There can be only one!

A Jorney Through the Intricacies of Machine Learning for Unravelling LoRa Frames under Collision

José I. Álamos.
HAW Hamburg
TU Dresden

INET-RG Christmas Talk
December 13, 2023

What is this talk about?

Journey and **lessons learned** from the development of **Machine Learning models** for unravelling **LoRa frames** under **collision**.

DALL-E. OpenAI

Disclaimer

- I'm not a Machine Learning expert.

Disclaimer

- I'm not a Machine Learning expert.
- Wild math formulas spotted in the area!

Disclaimer

- I'm not a Machine Learning expert.
- Wild math formulas spotted in the area!
  - Don't worry, they are like passing clouds.
  - Interesting to glance at, but you don't need to give them a second thought.

Overview of LoRa

Proprietary wireless modulation technique

Long range (up to 15 km)

Low power consumption (mJ)

Low data rate (bytes/s)

Use Cases

Elf Tracking

Naughty or Nice Monitor

Reindeer Health Tracking

DALL-E. OpenAI

Problem Statement

Region	Population ($\frac{1}{km^2}$)	10km-Radius Mean Arrival ($\frac{1}{s}$)
Paris	21000	18325
London	5518	4815
Berlin	4000	3490

Table: Mean arrival rate of LoRa frames in selected urban areas. Assumes 10 devices per person, 1 message per device per hour.

Draft IG LPWA Report (IEEE P802.15-17-0528-00-lpwa)

Problem Statement

Region	Population ($\frac{1}{km^2}$)	10km-Radius Mean Arrival ($\frac{1}{s}$)

Long range yields high collision probability

Table: Mean arrival rate of LoRa frames in selected urban areas. Assumes 10 devices per person, 1 message per device per hour.

Draft IG LPWA Report (IEEE P802.15-17-0528-00-lpwa)

Agenda

When LoRa Frames Collide

The Journey Begins

Towards Machine Learning

Light Ahead

Lessons Learned

When LoRa Frames Collide

DALL-E. OpenAI

LoRa Modulation

Chirp Spread Spectrum

$$Y_0(t) = e^{j\left(\pi B t^2\right)}$$

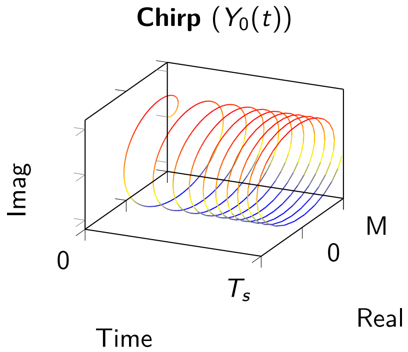

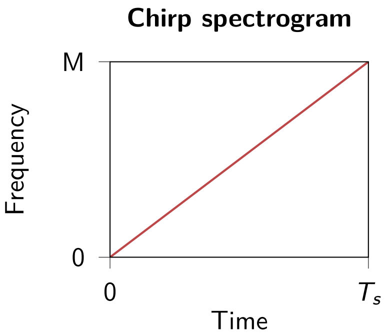

Decoding LoRa

LoRa symbol $(X(t))$

Downchirp $(\overline{X_0(t)})$

Dechirped symbol $(X(t) \cdot \overline{X_0(t)})$

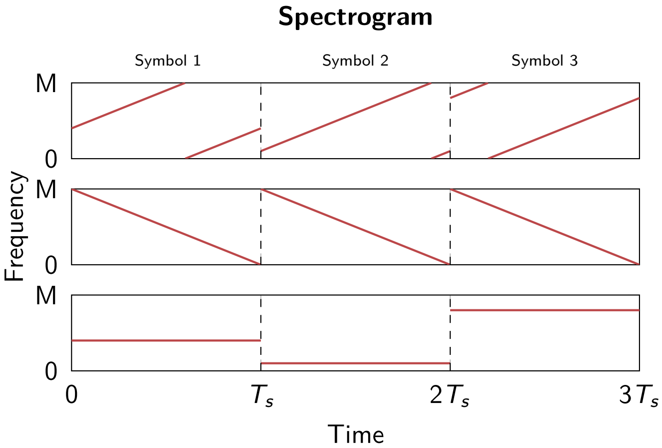

Decoding LoRa

Frequency bin with the highest magnitude.

