

A scenic sunset over a lake with mountains in the background. The sky is a mix of dark blue and orange, with the sun low on the horizon. The water reflects the colors of the sky. The mountains are silhouetted against the bright part of the sky.

doing more with *less*


effective use of limited acoustic bandwidth

Why acoustics?

Data Rate \leq Capacity

$$\text{Capacity} = C_{\text{monotonic}}(\Delta f, SNR)$$

$$\text{SNR}_{\text{dB}}(f, R) = \text{SL} + G - \alpha \log_{10} R - A(f)R - N(f)$$



More info: Stojanovic, "On the relationship between capacity and distance in an underwater acoustic communication channel", WUWNet 2006.

Bits are a precious resource!

Reducing Delivery Overheads

- Frugal headers
- "Short-circuiting"
- Time-to-live (TTL)
- Mailboxes
- "Juggling" and "super-TDMA"
- Erasure control coding (to reduce need for ACKs)
- ⋮

Payload Compression

Compressed Payload Size \geq Entropy $\mathcal{S}(\mathbf{x})$

$$\mathcal{S}(\mathbf{x}) = - \sum_{\mathbf{x}} \mathbb{P}(\mathbf{x}) \log_2 \mathbb{P}(\mathbf{x}) \text{ bits}$$

Entropy Example

x	P(x)
000	0.125
001	0.125
010	0.125
011	0.125
100	0.125
101	0.125
110	0.125
111	0.125

$$S = -8 \times 0.125 \log_2(0.125) = 3 \text{ bits}$$



x	P(x)
000	0.0
001	0.0
010	0.0
011	0.5
100	0.0
101	0.0
110	0.5
111	0.0

$$S = -2 \times 0.5 \log_2(0.5) = 1 \text{ bit}$$


Isn't it reasonable to assume that application data is already high-entropy?

