Diego Rodriguez.

Mechanical engineer drawn to physical products that have to work in the place they're used — a teaching otoscope in a vet lab, a rover in the Utah desert, a CO_2 alert in a hibernation cold room. **User testing isn't a gate I pass — it's how I figure out what to build next**; most of what I ship was shaped by something I learned holding the prototype.

DISCIPLINE	LOCATION	CONTACT
Mechanical Engineering B.S. May 2026 University of Wisconsin-Madison	Minnetonka, MN Open to Minneapolis-Saint Paul	diegov.rodriguezeng@gmail.com (952) 334-3763

P · 01	Veterinary Video Otoscope Housing MEDICAL DEVICE · PRODUCT DEVELOPMENT	6 UNITS DELIVERED	02 / 09
P · 02	Wisconsin Robotics Swerve-Drive System ROBOTICS · MECHANICAL DESIGN	12TH OF 34 · URC	04 / 09
P · 03	Two-Stage Gear Reducer + ML Optimization ANALYTICAL DESIGN · MACHINE LEARNING	12% MASS REDUCTION	06 / 09
P · 04	CO_2 Alert System for Hibernation Research RESEARCH TOOL · CLIENT-FACING	50% UNDER BUDGET	08 / 09

01 / 04 UW-MADISON SCHOOL OF VETERINARY MEDICINE

Veterinary Video Otoscope Housing.

Redesigned a flimsy training prototype into a durable, weighted teaching tool for UW–Madison veterinary labs — approaching the handfeel of the Welch Allyn clinical instrument students will use in practice. **Six production units delivered** across three prototype iterations.

TEAM	MY ROLE	OUTPUT	USERS
4 students + faculty	Mechanical designer · machinist · tester	6 units · ~$180 / unit	40-50 students per lab

DEVICE WEIGHT	COST / UNIT
1.1 lb From 0.31 lb → near Welch Allyn target (0.87 lb)	**$180** vs. ≈ $500 commercial otoscope

PROBLEM 01

Plastic prototype felt nothing like a clinical otoscope.

The existing BME prototype could transmit video, but it did not match the weight, balance, grip, or durability of the **Welch Allyn** instrument students use in real exams. It was **too bulky** in head and handle diameter, **too lightweight** at 0.31 lb (vs. 0.87 lb), **too fragile** — plastic split-shell, delicate internals — and ran a complicated beam-splitter + camera optical path.

ACTION 02

Pivoted to a sourced camera and a stock steel handle.

Replaced the team's first attempt (custom ESP32 + beam splitter) with the compact Bebird camera. Used knurled steel dumbbell handles — hollowed on the lathe, threads preserved — as the housing. Contributed to the 3D-printed head and **owned the rotating insert** that lets the charging cable spin freely while the end cap threads on.

SOLUTION 03

A 1.1 lb knurled instrument shipping in lab.

Three handle iterations: **(1)** 3D-printed knurled shell housing the ESP32 + beam splitter — camera overheated and barely fit; **(2)** spare tube stock after finding the Bebird — hand-knurling the stock proved unfeasible; **(3)** the final dumbbell-handle design. Six production units at ≈ $180 each. A supplemental LED was added after cadaver testing revealed a brightness gap CAD couldn't.

FIG 01.1 **BEFORE** Original BME video prototype — 0.31 lb, plastic split shell

FIG 01.2 **AFTER** Final trainer — 1.1 lb knurled steel housing, external charging

FIG 01.3 **EXPLODED VIEW** 3D-printed head, Bebird camera, machined knurled steel handle, rotating insert, threaded end cap

