Multi-View Channels and Applications

Arvind Rameshwar

Dept. of EE, IIT Madras

JTG ITSoc Summer School 2026

Joint work with

Nir Weinberger

Assistant Professor
Technion, Israel

Itay Geva

M.Sc. Candidate
Technion, Israel

What is the talk about?

The decoder obtains d noisy views of an input sequence.

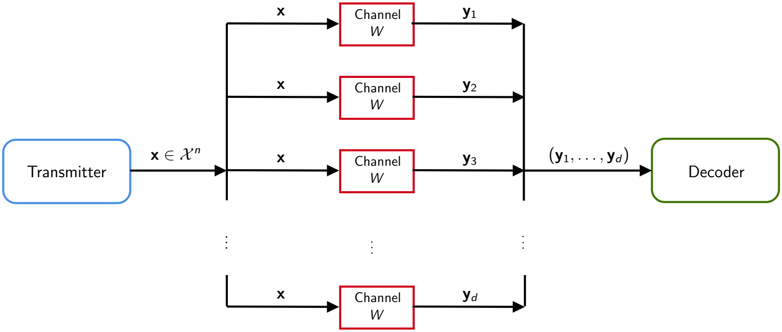

d views

What is the talk about?

The decoder obtains d noisy views of an input sequence.

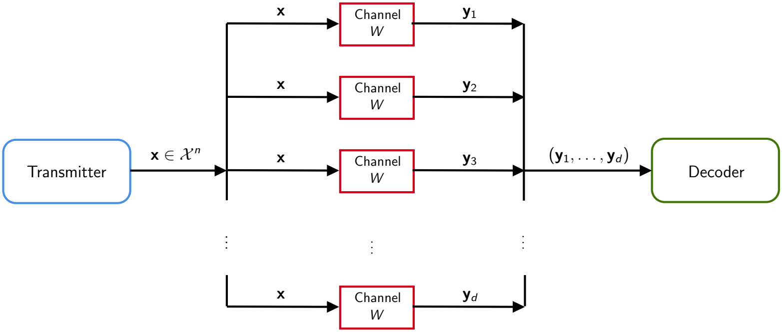

d views

Q1: What rates are achievable with vanishing reconstruction error probability?

Q2: Can we use **A1** to obtain achievable rates over practical channel models for modern storage media?

Where does this problem arise? - I

- Short molecules in a DNA pool are “amplified” by Polymerase Chain Reaction (PCR) and sampled multiple times [Shomorony and Heckel (2022)].

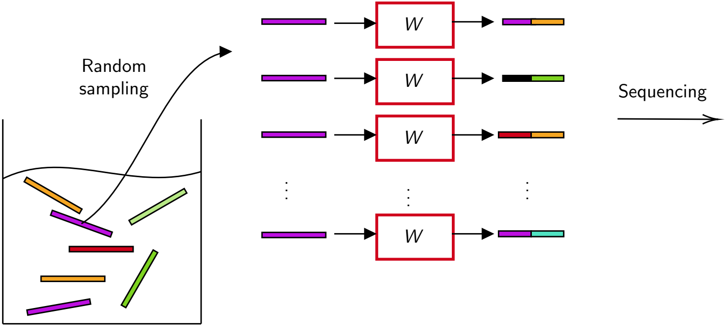

Where does this problem arise? - I

- Short molecules in a DNA pool are "amplified" by Polymerase Chain Reaction (PCR) and sampled multiple times [Shomorony and Heckel (2022)].

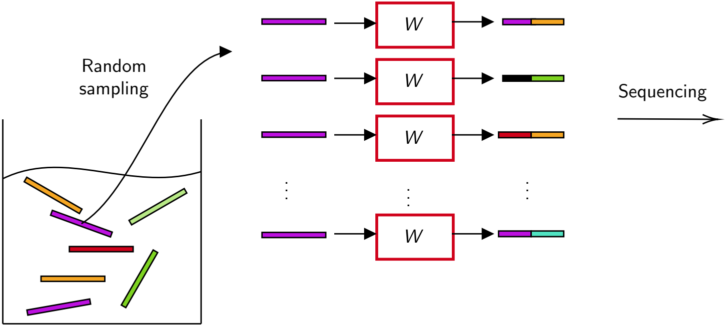

- The capacity is determined by mutual information terms corresponding to different # of views of the input sequence; see [Lenz et al. (2019, 2020), Shomorony and Heckel (2022), Weinberger and Merhav (2022)].

