AUTONOMOUS
INTELLIGENCE GROUP

LECTURE 10 — OBSERVATION MODELS (II): COMPLEMENTARY FILTER AND TIME OF FLIGHT

INTRODUCTION TO AUTONOMOUS SYSTEMS · MODULE III

Topics: Complementary filter · CF → KF cascade · Time of flight · SNR & detection · Matched filter · Resolution & pulse compression

AGENDA

§	TOPIC	KEY IDEAS
9.3	**The Complementary Filter**	Gyro + accelerometer · = steady-state KF · cascade vs. the 15-state filter
10.1	**Time of Flight and the Matched Filter**	SNR · detection thresholds · matched filter · sufficiency · range resolution · pulse compression, FMCW

> *Lecture 9 derived the pipelines of proprioceptive sensors — and they drift. Today: fuse the IMU's sensors, then derive the pipeline of the sensors that measure the world.*

LITERATURE

TOPIC	REFERENCE
Complementary filter on $SO(3)$	Mahony, Hamel & Pflimlin — "Nonlinear Complementary Filters on the Special Orthogonal Group", *IEEE TAC* (2008)
SNR, matched filter, thresholds	Kay — *Fundamentals of Statistical Signal Processing, Vol. II: Detection Theory* (1998), Ch. 3–4
Radar, Doppler, CFAR	Richards — *Fundamentals of Radar Signal Processing* (2nd ed., 2014), Ch. 4–6
Measurement models, SLAM	Thrun, Burgard & Fox — *Probabilistic Robotics* (2005), Ch. 6, 10
Information theory, DPI	Cover & Thomas — *Elements of Information Theory* (2006), Ch. 2

RECAP — LECTURE 9

- **The pipeline.** $x \to s \to o = T(s)$. Processing can only lose information: $I(x; T(s)) \le I(x; s)$ (**DPI**).
- **Lossless** $\iff$ $T(s)$ is a **sufficient statistic** $\iff$ $p(s \mid x) = g(T(s), x)\, k(s)$ (factorisation).
- **Proprioceptive sensors** measure *derivatives* of the pose — encoders, gyro, accelerometer — with derived noise models ($\Delta^2/12$, white noise + bias random walk from the Allan plot).
- **Integration drifts.** Gyro bias → ramp, white noise → random walk. The accelerometer's gravity reading is absolute in roll and pitch, but noisy.

Today: fuse gyro and accelerometer (§9.3), then the sensors that measure the pose relative to the world — derive their pipeline from signal to range (§10.1).

§9.3 · THE COMPLEMENTARY FILTER

TWO DRIFTING SENSORS, ONE STEADY-STATE KALMAN FILTER

THE COMPLEMENTARY FILTER

Estimate a **tilt** angle θ — **roll or pitch**, one filter each; not yaw, which gravity cannot give — from two sensors:

- **Gyro:** integrate the rate, $\theta_g = \theta(0) + \int \omega_m \, dt$ — smooth and precise short-term, but **drifts**.
- **Accelerometer:** gravity's direction, roll $\theta_a = \operatorname{atan2}(f_y, f_z)$, pitch $\operatorname{atan2}(-f_x, \sqrt{f_y^2 + f_z^2})$ — **absolute**, but noisy and fooled by every linear acceleration. Yet **right over long horizons**: each reading is computed afresh, so nothing accumulates, and averaged over time its errors cancel — noise shrinks like $1/\sqrt{T}$.

They fail in **opposite** ways, so blend them: follow the gyro, and pull a little towards the accelerometer every step,

$$\hat{\theta}_t = \alpha \left(\hat{\theta}_{t-1} + \Delta t \, \omega_{m,t-1} \right) + (1 - \alpha) \, \theta_{a,t}, \qquad \alpha \in (0, 1).$$