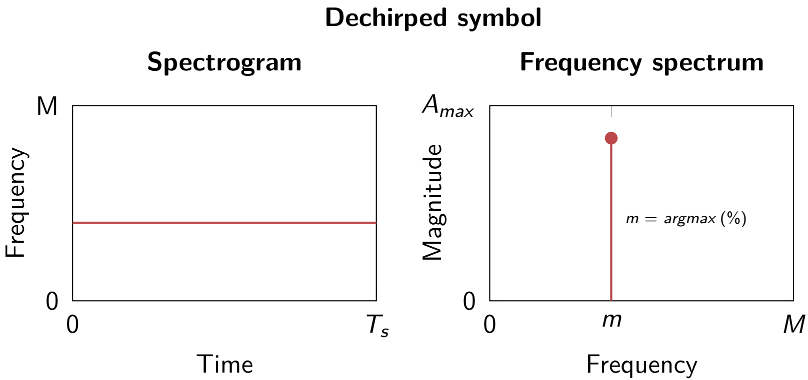

LoRa Collisions

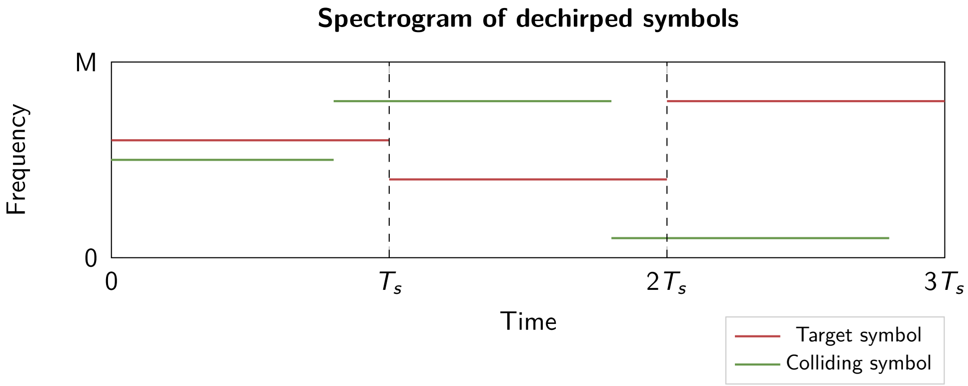

LoRa Collisions

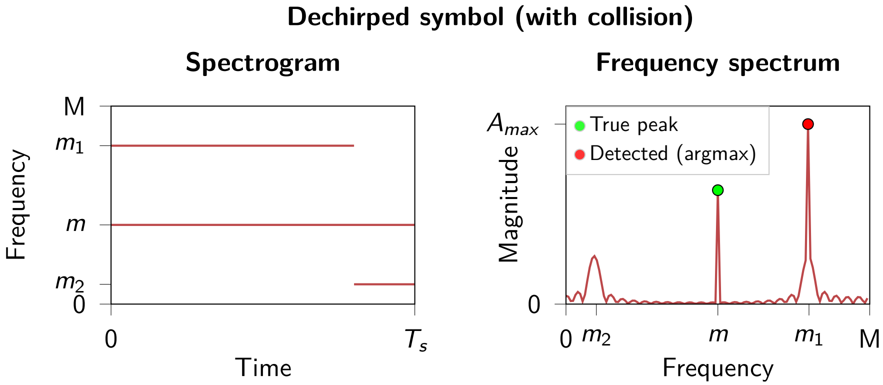

LoRa Collisions

Dechirped symbol (with collision)

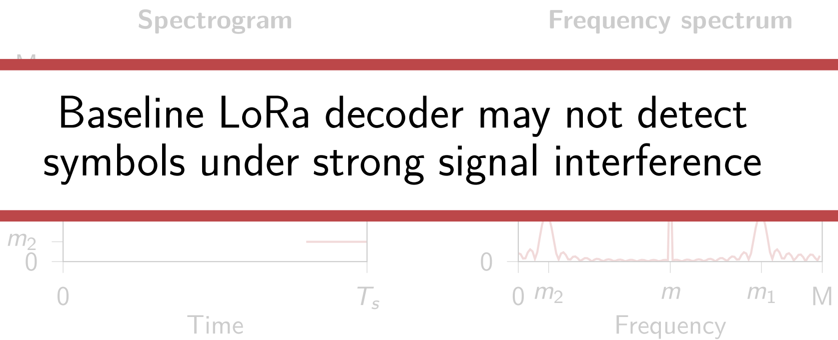

Baseline LoRa decoder may not detect symbols under strong signal interference

Related Work

Ranking of LoRa collision recovery algorithms

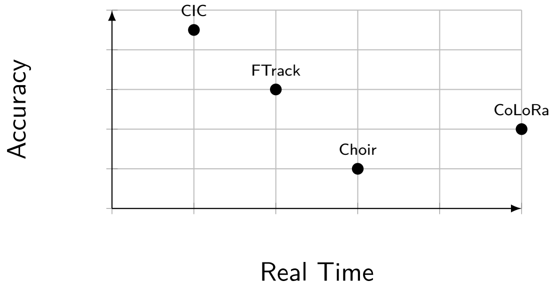

C. Shao et al., "Toward Ubiquitous Connectivity via LoRaWAN: An Overview of Signal Collision Resolving Solutions," in IEEE Internet of Things Magazine, vol. 4, no. 4, pp. 114-119, Dec. 2021.

Related Work

Ranking of LoRa collision recovery algorithms

Trade-off between accuracy and real-time performance.

Real Time

C. Shao et al., "Toward Ubiquitous Connectivity via LoRaWAN: An Overview of Signal Collision Resolving Solutions," in IEEE Internet of Things Magazine, vol. 4, no. 4, pp. 114-119, Dec. 2021.

Goal

Find the full sine wave in the dechirped symbol.

acgt.me

The Journey Begins

DALL-E. OpenAI

Gathering Symbols

Software Defined Radio (SDR)

- Reuse existing deployment.
- Capture real-world symbol data.