Where does this problem arise? - II

- Errors due to synchronization in packet-switched communications and in reading magneto-optical media

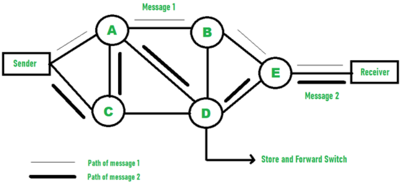

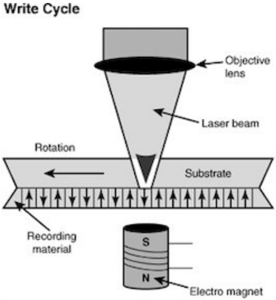

- Each output "run" at the end of (noisy) duplications corresponds to a multi-view channel [Mitzenmacher (2008), Cheraghchi and Ribeiro (2019)]

Where does this problem arise? - III

- More fundamentally, in the iterative decoding of LDPC codes [Gallager (1962), Richardson and Urbanke (2001)]

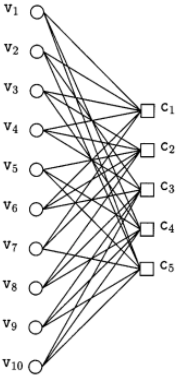

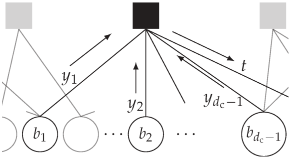

- A single variable node receives multiple "views"/estimates of its value, from different check nodes.

Answering **Q1** in the large view limit

A simpler setting: Multi-view DMCs

The setting, revisited

The decoder obtains d independent, noisy views of an input symbol.

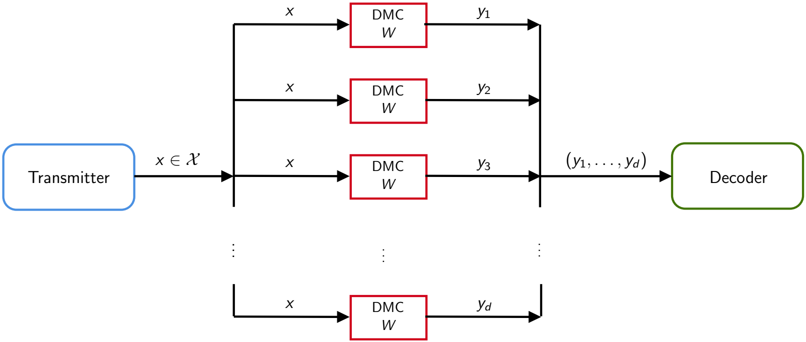

d views

The setting, revisited

The decoder obtains d independent, noisy views of an input symbol.

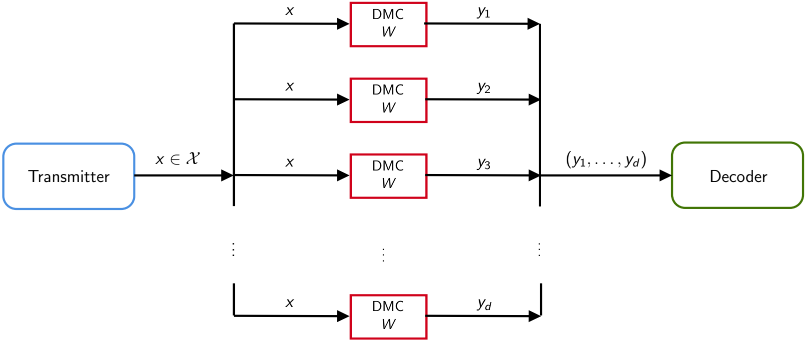

d views

Since a multi-view DMC is also a DMC, it suffices for us to focus on the transmission of a single symbol, for rate computations.

The setting, revisited

The decoder obtains d independent, noisy views of an input symbol.

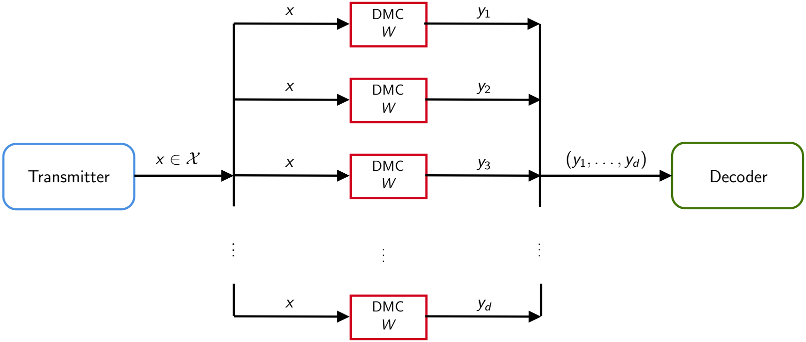

d views

Goal: Exact asymptotics of the mutual info. (+ capacity) and dispersion of such a multi-view channel, for arbitrary P_X.

What is known about this setting?

- Hellman and Raviv (1970), Kanaya and Han (1995): Exact asymptotics of information rates over DMCs with multiple views
- Levenshtein (2001): Characterization of # of views for exact reconstruction over comb. error channels and for reconstruction with decaying error prob. over multi-view DMCs

What is known about this setting?

