

CEN 5016: Software Engineering

Fall 2026



University of
Central Florida

Dr. Kevin Moran

Week 2 - Class 1: Measurement & Metrics





- Let me know if you are not on Ed Discussions
- Assignment 1, Getting started with Git, GitHub, and Typescript is due *Friday at 11:59 pm*
- Use Megathread on Ed Discussions to ask questions
- Assignment 2 out Thursday
- *No class on Thursday!! - UCF Football Game* 

Software Archeology & Anthropology




(Recap)

Today: How to tackle Codebases



- Goal: Develop and test a working model or set of working hypotheses about how (some part of) a system works
- Working model: an understanding of the pieces of the system (components), and the way they interact (connections)
- Focus: Observation, probes, and hypothesis testing
 - Helpful tools and techniques!



Where do Coding Agents Fit in?



Same three moves as before — the agent changes how fast you make them.

1. Observation

Reads far more of the repo in a minute than you can in an hour: file structure, naming conventions, candidate entry points.

But: it reports what the code says, not what it does.

2. Probes

Writes the log line, the throwaway test, the debugger config — the tedious parts you skip when you are in a hurry.

But: someone still has to run it.

3. Hypothesis Testing

Turns a vague question ("how does login work?") into a checkable claim with file:line citations.

But: it answers confidently whether or not it checked.

The agent accelerates the loop. It does not build the working model for you — that still has to end up in your head.

Agent Strategies for Code Comprehension



1 Ask for a map, not a summary

"List every entry point in this repo. For each, cite file:line and what triggers it."

A summary is unfalsifiable. A map is something you can go check.

2 Trace one path, end to end

"Trace what happens from <user action> to <observable effect>. Every function, in call order, with file:line."

Components are easy. Connections are the hard part — ask for those.

3 Make it write a probe, then run it

"Add logging that would prove or disprove that X is called on Y. Don't explain — write the diff."

Converts the agent's assertion into your observation.

4 Bound the question to a diff

"If I wanted to change <behavior>, which files would I touch, and why?"

Right-sized, verifiable, and it's the question you actually have.

Example: Hypothesis Testing Loop



Setup: unfamiliar repo, 200k lines. Ticket says “uploads over 10 MB fail silently.” You have never opened this code.

LOCATE

"Where is an upload size limit enforced? Cite file:line for every place a size is compared against a constant."

Agent returns 3 candidate sites in ~20 seconds.

PREDICT

You write the hypothesis down: "If site #2 is the live one, a 12 MB upload logs a rejection there."

This is yours, not the agent's..

PROBE

"Add a distinct log line at each of the three sites. Do not change behavior. Show me the diff."

Thirty seconds of typing you didn't have to do.

VERIFY

Run it. Upload a 12 MB file. Site #2 fires; sites #1 and #3 never do.

Two of the three were dead config and only running the code can tell you this.

The agent gave you candidates. Running the probe turned them into knowledge.

Trust, but Verify: Where Agents Fail



Fluency is not evidence.

It will summarize code it never ran, in the same confident register it uses for things it checked.



Dead code reads exactly like live code.

Nothing in the source text says "this branch has not executed since 2019." Only running it tells you.



It knows other people's repos.

Strong priors from a million codebases become confident wrong answers about the ways yours is unusual.



Tacit knowledge is still tacit.

Why the code is this way lives with people, in issue threads and hallway conversations — not in the files.

The rule: every agent claim should reduce to a file:line, a test, or a log line.
If it can't be checked, it isn't a finding — it's a hypothesis.

Metrics & Measurement



Goals for Today



- Use measurements as a decision tool to reduce uncertainty
- Understand difficulty of measurement; discuss validity of measurements
- Provide examples of metrics for software qualities and process
- Understand limitations and dangers of decisions and incentives based on measurements



- Software Engineering: Principles, practices (technical and non-technical) for confidently building high-quality software.

What does this mean?
How do we know?
-> Measurement & Metrics
are key concerns

Case Study: Autonomous Vehicles



Case Study: Autonomous Vehicles



- By what methods can we judge AV software quality (e.g., safety)?



