METEOR

Steganography for Realistic Distributions

Gabriel Kaptchuk (Boston University)
Tushar Jois, Matthew Green, Aviel Rubin (Johns Hopkins University)

Widespread Success of Encrypted Systems

Encrypted Messengers

~2 Billion Monthly Users

Encrypted Browsing

Ubiquitous Adoption and Significant Usability Progress

Censorship Resistance

>2 Million Daily Connections

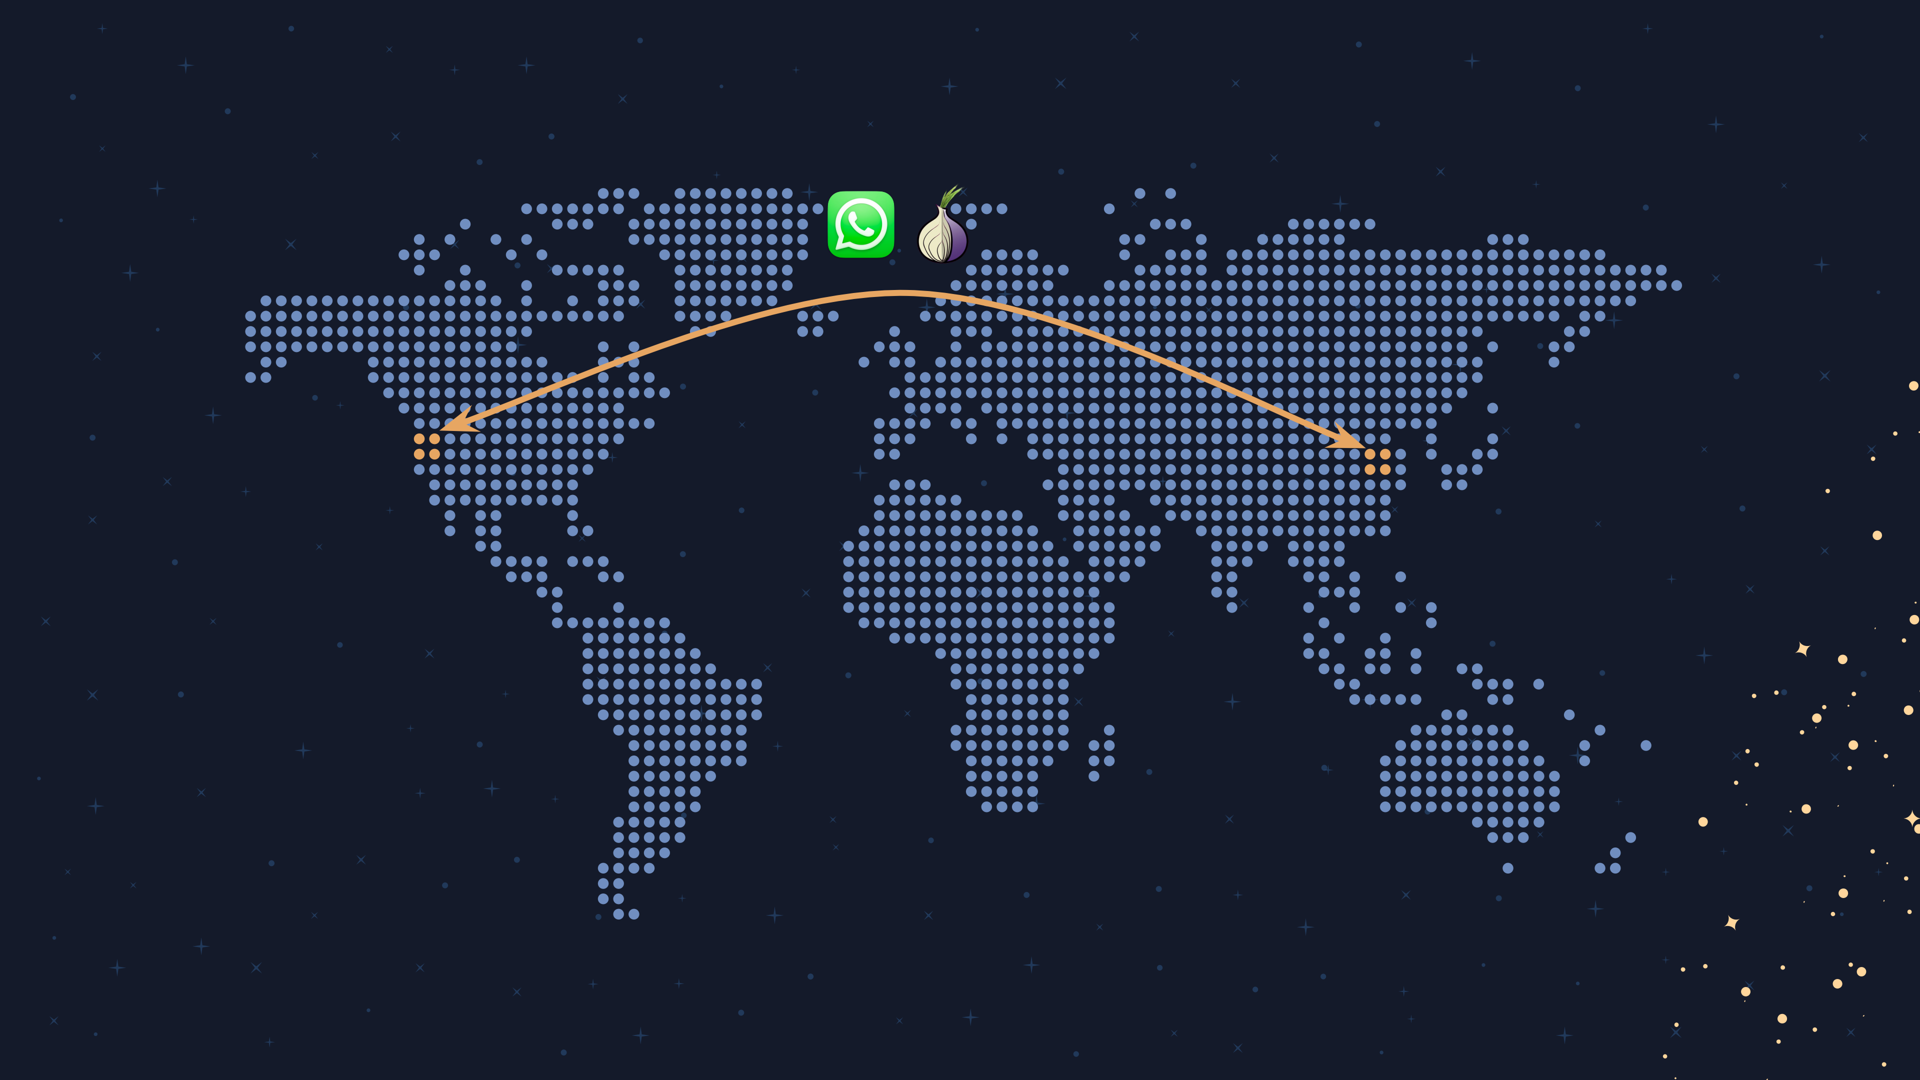

Perfect Tool: Universal Steganography!

Perfect Tool: Universal Steganography!