- Hellman and Raviv (1970), Kanaya and Han (1995): Exact asymptotics of information rates over DMCs with multiple views
- Levenshtein (2001): Characterization of # of views for exact reconstruction over comb. error channels and for reconstruction with decaying error prob. over multi-view DMCs
- Land et al. (2005), Land and Huber (2006): Bounds on the mutual information rates over multi-view DMCs via information combining
- Mitzenmacher (2006): Calculation of the capacity of a multi-view binary symmetric channel (BSC)
- Haeupler and Mitzenmacher (2014): Information rates over multi-view deletion channels when Pr[deletion] → 0.
- Batu et al. (2004), Cheraghchi et al. (2020), Peres and Zhai (2017): Trace reconstruction over the deletion channel

What is known about this setting?

- Hellman and Raviv (1970), Kanaya and Han (1995): Exact asymptotics of information rates over DMCs with multiple views
- Levenshtein (2001): Characterization of # of views for exact reconstruction over comb. error channels and for reconstruction with decaying error prob. over multi-view DMCs
- Land et al. (2005), Land and Huber (2006): Bounds on the mutual information rates over multi-view DMCs via information combining
- Mitzenmacher (2006): Calculation of the capacity of a multi-view binary symmetric channel (BSC)
- Haeupler and Mitzenmacher (2014): Information rates over multi-view deletion channels when $\Pr[\text{deletion}] \to 0$.
- Batu et al. (2004), Cheraghchi et al. (2020), Peres and Zhai (2017): Trace reconstruction over the deletion channel

Our contributions:

1. Unified treatment of info. rate and channel dispersion

$$\Downarrow$$

Finite-blocklength achievable rates with fixed error probability

What is known about this setting?

- Hellman and Raviv (1970), Kanaya and Han (1995): Exact asymptotics of information rates over DMCs with multiple views
- Levenshtein (2001): Characterization of # of views for exact reconstruction over comb. error channels and for reconstruction with decaying error prob. over multi-view DMCs
- Land et al. (2005), Land and Huber (2006): Bounds on the mutual information rates over multi-view DMCs via information combining
- Mitzenmacher (2006): Calculation of the capacity of a multi-view binary symmetric channel (BSC)
- Haeupler and Mitzenmacher (2014): Information rates over multi-view deletion channels when Pr[deletion] $\to 0$.
- Batu et al. (2004), Cheraghchi et al. (2020), Peres and Zhai (2017): Trace reconstruction over the deletion channel

Our contributions:

2. Directly extensible proofs for multi-letter channels

Background

Some formalism

- Consider a DMC W with input alphabet $\mathcal{X}$ and output alphabet $\mathcal{Y}$, both finite. Assume that $|\mathcal{X}|$, $|\mathcal{Y}|$ do not depend on d.
- The d-view DMC $W^{(d)}$ obeys the channel law

$$W^{(d)}((y_1, \ldots, y_d) \mid x) = \prod_{i=1}^{d} W(y_i \mid x),$$

for $(y_1, \ldots, y_d) \in \mathcal{Y}^d$ and $x \in \mathcal{X}$.

Some formalism

- Consider a DMC W with input alphabet $\mathcal{X}$ and output alphabet $\mathcal{Y}$, both finite. Assume that $|\mathcal{X}|$, $|\mathcal{Y}|$ do not depend on d.
- The d-view DMC $W^{(d)}$ obeys the channel law

  $$W^{(d)}((y_1, \ldots, y_d) \mid x) = \prod_{i=1}^{d} W(y_i \mid x),$$

  for $(y_1, \ldots, y_d) \in \mathcal{Y}^d$ and $x \in \mathcal{X}$.
- Fix a d-independent input distribution P_X. We are interested in:

  $$I^{(d)} = I(X; Y^d) = \mathrm{H}(X) - \mathrm{H}(X \mid Y^d) \qquad \text{[Mutual info.]}$$

  and

  $$V^{(d)} = \mathbb{E}\left[(\iota(X; Y^d) - I^{(d)})^2\right], \qquad \text{[Channel dispersion]}$$

  where

  $$\iota(X; Y^d) = \log \frac{P(Y^d \mid X)}{P(Y^d)}. \qquad \text{[Info. spectrum]}$$

Building intuition: Multi-view BSC$^{(d)}$

Consider the d-view BSC(p), with binary input $X \in \{0, 1\}$:

$$Y = X + Z \pmod 2,$$

where $Z \sim \text{Ber}(p)$.

Building intuition: Multi-view BSC$^{(d)}$

Consider the d-view BSC(p), with binary input $X \in \{0, 1\}$:

$$Y = X + Z \ (\text{mod } 2),$$

where $Z \sim \text{Ber}(p)$. Suppose that $X \sim \text{Ber}(1/2)$.