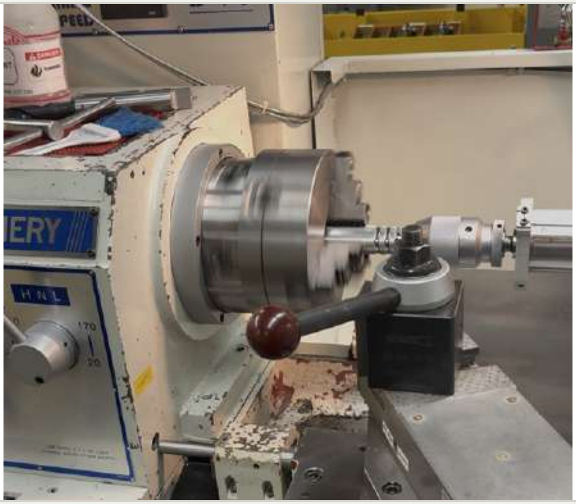

FIG 01.4 **FABRICATION** Lathe used to hollow each handle to 0.8125" ID for the camera + battery

KEY DECISION · № 01

Use a knurled steel dumbbell handle as the housing base.

Stock parts gave us the clinical grip texture, the weight, and reusable threads — without custom machining the most ergonomically important surface.

IMPACT

- **0.31 lb → 1.1 lb** — within 0.23 lb of the 0.87 lb clinical target
- Knurled grip **without custom CNC work**
- Threaded ends reused for end-cap assembly

KEY DECISION · № 02

A rotating 3D-printed insert that decouples cable from cap.

The charging cable had to pass through the end cap. The end cap had to rotate to thread. Those two facts together would have ruined the cable on every reassembly.

IMPACT

- No cable twist during assembly
- **External charging** — no disassembly required
- Improved serviceability and cable durability

IN HINDSIGHT

WHAT I'D CHANGE

I'd test the camera and lighting in realistic canine ear anatomy much earlier. We discovered late that the Bebird light wasn't strong enough, so the supplemental LED and battery had to be packaged near the end. Earlier cadaver testing would have let us design **lighting, battery placement, and wire routing into the housing from the start** — not retrofitted into a head that wasn't designed for them.

I'd also redesign the handle for **tunable weight** — a thinner-wall tube or aluminum handle with internal ballast — so we could dial in the 0.87 lb clinical target without being locked to the geometry of stock dumbbell stock.

02 / 04 WISCONSIN ROBOTICS · UNIVERSITY ROVER CHALLENGE

Wisconsin Robotics Swerve-Drive System.

Designed an integrated tube-and-gearbox connection for the swerve modules on Wisconsin Robotics' ~50 lb URC competition rover. **Eliminated module flex, raised the swerve motor above terrain, and enabled reliable in-place turning.** The rover placed **12th of 34 teams** at the University Rover Challenge.

TEAM	MY CONTRIBUTION	VEHICLE	STATUS
Mechanical · Swerve-Drive	Owned tube-to-gearbox interface	~50 lb URC rover	Installed · Competed

URC FINISH	FINAL SAFETY FACTOR
12 / 34 teams	**2.08**
University Rover Challenge, 2024	6061 Al shaft, vs. 1.80 target

PROBLEM 01

The previous rover needed a whole new turning system.

The prior rover used a six-wheel skid-steer drivetrain — fixed wheel modules, steering by varying left/right wheel speeds. Mechanically simple, but it scrubbed tires through every turn, fed extra load into the wheel mounts, gearbox connections, and frame, and sat the module low enough to take hits on climbs. The rover needed swerve.

ACTION 02

Integrate the shaft receiver into the gearbox housing.

Proposed two tube-to-gearbox concepts for the new swerve module: a clamp-style connector and an integrated receiver. The clamp flexed under load — modules wobbled, alignment stacked up. Selected the integrated version, modified the CAD, ran EES on the connecting shaft across three candidate materials against a SF > 1.8 target, prototyped in 3D print, and presented to the drivetrain team.

SOLUTION 03

6061 aluminum, SF 2.08, 12th at URC.

Installed on every swerve module of the competition rover. In-place turning held. Modules stopped wobbling. The motor stayed clean. The team finished **12th of 34** at the University Rover Challenge.