Simulated data

- Generate symbol data with known parameters.
- Model symbol as complex chirp with white gaussian noise.

Naive Approach

Use conventional time-frequency analysis techniques.

- Fast Fourier Transform (FFT).
- Short-time Fourier transform (STFT).

Naive Approach (cont.)

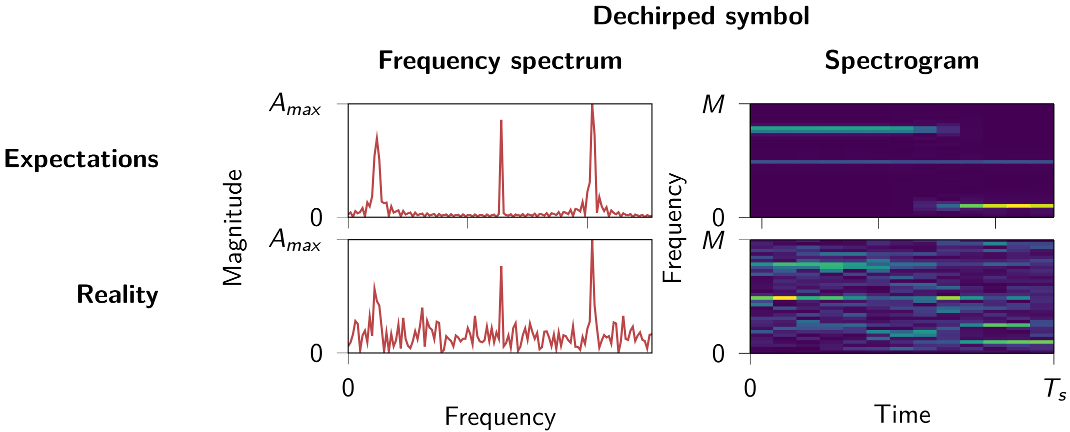

Towards Machine Learning

DALL-E. OpenAI

First Insights

- Finding the full sine wave is harder than expected.
  - LoRa collisions yield a complex frequency spectrum.
  - Hard to unravel using conventional signal processing techniques.

First Insights

- Finding the full sine wave is harder than expected.
  - LoRa collisions yield a complex frequency spectrum.
  - Hard to unravel using conventional signal processing techniques.
- Potential of Machine Learning techniques for decoding LoRa frames.

First Insights

- Finding the full sine wave is harder than expected.

Use Machine Learning to identify the full sine wave

frames.

Machine Learning Terminology

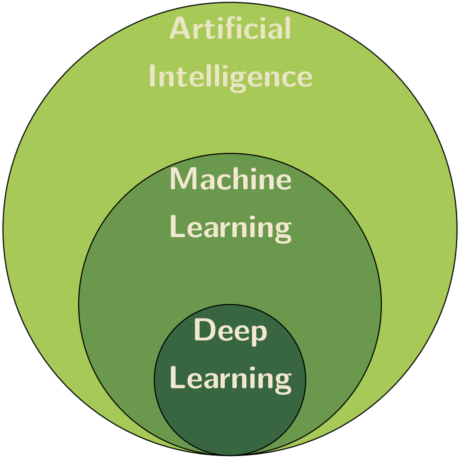

Choice of Deep Learning

There seem to be two types of Machine Learning researchers:

Choice of Deep Learning

There seem to be two types of Machine Learning researchers:

- The Deep Learning users.

Choice of Deep Learning

There seem to be two types of Machine Learning researchers:

- The Deep Learning users.
- The moving-towards-Deep-Learning users.

Deep Learning

- Neural Networks.

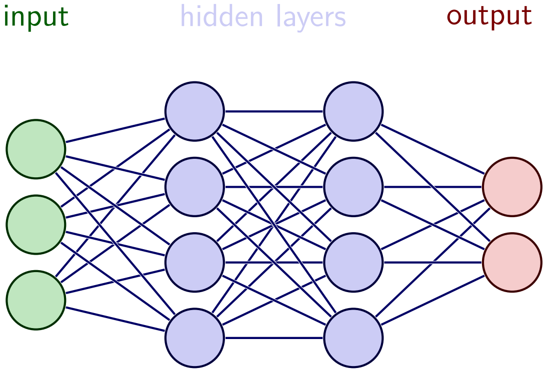

TiKZ.net

Deep Learning

- Neural Networks.
- Ability to learn from large and complex data.

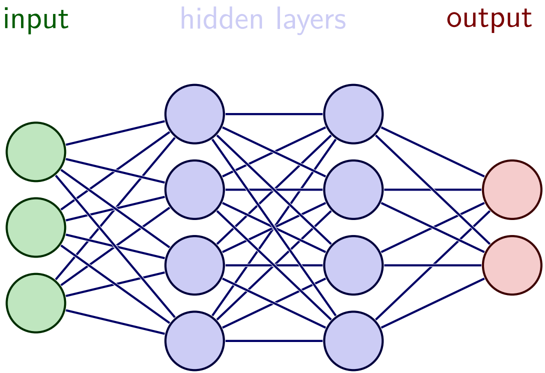

TiKZ.net

Deep Learning

- Neural Networks.
- Ability to learn from large and complex data.
- Automatic Feature Extraction.