Problem: Univ. Stegano. For Realistic Distributions Has Never Been Deployed

Our Contributions

Identify and overcome main barriers to realistic steganography

Analyze prior public key steganography protocols

Propose new symmetric key construction with better performance

Related Work

Classical Steganography and FTE

- Seminal work of Simmons [Sim83]
- MANY follow ups [AP98, ZFK+98, Mit99, Cac00, HLv02, RR03, Le03, LK03, vH04, BC05]

Related Work

Classical Steganography and FTE

- Seminal work of Simmons [Sim83]
- MANY follow ups [AP98, ZFK+98, Mit99, Cac00, HLv02, RR03, Le03, LK03, vH04, BC05]
- Keyless Steganography [ACI+20]

Related Work

Classical Steganography and FTE

- Seminal work of Simmons [Sim83]
- MANY follow ups [AP98, ZFK+98, Mit99, Cac00, HLv02, RR03, Le03, LK03, vH04, BC05]
- Keyless Steganography [ACI+20]
- Lysyanskaya and Meyerovich look at limits of using Markov Models [LM06]

Related Work

Classical Steganography and FTE

- Seminal work of Simmons [Sim83]
- MANY follow ups [AP98, ZFK+98, Mit99, Cac00, HLv02, RR03, Le03, LK03, vH04, BC05]
- Keyless Steganography [ACI+20]
- Lysyanskaya and Meyerovich look at limits of using Markov Models [LM06]
- Format Transforming Encryption [LDJ+14, DCRS13b, DCS15, OYZ+20]

Related Work

Classical Steganography and FTE

- Seminal work of Simmons [Sim83]
- MANY follow ups [AP98, ZFK+98, Mit99, Cac00, HLv02, RR03, Le03, LK03, vH04, BC05]
- Keyless Steganography [ACI+20]
- Lysyanskaya and Meyerovich look at limits of using Markov Models [LM06]
- Format Transforming Encryption [LDJ+14, DCRS13b, DCS15, OYZ+20]

Censorship Avoidance Tools

- obfs4/ScrambleSuit [WPF13]
- Domain Fronting [FLH+15]
- Skypemorph [MLDG12]
- FTEProxy [DCRS13a]
- StegoTorus [WWY+12]
- CensorProofer [WGN+12]
- FreeWave [HRBS13]

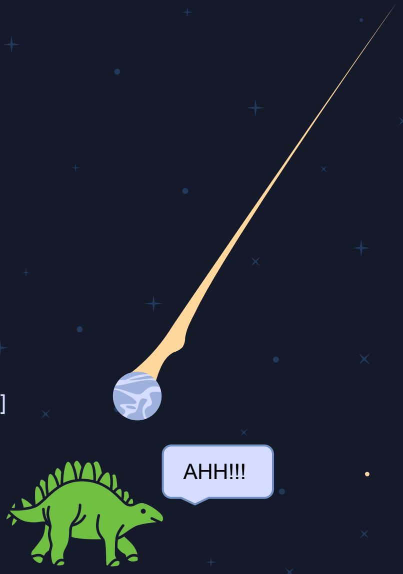

Related Work

Classical Steganography and FTE

- Seminal work of Simmons [Sim83]
- MANY follow ups [AP98, ZFK+98, Mit99, Cac00, HLv02, RR03, Le03, LK03, vH04, BC05]
- Keyless Steganography [ACI+20]
- Lysyanskaya and Meyerovich look at limits of using Markov Models [LM06]
- Format Transforming Encryption [LDJ+14, DCRS13b, DCS15, OYZ+20]

Censorship Avoidance Tools

- obfs4/ScrambleSuit [WPF13]
- Domain Fronting [FLH+15]
- Skypemorph [MLDG12]
- FTEProxy [DCRS13a]
- StegoTorus [WWY+12]
- CensorProofer [WGN+12]
- FreeWave [HRBS13]

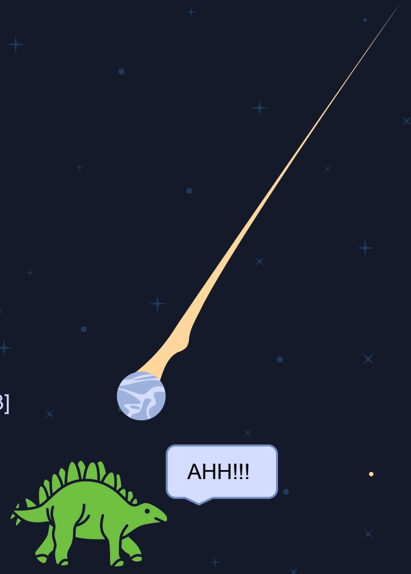

Ad-Hoc Steganography + Generative Models

- ML Steganography constructions [GGA+05, SSSS07, YHC+09, CC10, CC14, FJA17, VNBB17, YJH+18, Xia18, YGC+19, HH19, DC19, ZDR19]
- Attacking constructions [YHZ19, YWL+19, YWS+18, WBK15, KFH12, MHC+08]

Talk Outline

01 Steganography Refresher

02 Classical Schemes + Generative Models

03 METEOR: Dealing with Low Entropy

Talk Outline

01 Steganography Refresher

02 Classical Schemes + Generative Models

03 METEOR: Dealing with Low Entropy

Universal Steganography Refresher [Hop04]

Encode Message

1. Encrypt message $\mathtt{m}$ as $\mathtt{x}$ with IND$-CPA scheme

Universal Steganography Refresher [Hop04]

Encode Message

1. Encrypt message $\mathtt{m}$ as $\mathtt{x}$ with IND$-CPA scheme

2. For each bit $\mathtt{x_i}$ of the ciphertext:

Universal Steganography Refresher [Hop04]

Encode Message

1. Encrypt message m as x with IND$-CPA scheme

2. For each bit x_i of the ciphertext:

   a. Sample random c_i from covertext distribution

   b. If $h(c_i)=x_i$ (where h is an unbiased hash function):

      – Yes: append c_i to the stegotext, and proceed to next x_i

      – No: return to (a)

Universal Steganography Refresher [Hop04]

Encode Message

1. Encrypt message $\mathtt{m}$ as $\mathtt{x}$ with IND$-CPA scheme
2. For each bit $\mathtt{x}_\mathtt{i}$ of the ciphertext:
   a. Sample random $\mathtt{c}_\mathtt{i}$ from covertext distribution
   b. If $\mathtt{h(c}_\mathtt{i}\mathtt{)=x}_\mathtt{i}$ (where $\mathtt{h}$ is an <u>unbiased</u> hash function) :
      – Yes: append $\mathtt{c}_\mathtt{i}$ to the stegotext, and proceed to next $\mathtt{x}_\mathtt{i}$
      – No: return to (a)

Decode Message

1. Recover $\mathtt{x}_\mathtt{i}$ as $\mathtt{h(c}_\mathtt{i}\mathtt{)}$
2. Decrypt $\mathtt{x}$ to recover $\mathtt{m}$

Universal Steganography Refresher [Hop04]

Encode Message

1. Encrypt message m as x with IND$-CPA scheme
2. For each bit x_i of the ciphertext:
   a. Sample random c_i from covertext distribution
   b. If $h(c_i)=x_i$ (where h is an <u>unbiased</u> hash function):
      – Yes: append c_i to the stegotext, and proceed to next x_i
      – No: return to (a)

Decode Message

1. Recover x_i as $h(c_i)$
2. Decrypt x to recover m

Security Intuition

1. x_i are all random
2. h introduces no bias
3. Therefore, c_i are distributed as the covertext distribution

Barriers To Practical Universal Steganography

1. Lack of Appropriate Samplers

Barriers To Practical Universal Steganography

1. Lack of Appropriate Samplers

- Covertext distribution too complex
- Covertext distribution fundamentally unknowable (eg. human text)
- Best option: good approximation

Barriers To Practical Universal Steganography

1. Lack of Appropriate Samplers

- Covertext distribution too complex
- Covertext distribution fundamentally unknowable (eg. human text)
- Best option: good approximation

2. Unrealistic Entropy Requirements

Barriers To Practical Universal Steganography

1. Lack of Appropriate Samplers

- Covertext distribution too complex
- Covertext distribution fundamentally unknowable (eg. human text)
- Best option: good approximation

2. Unrealistic Entropy Requirements

- Low entropy means hash function likely must be biased
- Two potential outcomes:
  1) Sampler never finds “good” sample
  2) Resampling amplifies bias