THE COMPLEMENTARY FILTER IS A STEADY-STATE KALMAN FILTER

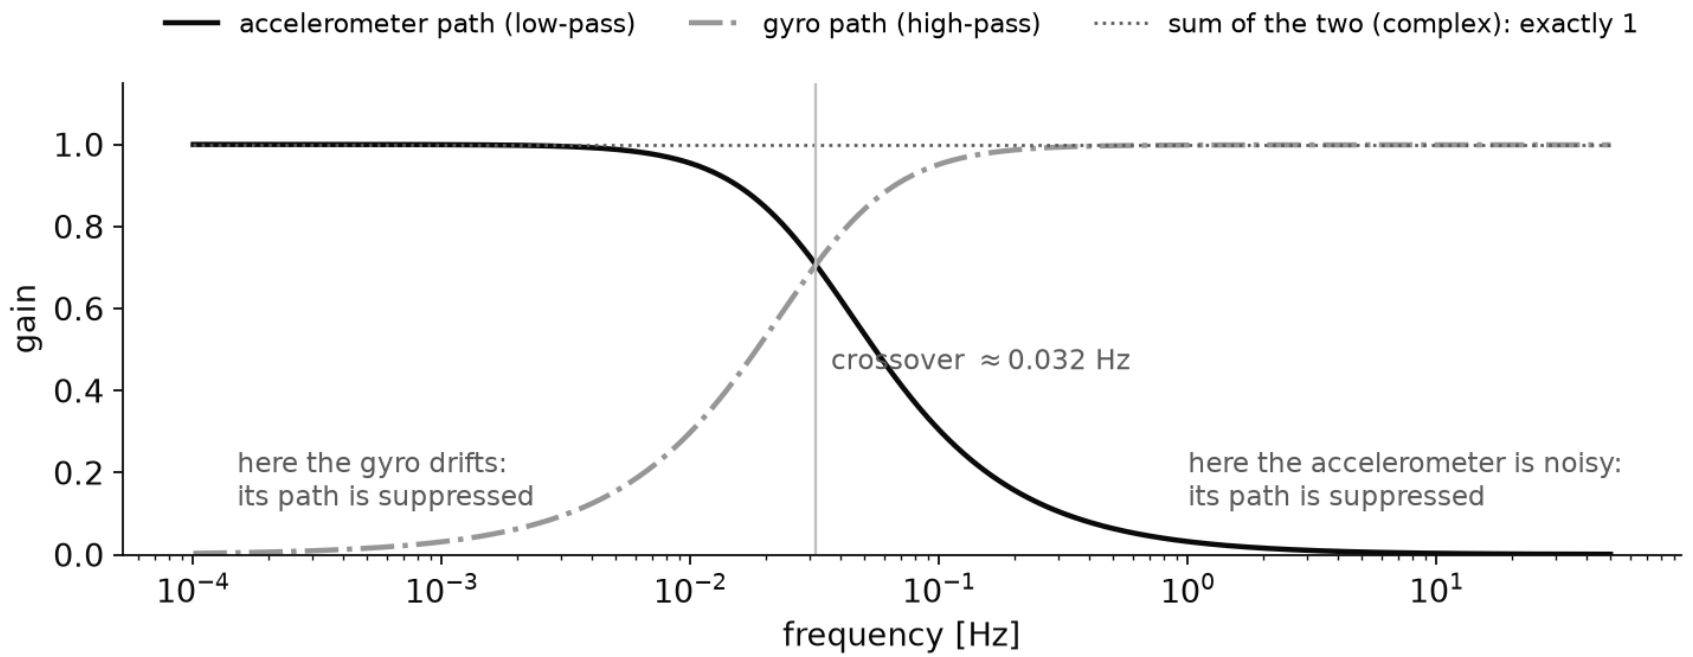
Proposition. Gyro as the action, $a_{t-1} = \omega_{m,t-1}$: $\theta_t = \theta_{t-1} + \Delta t\, \omega_{m,t-1} + w_t$, $w_t \sim \mathcal{N}(0, q)$; accelerometer as the sensor: $\theta_{a,t} = \theta_t + v_t$, $v_t \sim \mathcal{N}(0, r)$. The steady-state KF is the complementary filter with $\alpha = 1 - K_\infty$:

$$K_\infty = \frac{P_\infty}{P_\infty + r}, \qquad P_\infty = \frac{q + \sqrt{q^2 + 4qr}}{2},$$

with P_∞ the steady-state *predicted* variance (Lec 7, §4.6).
For $q \ll r$: $1 - \alpha = K_\infty \approx \sqrt{q/r}$.

Proof. KF of Lec 6, §4 with $A = H = 1$, $B = \Delta t$. Update: $\mu_t = (1 - K_t)\,\bar{\mu}_t + K_t\,\theta_{a,t}$, $\bar{\mu}_t = \mu_{t-1} + \Delta t\, \omega_{m,t-1}$ → complementary filter with $\alpha = 1 - K_t$.
$(A, H) = (1, 1)$ observable; the mode $\lambda = 1$ is excited by process noise ($q > 0$) → Property 4 of Lec 7, §4.6: $K_t \to K_\infty$.
Steady state, P predicted, Σ updated: $P = \Sigma + q$, $\Sigma = Pr/(P + r)$ → $P^2 - qP - qr = 0$ → positive root P_∞.
$q \ll r$: $P_\infty \approx \sqrt{qr}$, $K_\infty \approx \sqrt{q/r}$. □

COMPLEMENTARY FILTER — FREQUENCY VIEW

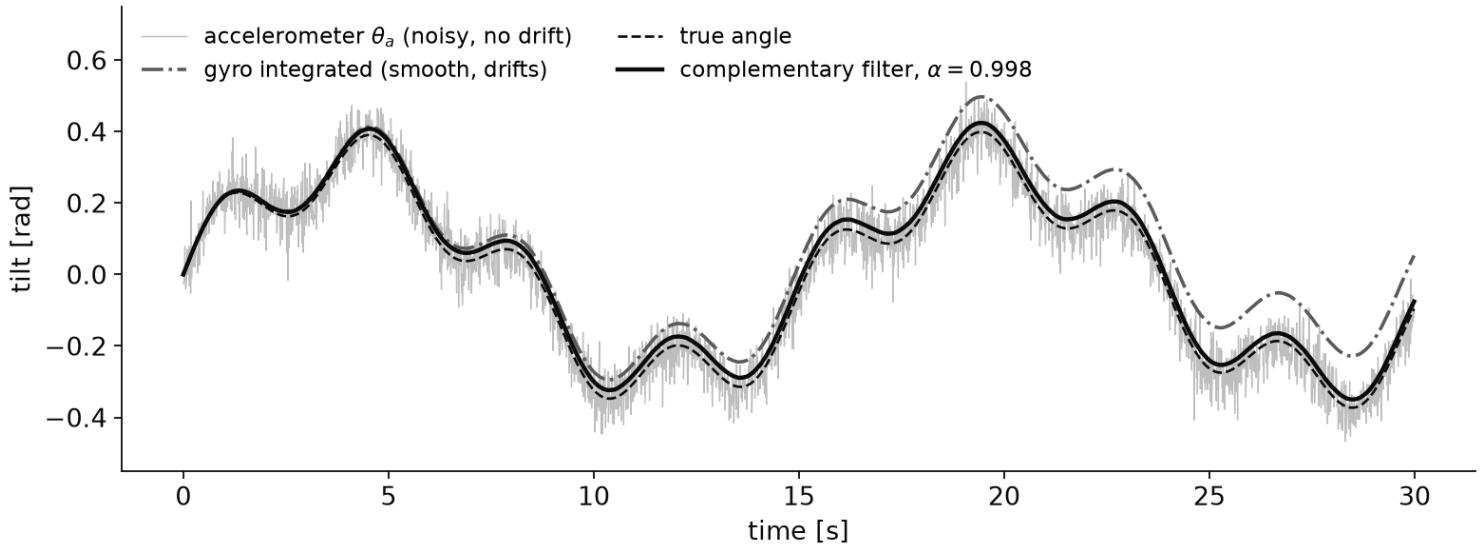


- **Low-pass** on the accelerometer angle, **high-pass** on the gyro-integrated angle. The two transfer functions **sum to one** ("complementary"): the true angle passes undistorted, each sensor is suppressed where its error lives. Plotted: $\alpha = 0.998,\ \Delta t = 10\ \text{ms}$.