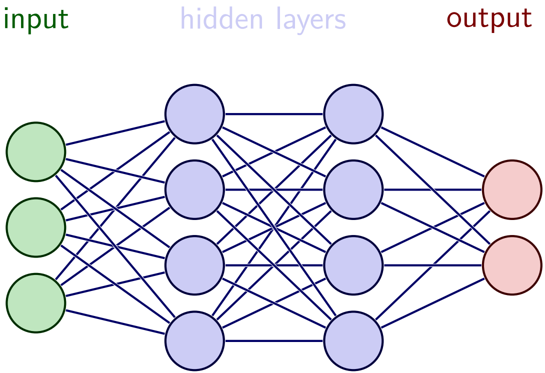

TiKZ.net

Neural Networks

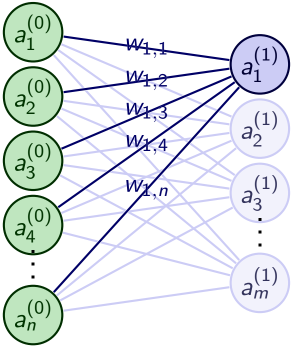

$$\begin{pmatrix} a_1^{(1)} \\ a_2^{(1)} \\ \vdots \\ a_m^{(1)} \end{pmatrix} = \sigma \left[\begin{pmatrix} w_{1,1} & w_{1,2} & \dots & w_{1,n} \\ w_{2,1} & w_{2,2} & \dots & w_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ w_{m,1} & w_{m,2} & \dots & w_{m,n} \end{pmatrix} \begin{pmatrix} a_1^{(0)} \\ a_2^{(0)} \\ \vdots \\ a_n^{(0)} \end{pmatrix} + \begin{pmatrix} b_1^{(0)} \\ b_2^{(0)} \\ \vdots \\ b_m^{(0)} \end{pmatrix} \right]$$

$$Input = \mathbf{a}^{(0)}$$

$$\mathbf{a}^{(1)} = \sigma\left(\mathbf{W}^{(0)}\mathbf{a}^{(0)} + \mathbf{b}^{(0)}\right)$$

$$Output = \mathbf{a}^{(K)} = F\left(\mathbf{W}, \mathbf{b}, \mathbf{a}^{(0)}\right)$$

$$Goal : min.F_{loss}$$

TiKZ.net

Deep Learning Architectures

Many! Some examples:

Deep Learning Architectures

Many! Some examples:

- Fully Connected Layers (FC).
  - Good at classification and regression.

Deep Learning Architectures

Many! Some examples:

- Fully Connected Layers (FC).
  - Good at classification and regression.
- Convolutional Neural Networks (CNN).
  - Good at finding patterns in data.

Deep Learning Architectures

Many! Some examples:

- Fully Connected Layers (FC).
  - Good at classification and regression.
- Convolutional Neural Networks (CNN).
  - Good at finding patterns in data.
- Recurrent Neural Networks (RNN).
  - Good at predicting sequences.

Deep Learning Architectures

Many! Some examples:

- Fully Connected Layers (FC).
  - Good at classification and regression.
- Convolutional Neural Networks (CNN).
  - Good at finding patterns in data.
- Recurrent Neural Networks (RNN).
  - Good at predicting sequences.
- Autoencoders (AE).
  - Good at dimensionality reduction.

It looks promising

Due to the Deep Learning hype, it is easy to believe the process is:

It looks promising

Due to the Deep Learning hype, it is easy to believe the process is:

1. Gather data.

It looks promising

Due to the Deep Learning hype, it is easy to believe the process is:

1. Gather data.
2. Build and train model.

It looks promising

Due to the Deep Learning hype, it is easy to believe the process is:

1. Gather data.
2. Build and train model.
3. Evaluate model.

It looks promising

Due to the Deep Learning hype, it is easy to believe the process is:

1. Gather data.
2. Build and train model.
3. Evaluate model.
4. Tune hyperparameters.

It looks promising

Due to the Deep Learning hype, it is easy to believe the process is:

1. Gather data.
2. Build and train model.
3. Evaluate model.
4. Tune hyperparameters.
5. Get state-of-the-art results.

CNN Symbol Classifier

- Train a CNN to find the longest sine wave in the dechirped symbol.

CNN Symbol Classifier

- Train a CNN to find the longest sine wave in the dechirped symbol.
- Use signal in time and frequency domain as input.

CNN Symbol Classifier

- Train a CNN to find the longest sine wave in the dechirped symbol.
- Use signal in time and frequency domain as input.
- Train from simulated data.

