



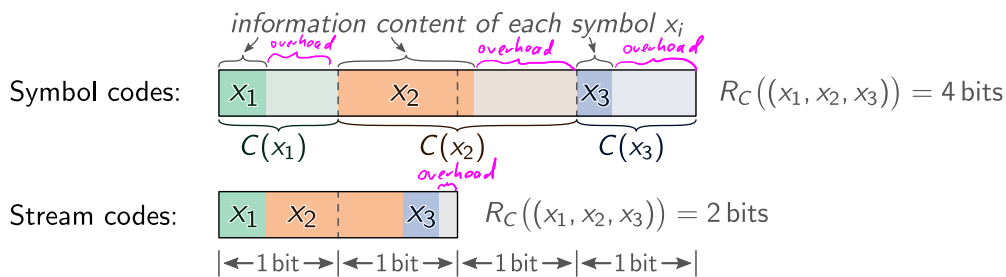
Lecture 6, Part 1:

# Asymmetric Numeral Systems (ANS)

Robert Bamler · Summer Term of 2023

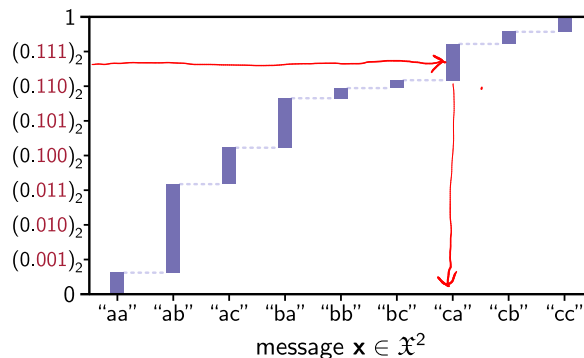
These slides are part of the course “Data Compression With and Without Deep Probabilistic Models” taught at University of Tübingen. More course materials—including video recordings, lecture notes, and problem sets with solutions—are publicly available at <https://robamler.github.io/teaching/compress23/>.

## Recall From Last Week: Stream Codes



- ▶ Reduces overhead from  $\leq 1$  bit *per symbol* to  $\leq 1$  bit *per message* (in theory)
- ▶ In practice: larger overhead due to finite precision & technical limitations (more on this today), but much smaller than for symbol codes

## Recall: Arithmetic Coding & Range Coding




- ▶ Simple algorithm in principle: iteratively refine interval  $[P(\mathbf{X} < \mathbf{x}), P(\mathbf{X} \leq \mathbf{x})]$
- ▶ Tricky to implement in practice (finite precision arithmetic, edge cases)
- ▶ Operates as a *queue*: “first in first out”


## (1) Markov Process




## (2) Hidden Markov Model



## (3) Autoregressive Model



## (4) Latent Variable Model



next week, using entropy coder that we discuss now

► Problem 7.3: use range coding for our autoregressive model of natural language

# Positional Numeral Systems

optimal symbol code (Huffman code)  $k(x_i) = 3, 3, 3, 3, 4, 4, 4, 4, 3, 3 \Rightarrow \text{avg} = 3.4$

Consider a data source that generates a random message  $\mathbf{X} \equiv (X_1, X_2, \dots, X_k)$  of length  $k$ ; symbols are i.i.d., with each  $X_i$  uniformly distributed over  $\mathcal{X} = \{0, 1, 2, \dots, 9\}$ .

(a) What is the entropy per symbol?  $\frac{1}{k} H_P[\mathbf{X}] = H_P[X_i] = \mathbb{E}_P[-\log_2 P(X_i = x_i)] = \mathbb{E}_P[-\log_2 \frac{1}{10}] = \log_2(10) \approx 3.32$

(b) What is the expected code word length of an optimal symbol code for this data source?

$L_{\text{Opt. symbol code}} := \mathbb{E}_P[|C_{\text{Opt. symbol code}}(X_i)|] = 3.4$  (see Huffman tree above)

(c) Can you do better than an optimal symbol code? Describe your approach first **in words** then **implement it in Python** or in **pseudo code**

next slide

- Don't think about arithmetic/range coding; it's much simpler.
- About 4 lines of code for encoding + 4 lines of code for decoding; no library function calls necessary.

encode(x): allows us to delete  
initialize n = 1  
for i in k down to 1:  
  update n ← n \* 10 + x;  
return binary\_rep(n)

decode(n):  
while n ≠ 1:  
  emit symbol n mod 10  
  update n ← ⌊n/10⌋

interpret x<sub>i</sub> as digits of a decimal number & then convert to binary

(d) What is the expected bit rate per symbol of your method from part (c) in the limit of long messages?  $\lim_{k \rightarrow \infty} \frac{1}{k} \mathbb{E}_P[R_{C_{\text{part (c)}}}(\mathbf{X})] \leq \lim_{k \rightarrow \infty} \frac{1}{k} R_{C_{\text{part (c)}}}((999\dots 9)) = \lim_{k \rightarrow \infty} \frac{\text{length}(\text{binary\_rep}(1999\dots 9))}{k} =$