Building intuition: Multi-view BSC$^{(d)}$

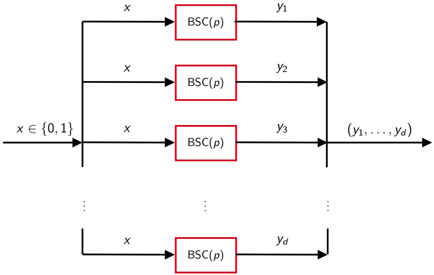

- Given outputs $Y^d = (Y_1, \ldots, Y_d)$, compute

$$M = \texttt{Majority}(Y^d).$$

Building intuition: Multi-view BSC$^{(d)}$

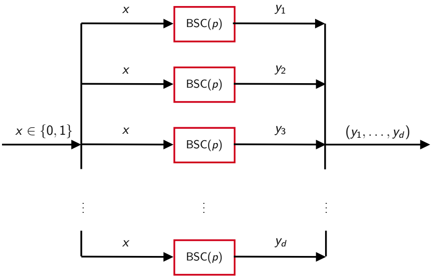

- Given outputs $Y^d = (Y_1, \ldots, Y_d)$, compute

$$M = \texttt{Majority}(Y^d).$$

- Then,

$$\begin{aligned} I(X; Y^d) &\geq I(X; M) \\ &= \log 2 - h_b(\Pr[X \neq M]). \end{aligned}$$

Building intuition: Multi-view BSC$^{(d)}$

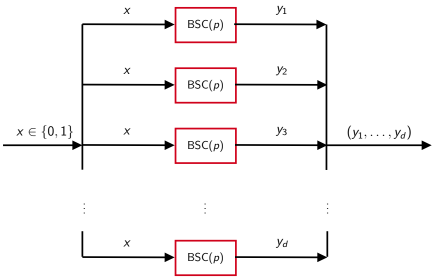

- Given outputs $Y^d = (Y_1, \ldots, Y_d)$, compute

$$M = \texttt{Majority}(Y^d).$$

- Then,

$$\begin{aligned} I(X; Y^d) &\geq I(X; M) \\ &= \log 2 - h_b(\Pr[X \neq M]). \end{aligned}$$

- By the Chernoff bound,

$$\Pr[X \neq M] \leq \exp(-dZ(p)),$$

where $Z(p) = 2\sqrt{p(1-p)}$.

Building intuition: Multi-view BSC$^{(d)}$

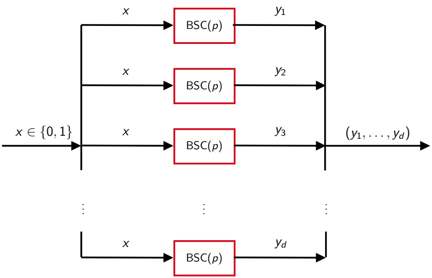

- Given outputs $Y^d = (Y_1, \ldots, Y_d)$, compute

$$M = \texttt{Majority}(Y^d).$$

- Then,

$$\begin{aligned} I(X; Y^d) &\geq I(X; M) \\ &= \log 2 - h_b(\Pr[X \neq M]). \end{aligned}$$

- By the Chernoff bound,

$$\Pr[X \neq M] \leq \exp(-dZ(p)),$$

where $Z(p) = 2\sqrt{p(1-p)}$.

Hence, for large d:

$$\begin{aligned} h_b(\Pr[X \neq M]) &\leq -2 \Pr[X \neq M] \cdot \log \Pr[X \neq M] \\ &= 2dZ(p) \cdot \exp(-dZ(p)). \end{aligned}$$

Building intuition: Multi-view BSC$^{(d)}$

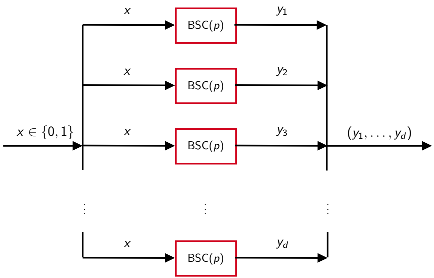

- Given outputs $Y^d = (Y_1, \dots, Y_d)$, compute

$$M = \texttt{Majority}(Y^d).$$

- Then,

$$\begin{aligned} I(X; Y^d) &\geq I(X; M) \\ &= \log 2 - h_b(\Pr[X \neq M]). \end{aligned}$$

- By the Chernoff bound,

$$\Pr[X \neq M] \leq \exp(-dZ(p)),$$

where $Z(p) = 2\sqrt{p(1-p)}$.

This gives us

$$I(X; Y^d) \geq \log 2 - 2dZ(p) \cdot \exp(-dZ(p)).$$

Exponentially fast convergence to $H(X) = \log 2$

Our Results

Our main result

- Intuitively, as d becomes large, we expect $I^{(d)} \approx \mathsf{H}(X)$ and

$$V^{(d)} \approx \underbrace{\mathsf{V}(X)}_{\text{i/p varentropy}} = \mathbb{E}\left[(\log P(X) + \mathsf{H}(X))^2\right].$$

Our main result

- Intuitively, as d becomes large, we expect $I^{(d)} \approx \mathsf{H}(X)$ and

$$V^{(d)} \approx \underbrace{\mathsf{V}(X)}_{\text{i/p varentropy}} = \mathbb{E}\left[(\log P(X) + \mathsf{H}(X))^2\right].$$