Barriers To Practical Universal Steganography

1. Lack of Appropriate Samplers

↓

Use (Public) Generative Models

2. Unrealistic Entropy Requirements

Generative Models

Generative Models

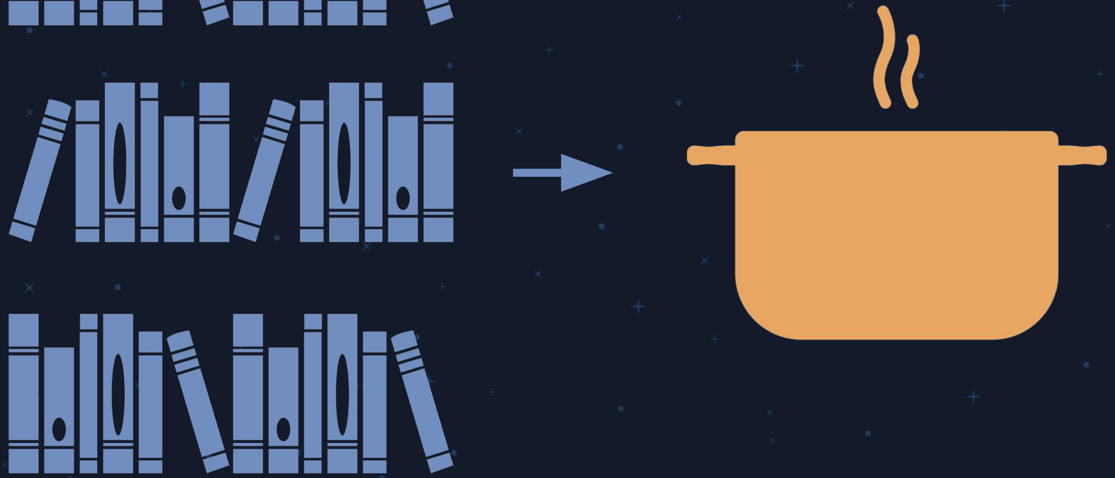

Generative Models

Generative Models

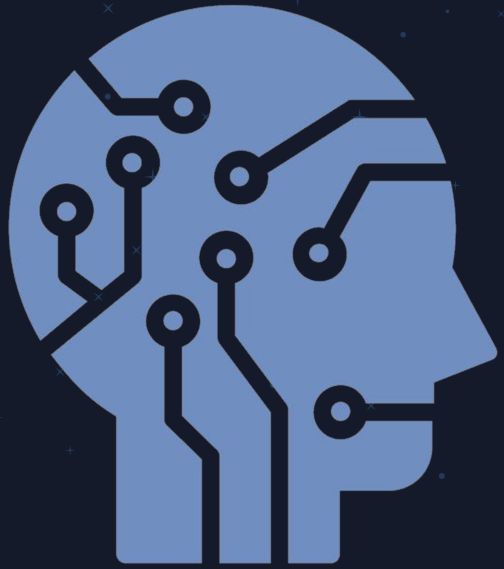

Generative Models

Context:

"Evidence indicates that the asteroid fell in the Yucatan Peninsula, at Chicxulub, Mexico."

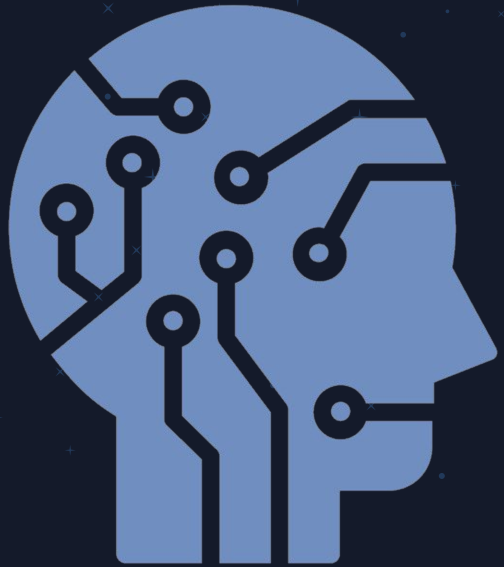

Generative Models

Context:
"Evidence indicates that the asteroid fell in the Yucatan Peninsula, at Chicxulub, Mexico."

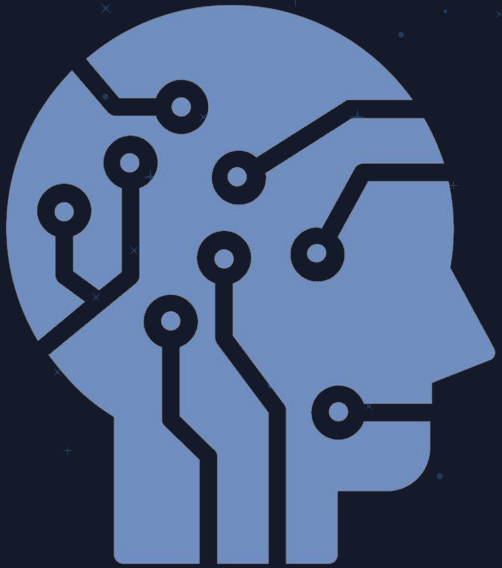

Next Word Prediction:
32% - "An"
17% - "The"
12% - "A"
23% - "However"
15% - "Since"
1% - Other Options

Generative Models

Context:

“Evidence indicates that the asteroid fell in the Yucatan Peninsula, at Chicxulub, Mexico.”

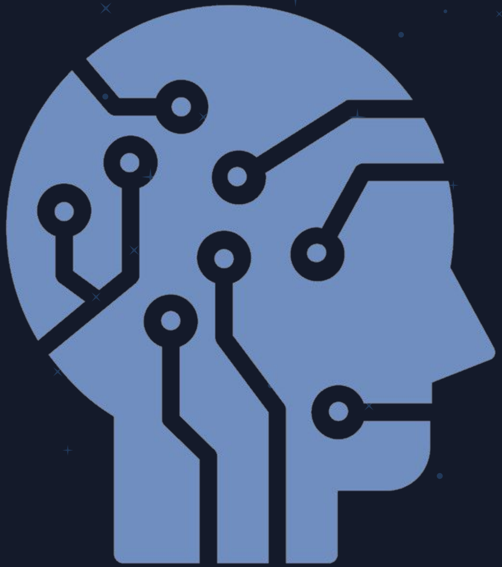

Next Word Prediction:

32% - “An”
17% - “The”
12% - “A”
23% - “However”
15% - “Since”
1% - Other Options

Generative Models

Context:
"Evidence indicates that the asteroid fell in the Yucatan Peninsula, at Chicxulub, Mexico. The"

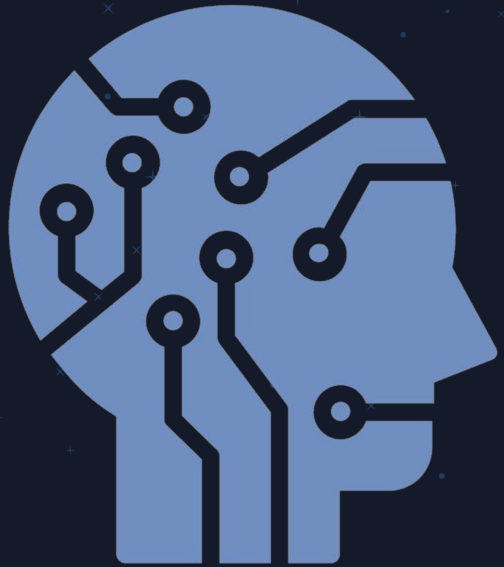

Generative Models

Context:
“Evidence indicates that the asteroid fell in the Yucatan Peninsula, at Chicxulub, Mexico. The”