COMPLEMENTARY FILTER — THE NUMBERS

Gyro noise 0.01 rad/s at $\Delta t = 0.01$ s ($q = 10^{-8}$), accelerometer tilt noise 0.05 rad ($r = 2.5 \times 10^{-3}$):

$$1 - \alpha \approx \sqrt{q/r} = 2 \times 10^{-3}, \qquad \alpha \approx 0.998.$$



α is the gyro/accelerometer noise ratio, not a magic constant. Gyro bias 0.005 rad/s: integrated, it drifts 0.15 rad in 30 s; the CF keeps a constant offset of ≈ 0.025 rad. The KF adds a covariance, a transient gain, and **bias estimation**.

CASCADE: COMPLEMENTARY FILTER → KALMAN FILTER

Split the 15 states between two filters run in series:

1. **Attitude — complementary filter** at IMU rate: roll and pitch from the gyro + accelerometer-tilt blend; yaw from the same blend with the magnetometer heading. Gyro bias b_g: the mean of ω_m while standing still.
2. **Position — linear Kalman filter** on $x = (p, v, b_a) \in \mathbb{R}^9$, with $\hat{R}$ from stage 1 treated as **known**:

$$\underbrace{\begin{bmatrix} p \\ v \\ b_a \end{bmatrix}_t}_{x_t} = \underbrace{\begin{bmatrix} I & \Delta t\, I & 0 \\ 0 & I & -\hat{R}_{t-1}\Delta t \\ 0 & 0 & I \end{bmatrix}}_{A_t} x_{t-1} + \underbrace{\begin{bmatrix} 0 \\ (\hat{R}_{t-1} f_m + g)\,\Delta t \\ 0 \end{bmatrix}}_{B\, a_{t-1}} + w_t$$

$$Q = \Delta t \begin{bmatrix} 0 & 0 & 0 \\ 0 & N_a^2 I_3 & 0 \\ 0 & 0 & K_a^2 I_3 \end{bmatrix}, \qquad \text{GNSS:} \quad o_t = \begin{bmatrix} I & 0 & 0 \end{bmatrix} x_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_p^2 I_3)$$

- **Linear**, because the rotation is an input, not a state: the Kalman filter of Lecture 6 as it stands — no Jacobians, no $SO(3)$.

CASCADE VS THE 15-STATE KALMAN FILTER

	CASCADE (CF → KF)	15-STATE ERROR-STATE KF
States	3 angles (CF) + 9 (KF); b_g calibrated at rest	all 15, jointly
Filter	fixed α + linear KF	EKF, linearised on $SO(3)$
Cost	tiny; 9×9 covariance	15×15 covariance + Jacobians
Tilt uncertainty	none — $\hat{R}$ assumed exact	in Σ, correlated with v
GNSS corrects attitude, b_g	never: information flows one way	yes, through the correlations
Linear acceleration	fools the tilt blend	accelerometer is not a tilt sensor; tilt is observed via v

- **Why the coupling matters:** a tilt error $\delta\theta$ puts a fake acceleration $g\,\delta\theta$ into v. At $\delta\theta = 0.01$ rad that is 0.1 m/s², so $\frac{1}{2}at^2 = 5$ m of position error after 10 s without a GNSS fix — which the cascade's KF can neither predict nor trace back to the tilt.
- **Rule of thumb:** cascade for gentle motion and an accurate tilt (simple flight controllers, early PX4); the joint filter for aggressive manoeuvres, or when GNSS must calibrate the IMU (PX4 EKF2, ArduPilot EKF3).

§10 · RANGING SENSORS

DRIFT IS BOUNDED ONLY BY MEASURING THE WORLD — AND THE FILTER FALLS OUT OF THE LIKELIHOOD

WHAT ARE RANGING SENSORS?

Ranging sensors are exteroceptive: they measure **distance** (often also direction or velocity) to things in the world by **time of flight** — time a known signal: an **echo** goes out and back, range $= c\,\Delta t/2$; **GNSS** is one-way (satellite → receiver, passive), range $= c\,\Delta t$.

SENSOR	SIGNAL	MEASURES	RANGE · PRECISION	RATE	FAILS ON
Sonar	ultrasound	range, wide beam ($\sim 25^\circ$)	$\lesssim 5$ m · 1 cm	10–40 Hz	soft, angled surfaces
Lidar	infrared laser	ranges, 10^3–10^5 beams/scan	~ 100 m · 2 cm	10–20 Hz	glass, dark surfaces, fog
Radar	77 GHz radio	range, azimuth, **radial velocity**	~ 200 m · 10 cm	10–20 Hz	coarse angle, clutter
GNSS	satellite radio	pseudoranges → global position	global · m (cm RTK)	1–10 Hz	indoors, urban canyons

- **Why they matter:** they measure the pose **relative to the world** (map, satellites, landmarks), so they bound the drift of §9.
- **Speed sets the timing:** sound 343 m/s (a 1 m echo: 6 ms); light 3×10^8 m/s (1 cm of range: 67 ps). Always the same problem: find the delay of a known signal in noise.

RTK (real-time kinematic): cm-level GNSS with a base station — §10.3, Lecture 11.