Test Coverage




- Amount of code executed during testing.
- Statement coverage, line coverage, branch coverage, etc.
- E.g., 75% branch coverage -> 3/4 if-else outcomes have been executed

```
:  
: 1698 : const TrajectoryPoint& StGraphData::init_point() const { return init_point; }  
:  
: 2264 : const SpeedLimit& StGraphData::speed_limit() const { return speed_limit; }  
:  
: 212736 : double StGraphData::cruise_speed() const {  
[ - + ]: 212736 :     return cruise_speed_ > 0.0 ? cruise_speed_ : FLAGS_default_cruise_speed;  
:  
:     }  
:  
: 1698 : double StGraphData::path_length() const { return path_data_length; }  
:  
: 1698 : double StGraphData::total_time_by_conf() const { return total_time_by_conf; }  
:  
: 1698 : planning_internal::STGraphDebug* StGraphData::mutable_st_graph_debug() {  
: 1698 :     return st_graph_debug;  
:  
:     }  
:  
: 566 : bool StGraphData::SetSTDrivableBoundary(  
:     const std::vector<std::tuple<double, double, double>>& s_boundary,  
:     const std::vector<std::tuple<double, double, double>>& v_obs_info) {  
[ + - ]: 566 :     if (s_boundary.size() != v_obs_info.size()) {  
:         return false;  
:     }  
:  
[ + + ]: 40752 :     for (size_t i = 0; i < s_boundary.size(); ++i) {  
:         80372 :         auto st_bound_instance = st_drivable_boundary_.add_st_boundary();  
:         160744 :         st_bound_instance->set_t(std::get<0>(s_boundary[i]));  
:         120558 :         st_bound_instance->set_s_lower(std::get<1>(s_boundary[i]));  
:         120558 :         st_bound_instance->set_s_upper(std::get<2>(s_boundary[i]));  
[ - + ]: 40186 :         if (std::get<1>(v_obs_info[i]) > -kObsSpeedIgnoreThreshold) {  
:             0 :             st_bound_instance->set_v_obs_lower(std::get<1>(v_obs_info[i]));  
:  
:             }  
[ + + ]: 40186 :         if (std::get<2>(v_obs_info[i]) < kObsSpeedIgnoreThreshold) {  
:         50254 :             st_bound_instance->set_v_obs_upper(std::get<2>(v_obs_info[i]));  
:  
:         }  
:  
:     }  
:  
: }
```

Model Accuracy



- Train machine-learning models on labelled data (sensor data + ground truth).
- Compute accuracy on a separate labelled test set.
- E.g., 90% accuracy implies that object recognition is right for 90% of the test inputs.





- Frequency of crashes / fatalities
- Per 1,000 rides, per million miles, per month (in the news)






Driving to Safety

How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?


Nidhi Kalra, Susan M. Paddock

Figure 3. Miles Needed to Demonstrate with 95% Confidence that the Autonomous Vehicle Failure Rate Is Lower than the Human Driver Failure Rate



Building the World's Most Experienced Driver™

The Waymo Driver gains experience with every mile, in each car. ↗



10+

More than a Decade of Autonomous Driving in More than 10 States

5

Generations of Autonomously Driven Vehicles

15+

Billion Autonomously Driven Miles in Simulation

20+

Million Real-World Miles on Public Roads

Source: waymo.com/safety (September 2021)

What is Measurement?



- Measurement is the empirical, objective assignment of numbers, according to a rule derived from a model or theory, to attributes of objects or events with the intent of describing them. – Craner, Bond, “Software Engineering Metrics: What Do They Measure and How Do We Know?”
- A quantitatively expressed reduction of uncertainty based on one or more observations. – Hubbard, “How to Measure Anything ...”



- IEEE 1061 definition: “A software quality metric is a function whose inputs are software data and whose output is a single numerical value that can be interpreted as the degree to which the software possesses a given attribute that affects its quality.”
- Metrics have been proposed for many quality attributes; may define own metrics

What Software Qualities Do We Care About?



- Functionality (e.g., data integrity)
- Scalability
- Security
- Extensibility
- Bugginess
- Documentation
- Performance
- Installability
- Availability
- Consistency
- Portability
- Regulatory compliance

What Process Qualities Do We Care About?



- On-time release
- Development speed
- Meeting efficiency
- Conformance to processes
- Time spent on rework
- Reliability of predictions
- Fairness in decision making
- Number of builds
- Code review acceptance rate
- Regulatory compliance
- Measure time, costs, actions, resources, and quality of work packages; compare with predictions
- Use information from issue trackers, communication networks, team structures, etc...

What People Qualities Do We Care About?



- **Developers**
 - Maintainability
 - Performance
 - Employee satisfaction and well-being
 - Communication and collaboration
 - Efficiency and flow
 - Satisfaction with engineering system
 - Regulatory compliance
- **Customers**
 - Satisfaction
 - Ease of use
 - Feature usage
 - Regulatory compliance

Everything is Measurable




- If X is something we care about, then X, by definition, must be detectable.
 - How could we care about things like “quality,” “risk,” “security,” or “public image” if these things were totally undetectable, directly or indirectly?
 - If we have reason to care about some unknown quantity, it is because we think it corresponds to desirable or undesirable results in some way.
- If X is detectable, then it must be detectable in some amount.
 - If you can observe a thing at all, you can observe more of it or less of it 21
- If we can observe it in some amount, then it must be measurable.

Why Measure?