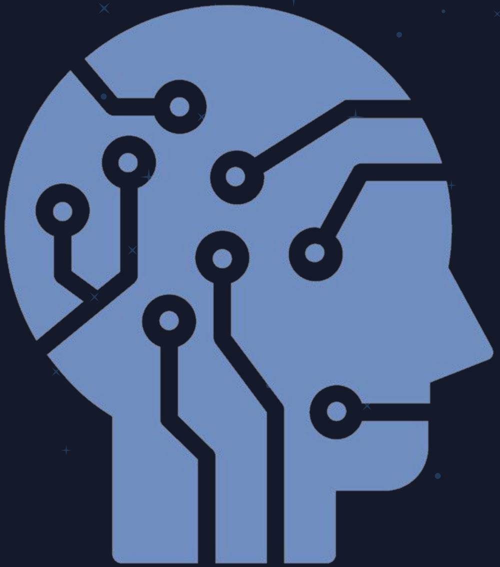

“first importance of Yucatan Peninsula is demonstrated with the following conclusion: the Pliocene Earth has lost about seven times as much vegetation as the Jurassic in regular parts of the globe, from northern India to Siberia…”

Barriers To Practical Universal Steganography

1. Lack of Appropriate Samplers

↓

Use (Public) Generative Models

2. Unrealistic Entropy Requirements

↓

Naturally Adapt Encoding Rate To Entropy

Talk Outline

01 Steganography Refresher

02 Classical Schemes + Generative Models

03 METEOR: Dealing with Low Entropy

Universal Steganography Barriers

Encode Message

1. Encrypt message $\mathtt{m}$ as $\mathtt{x}$ with IND$-CPA scheme

2. For each bit $\mathtt{x}_\mathtt{i}$ of the ciphertext:

    a. Sample random $\mathtt{c}_\mathtt{i}$ from covertext distribution

    b. If $\mathtt{h(c_i)=x_i}$ (where $\mathtt{h}$ is an unbiased hash function):

    - Yes: append $\mathtt{c}_\mathtt{i}$ to the stegotext, and proceed to next $\mathtt{x}_\mathtt{i}$

    - No: return to (a)

Universal Steganography Barriers

Encode Message

1. Encrypt message m as x with IND$-CPA scheme

2. For each bit x_i of the ciphertext:

   a. Sample random c_i using <u>GENERATIVE MODEL</u>

   b. If $h(c_i)=x_i$ (where h is an <u>unbiased</u> hash function):

      – Yes: append c_i to the stegotext, and proceed to next x_i

      – No: return to (a)

Universal Steganography Barriers

Encode Message

1. Encrypt message m as x with IND\$-CPA scheme

2. For each bit x_i of the ciphertext:

   a. Sample random c_i using GENERATIVE MODEL

   b. If $h(c_i)=x_i$ (where h is an unbiased hash function):

   - Yes: append c_i to the stegotext, and proceed to next x_i

   - No: return to (a)

Universal Steganography Barriers

Encode Message

1. Encrypt message m as x with IND$-CPA scheme
2. For each bit x_i of the ciphertext:
   a. Sample random c_i using <u>GENERATIVE MODEL</u>
   b. If $h(c_i)=x_i$ (where h is a <u>cryptographic</u> hash function):
      — Yes: append c_i to the stegotext, and proceed to next x_i
      — No: return to (a)

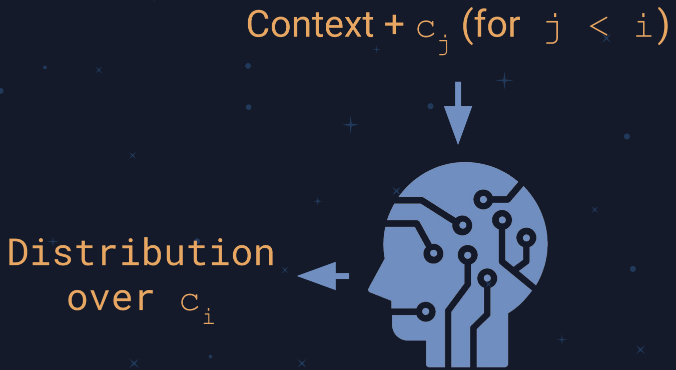

Universal Steganography Barriers

Encode Message

1. Encrypt message m as x with IND$-CPA scheme
2. For each bit x_i of the ciphertext:
   a. Sample random c_i using GENERATIVE MODEL
   b. If $h(c_i)=x_i$ (where h is a cryptographic hash function):
      – Yes: append c_i to the stegotext, and proceed to next x_i
      – No: return to (a)

Context + c_j (for $j < i$)

Distribution over c_i

Might introduce bias over low entropy distributions of c_i

Instantaneous Entropy Over GPT-2

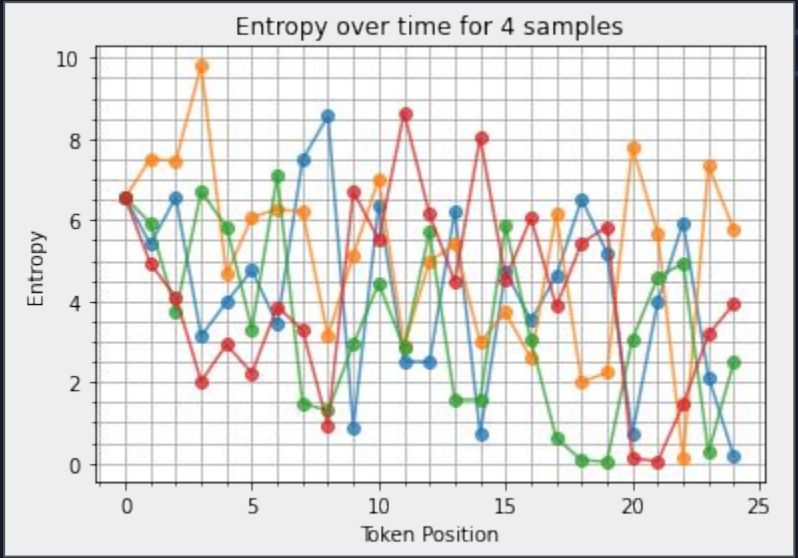

Adaptation Options

1. Skip Low Distribution Moments

Adaptation Options

1. Skip Low Distribution Moments

- Model is public information
- Entropy is public information
- Skip all low entropy sampling events (eg. Entropy < 4.5)

Adaptation Options

1. Skip Low Distribution Moments

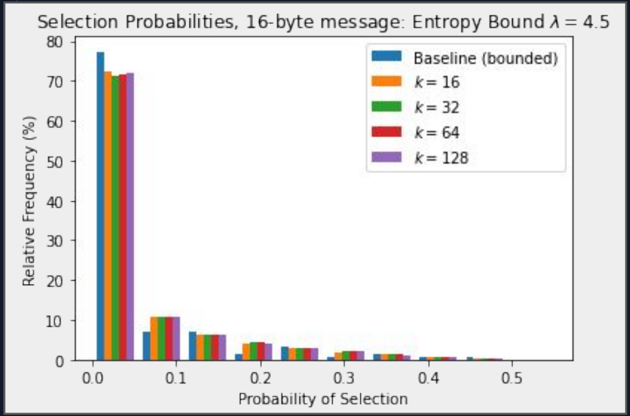

Adaptation Options

1. Skip Low Distribution Moments

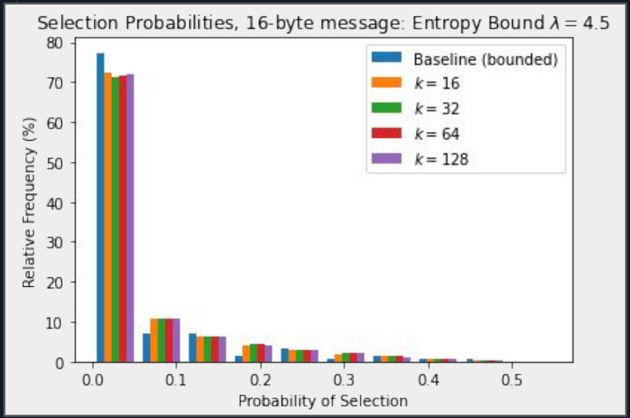

2. Accumulate Entropy

- Compile channel such that it has sufficient entropy
- Sample many tokens together