(similar calculation for lower bound:  $\lim_{k \rightarrow \infty} \frac{1}{k} \mathbb{E}_P[R_C(\mathbf{X})] \geq \log_2 10$ )

$= \lim_{k \rightarrow \infty} \frac{1}{k} \log_2 [2 \times 10^k - 1] \leq \lim_{k \rightarrow \infty} \frac{1}{k} (k \log_2 10 + 2) = \log_2 10$

# Positional Numeral Systems: Implementation

## ► Implementation in Python:

```
def encode_uniform(message, base=10):
    number = 1
    for symbol in reversed(message):
        number = number * base + symbol
    return number

def decode_uniform(number, base=10):
    while number != 1:
        yield number % base
        number //= base
```

we encode symbols in reverse order here so that the decoder below reconstructs them in normal order

→ positional numeral systems behave like a stack ("last in first out")

## ► Usage example:

```
compressed = encode_uniform([3, 5, 6])
print(bin(compressed)) # Prints: "0b11001110101"
print(list(decode_uniform(compressed))) # Prints: "[3, 5, 6]"
```

# Observations About Positional Numeral Systems

- ▶ encoding & decoding operates as a *stack* ("last in first out")
- ▶ encoding *amortizes* bit rate over several symbols ( $\neq$  symbol codes)

we changed the 3<sup>rd</sup> symbol  
we changed the 2<sup>nd</sup> symbol

```
print(bin(encode_uniform([3, 5, 7]))) # Prints: "0b11011011001"
print(bin(encode_uniform([3, 5, 6]))) # Prints: "0b11001110101"
print(bin(encode_uniform([3, 4, 6]))) # Prints: "0b11001101011"
```

those bits seem to depend on both the second and the third symbol  
⇒ we can no longer assign each bit to a single symbol, as we were able to do for symbol codes

- ▶ positional numeral systems are an *optimal lossless compression methods* if:
  - all symbols are from the same (finite) alphabet  $\mathcal{X}$
  - all symbols are uniformly distributed over  $\mathcal{X}$
  - all symbols are statistically independent
- ▶ **Goal today:** remove constraints (i) and (ii)
- ▶ **Goal next lecture:** remove constraint (iii) → "bits-back trick"

## Limitation (i): symbols from different alphabets

Not a real limitation: just make base position-dependent:

Usage example:

```
class UniformCoder:
    def __init__(self, compressed=0):
        self.compressed = compressed

    def push(self, symbol, base):
        """Encodes a symbol from alphabet {0, ..., base-1}."""
        self.compressed = self.compressed * base + symbol

    def pop(self, base):
        """Decodes a symbol from alphabet {0, ..., base-1}."""
        symbol = self.compressed % base
        self.compressed //= base
        return symbol
```

```
coder = UniformCoder()
coder.push(6, base=10)
coder.push(13, base=16)
coder.push(7, base=8)
print(bin(coder.compressed))
# Prints: "0b1101101111"

print(coder.pop(base=8))
# Prints: "7"
print(coder.pop(base=16))
# Prints: "13"
print(coder.pop(base=10))
# Prints: "6"
```

reconstructs original message (in reverse order)

## Limitation (ii): Non-uniformly Distributed Symbols

- ▶ Consider a single symbol  $x_i$
- ▶ **Step 1:** approximate  $P(X_i)$  in fixed point arithmetic:

$$Q(x_i) = \frac{m_i(x_i)}{n} \quad \text{where } n = 2^{\text{precision}} \text{ is a power of 2 (will become useful later)} \\ \text{and the integers } m_i(x_i) \text{ satisfy } \sum_{x_i \in \mathcal{X}_i} m_i(x_i) = n \text{ and are chosen such that } Q(x_i) \approx P(X_i)$$

- ▶ compression overhead:  $D_{KL}(P(X_i) \parallel Q(x_i))$  (asymptotically decays exponentially in precision)

- ▶ **Step 2:** interpret  $Q(X_i)$  as the marginal distribution of a *latent variable model*:

$$Q(x_i) = \sum_{z_i=0}^{n-1} Q(z_i) Q(x_i | z_i)$$

*uniform distribution over  $\{0, 1, \dots, n-1\}$  i.e.,  $Q(z_i) = \frac{1}{n} \forall z_i \in \{0, \dots, n-1\}$*


$$Q(x_i = x_i | z_i = z_i) = \begin{cases} 1 & \text{if } z_i \in \mathcal{Z}_i(x_i) \\ 0 & \text{otherwise} \end{cases}$$