CNN Architecture

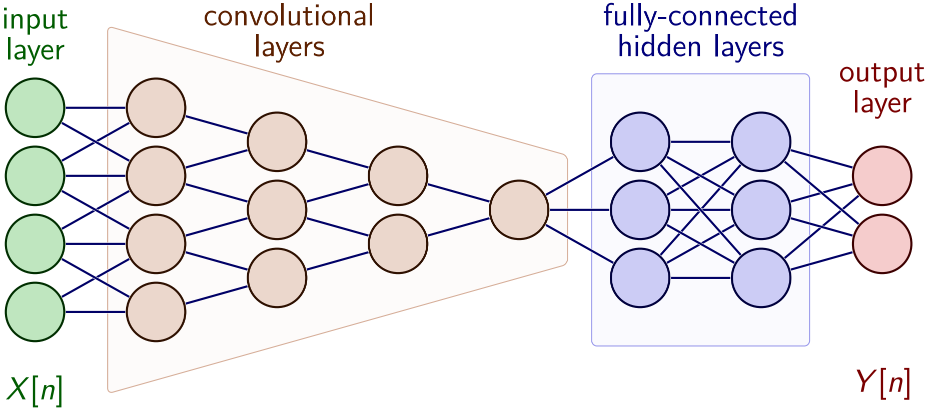

TiKZ.net

CNN Symbol Classifier Evaluation

- Time domain as input.
  - Does not converge

CNN Symbol Classifier Evaluation

- Time domain as input.
  - Does not converge
- Spectrogram (STFT) as input.
  - Worse than baseline decoder.

CNN Symbol Classifier Evaluation

- Time domain as input.
  - Does not converge
- Spectrogram (STFT) as input.
  - Worse than baseline decoder.
- Frequency spectrum (FFT) as input.
  - Does not detect symbols with collisions.

CNN Symbol Classifier Evaluation

- Time domain as input.
  - Does not converge
- Spectrogram (STFT) as input.
  - Worse than baseline decoder.
- Frequency spectrum (FFT) as input.
  - Does not detect symbols with collisions.
  - But performs slightly better than baseline decoder.

We need more signal features!

Wavelet Transform

We need more signal features!

Wavelet Transform

Wigner-Ville Distribution

We need more signal features!

Wavelet Transform

Wigner-Ville Distribution

Synchrosqueezing Transform

We need more signal features!

Wavelet Transform

Wigner-Ville Distribution

Synchrosqueezing Transform

Fractional Fourier Transform

We need more signal features!

Wavelet Transform

Wigner-Ville Distribution

Synchrosqueezing Transform

Hilbert-Huang Transform

Fractional Fourier Transform

We need more signal features!

Wavelet Transform

Wigner-Ville Distribution

Synchrosqueezing Transform

Hilbert-Huang Transform

Empirical Mode Decomposition

Fractional Fourier Transform

We need more signal features!

Wavelet Transform

Wigner-Ville Distribution

Synchrosqueezing Transform

Hilbert-Huang Transform

Empirical Mode Decomposition

Fractional Fourier Transform

etc.

Aftermath of CNN Classifier

- Some features yield slightly better accuracy than the baseline decoder.
  - At the cost of high computational complexity.

Aftermath of CNN Classifier

- Some features yield slightly better accuracy than the baseline decoder.
  - At the cost of high computational complexity.
- The classifier works best for symbols without collision.

Aftermath of CNN Classifier

Some features yield slightly better accuracy than the baseline

Gains are not enough to justify the complexity

Next attempt: Peak Classification

1. Find location of peaks.
   - Reuse CNN architecture from previous attempt.

Next attempt: Peak Classification

1. Find location of peaks.
   - Reuse CNN architecture from previous attempt.
2. Find probability of each detected peak being a full sine wave.
   - Based on peak features (power, frequency offset).
   - Utilize a simpler Machine Learning model (Gradient Boosting).
   - Compare against simple Bayesian (Zscore) method.

Next attempt: Peak Classification

1. Find location of peaks.
   - Reuse CNN architecture from previous attempt.
2. Find probability of each detected peak being a full sine wave.
   - Based on peak features (power, frequency offset).
   - Utilize a simpler Machine Learning model (Gradient Boosting).
   - Compare against simple Bayesian (Zscore) method.
3. Choose the candidate with highest probability.

CNN Peak Detection Evaluation

It works!