Adaptation Options

1. Skip Low Distribution Moments

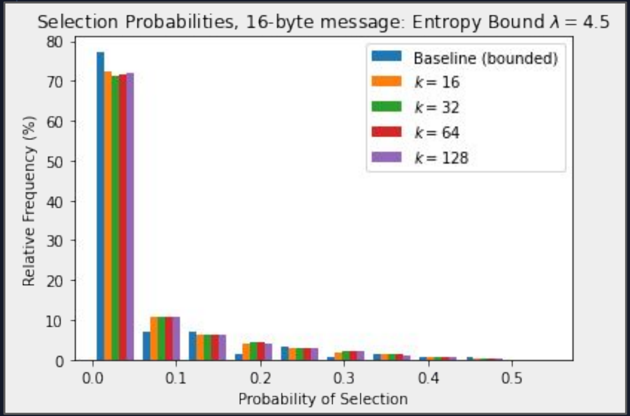

2. Accumulate Entropy

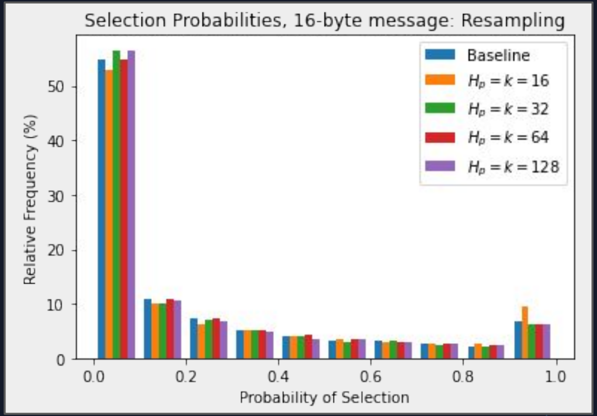

Performance When Accumulating Entropy

Parameters	Samples (Tokens)	Time (Sec)	Stegotext Len. (KiB)	Overhead (Length)
$H_p = k = 16$	502.8	42.69	2.3	149.4x
$H_p = k = 32$	880.4	128.41	4.1	261.8x
$H_p = k = 64$	1645.0	361.28	7.5	482.1x
$H_p = k = 128$	2994.6	765.40	13.6	870.7x

Talk Outline

01 Steganography Refresher

02 Classical Schemes + Generative Models

03 METEOR: Dealing with Low Entropy

Can We Do Better In The Symmetric Key Setting?

