

# Showing, Not Telling

Drawing on the learner's screen as tutoring feedback: a design note

The DuckyHelper team, MingLLM, Inc. October 2026 • Version 1.0 • [hello@duckyhelper.com](mailto:hello@duckyhelper.com)

## Abstract

Human tutors point. They circle the wrong step, underline the word that matters and sketch the idea in the margin. Most AI tutors cannot: they answer in a chat window, away from the learner's work. DuckyHelper's Mac app gives its voice tutor a pointer and a pen that work on the learner's real screen, over any app. This design note describes what the screen pen does, how it decides where a mark goes, and how each design choice follows research on multimedia learning, gesture and feedback. We also describe the web app, which cannot see the learner's screen and draws on its own blackboard instead, and the privacy limits on both. DuckyHelper has not been evaluated in a controlled study. This paper reports no learning outcomes, no usage numbers and no benchmark scores of our own. Every number in it belongs to the study that reported it.

Keywords: multimedia learning, gesture, signaling, feedback, screen annotation, voice tutor

## 1 The problem with telling

A learner stuck on a geometry worksheet asks for help. A chat assistant replies in a separate window: "In line 3, you squared 3 and wrote 6." The learner now has to find line 3, find the 3, and connect the sentence to the place. A tutor sitting next to them would just put a finger on it.

That small act is backed by research. Learning improves when words are tied to what they describe (Mayer & Moreno, 2003), when the learner can see an explanation being drawn rather than finished (Fiorella & Mayer, 2016), and when gesture carries part of the message (Goldin-Meadow & Alibali, 2013). Feedback works best when it is specific and about the task (Shute, 2008). A mark on the learner's own page does all four.

## 2 What the screen pen does

In DuckyHelper's Mac app, the learner presses fn and talks. While the pill is open the tutor keeps a small current view of the screen. It can then act on that screen in five ways, all drawn in a layer above the learner's apps:

- **Point.** A round blue pointer flies to a spot and shows a 2 to 5 word label.
- **Draw.** Circles, arrows, boxes, highlights, underlines, labels and freehand paths, drawn one mark at a time like a pen, and numbered for walkthroughs.
- **Write.** Handwriting in the margin beside the learner's line: what comes next in their working, or a correction beside a line that went wrong.

- **Spotlight.** On a busy screen, everything but one region dims, with a short callout. The learner can still click inside the lit region, and clicking anywhere brings the screen back.
- **Walk through.** Step-by-step tutorials in the learner's own apps, where the pointer shows where to click and the step completes when the learner really does it.




Figure 1. The pen circles a slip on the learner's own page ("3<sup>2</sup> is 9, not 6"). Frame from our tutorial film, real app interface.

The learner can point too. Holding fn and cmd while moving the mouse leaves a glowing blue trail and shows their words in a capsule; on release the tutor receives a picture with the path drawn on it, so "what is this?" and "merge those" refer to exactly what was traced.




Figure 2. The learner holds fn + cmd and traces the square root they are asking about. Frame from our tutorial film, real app interface.

### 3 Where a mark goes

A mark in the wrong place is worse than no mark. Three rules keep marks where they belong. First, the tutor may not describe what it has not seen:

#### IN DUCKYHELPER

The tutor does not describe what is on the learner's screen until it has actually looked at it.

Second, when the target has words on it, the mark is anchored to those words rather than to coordinates. The Mac reads the text on the screen, and the tutor names the words:

#### IN DUCKYHELPER

When the target has words on it, the mark goes on those words as the Mac reads them from the screen, not on an estimated position.

Third, marks stay attached to what they were drawn on when the learner scrolls, and clear when the lesson moves on. We have not measured how often a mark lands on the intended spot; like a human tutor's finger, it is only as good as the tutor's reading of the page.

### 4 Design choices and the research behind them

Table 1. Each design choice and the finding it follows.

| Design choice                                       | Finding                                                            | Source                                              |
|-----------------------------------------------------|--------------------------------------------------------------------|-----------------------------------------------------|
| One mark at a time, in the order of the explanation | Watching a drawing appear beat a finished diagram for beginners    | (Fiorella & Mayer, 2016)                            |
| Point whenever the tutor says "this"                | Gesture is part of how speakers communicate and how learners learn | (Goldin-Meadow & Alibali, 2013; Cook et al., 2008)  |
| Labels beside the part, joined by a line            | Put words next to the picture they describe                        | (Mayer & Moreno, 2003)                              |
| The spoken sentence starts with its mark            | Connect pictures and words in time                                 | (Mayer & Moreno, 2003; Clark & Paivio, 1991)        |
| Spotlight on busy screens                           | Reduce what the learner has to search through                      | (Mayer & Moreno, 2003; Hegarty, 2011)               |
| Circle the wrong step, not the whole page           | Feedback about the task, specific and timely                       | (Shute, 2008; Kluger & DeNisi, 1996)                |
| The learner can point back                          | Learners' own gestures and drawings help them learn                | (Goldin-Meadow et al., 2009; Bobek & Tversky, 2016) |

Two choices deserve a note. **Order.** The pen draws its marks in the order they are explained, because the rules tell the tutor to order them that way:

#### IN DUCKYHELPER

Marks appear one at a time, in the order the explanation needs them.