10 kB JPG



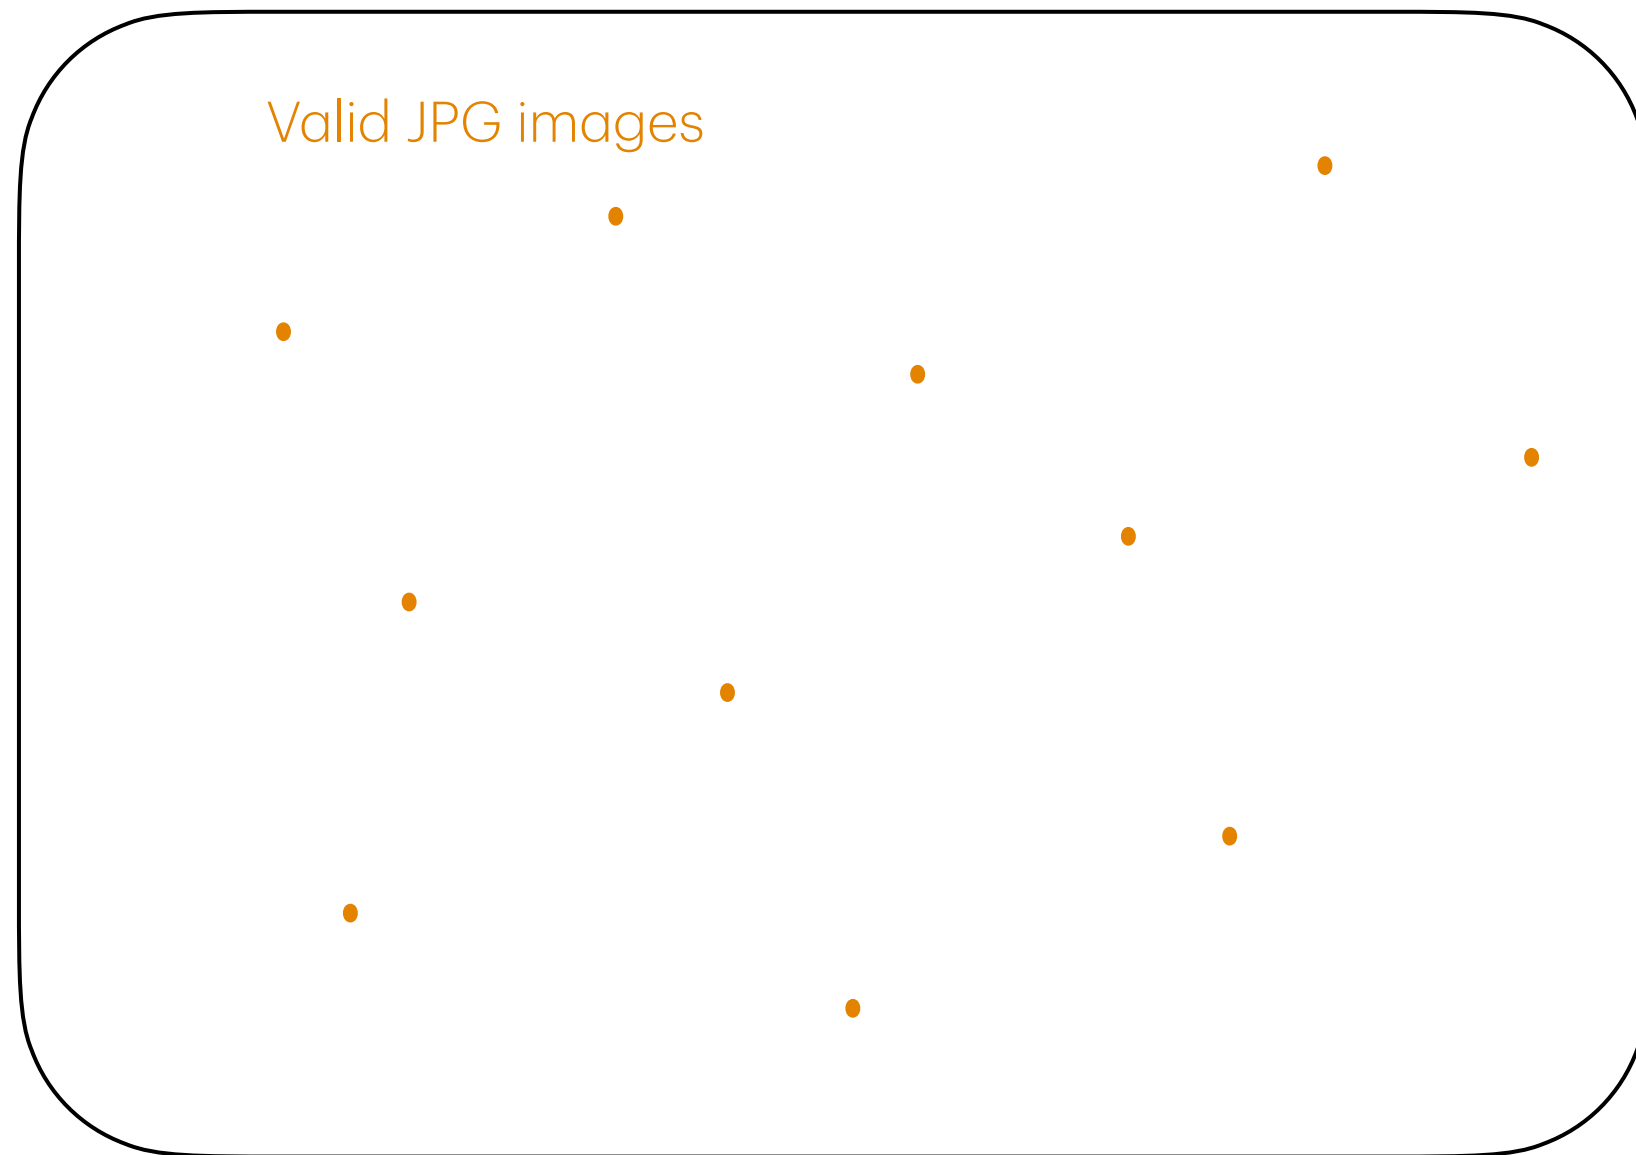


Space of all possible bit sequences of length N



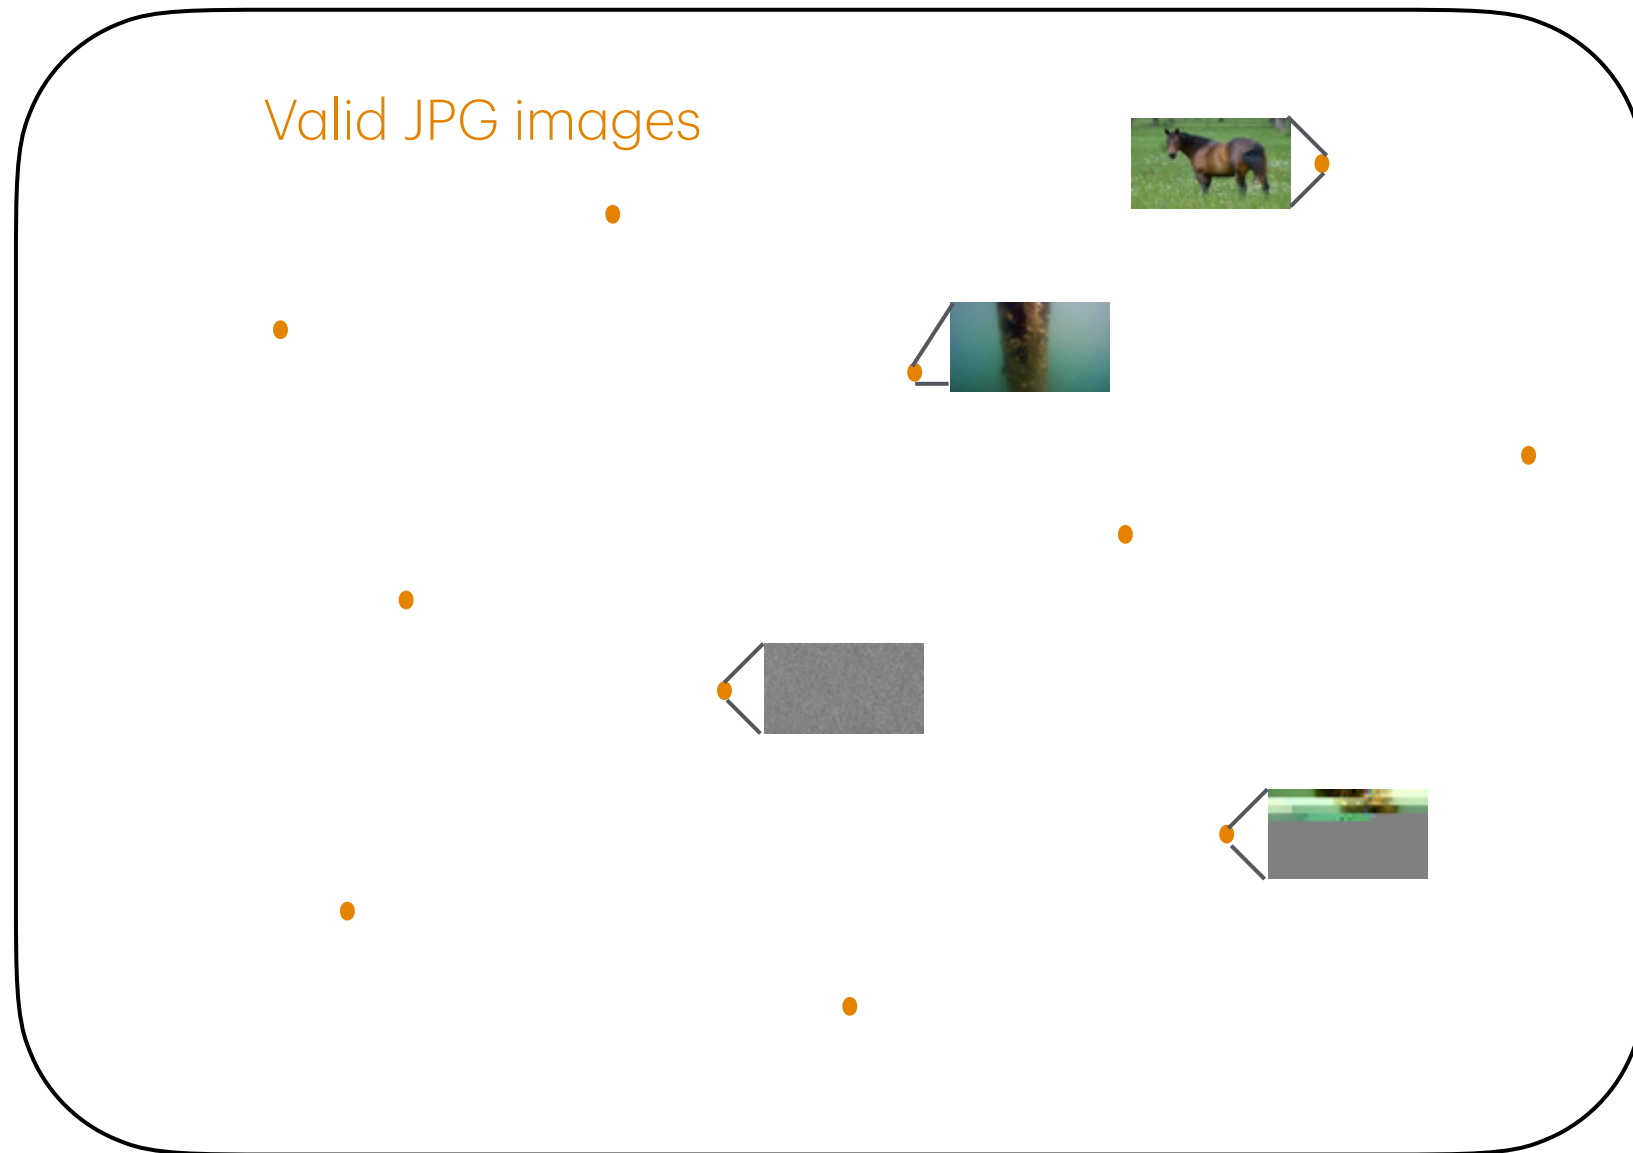
$$2^N$$

Space of all possible bit sequences of length N

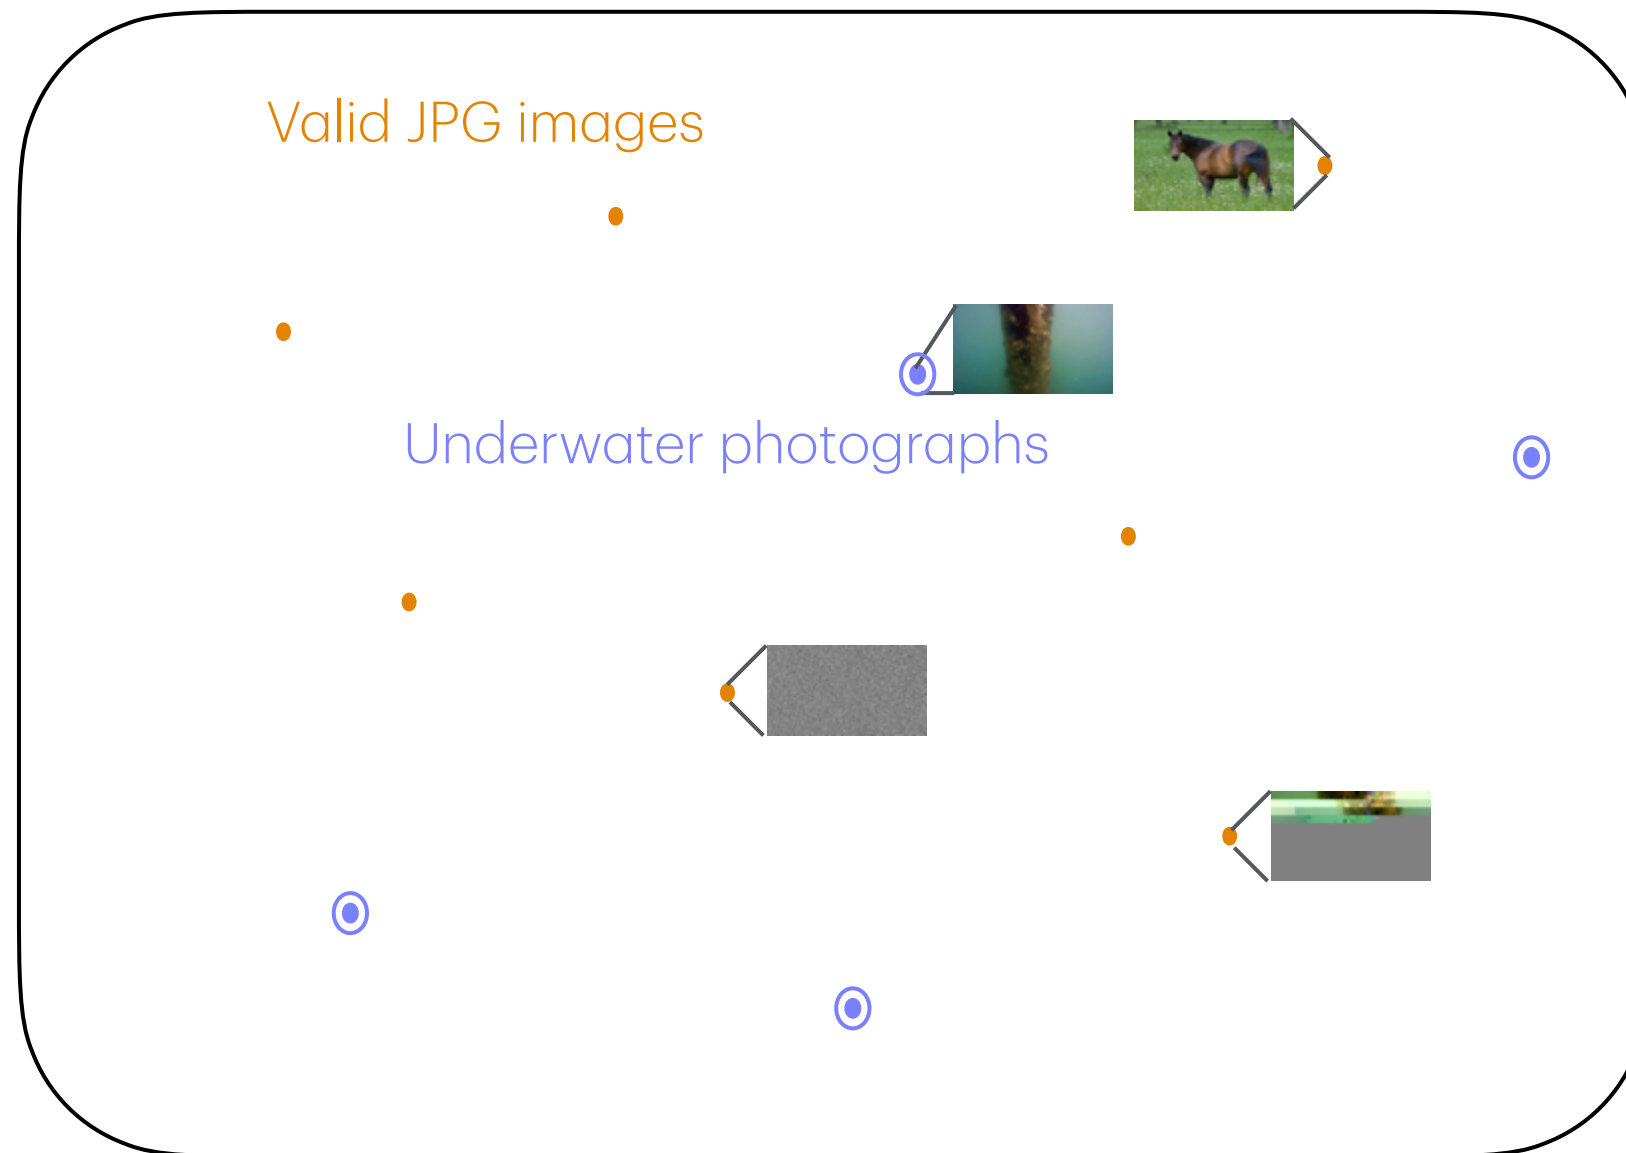


Space of all possible bit sequences of length N

Valid JPG images




Space of all possible bit sequences of length N



Space of all possible bit sequences of length N

Valid JPG images



Underwater photographs

Space of all possible bit sequences of length N

Valid JPG images



Underwater photographs

2^N

Space of all possible bit sequences of length N

Valid JPG images



Underwater photographs

2^N

Compression ratio: N/k

Energy / time usage

Typical terrestrial wireless:

- 4G LTE: 10 $\mu\text{J}/\text{bit}$, 100 Mbps
- 5G: 0.3 $\mu\text{J}/\text{bit}$, 1 Gbps

Energy / time usage

Typical terrestrial wireless:

- 4G LTE: 10 $\mu\text{J}/\text{bit}$, 100 Mbps
- 5G: 0.3 $\mu\text{J}/\text{bit}$, 1 Gbps

Typical acoustic modem:

- 185 dB 1 μPa @ 1m, 5 kbps
- 5 mJ/bit
- $\sim 10^3 - 10^4 \times$ larger!