§10.1 · TIME OF FLIGHT AND THE MATCHED FILTER

FIND A KNOWN PULSE BURIED IN NOISE

TIME OF FLIGHT — THE MODEL

Setup. The sensor emits a known short pulse $\psi(t)$ (duration T_p), then **records**: the analogue-to-digital converter samples the received voltage every T_s, keeping L samples — listening time $LT_s \geq 2r_{\max}/c + T_p$, long enough for an echo from the largest range of interest $r_{\max}$.

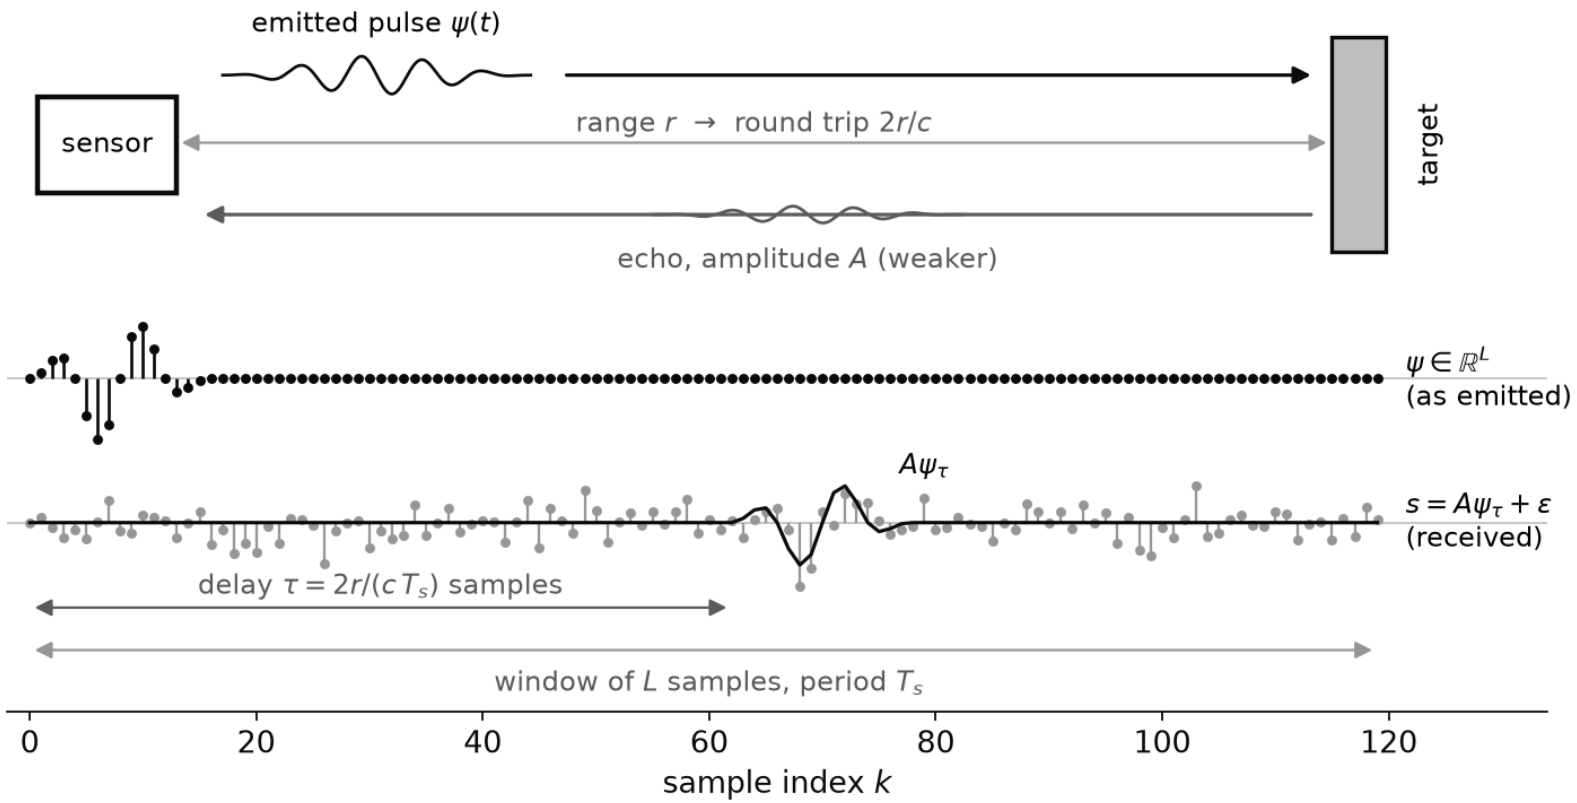
- **Pulse:** $\psi_k = \psi(kT_s)$ — nonzero only for the first T_p/T_s samples, then **zero-padded** to $\psi \in \mathbb{R}^L$ so it can be shifted and compared within the window. **Energy** $\|\psi\|^2 = \sum_k \psi_k^2$.
- **Delay:** round-trip time $2r/c$ (c: light for lidar/radar, sound for sonar). In samples: $\tau = 2r/(cT_s)$, taken integer (finer: interpolate the peak).
- **Echo** = the round-tripped pulse itself, **one target assumed**: $(\psi_\tau)_k = \psi_{k-\tau}$, scaled by an unknown amplitude A (range, reflectivity). The window holds it, so a delay only moves samples: $\|\psi_\tau\| = \|\psi\|$.
- **Noise:** white Gaussian, independent across samples, variance σ^2.

$$s = A\,\psi_\tau + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma^2 I_L).$$

Known: ψ, σ^2. **Unknown:** τ, A. Estimating the range = estimating the delay. The pulse is buried in noise — how do we find it?

Real beams: **no** echo (nothing in range, glass) or **several** (branch + wall, multipath), $s = \sum_i A_i \psi_{\tau_i} + \varepsilon$.

TIME OF FLIGHT — THE PICTURE



Top: the physics — the pulse travels to the target and back, $2r/c$, and returns weaker. **Bottom:** what the receiver stores — the window $s \in \mathbb{R}^L$: the known pulse shifted by τ samples, scaled by A, plus white noise.