- For distributions P, Q on $\mathcal{X}$, define the Chernoff information

$$\mathsf{C}(P, Q) = -\min_{\lambda \in [0,1]} \log \left(\sum_{x \in \mathcal{X}} P(x)^{1-\lambda} Q(x)^{\lambda} \right).$$

Our main result

- Intuitively, as d becomes large, we expect $I^{(d)} \approx \mathsf{H}(X)$ and

$$V^{(d)} \approx \underbrace{\mathsf{V}(X)}_{\text{i/p varentropy}} = \mathbb{E}\left[(\log P(X) + \mathsf{H}(X))^2\right].$$

- For distributions P, Q on $\mathcal{X}$, define the Chernoff information

$$\mathsf{C}(P,Q) = -\min_{\lambda \in [0,1]} \log\left(\sum_{x \in \mathcal{X}} P(x)^{1-\lambda} Q(x)^{\lambda}\right).$$

Theorem

When $\mathcal{X}, \mathcal{Y}, P_X$ do not depend on d,

$$I^{(d)} = \mathsf{H}(X) - \exp\left(-d\rho + \Theta(\log d|\mathcal{X}|)\right), \text{ and}$$

$$\left|V^{(d)} - \mathsf{V}(X)\right| = \exp\left(-d\rho + \Theta(\log d|\mathcal{X}|)\right),$$

where

$$\rho = \min_{x,x' :\ x \neq x'} \mathsf{C}(P_{Y|x}, P_{Y|x'}).$$

Interpreting the result

- The rate of convergence ρ of the mutual information and channel dispersion to their limits

  $$\rho = \min_{x, x' :\ x \neq x'} \mathsf{C}(P_{Y|x}, P_{Y|x'})$$

  is independent of the input distribution (except via its support)!

Interpreting the result

- The rate of convergence ρ of the mutual information and channel dispersion to their limits

  $$\rho = \min_{x,x' \colon\, x \neq x'} \mathsf{C}(P_{Y|x}, P_{Y|x'})$$

  is independent of the input distribution (except via its support)!

- For a binary-input memoryless symmetric (BIMS) channel W, the rate

  $$\begin{aligned} \rho &= -\log \sum_{y \in \mathcal{Y}} \sqrt{P_{Y|0}(y \mid 0) P_{Y|1}(y \mid 1)} \\ &= -\log Z_b(W), \end{aligned}$$

  where $Z_b(W)$ is the Bhattacharya parameter of the BIMS W.

Interpreting the result

- The rate of convergence ρ of the mutual information and channel dispersion to their limits

  $$\rho = \min_{x,x'\colon\ x \neq x'} \mathsf{C}(P_{Y|x}, P_{Y|x'})$$

  is independent of the input distribution (except via its support)!

- For a binary-input memoryless symmetric (BIMS) channel W, the rate

  $$\begin{aligned}\rho &= -\log \sum_{y \in \mathcal{Y}} \sqrt{P_{Y|0}(y \mid 0) P_{Y|1}(y \mid 1)} \\ &= -\log Z_b(W),\end{aligned}$$

  where $Z_b(W)$ is the Bhattacharya parameter of the BIMS W.

Hence, our earlier speed of convergence for the $\mathrm{BSC}^{(d)}$ is tight!

A finite blocklength corollary

- A characterization of finite-blocklength rates achievable over $W^{(d)}$, thus follows.
- For a fixed $\epsilon \in (0,1)$ and blocklength $n \geq 1$, let

$$M^{\star}(n,\epsilon) \leftarrow \text{largest } M \text{ s.t. } \exists \text{ length-}n \text{ code over } W^{(d)} \\ \text{with max. error } \epsilon.$$

A finite blocklength corollary

- A characterization of finite-blocklength rates achievable over $W^{(d)}$, thus follows.
- For a fixed $\epsilon \in (0,1)$ and blocklength $n \geq 1$, let

$$M^\star(n,\epsilon) \leftarrow \text{largest } M \text{ s.t. } \exists \text{ length-}n \text{ code over } W^{(d)} \text{ with max. error } \epsilon.$$

Theorem

If W is "non-singular", we have

$$\log|\mathcal{X}| - \frac{\log M^\star(n,\epsilon)}{n} \leq e^{-d\rho+\Theta(\log d|\mathcal{X}|)} - \Phi^{-1}(\epsilon)\cdot\frac{e^{-d\rho/2+\Theta(\log d|\mathcal{X}|)}}{\sqrt{n}} + \Theta\left(\frac{\log n}{n}\right).$$

A finite blocklength corollary

- A characterization of finite-blocklength rates achievable over $W^{(d)}$, thus follows.
- For a fixed $\epsilon \in (0,1)$ and blocklength $n \geq 1$, let

$$M^\star(n,\epsilon) \leftarrow \text{largest } M \text{ s.t. } \exists \text{ length-}n \text{ code over } W^{(d)}$$
$$\text{with max. error } \epsilon.$$