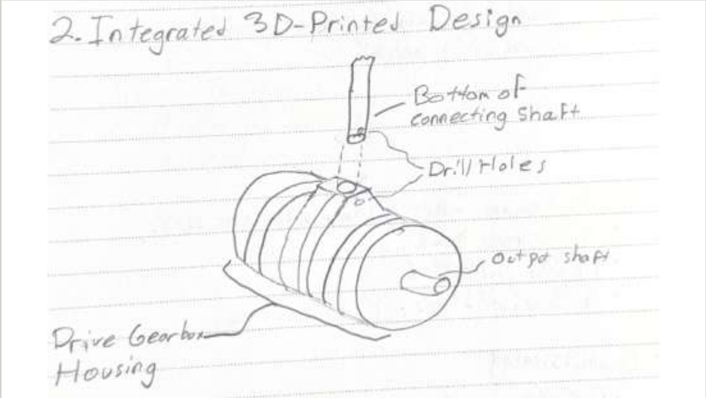

FIG 02.1 **CONCEPT** Integrated shaft receiver — selected over the clamp bracket

FIG 02.2 **FINAL** Rover with integrated swerve modules — 12th at URC

MATERIAL STUDY

EES · VON MISES

Picking the connecting shaft material.

Used EES to evaluate the shaft under combined rover weight, motor torque, impact, and shear loading. Compared three candidate materials against a Von Mises safety-factor criterion (target > 1.8). 6061 aluminum cleared the target with the lowest mass and cost — the right ship.

MATERIAL	SAFETY FACTOR	DECISION
● 6061 Aluminum	2.08	*Selected — met target, lowest weight + cost*
4130 Steel	3.45	*Stronger — but heavier, more expensive*
Titanium	6.59	*Highest strength — cost-prohibitive*

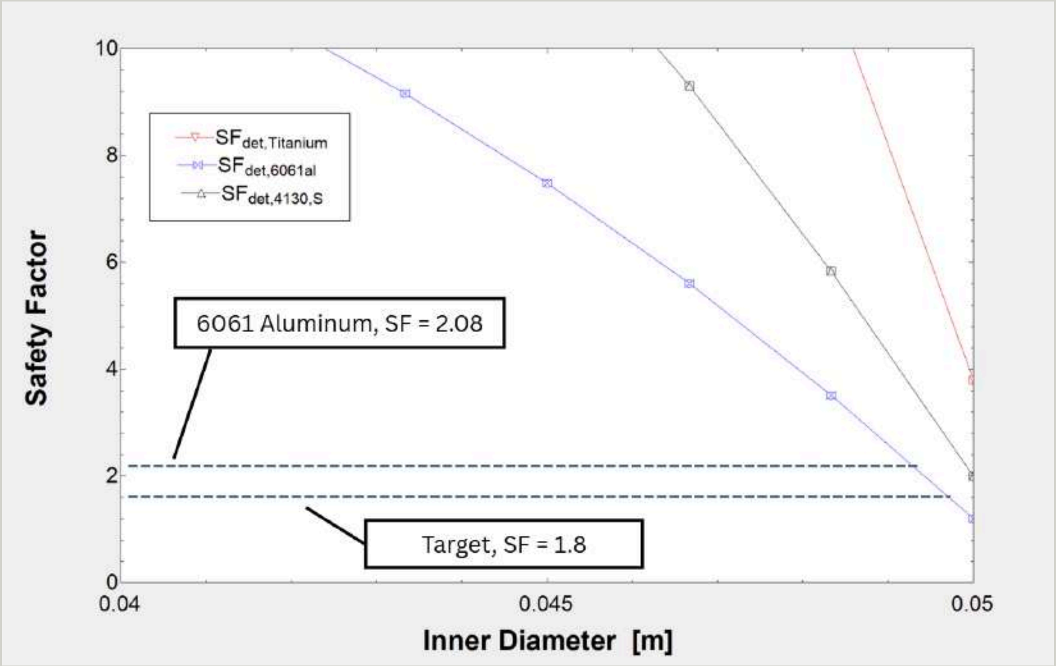

FIG 02.3

EES PLOT Safety factor vs. shaft inner diameter across three materials — 6061 Al selected at SF 2.08, above the 1.80 target

IN HINDSIGHT

WHAT I'D CHANGE

Next time I'd add **more structured validation before competition**. The redesign clearly improved alignment and let the rover turn in place, but I'd quantify that improvement with before/after wobble measurements — something a judge or a teammate could read off a chart, not just feel through the drive sticks.