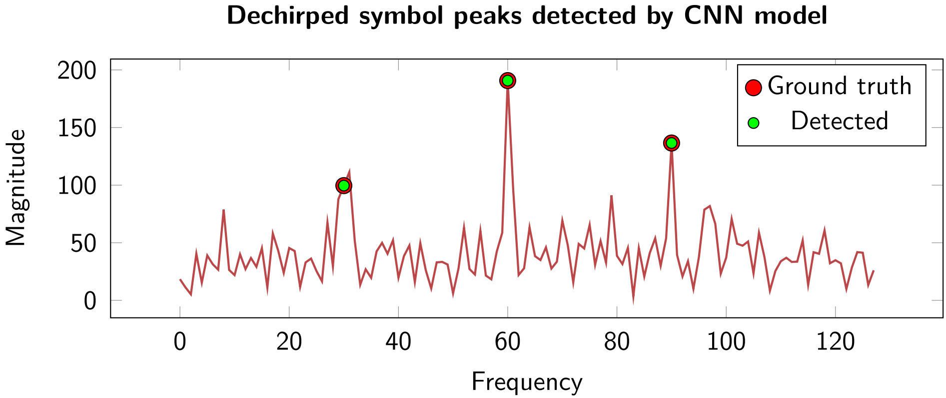

Peak Classification Evaluation

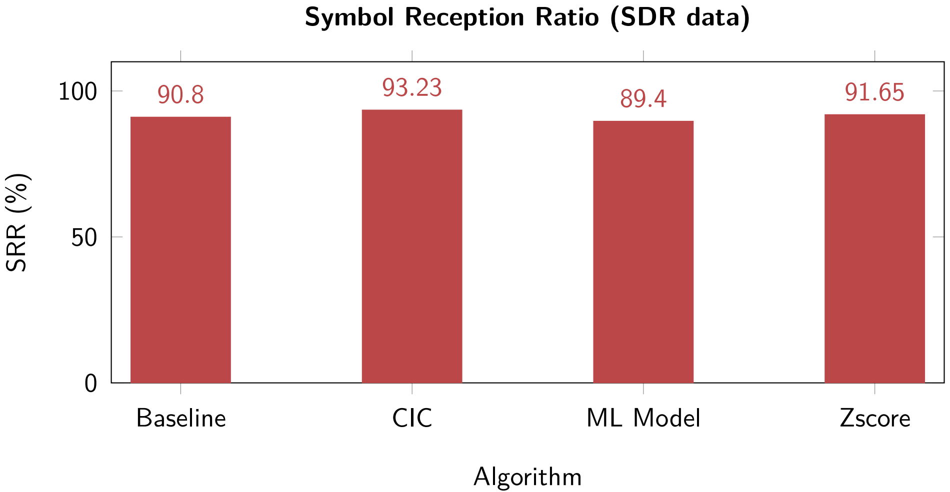

Aftermath of Peak Classification

- Peak detection works well, but computationally expensive.

Aftermath of Peak Classification

- Peak detection works well, but computationally expensive.
- Peak classification does not bring any gains over the state of the art.

Aftermath of Peak Classification

- Peak detection works well, but computationally expensive.
- Peak classification does not bring any gains over the state of the art.
  - The Machine Learning model performs worse than the simple Zscore method.

Aftermath of Peak Classification

- Peak detection works well, but computationally expensive.
- Peak classification does not bring any gains over the state of the art.
  - The Machine Learning model performs worse than the simple Zscore method.
  - ... and worse than the baseline decoder.

Aftermath of Peak Classification

- Peak detection works well, but computationally expensive.

Peak classification strategy does not pay off

- ... and worse than the baseline decoder.

Next attempt: Denoiser Autoencoder

- Noise remains as one of the main challenges.
  - Hard to remove with simple math.

Next attempt: Denoiser Autoencoder

- Noise remains as one of the main challenges.
  - Hard to remove with simple math.
- Train a neural network to remove noise from frequency domain.

Next attempt: Denoiser Autoencoder

- Noise remains as one of the main challenges.
  - Hard to remove with simple math.
- Train a neural network to remove noise from frequency domain.
- Then use simple math techniques to find the full sine wave.