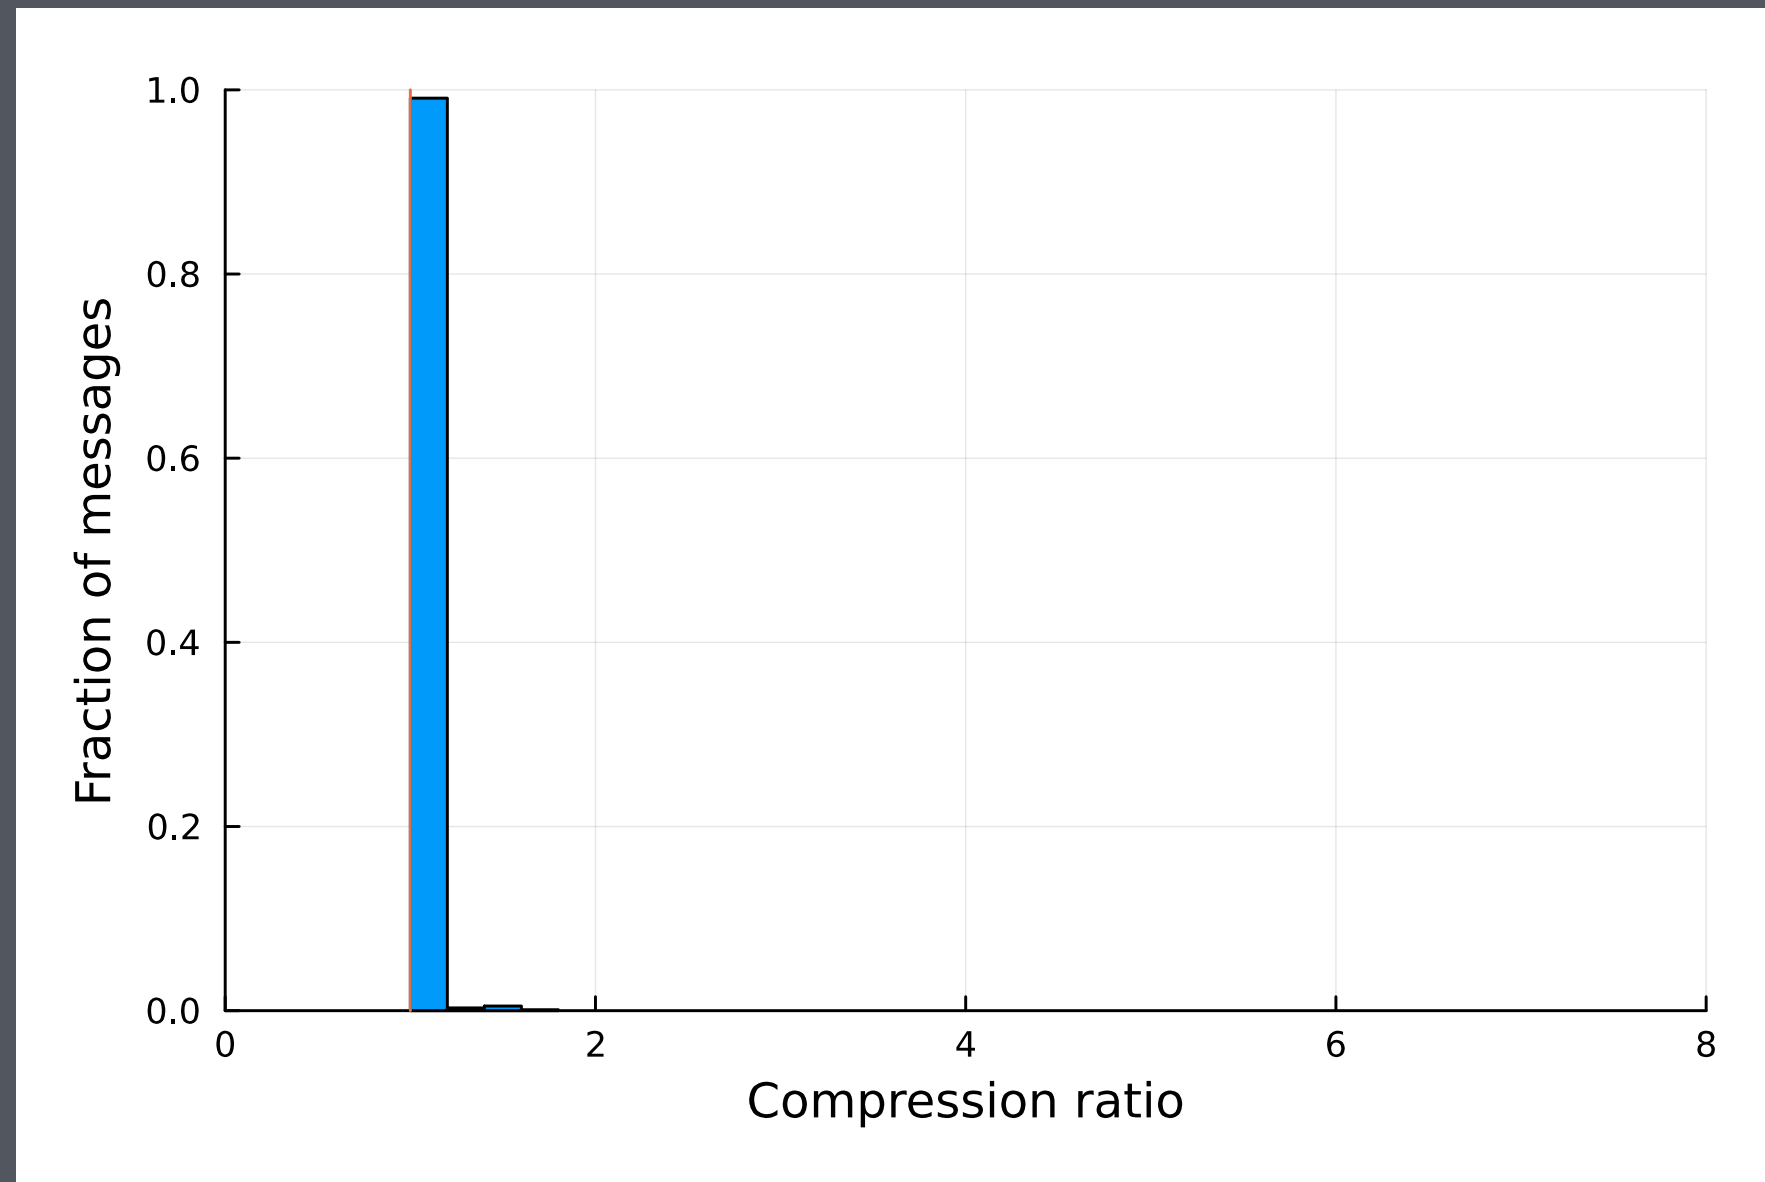
We can afford to spend
a lot more time & energy on computation
to reduce the number of bits to transmit!