I'd also run **FEA on the integrated gearbox extension**, especially around the bolt holes and shaft receiver, to confirm that simplifying the interface didn't introduce new stress concentrations the EES shaft study wouldn't have caught.

Two-Stage Gear Reducer + ML Optimization.

A 4.5:1 reverted reducer designed by hand to AGMA / Buckingham — then re-evaluated against 2,500 candidates by a Random Forest surrogate I built. **Found a feasible variant 12% lighter** — at the cost of fatigue-life margin (SF_gear 1.31 → 1.22).

COURSE	TEAM	MY CONTRIBUTION	EXTENSION
ME 342	4 students	Shaft + gear analysis · Python pipeline (solo)	Self-initiated

MASS REDUCTION

12 %

7.19 → 6.34 kg · trade: SF_gear 1.31 → 1.22

DESIGNS EVALUATED

2,500

Random Forest surrogate sweep

PROBLEM 01

Design a 4.5:1 reducer that meets strength, fatigue, and manufacturing.

Team brief: hit the reduction ratio, meet 99.9% reliability, balance reliability against weight and manufacturability, produce ready-to-fab CAD and drawings.

ACTION 02

Hand-analyzed shafts and gears. Then automated the sweep.

Sized shafts (EES) and gears (AGMA + Buckingham). Selected 1045 HR steel shafts and AISI 4140 gears. Then converted the EES design into a Python analytical model and trained a Random Forest surrogate to predict performance across the design space.

SOLUTION 03

Baseline shipped. Then found 12% lighter still feasible.

Baseline met every spec. Surrogate-driven sweep across 2,500 candidates surfaced a variant at **6.34 kg** (vs. 7.19 kg baseline), holding SF_shaft ≥ 1.5 and SF_gear ≥ 1.2.

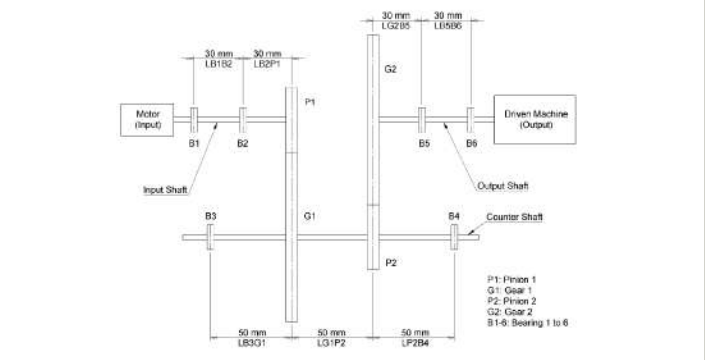

FIG 03.1 **ASSEMBLY** Two-stage reverted reducer — motor, input, counter, output shafts

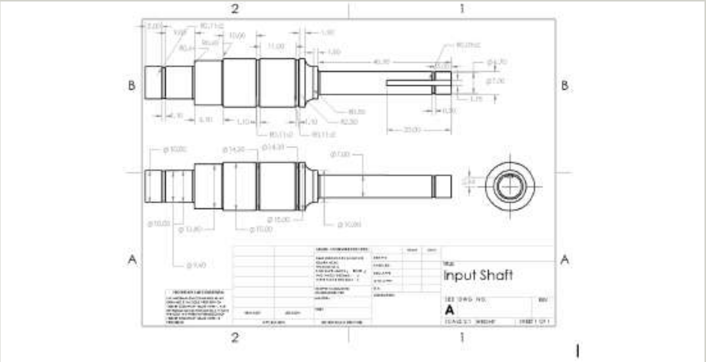

FIG 03.2 **DRAWING** Manufacturing-ready input shaft — dimensions, tolerances, keyway

SELF-INITIATED

The team's design met strength — it wasn't optimized for weight. **I built a Python pipeline to find out how much we left on the table.**

OPTIMIZATION

PYTHON · ML

Converted the EES design into a Python analytical model, performed an automated sweep across 2,500 candidate reducers, and trained a Random Forest surrogate to predict performance from the five geometric inputs (`d_in, d_co, d_out, b_p, b_g`) without re-running the full analysis on every candidate.