Denoiser Autoencoder Architecture

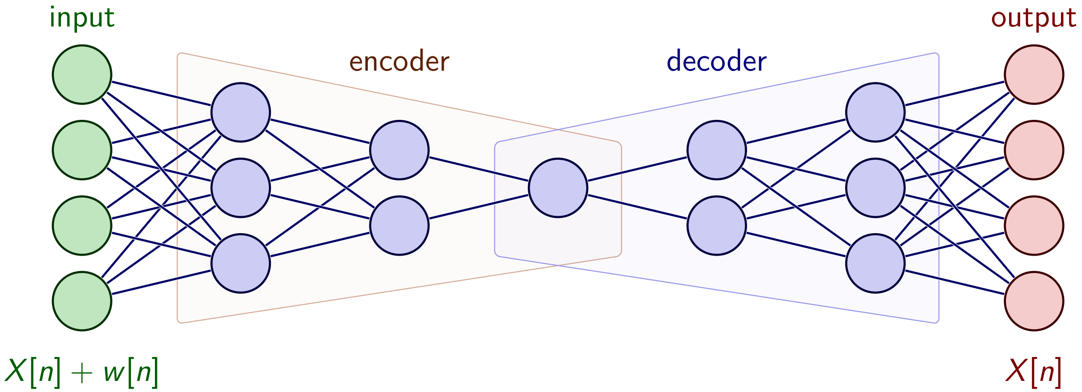

TiKZ.net

Denoiser Autoencoder Evaluation

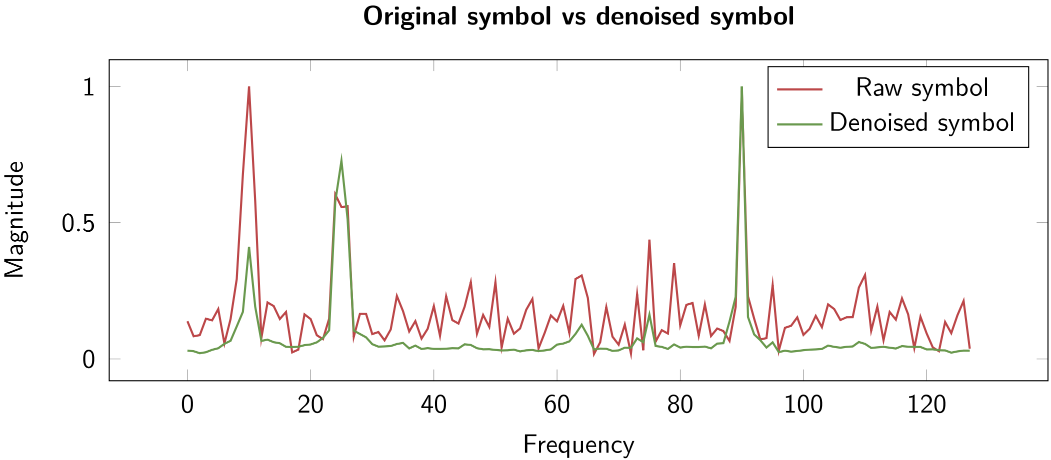

Aftermath of Denoiser Autoencoder

- SNR of the signal is improved.

Aftermath of Denoiser Autoencoder

- SNR of the signal is improved.
- Phase distortion introduces a new problem.

Phase matters!

- Location of sine waves is encoded in the phase.

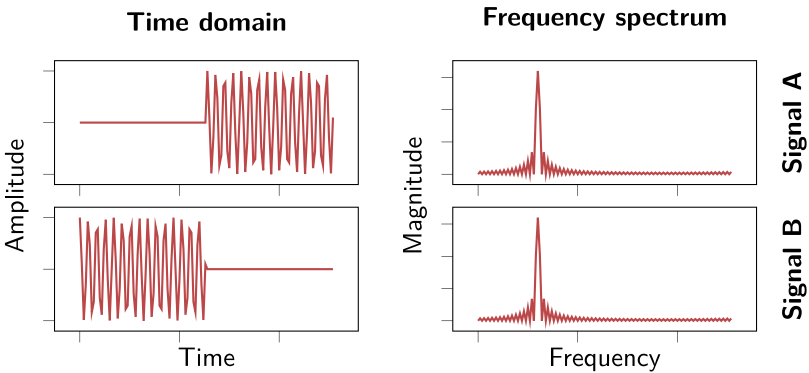

Aftermath of Denoiser Autoencoder (cont.)

- Different architectures and training methods improve results.

Aftermath of Denoiser Autoencoder (cont.)

- Different architectures and training methods improve results.
- Phase reconstruction still not good enough.

Aftermath of Denoiser Autoencoder (cont.)

Incorrect phase distortion rules out the use of time-frequency techniques

Light Ahead

DALL-E. OpenAI

Recap

- Addressed the problem from many different angles.

imgflip.com

Recap

- Addressed the problem from many different angles.
- Machine Learning show good results for certain tasks.

imgflip.com

Recap

- Addressed the problem from many different angles.
- Machine Learning show good results for certain tasks.
  - Still, the original problem remains unsolved.

imgflip.com

Recap

- Addressed the problem from many different angles.
- Machine Learning show good results for certain tasks.
  - Still, the original problem remains unsolved.
- Each iteration takes time, without any guarantees.

imgflip.com

Last hope: Simple Math (revisited)

- Leverage the knowledge acquired from the previous approaches.
  - Tons of signal processing techniques.
  - Understanding the dynamics of LoRa collisions.

Last hope: Simple Math (revisited)

- Leverage the knowledge acquired from the previous approaches.
  - Tons of signal processing techniques.
  - Understanding the dynamics of LoRa collisions.
- After some weeks of research, we identified an analytic scheme.
  - Exploits the structure of symbols under collision.
  - Isolates the full sine wave from the rest of the signal.

Preliminary Evaluation

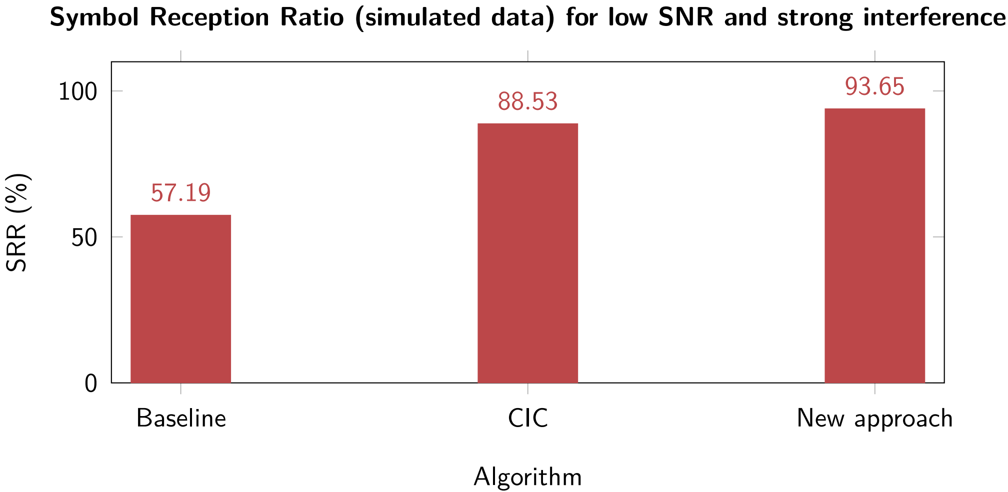

Lessons Learned

DALL-E. OpenAI

Post Mortem Analysis

- Spent 6 month trying to solve a problem that could be solved in 2 weeks with conventional math techniques.