- Fund project?
- More testing?
- Fast enough? Secure enough?
- Code quality sufficient?
- Which feature to focus on?
- Developer bonus?
- Time and cost estimation? Predictions reliable?



Benchmarking Against Standards




- Monitor many projects or many modules, get typical values for metrics
- Report deviations



Antipatterns in Effort Estimation




- IBM in the 60s: Would account in “person-months”
e.g. Team of 2 working 3 months = 6 person-months
- LoC ~ Person-months ~ \$\$
\$
- Brooks: “Adding manpower to a late software project [just] makes it later.”



Measurement is Difficult



The Streetlight Effect



The Streetlight Effect





- A known observational bias.
- People tend to look for something only where it's easiest to do so.
- If you drop your keys at night, you'll tend to look for it under streetlights.

What could Possibly go Wrong?




- Bad statistics: A basic misunderstanding of measurement theory and what is being measured.
- Bad decisions: The incorrect use of measurement data, leading to unintended side effects.
- Bad incentives: Disregard for the human factors, or how the cultural change of taking measurements will affect people.






- To infer causation:
 - Provide a theory (from domain knowledge, independent of data)
 - Show correlation
 - Demonstrate ability to predict new cases (replicate/validate)

Spurious Correlations



Confounding Variables



- If you look only at the coffee consumption → cancer relationship, you can get very misleading results
- Smoking is a confounder



RESEARCH-ARTICLE

Coverage is not strongly correlated with test suite effectiveness

Authors:  [Laura Inozemtseva](#),  [Reid Holmes](#) [Authors Info & Affiliations](#)

ICSE 2014: Proceedings of the 36th International Conference on Software Engineering • May 2014 • Pages 435–445 • <https://doi.org/10.1145/2568225.2568271>

“We found that there is a low to moderate correlation between coverage and effectiveness when the number of test cases in the suite is controlled for.”


5



- **Construct validity** – Are we measuring what we intended to measure?
- **Internal validity** – The extent to which the measurement can be used to explain some other characteristic of the entity being measured
- **External validity** – Concerns the generalization of the findings to contexts and environments, other than the one studied



- Extent to which a measurement yields similar results when applied multiple times
- Goal is to reduce uncertainty, increase consistency
- Example: Performance
 - Time, memory usage
 - Cache misses, I/O operations, instruction execution count, etc.
- Law of large numbers
 - Taking multiple measurements to reduce error
- Trade-off with cost





- Measure whatever can be easily measured.
- Disregard that which cannot be measured easily.
- Presume that which cannot be measured easily is not important.
- Presume that which cannot be measured easily does not exist.



- There seems to be a general misunderstanding to the effect that a mathematical model cannot be undertaken until every constant and functional relationship is known to high accuracy. This often leads to the omission of admittedly highly significant factors (most of the “intangibles” influences on decisions) because these are unmeasured or unmeasurable. To omit such variables is equivalent to saying that they have zero effect... Probably the only value known to be wrong...
- J. W. Forrester, Industrial Dynamics, The MIT Press, 1961

- Goodhart's law: "When a measure becomes a target, it ceases to be a good measure."



Simplistic Productivity Measures



- Lines of code per day?
 - Industry average 10-50 lines/day
 - Debugging + rework ca. 50% of time
- ● Function/object/application points per month ● Bugs fixed?
 - Milestones reached?



- What happens when developer bonuses are based on
 - Lines of code per day?
 - Amount of documentation written?
 - Low number of reported bugs in their code?
 - Low number of open bugs in their code?
 - High number of fixed bugs?
 - Accuracy of time estimates?

Developer Productivity Myths



- Productivity is all about developer activity
- Productivity is only about individual performance
- One productivity metric can tell us everything
- Productivity measures are useful only for managers
- Productivity is only about engineering systems and developer tools

WARNING!!



- Most software metrics are controversial
 - Usually only plausibility arguments, rarely rigorously validated
 - Cyclomatic complexity was repeatedly refuted, yet is still used
 - “Similar to the attempt of measuring the intelligence of a person in terms of the weight or circumference of the brain”
- Use carefully!
- Code size dominates many metrics
- Avoid claims about human factors (e.g., readability) and quality, unless validated
- Calibrate metrics in project history and other projects
- Metrics can be gamed; you get what you measure



- Measurement is difficult but important for decision making
- Software metrics are easy to measure but hard to interpret,
validity often not established
- Many metrics exist, often composed; pick or design suitable metrics if needed
- Careful in use: monitoring vs incentives
- Strategies beyond metrics

Questions to Consider for Your Projects



- What properties do we care about and how do we measure them?
- What is being measured? Does it (to what degree) capture the thing you care about? What are its limitations?
- How should it be incorporated into process?
- What are potentially negative side effects or incentives?