*$(\mathcal{Z}_i(x_i))_{x_i \in \mathcal{X}_i}$  are pairwise disjoint subsets of  $\{0, \dots, n-1\}$  of size  $|\mathcal{Z}_i(x_i)| = m_i(x_i)$*

$$\Rightarrow Q(x_i = x_i) = \sum_{z_i \in \mathcal{Z}_i(x_i)} Q(z_i = z_i) Q(x_i = x_i | z_i = z_i) = \sum_{z_i \in \mathcal{Z}_i(x_i)} \frac{1}{n} = \frac{|\mathcal{Z}_i(x_i)|}{n} = \frac{m_i(x_i)}{n} \checkmark$$

# Limitation (ii): Non-uniformly Distributed Symbols

- ▶ Consider a single symbol  $x_i$
- ▶ **Step 1:** approximate  $P(X_i)$  in fixed point arithmetic:
- ▶ **Step 2:** interpret  $Q(X_i)$  as the marginal distribution of a *latent variable model*.
- ▶ **Step 3:** Bits-back trick:
  - ▶  $\{m_i(x_i)\}$  partition the range  $\{1, 2, \dots, n-1\}$  into  $|\mathcal{X}_i|$  non-overlapping subranges  $\mathfrak{Z}_i(x_i)$ :



- ▶ **Naive idea:** encode arbitrary  $z_i \in \mathfrak{Z}_i(x_i) \rightarrow$  bit rate:  $-\log_2 Q(z_i; x_i) = \log_2 n =$  precision bits per symbol 😞
- ▶ **Better idea:** piggy-back some information into choice of  $z_i$  by decoding  $z_i$  from previously encoded data, using a uniform model over  $\mathfrak{Z}_i(x_i) \rightarrow$  consumes  $\log_2 |\mathfrak{Z}_i(x_i)| = \log_2 m_i(x_i)$  bits  $\Rightarrow$  net bit rate:  $\log_2 n - \log_2 m_i(x_i) = -\log_2 Q(x_i; x_i)$  😊
- $\rightarrow$  "bits-back trick" (more general discussion next lecture)

## (Slow) Implementation of ANS in Python

Prove correctness  
 and analyze  
 bit rate on  
 problem set

```

class SlowAnsCoder:
    def __init__(self, precision, compressed=0):
        self.n = 2**precision # ("**" denotes exponentiation.)
        self.uniform_coder = UniformCoder(compressed) # See slide 7.

    def push(self, symbol, m): # Encodes one symbol.
        z = self.uniform_coder.pop(base=m[symbol]) + sum(m[0:symbol])
        self.uniform_coder.push(z, base=self.n)

    def pop(self, m): # Decodes one symbol.
        z = self.uniform_coder.pop(base=self.n)
        # Find the unique symbol that satisfies z in Z_i(symbol)
        # (using linear search just to simplify exposition):
        for symbol, m_symbol in enumerate(m):
            if z >= m_symbol:
                z -= m_symbol
            else:
                break
        self.uniform_coder.push(z, base=m_symbol)
        return symbol

    def get_compressed(self):
        return self.uniform_coder.compressed
  
```

decode  $z$  from  
 any data that's  
 already accumulated  
 on self.uniform\_coder  
 using uniform distrib.  
 over  $\mathfrak{Z}_i(x_i)$   
 $\rightarrow$  consumes  
 $\log_2 m_i(x_i)$  bits

reconstruct  
 the data that  
 we consumed  
 on the first  
 line of push.

### Usage example:

```

# We use a very low precision
# here for demonstration purpose;
# real deployments should use
# higher precision.
precision = 4
m = [7, 6, 3] # (Sums to 16 = 2^4.)
coder = SlowAnsCoder(precision)

coder.push(0, m)
coder.push(2, m)
coder.push(1, m)
print(bin(coder.get_compressed()))
# Prints: "0b101000"



print(coder.pop(m)) # Prints: 1
print(coder.pop(m)) # Prints: 2
print(coder.pop(m)) # Prints: 0
  
```

This means:  
 $Q(x_i=0) = \frac{7}{2^4} = \frac{7}{16}$   
 $Q(x_i=1) = \frac{6}{16}$   
 $Q(x_i=2) = \frac{3}{16}$

We could also use a different  
 model for each symbol as long  
 as we use matching models for  
 decoding below.