Post Mortem Analysis

- Spent 6 month trying to solve a problem that could be solved in 2 weeks with conventional math techniques.
- Was it worth it?

Post Mortem Analysis

Absolutely!

Post Mortem Analysis (cont.)

- Knowledge acquired led to final solution.

Post Mortem Analysis (cont.)

- Knowledge acquired led to final solution.
- Some Deep Learning tasks show promising results.
  - Denoising
  - Peak detection

Post Mortem Analysis (cont.)

- Knowledge acquired led to final solution.
- Some Deep Learning tasks show promising results.
  - Denoising
  - Peak detection
- These strategies have still interesting applications for LoRa decoding
  - Peak detection in extremely noisy environments.

Lessons Learned

Although very powerful:

Lessons Learned

Although very powerful:

1. Machine Learning is not a silver bullet.

Lessons Learned

Although very powerful:

1. Machine Learning is not a silver bullet.
2. Machine Learning is not magic.

Lessons Learned

Although very powerful:

1. Machine Learning is not a silver bullet.
2. Machine Learning is not magic.
3. Machine Learning is not straightforward.

Not a Silver Bullet

- Certain tasks are solved better with conventional math techniques.

Not a Silver Bullet

- Certain tasks are solved better with conventional math techniques.
- Training/tuning takes time
  - It is not always worth it.

Not a Silver Bullet

- Certain tasks are solved better with conventional math techniques.
- Training/tuning takes time
  - It is not always worth it.
- Complex models do not ensure better results.
  - Deep Learning models do not always perform better than simpler Machine Learning models.
  - And they are harder to train.

Not a Silver Bullet (cont.)

But for certain tasks, Deep Learning is likely the best approach.

DALL-E. OpenAI.

Not a Silver Bullet (cont.)

But for certain tasks, Deep Learning is likely the best approach.

- Large Language Models

DALL-E. OpenAI.

Not a Silver Bullet (cont.)

But for certain tasks, Deep Learning is likely the best approach.

- Large Language Models
- Denoising

DALL-E. OpenAI.

Not a Silver Bullet (cont.)

But for certain tasks, Deep Learning is likely the best approach.

- Large Language Models
- Denoising
- Sketching Santa Claus drinking Glühwein in Hamburg.

DALL-E. OpenAI.

Not Magic

- Underneath the fancy names, Machine Learning models are just a bunch of matrix multiplications.
  - Days of training may converge to a mathematical function that can be expressed in a single line of code.

Not Magic

- Underneath the fancy names, Machine Learning models are just a bunch of matrix multiplications.
  - Days of training may converge to a mathematical function that can be expressed in a single line of code.
- They are not able to learn anything that is not in the data.

Not Magic

- Underneath the fancy names, Machine Learning models are just a bunch of matrix multiplications.
    - Days of training may converge to a mathematical function that can be expressed in a single line of code.
- They are not able to learn anything that is not in the data.
- They do not replace domain knowledge

Not Magic (cont.)

After some time, one develops a sense of what is possible and what is not.

imgflip.com

Not Straightforward

- Model may learn something completely different from what you expect.
  - First iterations of autoencoder learned the identity function.

Not Straightforward

- Model may learn something completely different from what you expect.
  - First iterations of autoencoder learned the identity function.
  - If data is not balanced, model may learn to always predict the most common class.

Not Straightforward

- Model may learn something completely different from what you expect.
  - First iterations of autoencoder learned the identity function.
  - If data is not balanced, model may learn to always predict the most common class.
  - Requires some experience to identify and solve these problems.

Not Straightforward (cont.)

- Lack of interpretability.
  - It is not always clear why a model is not working.
  - It is not always clear how to improve a model.

Not Straightforward (cont.)

- Lack of interpretability.
  - It is not always clear why a model is not working.
  - It is not always clear how to improve a model.
- Architecture selection is not straightforward.
  - It is not always clear which architecture is the best for a given task.
  - Loss, activation functions and size of network make a huge difference.

Should you use Machine Learning?

Should you use Machine Learning?

Yes, but carefully.

Should you use Machine Learning?

Yes, but carefully.

- Try conventional techniques first.
  - Stay there if you are happy with the results.

Should you use Machine Learning?

Yes, but carefully.

- Try conventional techniques first.
  - Stay there if you are happy with the results.
- Identify tasks that may benefit from Machine Learning.
  - Pattern recognition
  - Prediction

Should you use Machine Learning?

Yes, but carefully.

- Try conventional techniques first.
  - Stay there if you are happy with the results.
- Identify tasks that may benefit from Machine Learning.
  - Pattern recognition
  - Prediction
- Start with simple models before moving to Deep Learning.
  - Specially if data is limited.

Thank you

for your attention!