SIGNAL-TO-NOISE RATIO

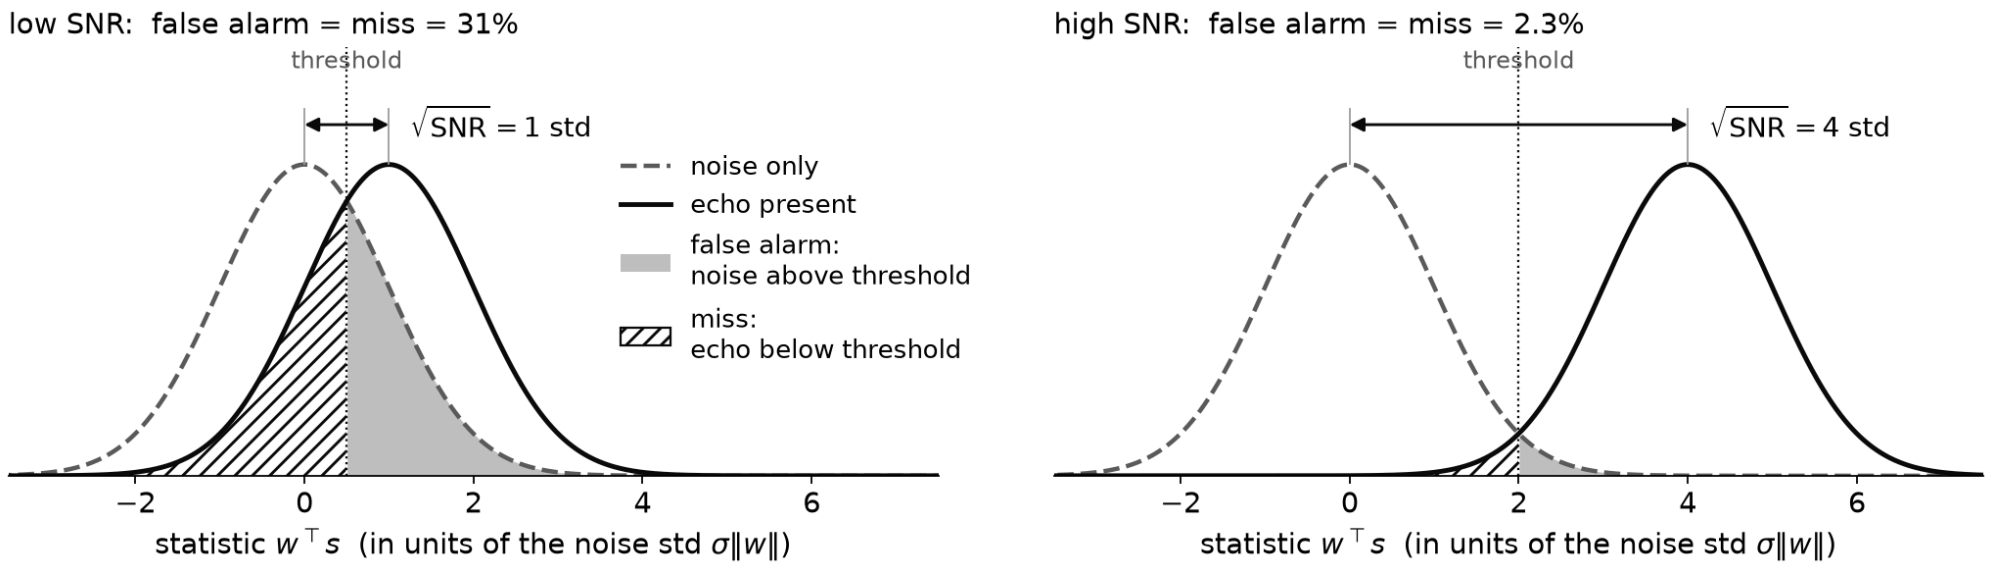
Suppose the echo is at delay τ. Which linear statistic $w^\top s$ reveals it best? Signal part $A\, w^\top \psi_\tau$, noise part $w^\top \varepsilon$ with variance $\sigma^2 \|w\|^2$.

> **Definition (Signal-to-noise ratio).** The SNR of a statistic is its squared signal part divided by the variance of its noise part. For $w^\top s$,
>
> $$\mathrm{SNR}(w) = \frac{\left(\mathbb{E}[w^\top s]\right)^2}{\mathrm{Var}(w^\top s)} = \frac{A^2\,(w^\top \psi_\tau)^2}{\sigma^2\,\|w\|^2}.$$

Meaning: with no pulse, $w^\top s$ has mean 0, std $\sigma\|w\|$; the pulse shifts the mean to $A\, w^\top \psi_\tau$. $\sqrt{\mathrm{SNR}}$ = this **shift in noise std's**. Rescaling w changes neither.

SNR AND DETECTION

Decide "echo or noise?" by thresholding $w^\top s$ — the two cases are two bells $\sqrt{\mathrm{SNR}}$ std's apart:



- **High SNR** ($\sqrt{\mathrm{SNR}} \gtrsim 4$): the two bells barely overlap — reliable detection, a sharp correlation peak, a precise range.
- **Low SNR** ($\lesssim 1$): they overlap — any threshold misses echoes or fires on noise; the range estimate jumps to noise peaks.
- **Raise it:** more pulse energy $A^2\|\psi\|^2$ (power, duration) or less noise (averaging n pulses multiplies SNR by n).

Which w gives the highest SNR? The definition leaves w free; Cauchy–Schwarz picks it.

> **Theorem (Matched filter).** Among all linear statistics $w^\top s$, $\mathrm{SNR}(w)$ is maximised by $w \propto \psi_\tau$, and its maximum is $A^2\|\psi\|^2/\sigma^2$.