The surrogate replaced a per-design analytical evaluation with a near-instant prediction — fast enough to sweep 2,500 candidates in seconds. The trade-off is explicit: gear safety factor dropped from **1.31 → 1.22**, closer to the 1.20 constraint. Both stay inside spec; the baseline is the safer ship.

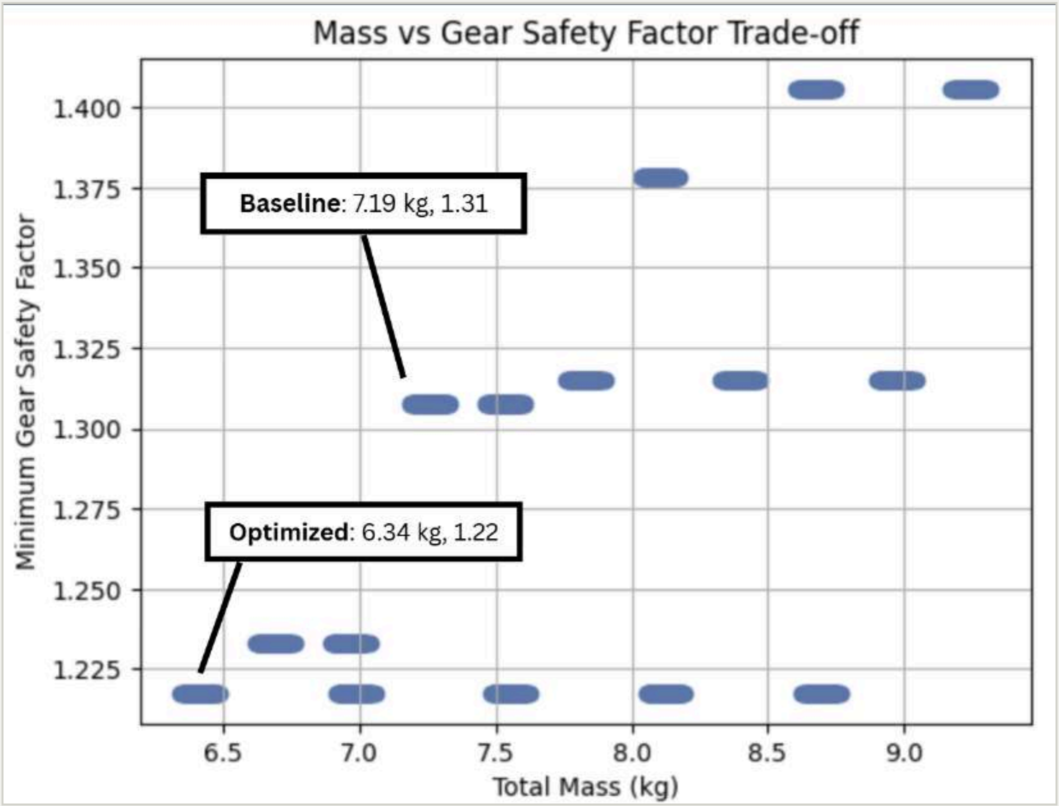

FIG 03.3

OPTIMIZATION PLOT Mass vs. minimum gear safety factor across the feasible sweep — the weight / durability trade-off frontier

IN HINDSIGHT

WHAT I'D CHANGE

I'd make the optimization **more manufacturing-aware**. The model reduced mass while holding safety factors, but next time I'd fold machining tolerances, assembly clearances, and cost into the objective — not just strength and weight. A 12% lighter design that's harder to fabricate isn't actually lighter on the shop floor.