Example

NMEA data from a tether-less ROV

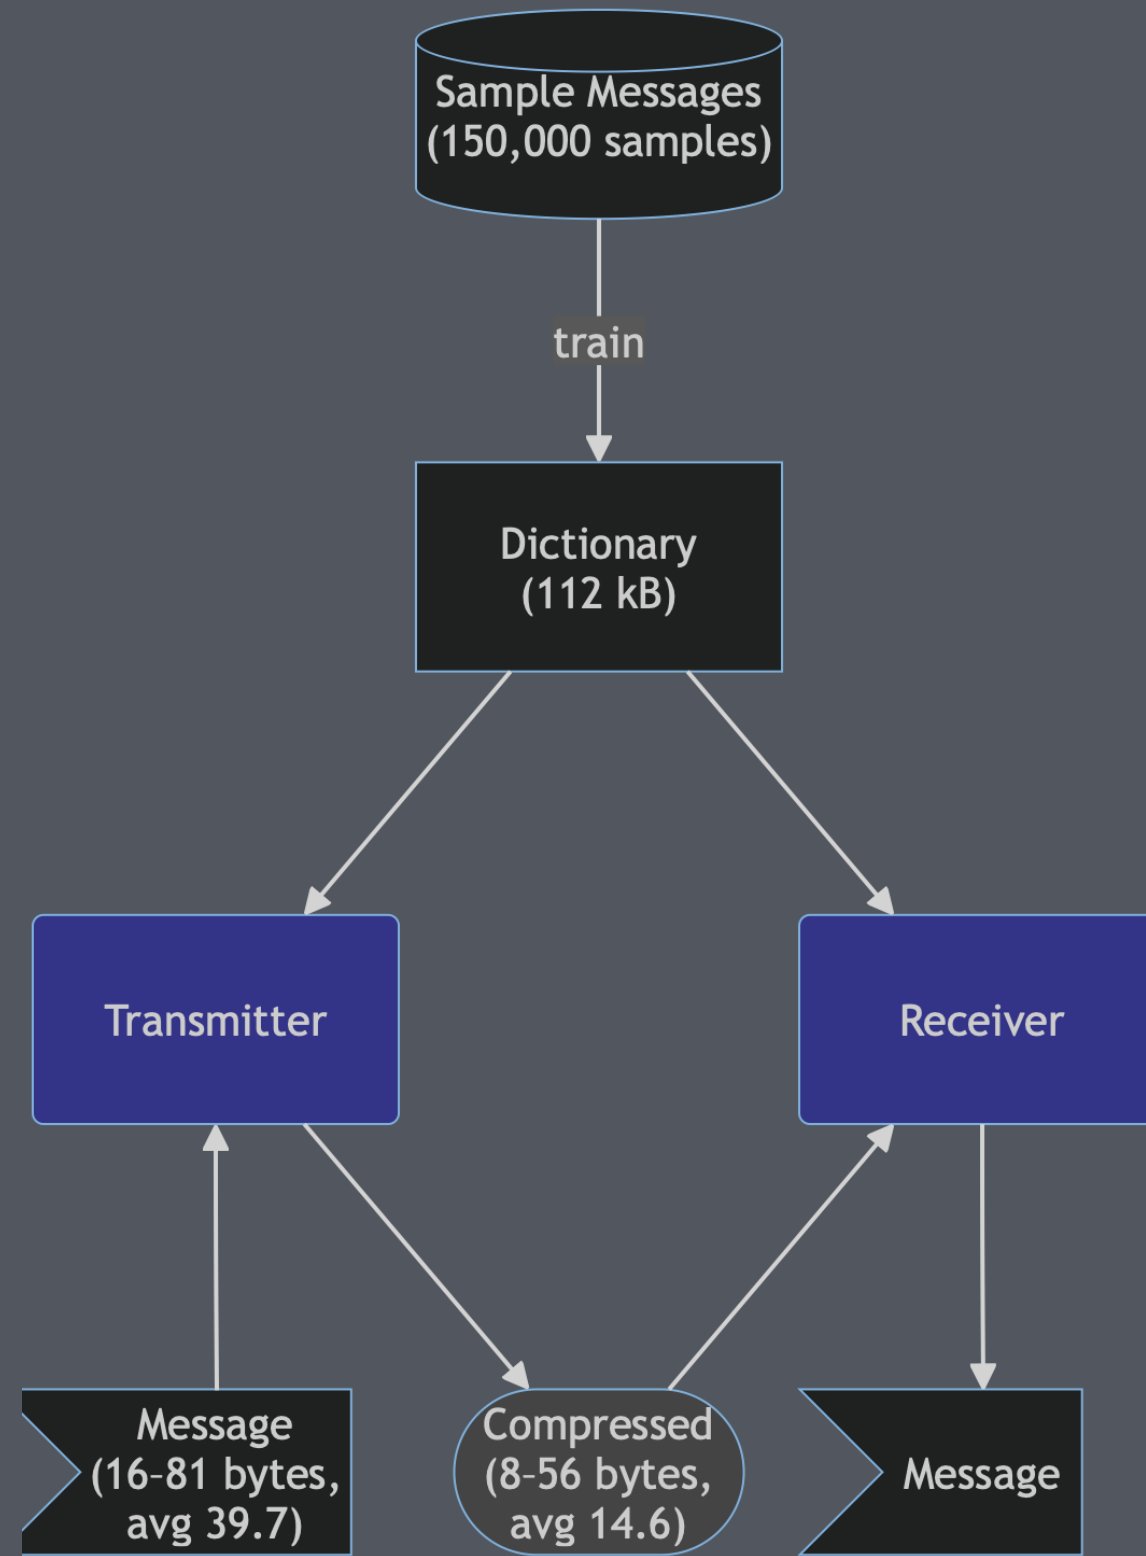
```
$PCDIN,01F200,271EFE80,38,00461FFFFFF7FFFFFF*50
$PCDIN,01F201,271EFC3,38,00FFFFFFFDB8D7305FF7F40C64400FFFFFFF7F00000007F7F*26
$MXPGN,01F200,2838,00461FFFFFF7FFFFFF*10
!AIVDM,1,1,,A,13P=Jk@0iN0PVn8Heh<;<pqp0000,0*2F
!AIVDM,1,1,,A,402FhgivMm7010eBbHHnGr700000,0*79
$STALK,84,16,59,66,4A,00,FC,00,06*1C
$YDRSA,-4.0,A,,V*4D
$YDVHW,,T,140.0,M,5.10,N,9.44,K*6E
$YDHDM,140.0,M*3A
$YDXDR,C,27.0,C,Air,P,102250,P,barometer*63
$YDMA,30.1944,I,1.0225,B,27.0,C,,C,,C,,T,,M,,N,,M*23
$YDXDR,C,90.0,C,Engine#0,U,13.9,V,Alternator#0,G,1252.00,,EngineHours#0*57
$YDMWV,41.9,R,4.1,M,A*1A
$STALK,90,00,03*67
$STALK,00,02,64,C4,04*1E
$STALK,65,00,00*6E
$STALK,20,01,31,00*40
$STALK,26,04,E6,01,3A,01,09*65
$GPGSV,3,1,12,05,11,036,33,10,07,165,38,16,47,307,43,18,76,077,43*7F
$GPGGA,070003.00,4240.856732,N,00837.913045,E,1,09,0.6,-3.6,M,48.0,M,,*4A
$GPVTG,131.5,T,131.7,M,7.2,N,13.4,K,A*12
$GPRMC,070003.00,A,4240.856732,N,00837.913045,E,7.2,131.5,100723,0.2,W,A*2C
$GPGSA,A,2,05,16,18,23,26,27,28,29,31,,0.9,0.6,0.7*33
$GPBOD,130.5,T,127.3,M,GoTo,From*62
$GPRMB,A,0.002,R,From,GoTo,4236.040,N,00845.293,E,7.3,131.5,7.146,V*3F
$GPXTE,A,A,0.002,R,N*42
$GPBWC,070002.00,4236.040,N,00845.293,E,131.71011670543692,T,128.4,M,7.3,N*32
$GPWPL,4337.638,N,00708.428,E,From*72
$GPWPL,4337.464,N,00709.442,E,Wp0*55
$GPWPL,4236.040,N,00845.293,E,Wp1*5B
$GPRTE,1,1,c,0,Wp0,Wp1*06
$YDVR,41.9,R,8.0,N,4.1,M,14.8,K*58
$PCDIN,01F200,271F0073,38,004D1FFFFFF7FFFFFF*2E
$PCDIN,01F201,271EFEC7,38,00FFFFFFFDB8D7305FF7F40C64400FFFFFFF7F00000007F7F*23
$MXPGN,01F200,2838,004D1FFFFFF7FFFFFF*62
!AIVDM,1,1,,A,33JrcV00h1PR:<@I27S<KI220000,0*68
$STALK,60,0C,01,12,22,00,00,00,00,00,00,00,00,00,00,01*1B
$STALK,9C,51,19,FC*32
$YDRSA,-4.0,A,,V*4D
$YDHDM,141.0,M*3B
$YDRPM,E,0,2000.5,,A*53
$YDMWV,40.9,R,4.4,M,A*1E
!AIVDM,1,1,,A,13VCT200PaP` :o0IB<d3b2mj0000,0*04
$PCDIN,01F200,271F0271,38,00401FFFFFF7FFFFFF*5A
$PCDIN,01F201,271F00BA,38,00FFFFFFFDB8D7305FF7F40C64400FFFFFFF7F00000007F7F*54
$PCDIN,01F214,271F00C5,38,007305FF7FFFFFFF*5F
$MXPGN,01F200,2838,00401FFFFFF7FFFFFF*16
$MXPGN,01F214,6838,007305FF7FFFFFFF*65
$STALK,23,01,17,49*4A
$YDRSA,-4.0,A,,V*4D
```

Generic per-message compression (gzip)