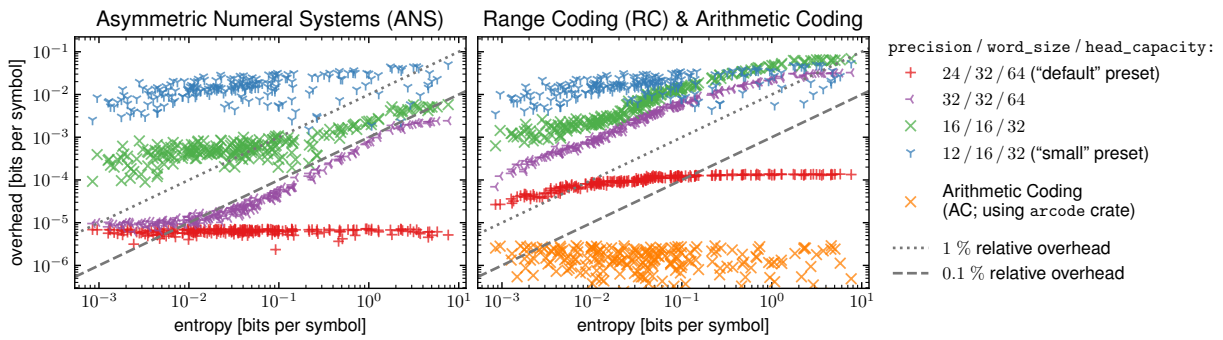
## Streaming ANS

- ▶ SlowAnsCoder is slow (run-time  $O(k^2)$ )
- ▶ **Idea ("streaming ANS"):** operate mostly on a compressed representation with *finite capacity*. If it would overflow, push an *integer number of bits* to a growable buffer.



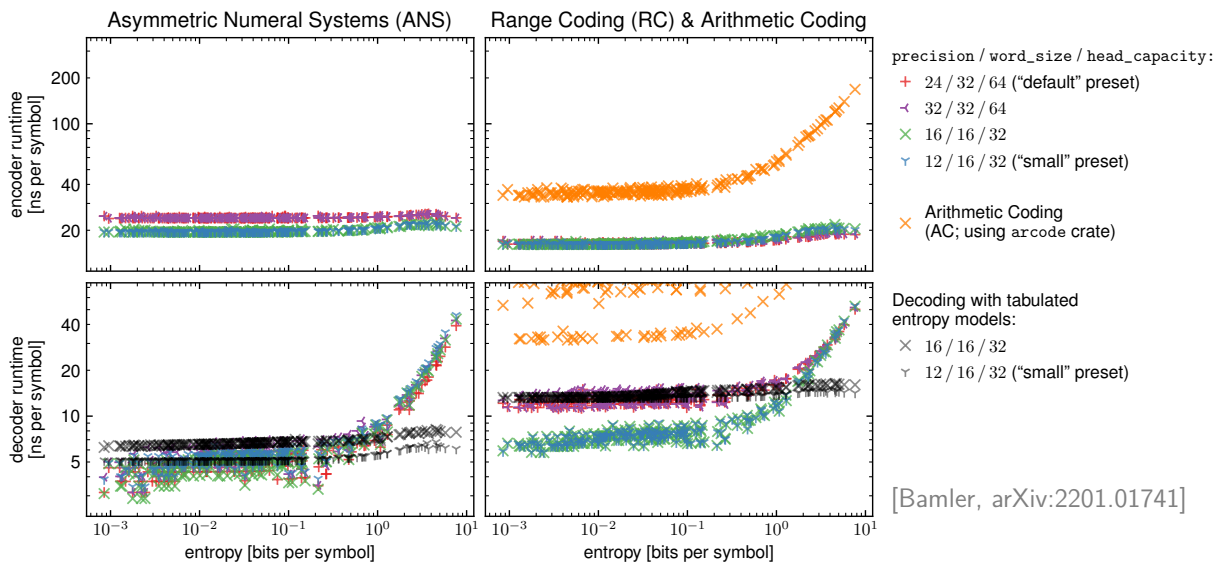
$\rightarrow$  You'll walk through an illustrated  
 example of this in Problem 6.3.

encode & decode mostly  
 on fixed-capacity head but  
 transfer least significant half  
 of it to/from bulk whenever  
 it over/underflows.



[Bamler, arXiv:2201.01741 (2022)]

*uses ANS & range coding implementation from "constriction" library, which you'll also use on later problem sets (starting with Problem 7.3).*



## Outlook

- ▶ **Problem Set:**
  - ▶ Prove correctness and analyze compression performance of our `SlowAnsCoder` implementation.
  - ▶ Illustrate an example of streaming ANS.
- ▶ **Next Lecture:** revisit & generalize the bits-back trick
- ▶ **Afterwards:** use ANS for (net) optimal lossless compression with latent variable models (e.g., variational autoencoders)