Proof. Bound the numerator of the SNR with Cauchy–Schwarz, $(w^\top \psi_\tau)^2 \le \|w\|^2\|\psi_\tau\|^2$:

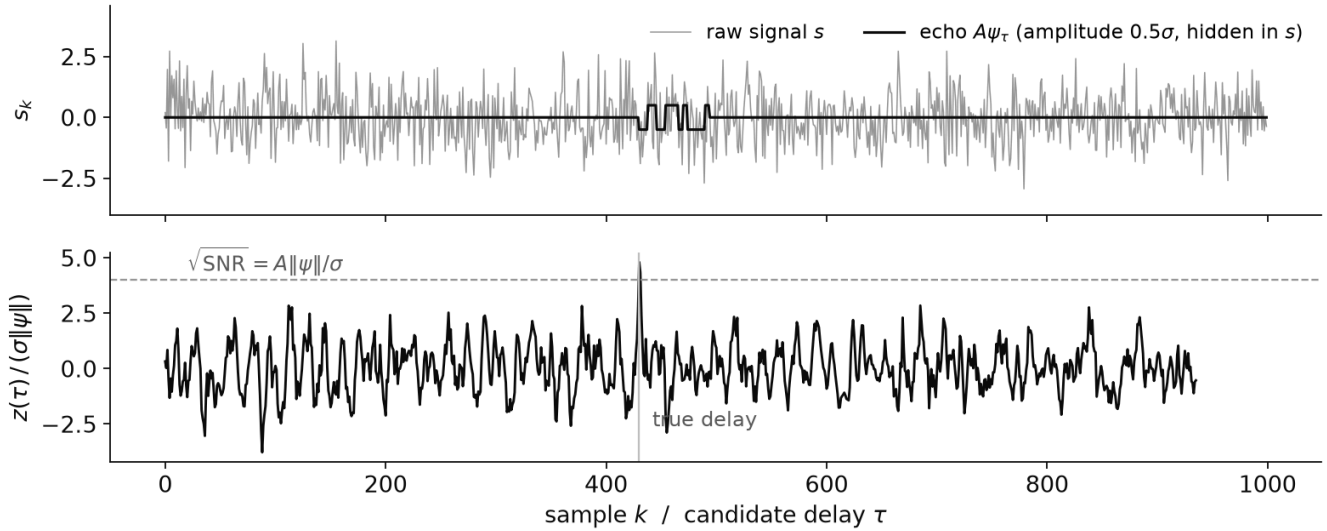
$$\mathrm{SNR}(w) = \frac{A^2\,(w^\top \psi_\tau)^2}{\sigma^2\,\|w\|^2} \;\le\; \frac{A^2\,\|w\|^2\,\|\psi_\tau\|^2}{\sigma^2\,\|w\|^2} = \frac{A^2\,\|\psi\|^2}{\sigma^2},$$

using $\|\psi_\tau\| = \|\psi\|$. Equality in Cauchy–Schwarz holds iff $w = c\,\psi_\tau$ (collinear, any $c \neq 0$). $\square$

Conclusion: take $w = \psi_\tau$, so the best statistic is $w^\top s = \psi_\tau^\top s = \sum_k s_k\,\psi_{k-\tau}$ — **correlate the received signal with a copy of the pulse** shifted to the candidate delay. "Matched": the weights match the pulse. The best SNR depends only on the pulse **energy** $A^2\|\psi\|^2$, not its shape.

THE CORRELATION FUNCTION

The true delay is unknown, so correlate with the pulse copy at every candidate delay: $z(\tau) = \sum_k s_k \, \psi_{k-\tau} = \psi_\tau^\top s$, the convolution of s with the time-reversed pulse ("filter"). **Range estimate = location of the peak.**



- **Bottom, in noise std's** $z/(\sigma\|\psi\|)$: noise $\sim \mathcal{N}(0,1)$ off the echo; mean $A\|\psi\|/\sigma = \sqrt{\text{SNR}}$ at the true delay (dashed).
- **Here:** the pulse is 64 samples, each $+1$ or -1 (a fixed pseudo-random pattern, black on top), so $\|\psi\| = 8$; with $A = 0.5\sigma$, $\sqrt{\text{SNR}} = 4$. One sample: SNR 0.25, invisible; correlating 64 multiplies it by 64.

THE DERIVED PIPELINE IS LOSSLESS

By the model, $s \mid \tau, A \sim \mathcal{N}(A\psi_\tau,\ \sigma^2 I_L)$ — a fixed echo plus white Gaussian noise. Its likelihood, expanding the square with $\psi_\tau^\top s = z(\tau)$ and $\|\psi_\tau\| = \|\psi\|$:

$$p(s \mid \tau, A) \propto \exp\Big(-\tfrac{1}{2\sigma^2}\|s - A\psi_\tau\|^2 \Big) = \underbrace{\exp\Big(-\tfrac{1}{2\sigma^2}\|s\|^2 \Big)}_{k(s)}\ \underbrace{\exp\Big(\tfrac{1}{\sigma^2}\big(A\, z(\tau) - \tfrac{A^2}{2}\|\psi\|^2\big) \Big)}_{g(z,\,(\tau,A))}.$$

- Factorisation theorem (§8) → **the correlation function $T(s) = z(\cdot)$ is a sufficient statistic for the range.**
- **Maximum likelihood (ML, Lec 4):** pick the $(\hat\tau, \hat A)$ that maximise $p(s \mid \tau, A)$ — only g depends on them. Over A: $\hat A = z(\tau)/\|\psi\|^2$, leaving $\log g = z(\tau)^2/(2\sigma^2\|\psi\|^2)$ → $\hat\tau$ maximises $|z(\tau)|$ ($\arg\max_\tau z(\tau)$ if $A > 0$ is known).