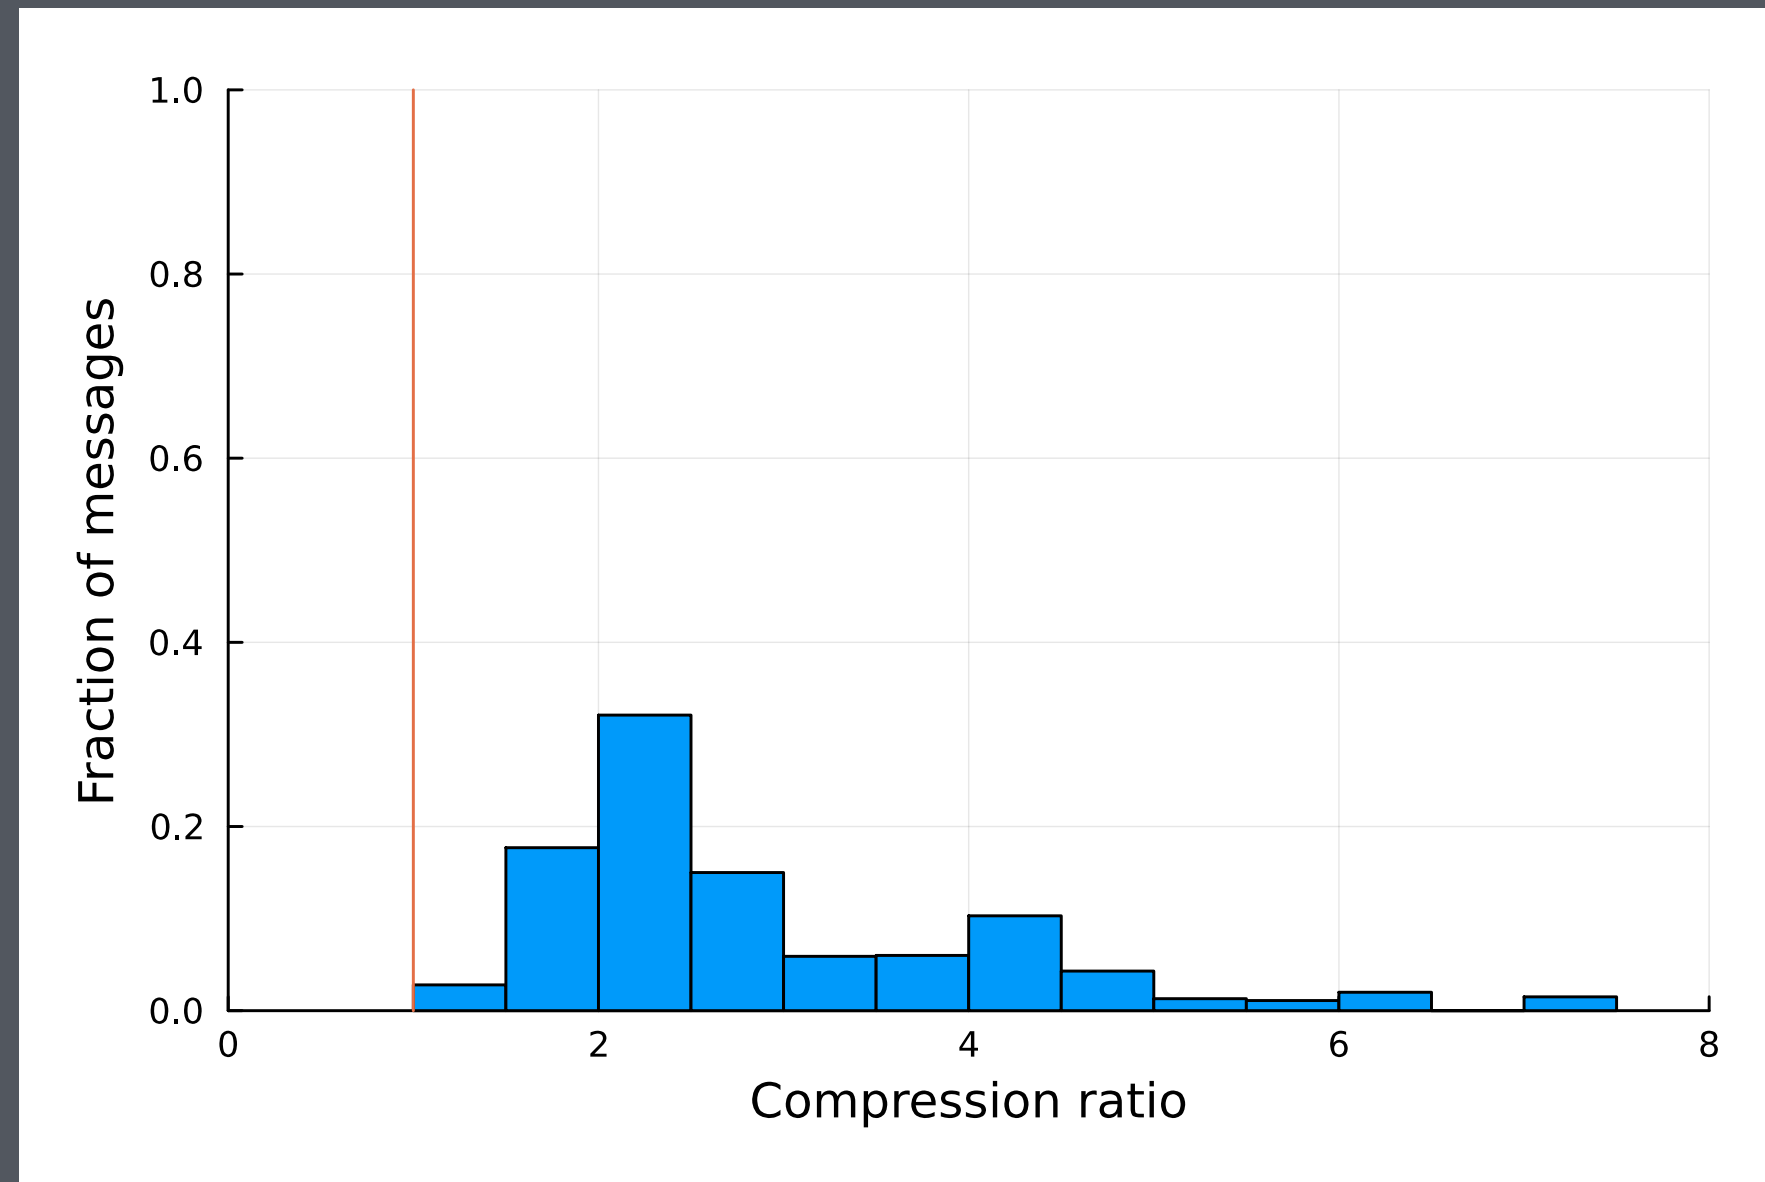


Application-specific compression

- Not all bit patterns are equally likely (compressible!)
- Generic compression performs poorly due to short messages
- Learn distribution from a large sample dataset of messages
- Use distribution (dictionary) as *prior* for message compression




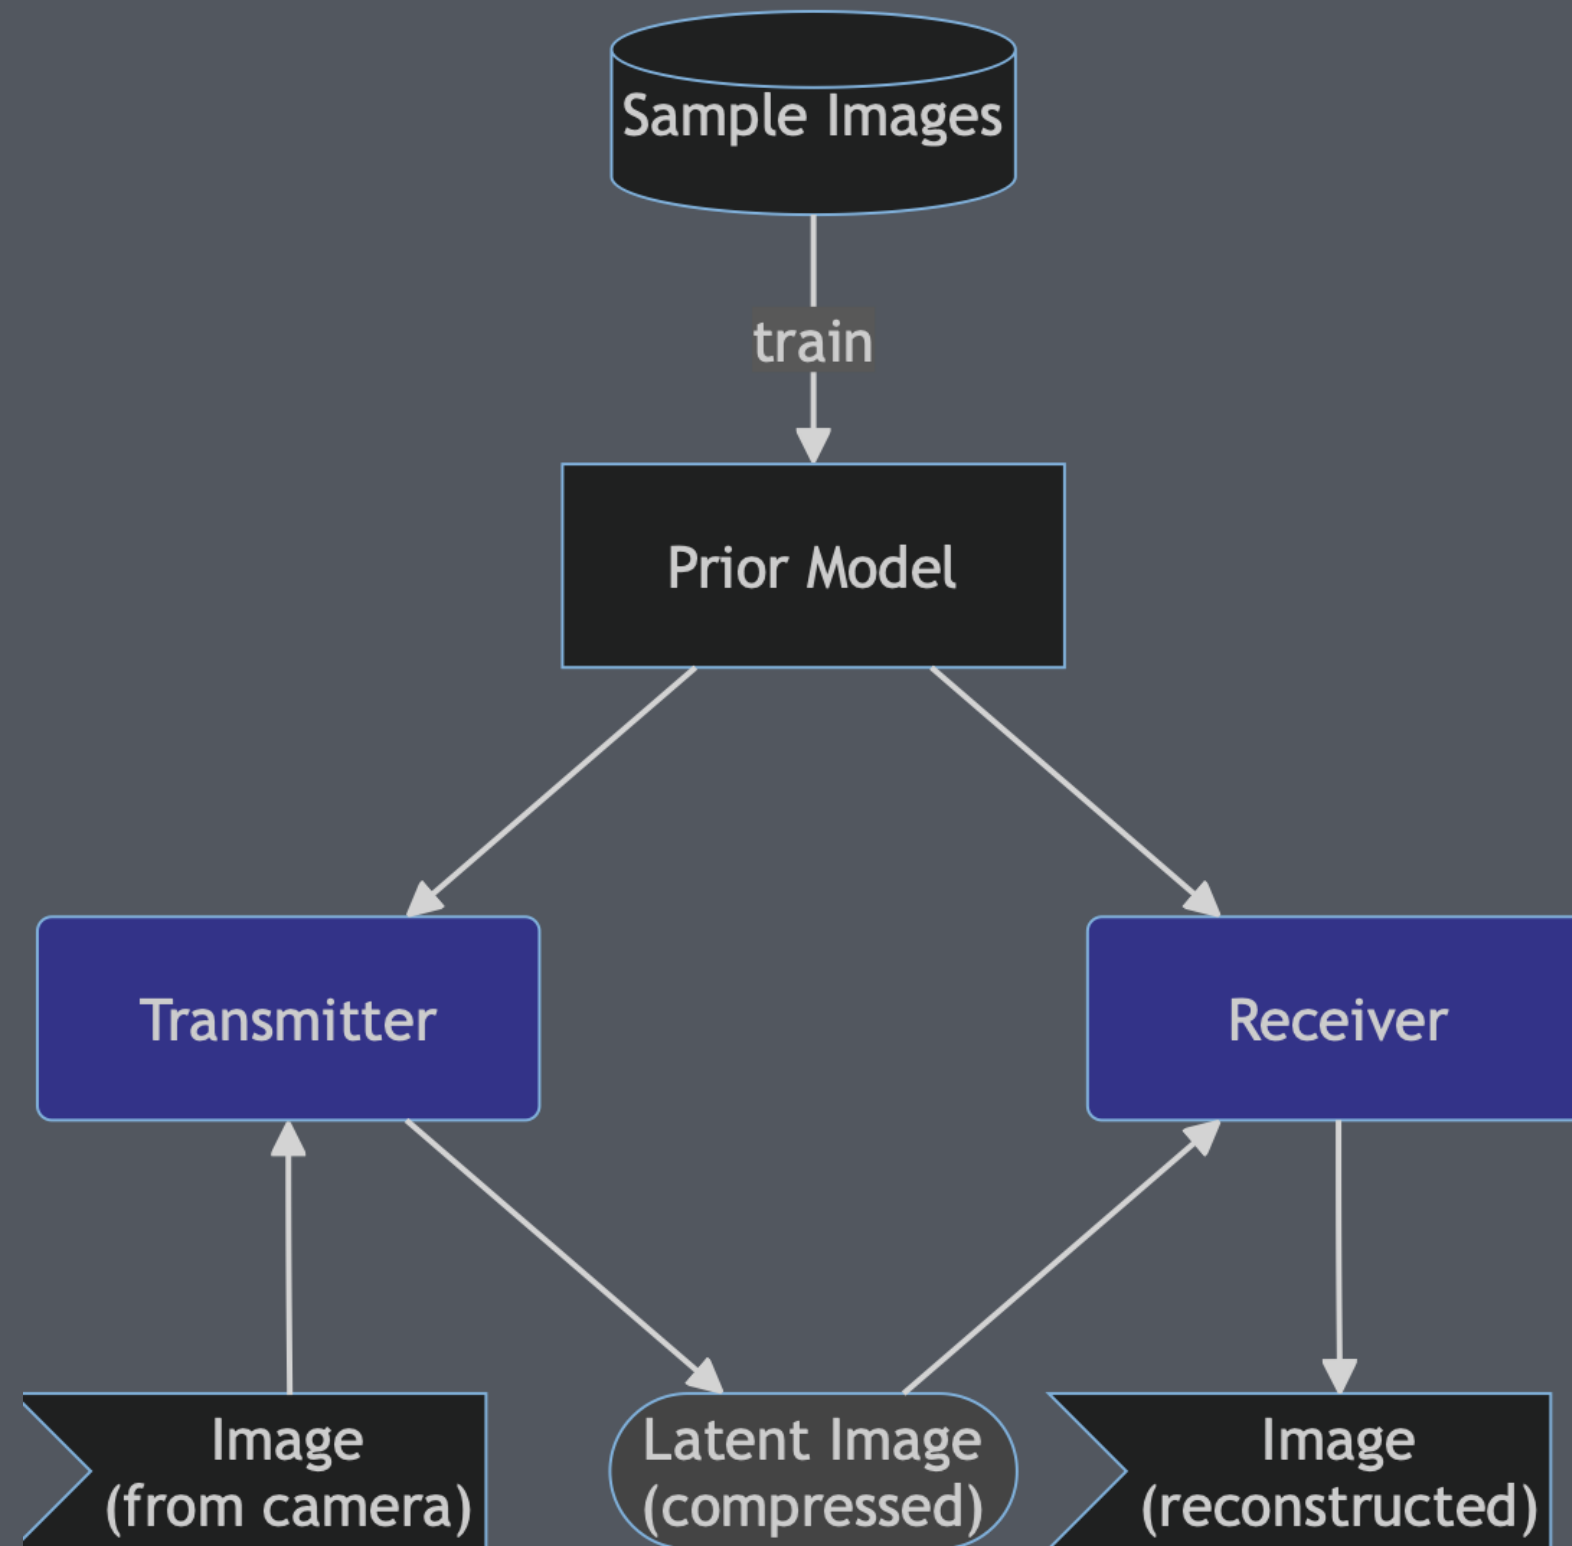
Application-specific compression



Example

Reinspection with a tether-less ROV







Prior model*

$$\text{pose} \equiv (x, y, z, q_1, q_2, q_3, q_4)$$




train
→

⟨ NVS Model (e.g. NeRF+, 3DGS#) ⟩



Synthesized image





* Novel View Synthesis (NVS) techniques


+ Neural Radiance Fields (NeRF)