Theorem

If W is "non-singular", we have

$$\log|\mathcal{X}| - \frac{\log M^\star(n,\epsilon)}{n}$$
$$\leq e^{-d\rho+\Theta(\log d|\mathcal{X}|)} - \Phi^{-1}(\epsilon)\cdot\frac{e^{-d\rho/2+\Theta(\log d|\mathcal{X}|)}}{\sqrt{n}} + \Theta\left(\frac{\log n}{n}\right).$$

- In particular, choosing $d = \rho^{-1}\log n$, we can achieve rates

$$\boxed{R_{n,\epsilon} \geq \log|\mathcal{X}| - O_\epsilon\left(\frac{\log n}{n}\right).}$$

An Application: Achievable Rates Over Noisy Nanopore Channels

What is the talk about?

- We are interested in nanopore-based DNA sequencing – a third generation technology in the sequencing of biopolymers.

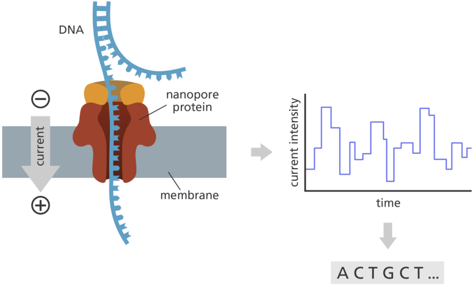

Image credit: Laura Olivares Boldú, Wellcome Connecting Science

- A nanopore sequencer can read strands of length 10–100 Kilo-bases, which is significantly larger than the read capability of other sequencers (≈ 100 bases).

How does a nanopore sequencer work?

- The bases of (single-stranded) DNA are "ratcheted" (translocated) through the nanopore, which is a microscopic pore.

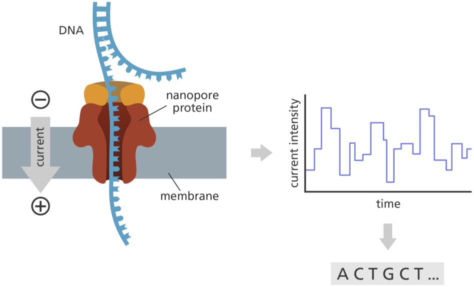

Image credit: Laura Olivares Boldú, Wellcome Connecting Science

How does a nanopore sequencer work?

- The bases of (single-stranded) DNA are “ratcheted” (translocated) through the nanopore, which is a microscopic pore.

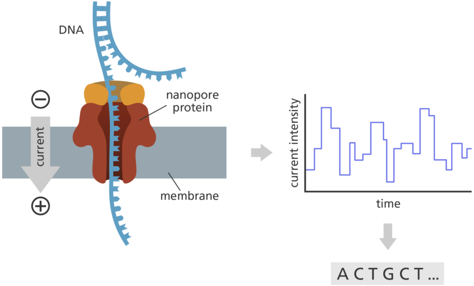

Image credit: Laura Olivares Boldú, Wellcome Connecting Science

- This traversal of bases creates electric current flow – a function of the bases in the nanopore – which can then be measured.

- Predictably, such measurements are error prone.

Error sources and modelling

The following types of errors corrupt nanopore reads:

[Mao, Diggavi, Kannan (2018); Hulett, Chandak, Wootters (2021); Hamoum et al. (2021); McBain, Viterbo, Saunderson (2024)]

1. Intersymbol interference (ISI): non-trivial nanopore width
2. "Backtracking" / "skipping": imperfect translocation speed
3. Measurement noise: random noise in current readings

Error sources and modelling

The following types of errors corrupt nanopore reads:

[Mao, Diggavi, Kannan (2018); Hulett, Chandak, Wootters (2021); Hamoum et al. (2021); McBain, Viterbo, Saunderson (2024)]

1. Intersymbol interference (ISI): non-trivial nanopore width
2. "Backtracking" / "skipping": imperfect translocation speed
3. Measurement noise: random noise in current readings

These errors can be modelled by the following abstractions:

1. Intersymbol interference (ISI): Block-Markov inputs
2. "Backtracking" / "skipping": Insertion/deletion errors
3. Measurement noise: Memoryless noise

Channel model

We work with the channel model in [McBain, Viterbo, Saunderson (2024)]:

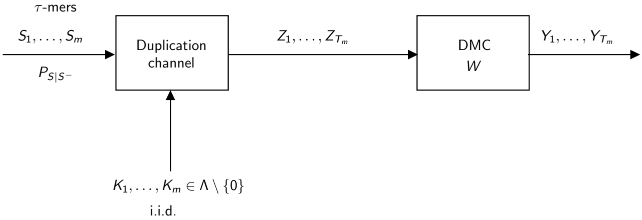

Channel model

We work with the channel model in [McBain, Viterbo, Saunderson (2024)]:

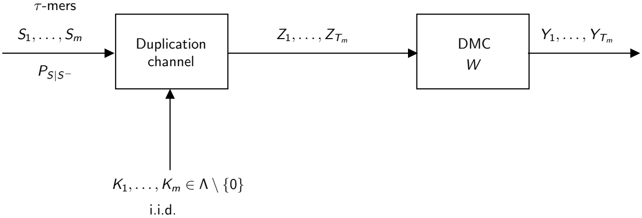

Quick description:

1. $(S_1, \dots, S_m)$: "de-Bruijn"-Markov input sequence of contiguous length-τ subsequences of input bases $B_1, \dots, B_{m+\tau-1} \in \mathcal{X}$
2. Each symbol S_i is repeated K_i times, leading to Z^{T_m}, where $T_m = \sum_{i \leq m} K_i$.
3. Each Z_i is corrupted (independently) by W, giving rise to Y^{T_m}.