> Under white Gaussian noise, the classical pipeline — correlate with the pulse, take the peak — throws away **nothing** about the range $r = c\,T_s\,\tau/2$. It was not designed by trial and error — **it fell out of the likelihood.**

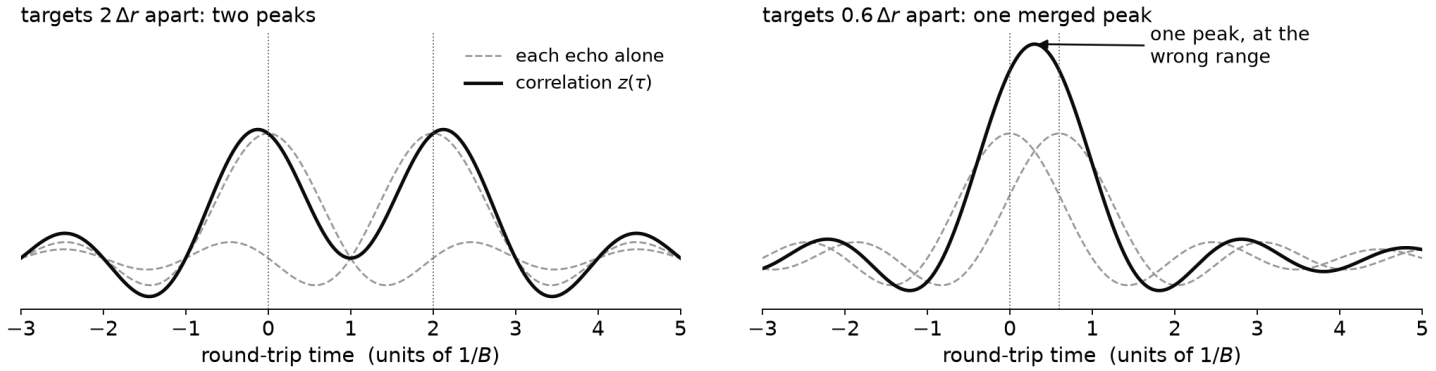
RANGE RESOLUTION

> **Definition (Range resolution).** The range resolution Δr is the smallest distance between two targets in the same beam at which their correlation peaks are still separate.

Peaks are $\sim 1/B$ wide in time (B = pulse bandwidth); echoes $2\Delta r/c$ apart separate when $2\Delta r/c \geq 1/B$:

$$\Delta r = \frac{c}{2B} \qquad (\text{radar with } B = 1\ \text{GHz: } 15\ \text{cm}).$$

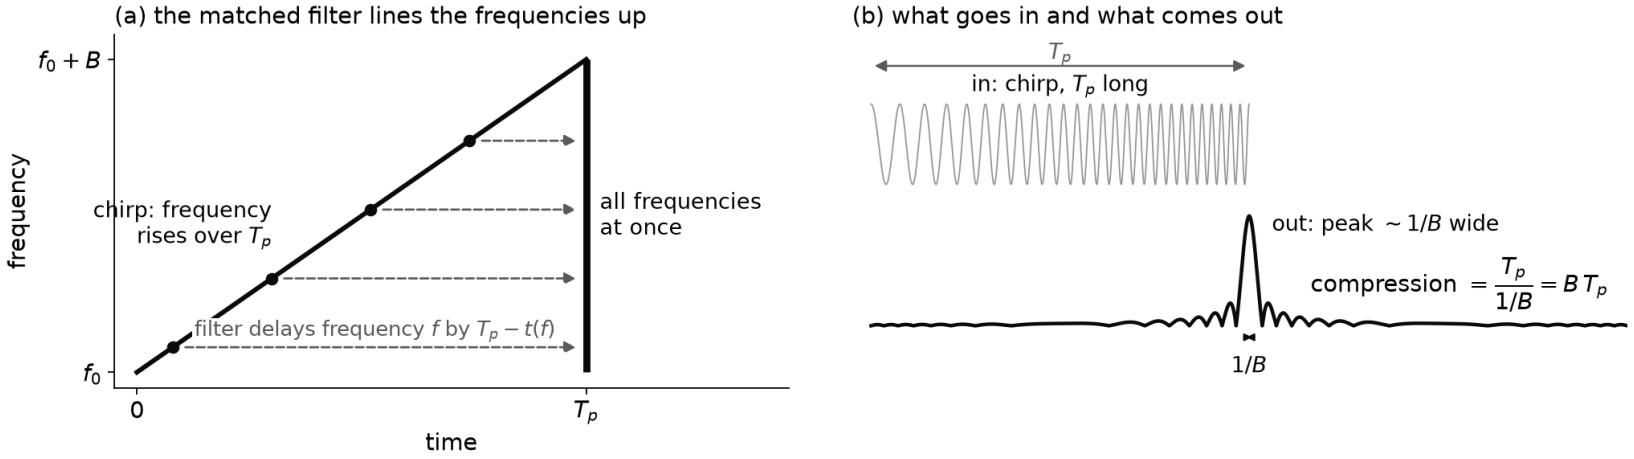
When it matters: several reflectors in one beam (a pedestrian beside a parked car — radar beams are wide; a branch before a wall). Merged peaks give **one wrong range**; a single target needs precision instead.