3D Gaussian Splatting (3DGS)

Prior model (3DGS)



Camera image



$\langle 3DGS^{-1} \text{ Model} \rangle$




$(x, y, z, q_1, q_2, q_3, q_4)$




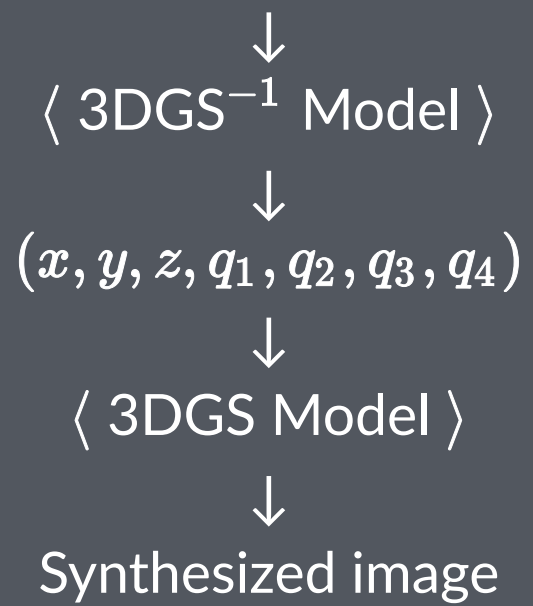

$\langle 3DGS \text{ Model} \rangle$



Synthesized image



Camera image


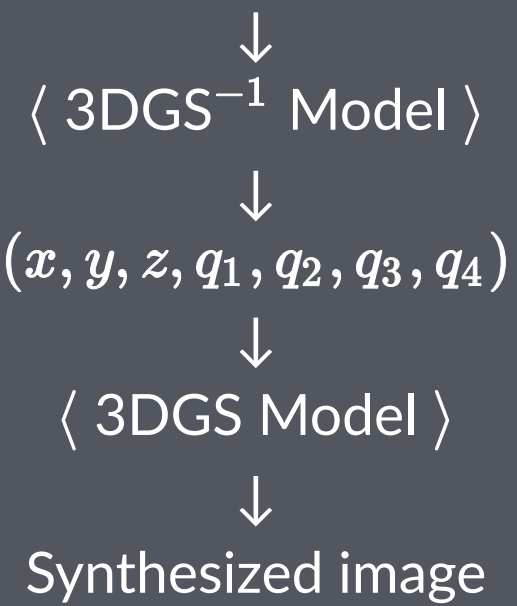



Latent image

$$(\mathbf{x}, \mathbf{y}, \mathbf{z}, q_1, q_2, q_3, q_4) + \text{Difference image}$$



Camera image



Latent image

$$(\mathbf{x}, \mathbf{y}, \mathbf{z}, q_1, q_2, q_3, q_4) + \text{Difference image}$$