Channel model

We work with the channel model in [McBain, Viterbo, Saunderson (2024)]:

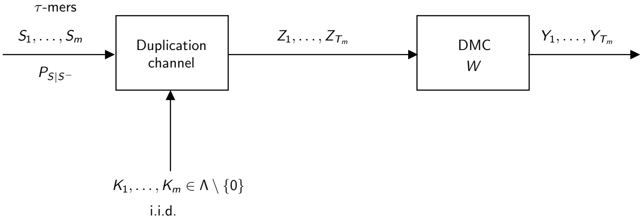

Quick description:

1. $(S_1, \ldots, S_m)$: "de-Bruijn"-Markov input sequence of contiguous length-τ subsequences of input bases $B_1, \ldots, B_{m+\tau-1} \in \mathcal{X}$
2. Each symbol S_i is repeated K_i times, leading to Z^{T_m}, where $T_m = \sum_{i \leq m} K_i$.
3. Each Z_i is corrupted (independently) by W, giving rise to Y^{T_m}.

Q: What information rates are achievable over such a channel W_{nn}?

Contributions in our work

Theorem (McBain, Viterbo, Saunderson (2024))

The capacity of W_{nn} is

$$C(W_{nn}) = \sup_{P_{S|S^-}} \lim_{m\to\infty} \frac{1}{m} I(S^m; Y^{T_m}),$$

where the sup. is over all stationary and ergodic transition kernels $P_{S|S^-}$.

Contributions in our work

Theorem (McBain, Viterbo, Saunderson (2024))

The capacity of W_{nn} is

$$C(W_{nn}) = \sup_{P_{S|S^-}} \lim_{m\to\infty} \frac{1}{m} I(S^m; Y^{T_m}),$$

where the sup. is over all stationary and ergodic transition kernels $P_{S|S^-}$.

Our contributions:

1. Explicit computation of the capacity without DMC noise
2. Computable bounds on capacity for general, noisy nanopore channels (NNCs)
3. Low-complexity, high-rate encoding/decoding algorithms when $W = \text{EC}$
4. Change-point detection-based decoder for high current measurement ("sampling") rates

Contributions in our work

Theorem (McBain, Viterbo, Saunderson (2024))

The capacity of W_{nn} is

$$C(W_{nn}) = \sup_{P_{S|S^-}} \lim_{m\to\infty} \frac{1}{m} I(S^m; Y^{T_m}),$$

where the sup. is over all stationary and ergodic transition kernels $P_{S|S^-}$.

Our contributions:

1. Explicit computation of the capacity without DMC noise
2. Computable bounds on capacity for general, noisy nanopore channels (NNCs)
3. Simple, high-rate encoding/decoding algorithms when $W = \mathrm{EC}$
4. Change-point detection-based decoder for high current measurement ("sampling") rates

NNCs Under High Sampling Rates

Setup and intuition

- We focus on the setting of high "sampling rates," or high rates of current measurements.
- The sampling rate is controlled by the system designer, e.g., a more recent Oxford Nanopore Technologies sequencer has a sequencing speed of 450 b/s.

  [Wick, Judd, Holt (2019), Zeng et al. (2020)]
- High sampling rates give rise to several duplications of each base.
- Intuitively, this leads to several noisy views of each τ-mer, allowing for reliable recovery.

High sampling rate regime

- We recommend that the designer set sampling rates high enough so that
$$P_K(\ell_m \leq K \leq h_m) \geq 1 - \frac{1}{m^{1+\eta}},$$
for some $\eta > 0$, where $\ell_m = m^2(\ln m)^3$ and $h_m = \gamma m^2(\ln m)^3$, for some $\gamma > 1$.

High sampling rate regime

- We recommend that the designer set sampling rates high enough so that
$$P_K(\ell_m \leq K \leq h_m) \geq 1 - \frac{1}{m^{1+\eta}},$$
for some $\eta > 0$, where $\ell_m = m^2(\ln m)^3$ and $h_m = \gamma m^2(\ln m)^3$, for some $\gamma > 1$.

High sampling rate regime

- We recommend that the designer set sampling rates high enough so that
$$P_K(\ell_m \leq K \leq h_m) \geq 1 - \frac{1}{m^{1+\eta}},$$
for some $\eta > 0$, where $\ell_m = m^2(\ln m)^3$ and $h_m = \gamma m^2(\ln m)^3$, for some $\gamma > 1$.