I'd also validate the optimized design with an **independent check** and treat the machine learning model as a fast surrogate, not final proof. The sweep is what surfaced the candidate; a hand recalculation (or FEA) is what should clear it to build.

CO_2 Alert System for Hibernation Research.

A noninvasive CO_2 monitoring enclosure that alerts researchers when hibernating ground squirrels re-enter interbout euthermia. **Delivered under budget, working in cold storage.**

CLIENT	MY ROLE	BUDGET	OPERATING TEMP
Dr. Sprenger · Dr. Sajdak	Mechanical designer · client lead	$300 → came in at $149.20	3–10 °C

ALERT LATENCY	FINAL COST
9.84 s	**$149**.20
Threshold → email · avg., n = 10 trials	vs. $300 budget · 50% under

PROBLEM 01

No way to know when squirrels woke.

Hibernating ground squirrels only enter **interbout euthermia every 6–15 days**, making the wake events hard to catch by hand. Animals were housed remotely in dark, refrigerated habitats; any sensor system had to be noninvasive and survive the cold without disturbing hibernation.

ACTION 02

Acrylic enclosure with a mesh-guarded CO_2 sensor.

Modeled the enclosure, lid, sensor placement, and air pump in SolidWorks. Selected acrylic and the wire mesh guard. Load-tested the mesh under 1 lb (≈ 3× squirrel weight). Wired the Arduino threshold-and-email logic.

SOLUTION 03

9.84 s avg. alert. $149.20 against a $300 budget.

Delivered a working prototype to Dr. Sprenger, Dr. Sajdak, and Fauna Bio. Threshold tested at 5,000 ppm across 10 trials. Mesh held its load with no visible stress.

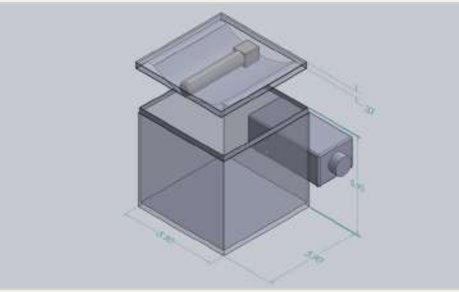

FIG 04.1 **CAD** Removable lid · sensor placement · mesh guard · air-pump connection

FIG 04.2 **PROTOTYPE** Acrylic enclosure · Arduino + WiFi alert · air pump

IN HINDSIGHT

WHAT I'D CHANGE

With the engineering skills I've built since this project, I'd take it further. I'd run **more refrigerator testing** at the actual operating temperature, and **package the electronics in a dedicated 3D-printed external enclosure** instead of leaving them exposed alongside the habitat.

I'd add **protection for the exposed tubing** and re-evaluate the air pump — quieter options, or revisit whether a pump is needed at all if passive diffusion proves sufficient at the alert threshold.

Let's build something.

The thread across these projects is the same: pick the right test, ask the next question, and ship something a real person can use. I'm open to full-time mechanical engineering and product development roles starting summer 2026.

STRENGTH · 01

Bias toward real-world testing.

The cadaver test on the otoscope and the CO_2 latency trials in a cold room both came out of the same instinct: **build it, then break it in the place it'll actually live.**

STRENGTH · 02

Comfortable owning analysis.

EES, AGMA, Buckingham, hand-derived FBDs, and Python sweeps — **I'll show my work** and defend the assumptions in review.

STRENGTH · 03

Translate between disciplines.

I've worked with veterinarians, animal physiologists, software-leaning roboticists, and ME peers; I try to learn the other side's vocabulary **before pushing a design.**

STRENGTH · 04

Take initiative beyond the brief.

The Python optimization on the gear reducer wasn't in the assignment. **It was the right next question to ask.**

EMAIL
diegov.rodriguezeng@gmail.com

PHONE
(952) 334-3763

LINKEDIN
linkedin.com/in/diego-rodriguez