Reconstruction

Latent image

$(x, y, z, q_1, q_2, q_3, q_4)$

+




Reconstruction


Latent image

$(x, y, z, q_1, q_2, q_3, q_4)$

+



+




Reconstruction


Latent image

$$(x, y, z, q_1, q_2, q_3, q_4)$$


+





+

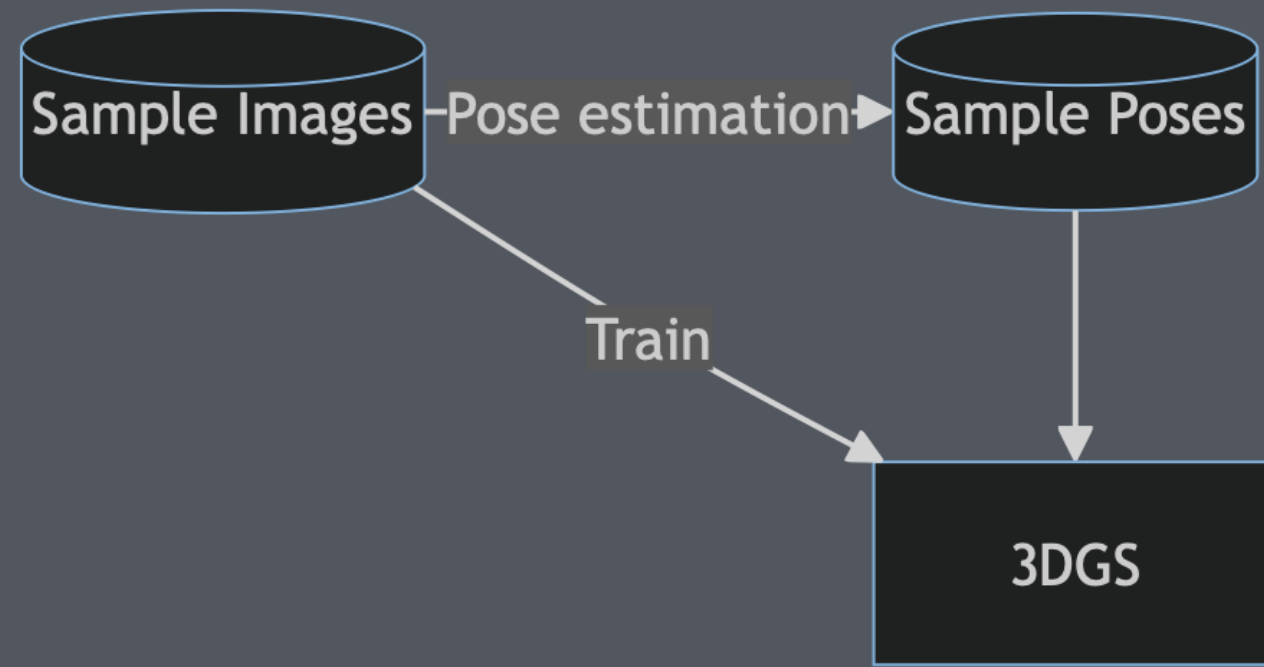



Reconstructed

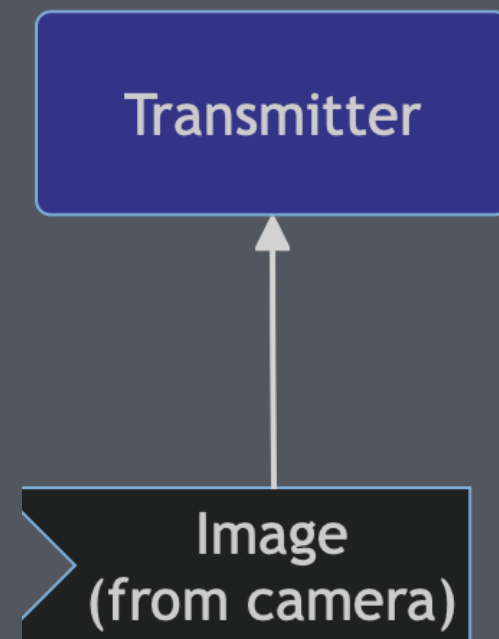



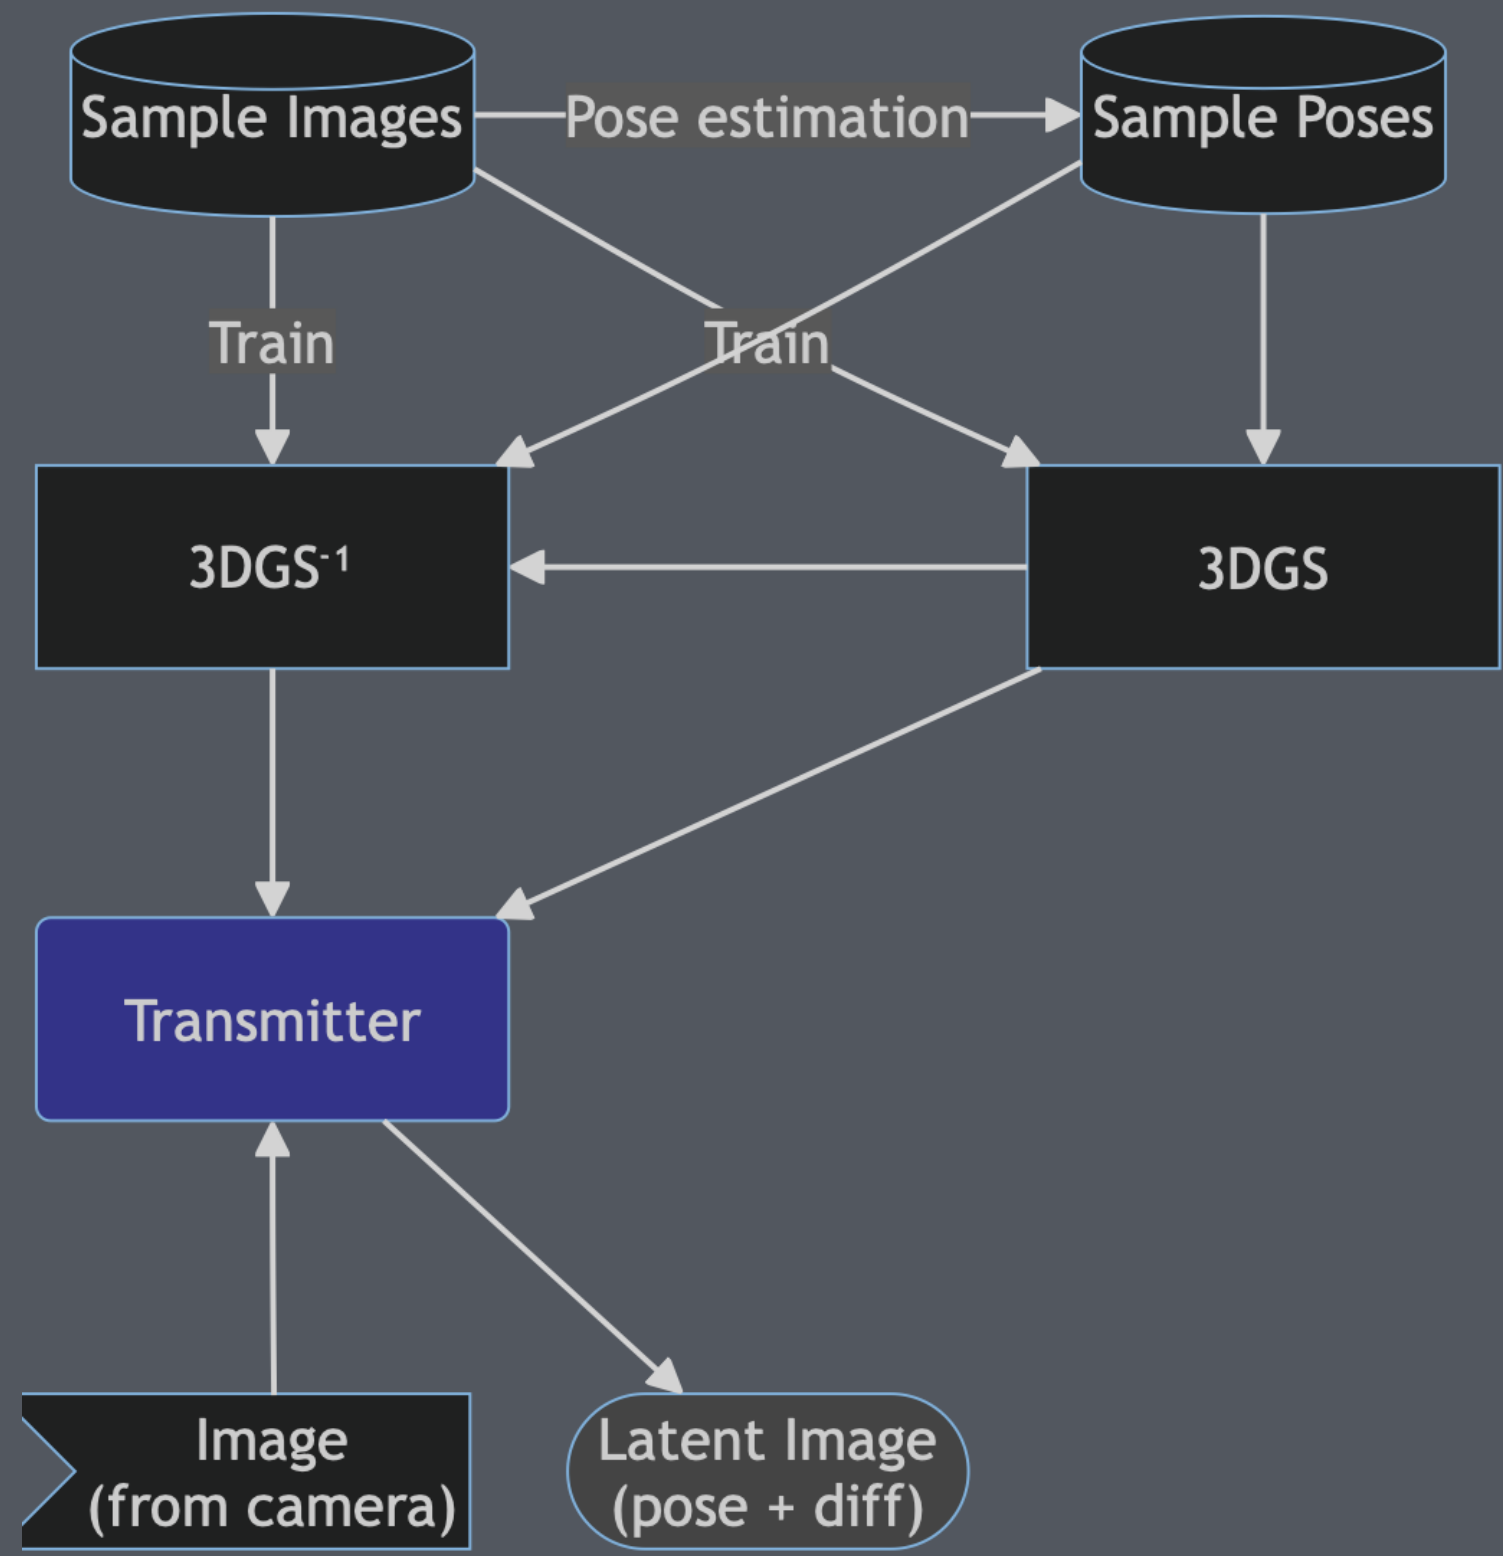


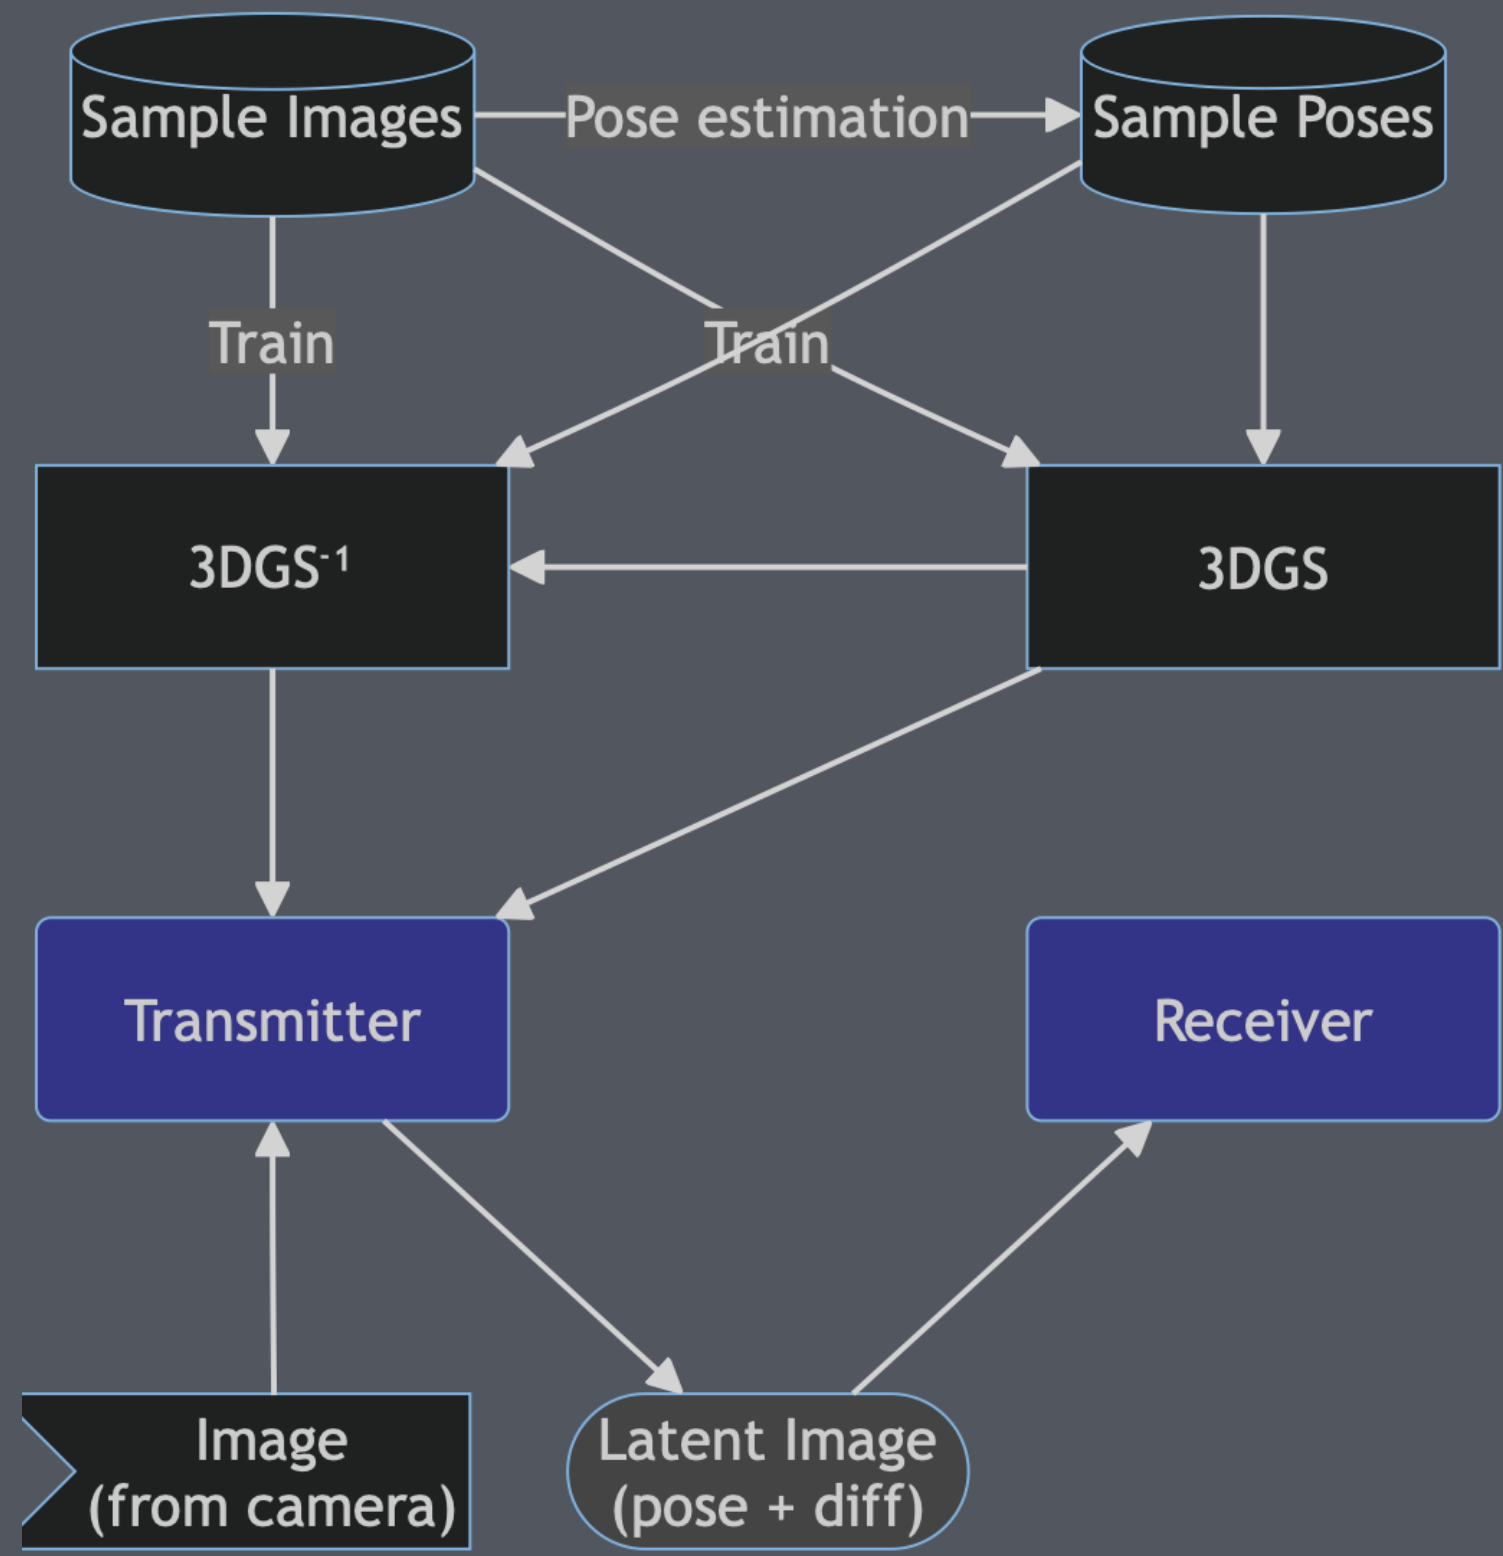


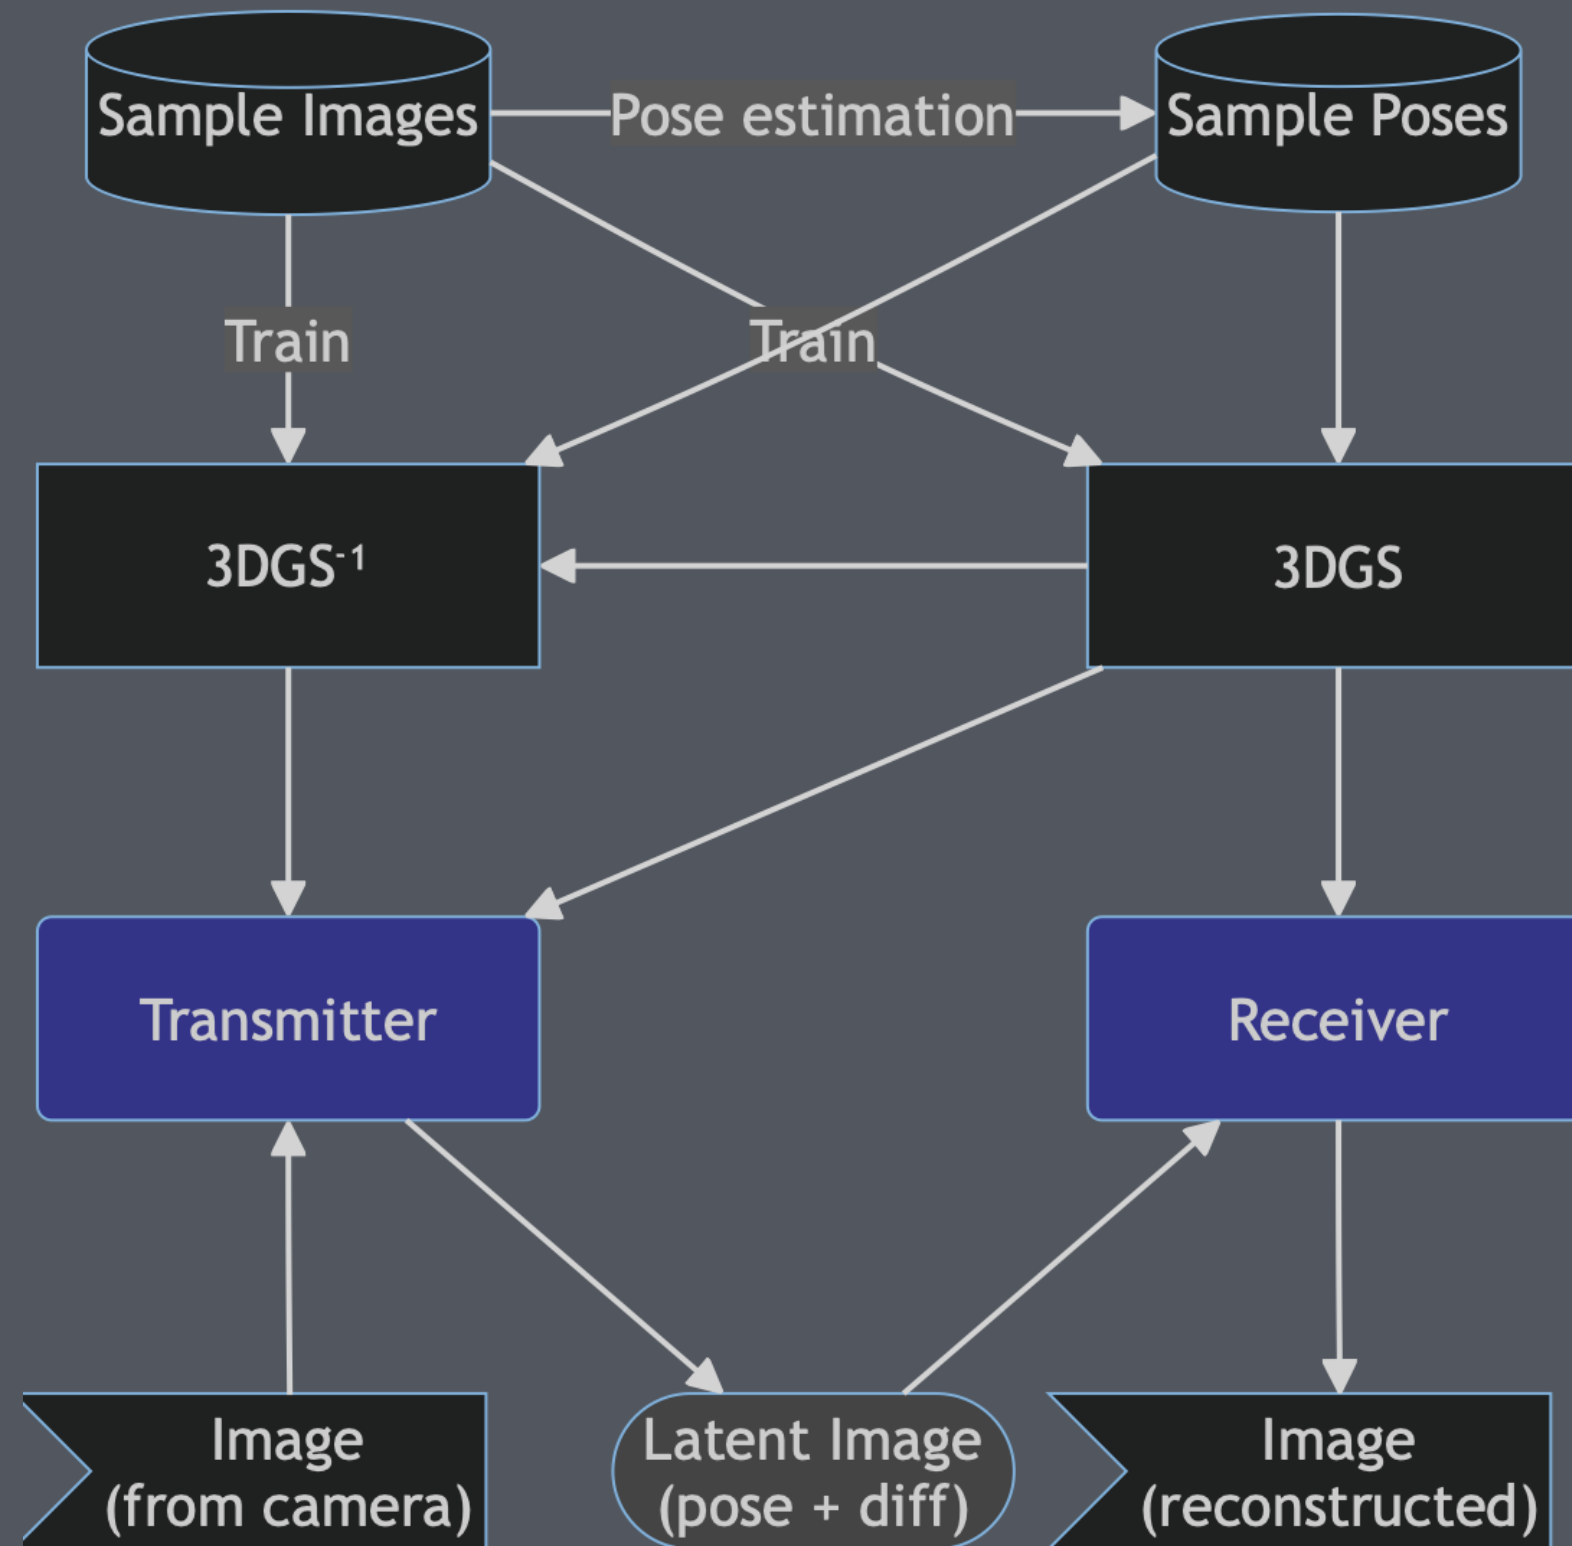















Camera

45807 bytes Reconstructed




3DGS





Difference

4688 bytes
ratio: 9.8x




Camera

43302 bytes Reconstructed



3DGS



Difference

7789 bytes
ratio: 5.6x



Camera

33297 bytes Reconstructed

PSNR: 37.9 dB

3DGS





Difference

3147 bytes
ratio: 10.6x

Acknowledgements[&]

... based on excellent work by:

Too Yuen Min, Peng Luyuan, Hari Vishnu, Bharath Kalyan, Rajat Mishra, Tan Soo Pieng



[&] This work was partially supported by A*STAR under its RIE2020 Advanced Manufacturing and Engineering (AME) Industry Alignment Fund - Pre-Positioning (IAF-PP) Grant No. A20H8a0241.

Takeaways

Use frugal communication protocols to minimize overheads

Do **more** with *less* bits by leveraging priors