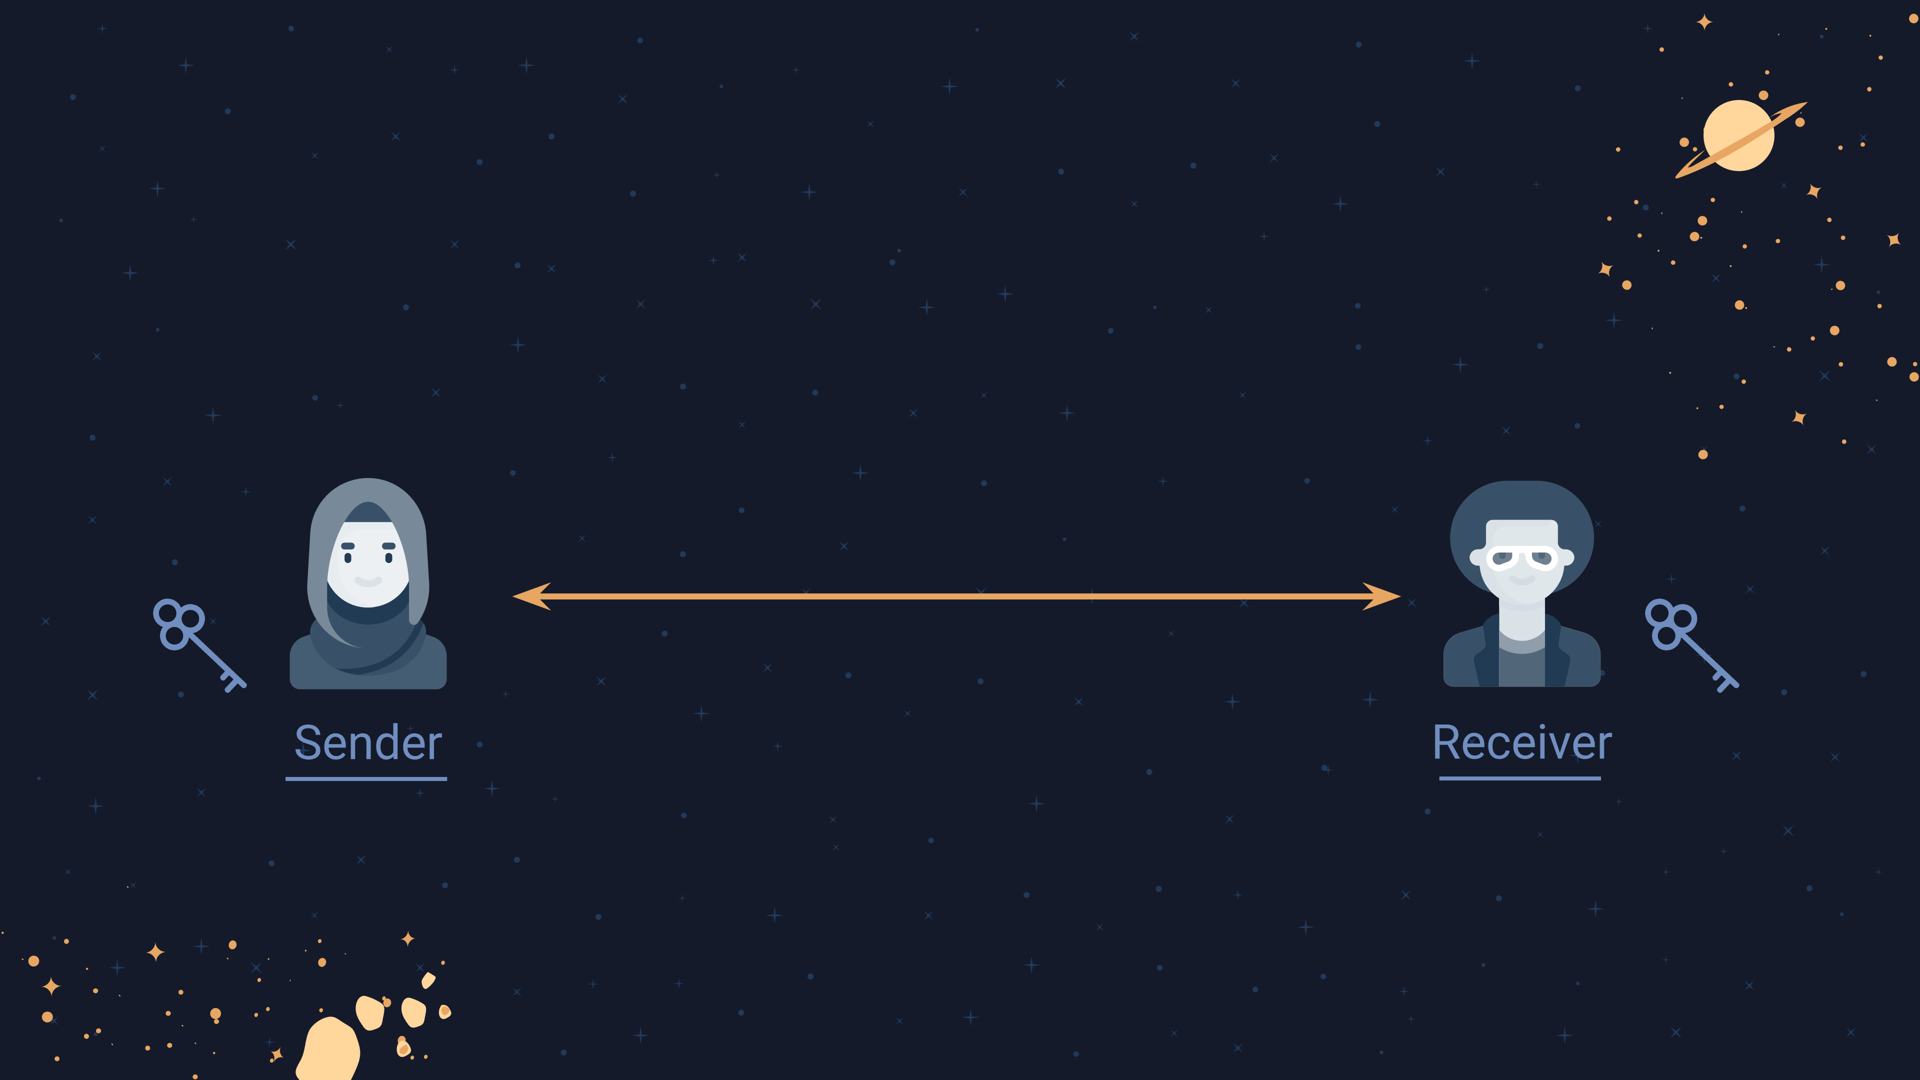
Sender
Receiver


Sender
Receiver

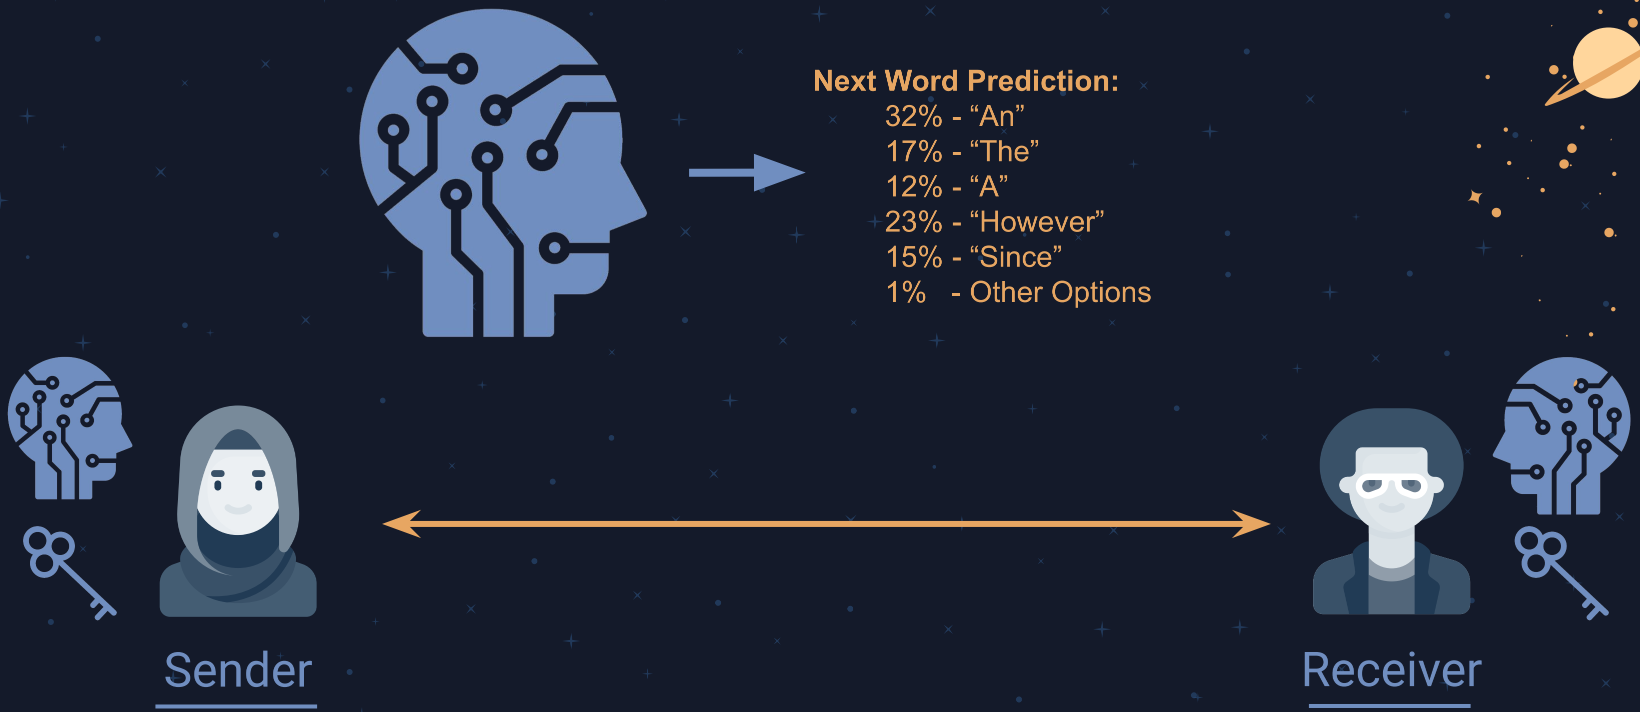
Next Word Prediction:
32% - “An”
17% - “The”
12% - “A”
23% - “However”
15% - “Since”
1% - Other Options
Sender
Receiver

Encoding Intuition

Context

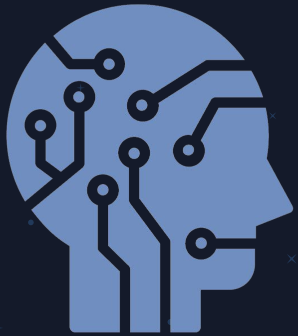

Next Word Prediction:

32% - “An”
17% - “The”
12% - “A”
23% - “However”
15% - “Since”
1% - Other Options

Encrypted Message
(as bits)

“The”

Decoding Intuition

Context

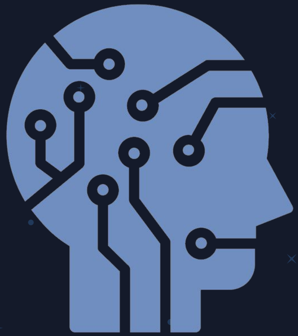

Next Word Prediction:

32% - “An”
17% - “The”
12% - “A”
23% - “However”
15% - “Since”
1% - Other Options

Encrypted Message
(as bits)

“The”