- Our decoder computes change-point instants and employs a **MAP decoder** between change-points.

  see [Laszlo et al. (2014)] for a change-point detection-based decoder in nanopore sequencing

- The existing analysis of multi-view information rates helps us obtain a handle on MAP error probability between change-points.

Our change-point detection-based decoder DECODE

Fix a false-alarm probability $\alpha_m = \frac{1}{m^3(\ln m)^4}$ and a trimming length $c_m = m(\ln m)^2$.

$$Y_1 \quad Y_2 \quad Y_3 \quad Y_4 \;\ldots\; Y_{T_1-1} \quad Y_{T_1} \quad Y_{T_1+1} \;\ldots\; Y_{T_2} \quad Y_{T_2+1} \ldots \; Y_{T_m}$$

Our change-point detection-based decoder Decode

Fix a false-alarm probability $\alpha_m = \frac{1}{m^3(\ln m)^4}$ and a trimming length $c_m = m(\ln m)^2$.

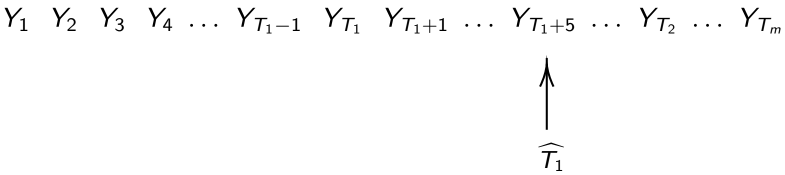

$$Y_1 \quad Y_2 \quad Y_3 \quad Y_4 \quad \dots \quad Y_{T_1-1} \quad Y_{T_1} \quad Y_{T_1+1} \quad \dots \quad Y_{T_1+5} \quad \dots \quad Y_{T_2} \quad \dots \quad Y_{T_m}$$

$\widehat{T_1}$

- Use the Shiryaev algorithm to obtain a change-point estimate with false-alarm prob. α_m.

Our change-point detection-based decoder DECODE

Fix a false-alarm probability $\alpha_m = \frac{1}{m^3(\ln m)^4}$ and a trimming length $c_m = m(\ln m)^2$.

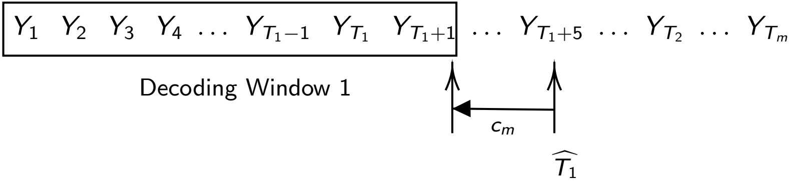

- Use the Shiryaev algorithm to obtain a change-point estimate with false-alarm prob. α_m.
- Trim by c_m to obtain the decoding window.
- Use the MAP decoder in the decoding window to obtain an estimate of the first τ-mer.

Our change-point detection-based decoder DECODE

Fix a false-alarm probability $\alpha_m = \frac{1}{m^3(\ln m)^4}$ and a trimming length $c_m = m(\ln m)^2$.

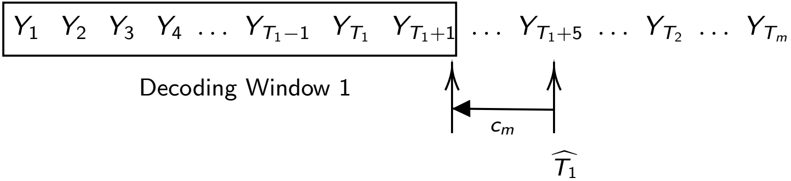

- Use the Shiryaev algorithm to obtain a change-point estimate with false-alarm prob. α_m.
- Trim by c_m to obtain the decoding window.
- Use the MAP decoder in the decoding window to obtain an estimate of the first τ-mer.

Recurse from the end of the window.

Main result

With the aid of properties of the Shiryaev algorithm, we obtain the following results:

Main result

With the aid of properties of the Shiryaev algorithm, we obtain the following results:

Theorem

For any no-self-loop de Bruijn Markov $P_{S|S^-}$, in the high sampling-rate regime,

$$\lim_{m\to\infty} \Pr\left[\textsc{Decode}(Y^{T_m}) \neq S^m\right] = 0.$$

Theorem

In the high sampling-rate regime, rates of up to $C_\tau^{\text{no-noise, no-loop}}$ are achievable using $\textsc{Decode}$*.*

Ongoing/Future work

- How many points can we pack into a (high-dimensional) probability simplex that are ϵ-separated in Chernoff distance?
  - Chernoff sep. $\equiv$ Rate of $\rightarrow H(X)$
  - Points in prob. simplex $\equiv$ Messages in composite DNA storage

- Concatenated multi-view channels? Feedback?

Thank You!