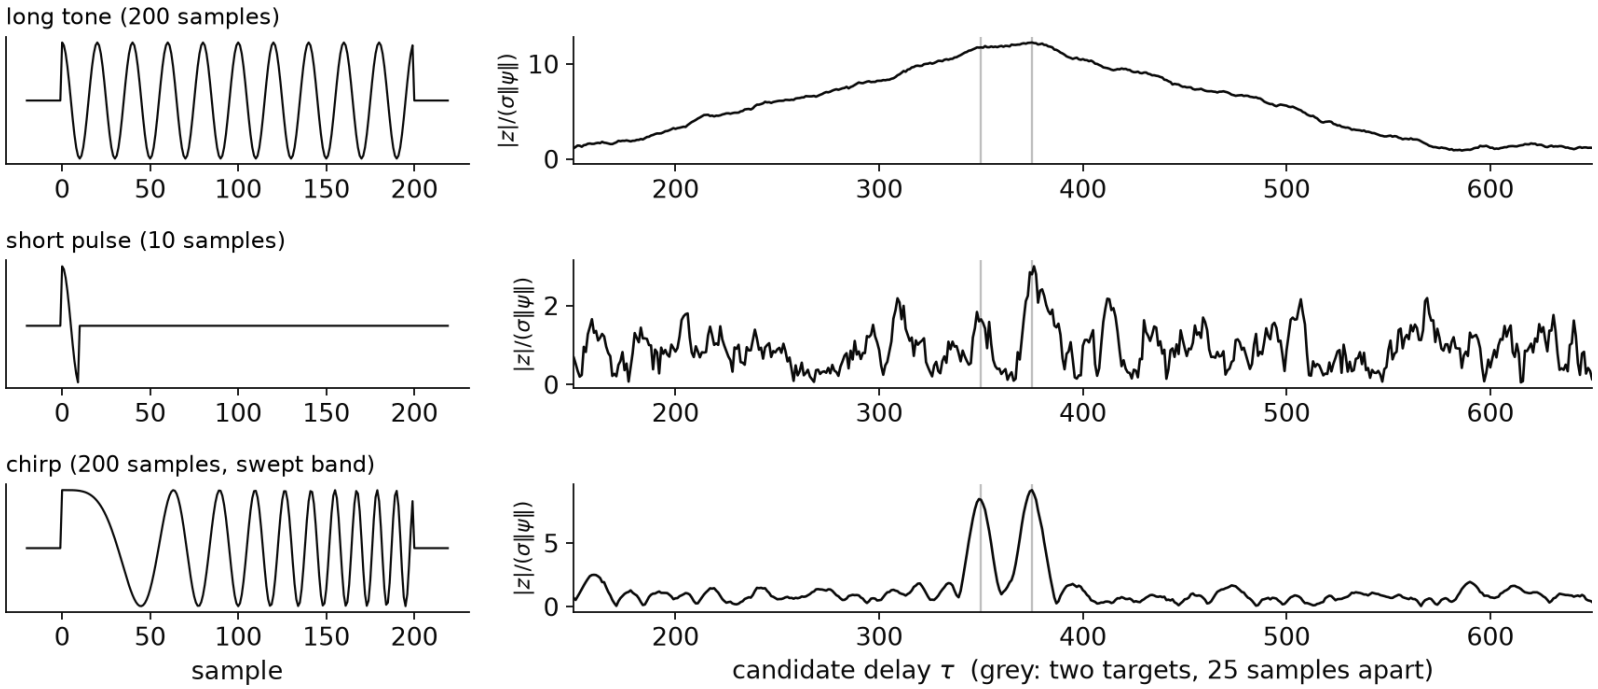
PULSE COMPRESSION

Two demands: detection and precision need **energy** (best SNR $= A^2\|\psi\|^2/\sigma^2$); separating targets needs **bandwidth** ($\Delta r = c/2B$).

- **The trade-off.** Peak power is limited (amplifiers, eye safety), so more energy means a **longer** pulse. A plain pulse of duration T_p has $B \approx 1/T_p$: longer → narrower band → targets merge. One knob, two opposite effects.
- **Pulse compression.** A long **chirp** sweeps its frequency across B: energy from the length, bandwidth from the sweep. The matched filter delays each frequency so all arrive at once (a): T_p long in, $1/B$ wide out — **compressed** by $B\,T_p$ (b). Automotive: $50\ \mu$s, 1 GHz → $5 \cdot 10^4$; a plain pulse has $B\,T_p \approx 1$.



PULSE COMPRESSION — PICTURED



Equal peak amplitude. **Long tone:** targets merge. **Short pulse:** lost in noise. **Chirp:** both resolved.

RESOLUTION VS. PRECISION

- **Resolution** $c/2B$: how far apart **two** targets must be to be separated.
- **Precision** on **one** target's range: much finer, improves with SNR — the std of the round-trip time $\hat{\tau}\, T_s$ scales as

$$\frac{1}{B\sqrt{\mathrm{SNR}}}.$$

> Common confusion: *"resolution and precision are the same."* No — $c/2B$ separates two targets; one target is located far more precisely at high SNR.

SUMMARY — LECTURE 10

§	QUESTION	ANSWER
9.3	How to fuse gyro and accelerometer?	**Complementary filter** = steady-state KF, $1-\alpha = K_\infty \approx \sqrt{q/r}$: high-pass the gyro, low-pass the accelerometer. CF → linear KF cascade vs. the joint 15-state filter.
10.1	How is a range measured?	Find the delay of a known pulse in noise. **Matched filter**: correlation with the pulse maximises SNR $= A^2\lvert\psi\rvert^2/\sigma^2$ and is **sufficient** for the delay. Threshold by Neyman–Pearson. Resolution $\Delta r = c/2B$; precision $\propto 1/(B\sqrt{\mathrm{SNR}})$; pulse compression (chirp, FMCW) gets energy and bandwidth at once.

> **The thread:** every pipeline so far was **derived** — and the matched filter is **exactly** sufficient.

WHAT'S NEXT

Lecture 11 — Observation Models (III): Ranging Sensors and Vision

The matched filter turns a signal into a range. Lecture 11 turns ranges into likelihoods: the **beam model** for lidar and sonar (Monte Carlo localisation), **GNSS** (pseudoranges, four satellites, RTK) and **radar** (Doppler). Then the camera — where geometry gets part of the way and the pipeline has to be **learned** — and the limits of what any observation model can recover.

> *Central question:* What does the likelihood $p(o \mid x)$ of a real ranging sensor look like, and what happens when it cannot be derived?