Encoding Intuition

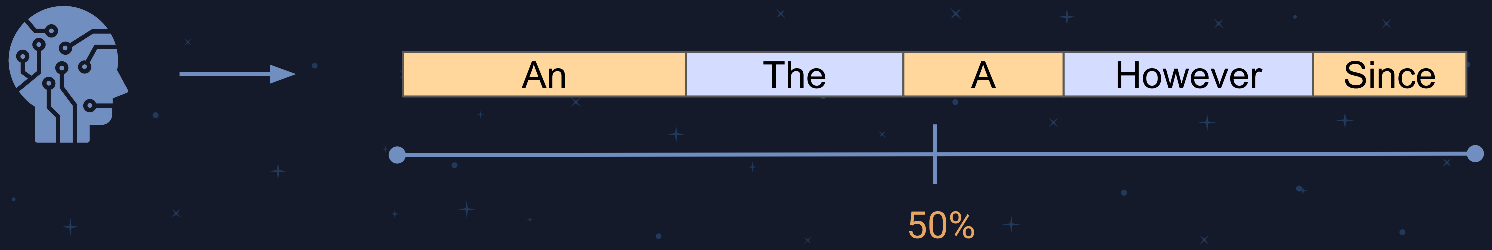

Encoding Intuition

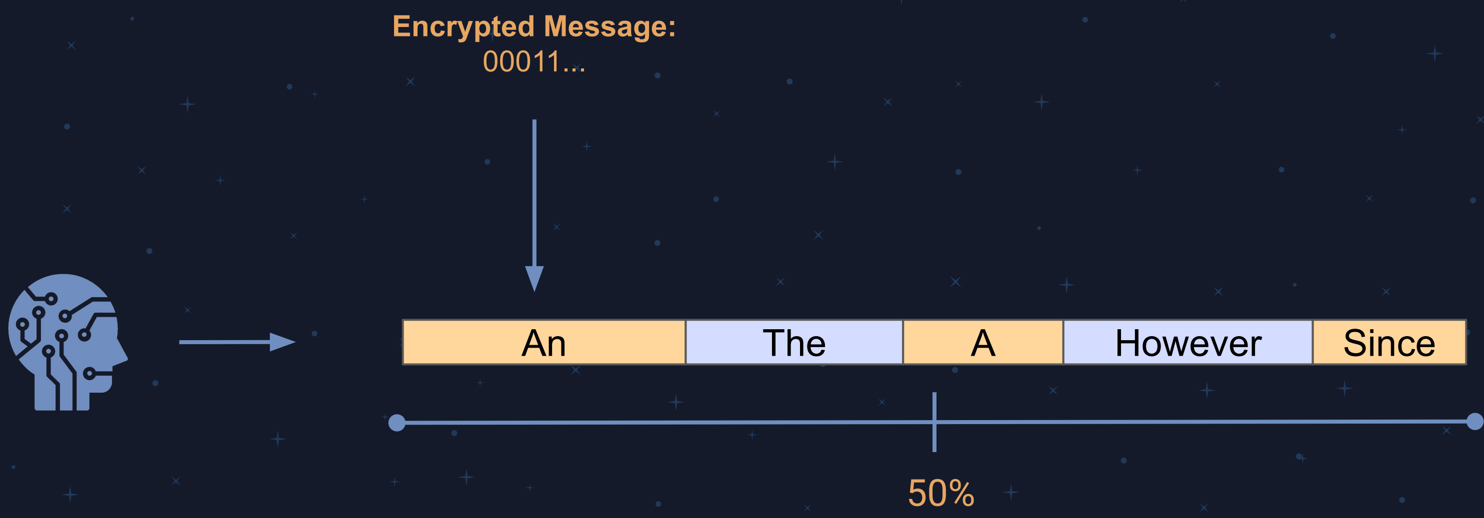

Encoding Intuition

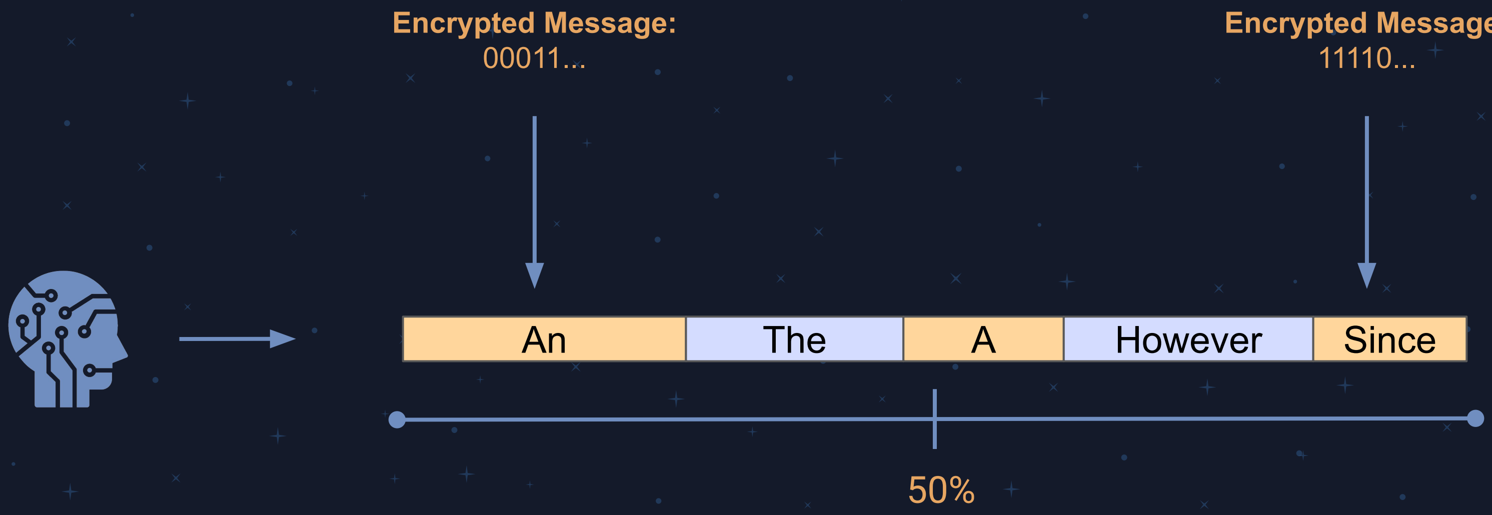

Encoding Intuition

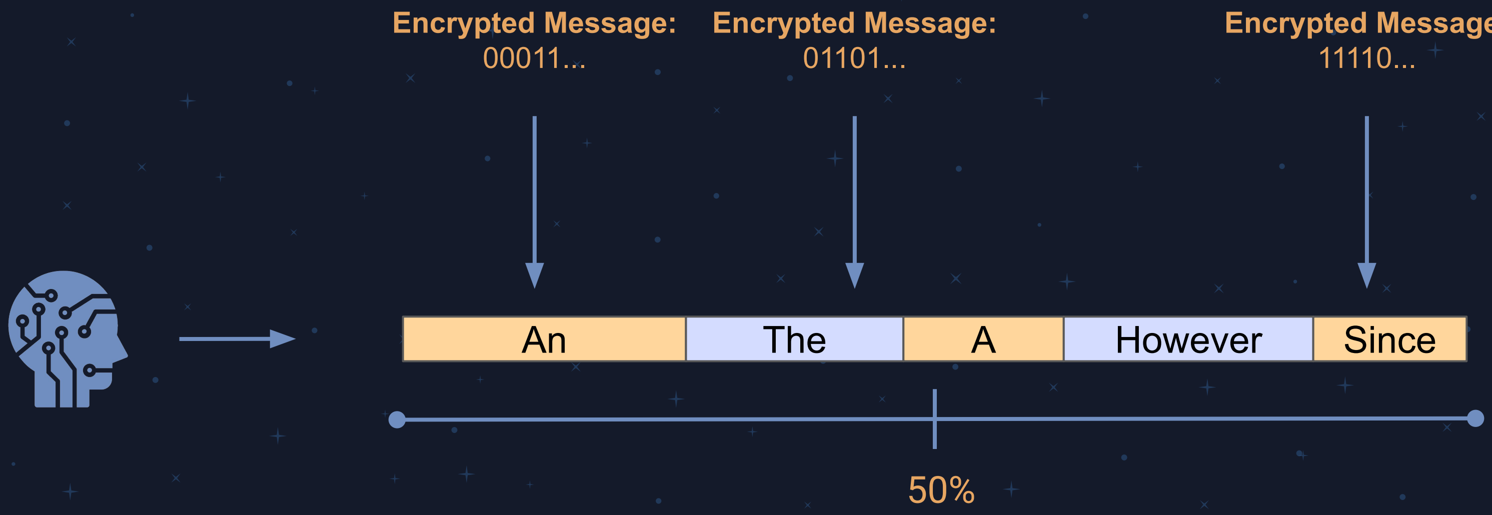

Decoding Intuition

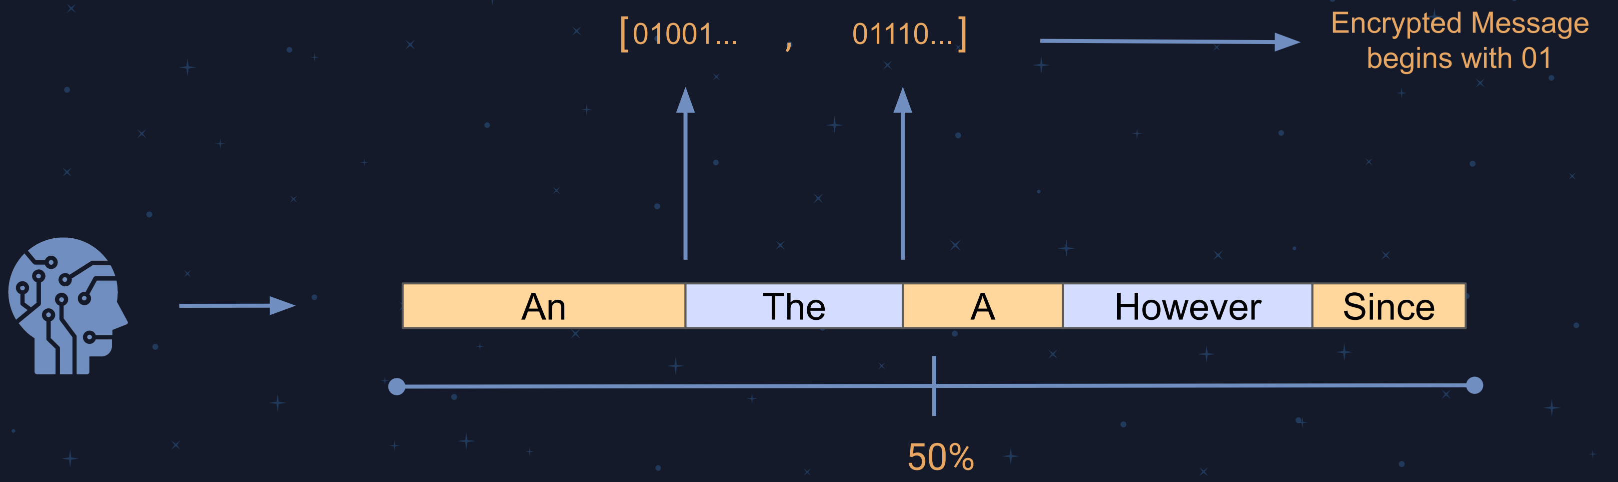

Decoding Intuition

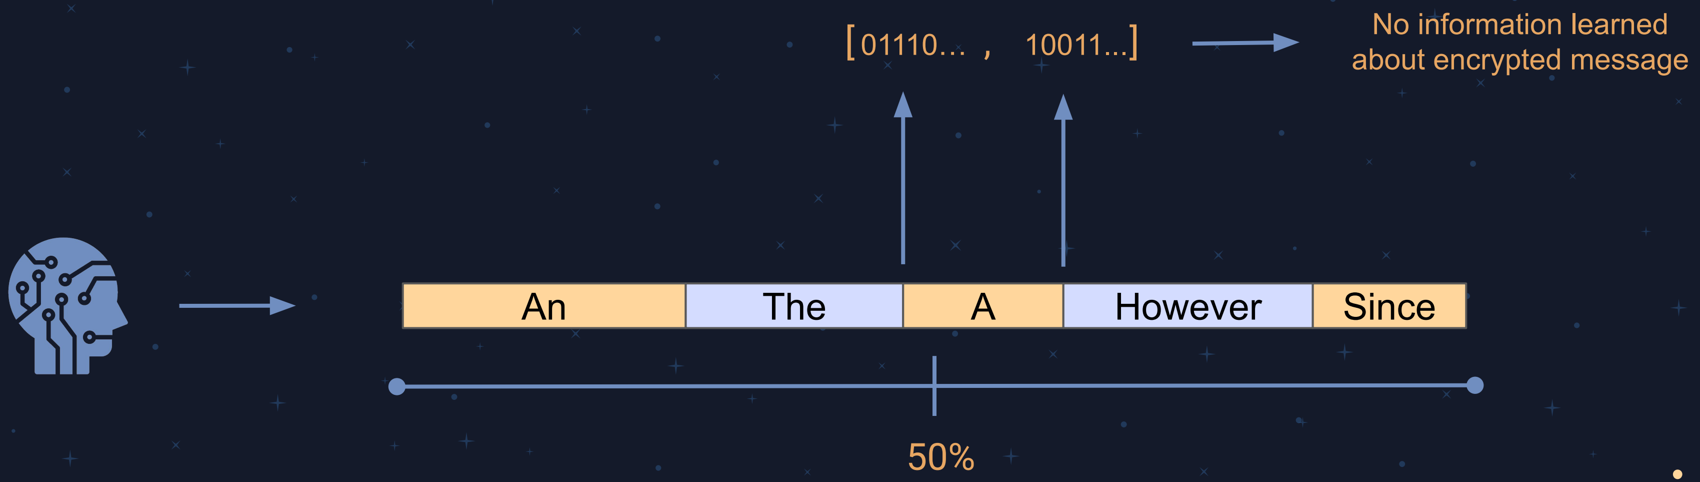

METEOR

Encode Message

1. While message not fully encoded:

METEOR

Encode Message

1. While message not fully encoded:

    a. Sample and apply random mask (from PRG)

METEOR

Encode Message

1. While message not fully encoded:
    a. Sample and apply random mask (from PRG)
    b. Sample distribution for next c_i from model

METEOR

Encode Message

1. While message not fully encoded:

    a. Sample and apply random mask (from PRG)

    b. Sample distribution for next c_i from model

    c. Use masked message to determine c_i

METEOR

Encode Message

1. While message not fully encoded:
    a. Sample and apply random mask (from PRG)
    b. Sample distribution for next c_i from model
    c. Use masked message to determine c_i
    d. Compute number of bits transferred

METEOR

Encode Message

1. While message not fully encoded:

    a. Sample and apply random mask (from PRG)

    b. Sample distribution for next c_i from model

    c. Use masked message to determine c_i

    d. Compute number of bits transferred

    e. Mark transferred bits as encoded and add c_i to message

METEOR

Encode Message

1. While message not fully encoded:
    a. Sample and apply random mask (from PRG)
    b. Sample distribution for next c_i from model
    c. Use masked message to determine c_i
    d. Compute number of bits transferred
    e. Mark transferred bits as encoded and add c_i to message

