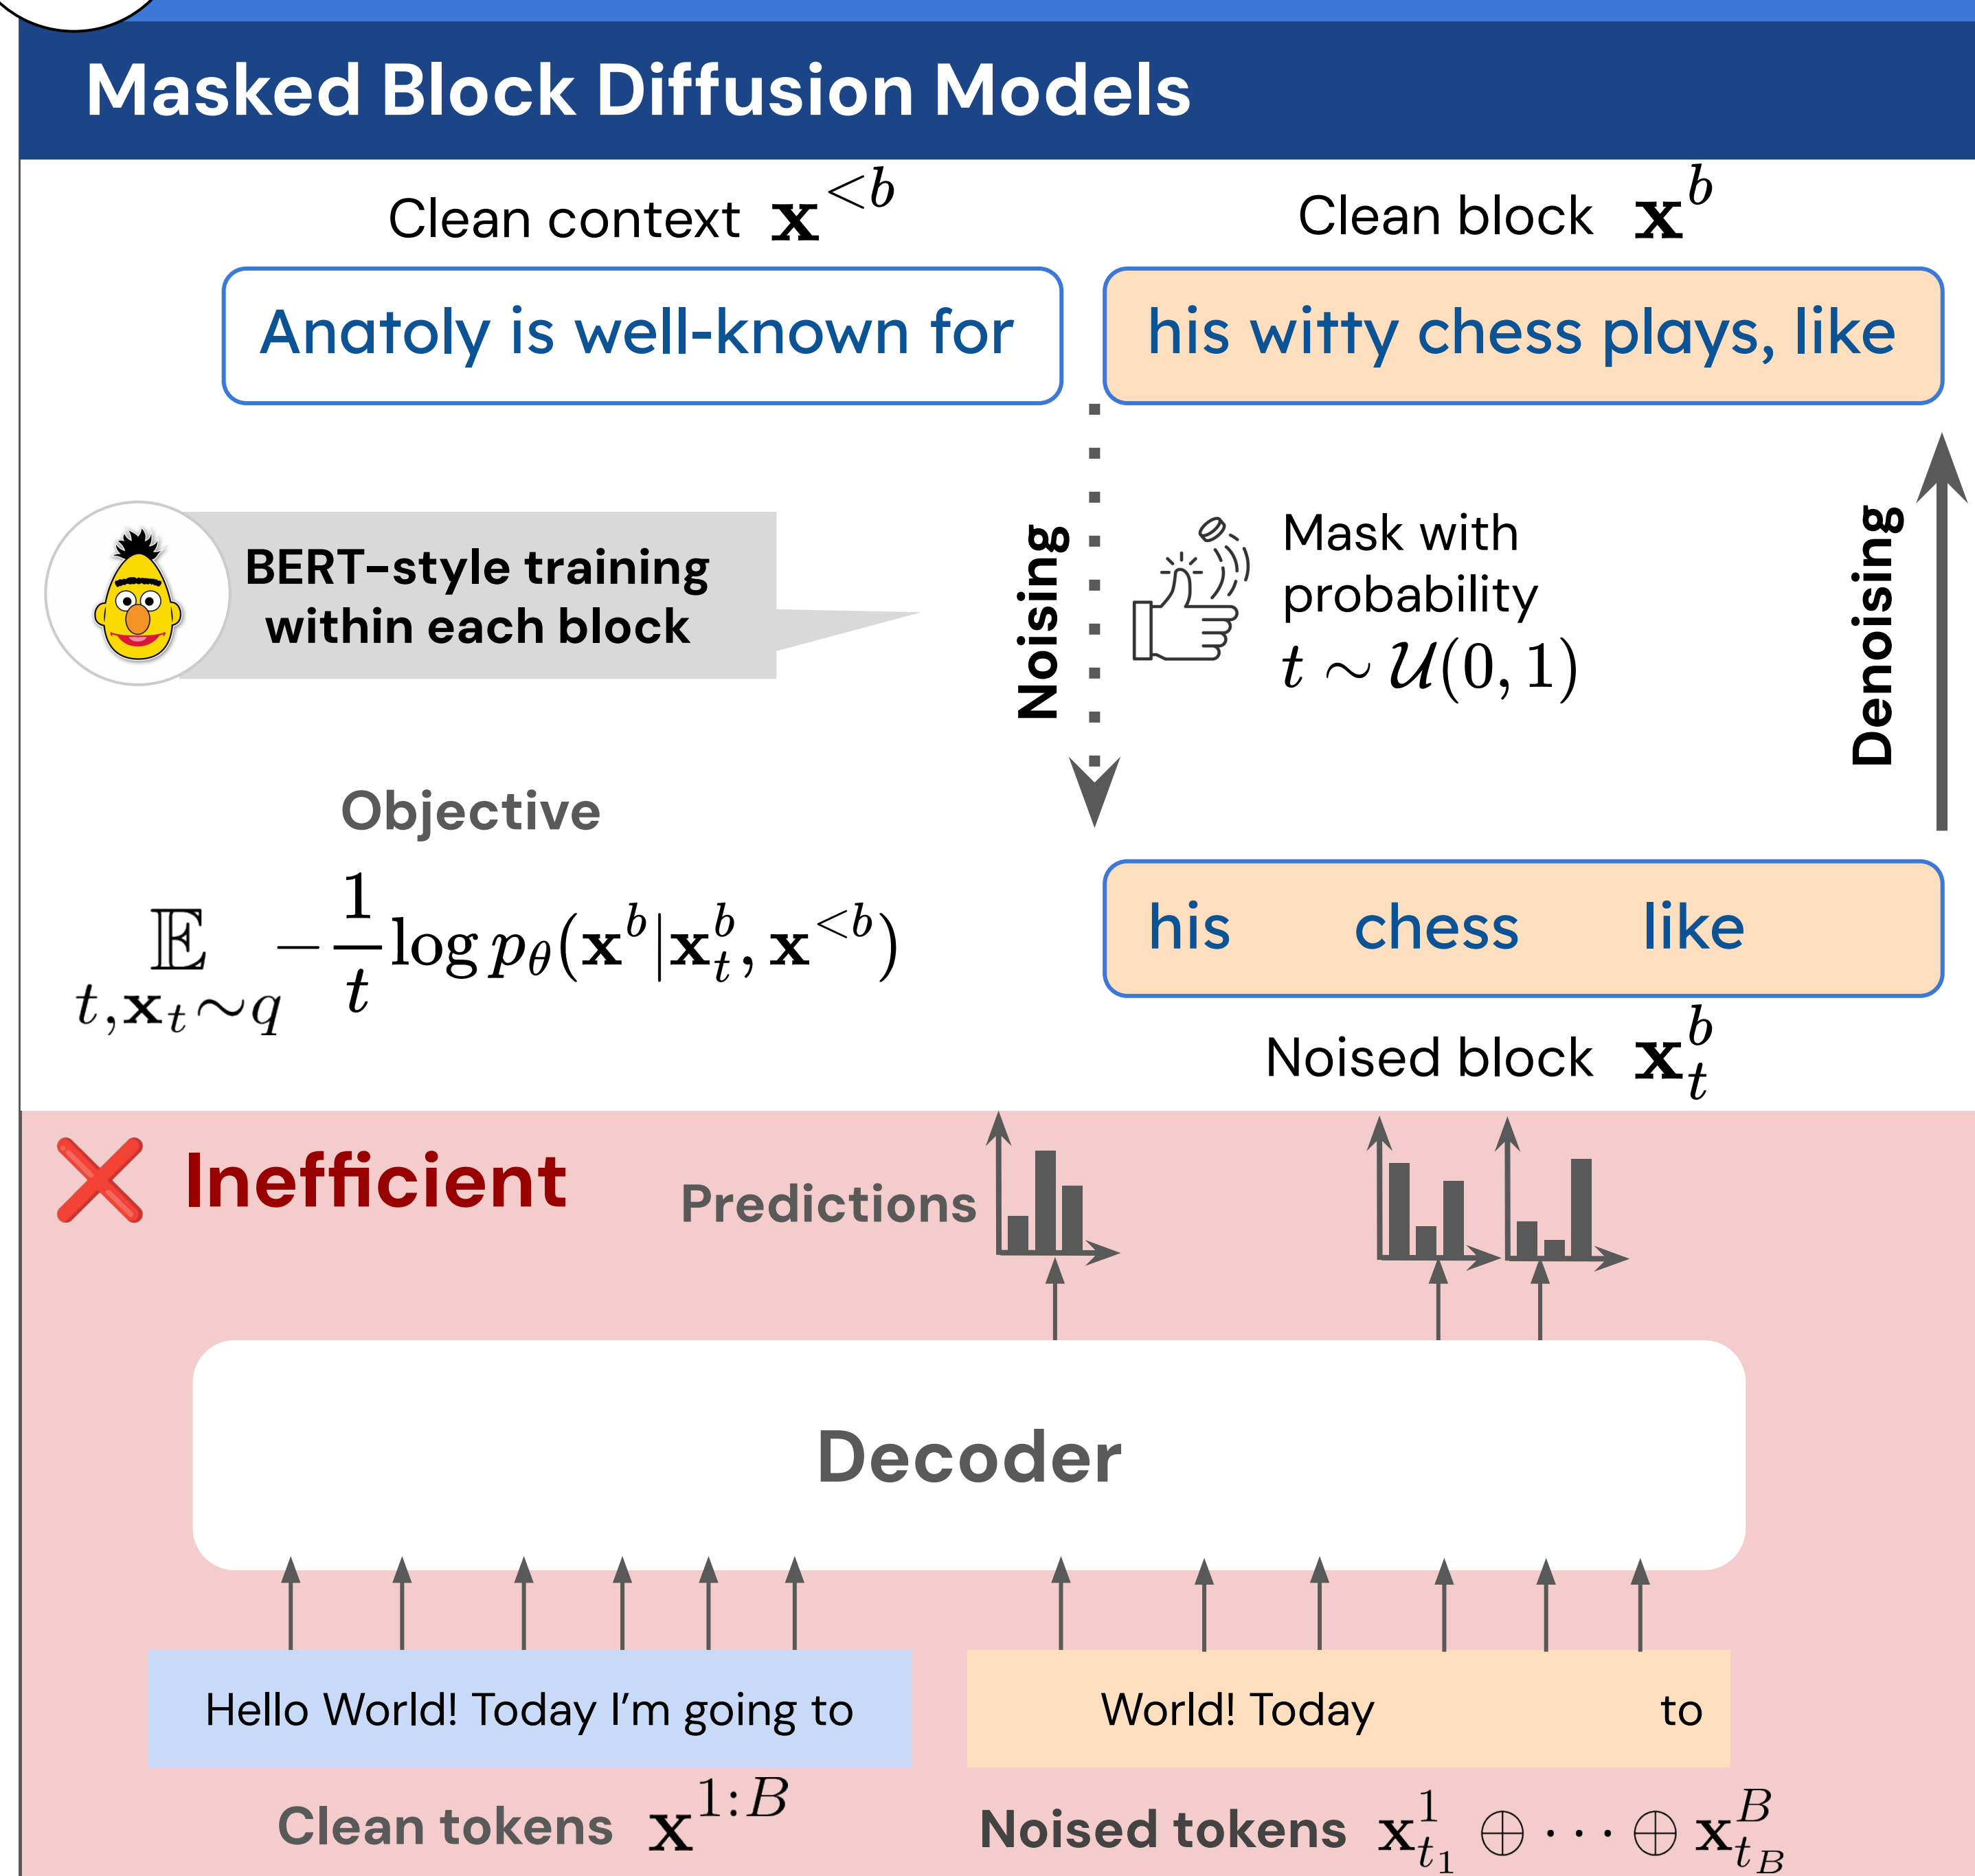
## Fast, encoder-decoder diffusion



## 2 Architecture design for faster inference



## 3 Efficient training splits computation between encoder/decoder



## 4 Improved quality-speed tradeoffs