Decode Message

1. While message not fully decoded:
    a. Sample distribution for next c_i from model
    b. Compute number of bits transferred by c_i
    c. Sample and apply random mask (from PRG)
    d. Mark transferred bits as encoded and add recovered bits to message

Mode	Desktop/GPU (sec)	Laptop/CPU (sec)	Stegotext Length (bytes)	Overhead (length)	Capacity (bits/token)
GPT-2	18.089	82.214	1976	12.36×	3.09
GPT-2 (Reorder)	30.570	82.638	1391	8.69×	4.11
GPT-2 (Compress)	11.070	42.942	938	3.39×	3.39
Wikipedia	19.791	46.583	2002	12.51×	0.64
Wikipedia (Reorder)	15.515	39.450	1547	9.67×	0.83
HTTP Headers	49.380	103.280	6144	38.4×	0.21
HTTP Headers (Reorder)	57.864	127.759	7237	45.23×	0.18

Device	Load	Encode	Decode	Overhead (time)
GPU	5.867	6.899	6.095	1×
CPU	5.234	41.221	40.334	4.6×
Mobile	1.830	473.58	457.57	49.5×

Benefits of Meteor’s Approach

1. Implicit Adjustment

Encoding rate is asymptotically equal to entropy

Benefits of Meteor’s Approach

1. Implicit Adjustment

Encoding rate is asymptotically equal to entropy

2. Concretely Efficient Enough to Really Run In Practice

Implemented and benchmarked run on GPU, CPU, and Mobile

Benefits of Meteor’s Approach

1. Implicit Adjustment

Encoding rate is asymptotically equal to entropy

2. Concretely Efficient Enough to Really Run In Practice

Implemented and benchmarked run on GPU, CPU, and Mobile

3. Clear Security Analysis

Straightforward reduction to security of PRG

Other Parts of Our Work

Comparison to Prior (Informal) Work

Ad-hoc Optimizations For Performance

Easy-to-use Code Demo on Google Co-Lab

Thanks!

ia.cr/2021/686 `meteorfrom.space`

Gabriel Kaptchuk (Boston University)

Tushar Jois, Matthew Green, Aviel Rubin

(Johns Hopkins University)