This follows Fiorella and Mayer (2016), where learners with low prior knowledge transferred more after watching the instructor draw:

*“resulted in significantly better transfer test performance than viewing already-drawn diagrams for learners with low prior knowledge”*

(Fiorella & Mayer, 2016)

**Pointing.** The rules ask for a pointer whenever the tutor refers to something on screen, the way a speaker gestures. Goldin-Meadow and Alibali (2013) review evidence that

*“the gestures speakers produce when they talk are integral to communication”*

(Goldin-Meadow & Alibali, 2013)

## 5 Feedback as a mark

The most common use of the pen in tutoring is feedback. A circle around the exact slip, with a few words, answers Hattie and Timperley’s question of where the learner is (Hattie & Timperley, 2007) without turning attention to the learner as a person, which Kluger and DeNisi (1996) found can make feedback backfire. On the shared whiteboard, the rules make the same move with the learner’s own lines:

### IN DUCKYHELPER

On the whiteboard, the tutor marks the learner’s first wrong step with a short note and leaves the fix to the learner.

The mark says *where*; the learner still does the fixing. This is the same line the tutor holds everywhere: hints before answers (Bastani et al., 2025).

## 6 The web app: drawing without seeing

On Chromebooks and other computers, DuckyHelper runs as a web app. Eddy, the tutor there, never sees the learner’s screen. It teaches on its own desk instead: on a blackboard it draws in time with its voice, building one picture piece by piece, and on pictures it labels one part at a time as it names them. The learner can click a part of a picture to ask about it. The principles are the same; the canvas is the tutor’s own.

## 7 Privacy and control

- The Mac app looks at the screen only while the pill is open, and watches continuously only when the learner asks it to.
- Nothing is captured while a password field has focus, and DuckyHelper’s own windows are never in the picture.
- The pen only draws. Real clicks and typing happen only when the learner clearly asks, or as a demonstration. Never in a password field, and never to complete or turn in the learner’s work, pay for something or delete anything.
- The web app has no screen capture and no camera.

## 8 Limits

- DuckyHelper has not been evaluated in a controlled study. This paper reports no learning outcomes, no usage numbers and no benchmark scores of our own. Every number in it belongs to the study that reported it.
- Drawing on the real screen is in the Mac app only, on the display with the mouse.
- The research cited here studied human instructors, recorded lessons and classroom gestures. Applying it to an AI's pointer and pen is a design bet that a study would need to test.

## 9 Conclusion

Telling is cheap for an AI. Showing is harder, and closer to what good tutors do. DuckyHelper's screen pen puts the tutor's explanation on the learner's own work, one mark at a time, in time with its voice, and its rules keep each mark about the task. Whether this helps learners more than telling them is an open question that we would like to test.

## References

- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- Bobek, E., & Tversky, B. (2016). Creating visual explanations improves learning. *Cognitive Research: Principles and Implications*, 1(1), 27. <https://doi.org/10.1186/s41235-016-0031-6>
- Clark, J. M., & Paivio, A. (1991). Dual coding theory and education. *Educational Psychology Review*, 3(3), 149-210. <https://doi.org/10.1007/BF01320076>
- Cook, S. W., Mitchell, Z., & Goldin-Meadow, S. (2008). Gesturing makes learning last. *Cognition*, 106(2), 1047-1058. <https://doi.org/10.1016/j.cognition.2007.04.010>
- Fiorella, L., & Mayer, R. E. (2016). Effects of observing the instructor draw diagrams on learning from multimedia messages. *Journal of Educational Psychology*, 108(4), 528-546. <https://doi.org/10.1037/edu0000065>
- Goldin-Meadow, S., & Alibali, M. W. (2013). Gesture's Role in Speaking, Learning, and Creating Language. *Annual Review of Psychology*, 64(1), 257-283. <https://doi.org/10.1146/annurev-psych-113011-143802>
- Goldin-Meadow, S., Cook, S. W., & Mitchell, Z. A. (2009). Gesturing Gives Children New Ideas About Math. *Psychological Science*, 20(3), 267-272. <https://doi.org/10.1111/j.1467-9280.2009.02297.x>
- Hattie, J., & Timperley, H. (2007). The Power of Feedback. *Review of Educational Research*, 77(1), 81-112. <https://doi.org/10.3102/003465430298487>
- Hegarty, M. (2011). The Cognitive Science of Visual-Spatial Displays: Implications for Design. *Topics in Cognitive Science*, 3(3), 446-474. <https://doi.org/10.1111/j.1756-8765.2011.01150.x>
- Kluger, A. N., & DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, 119(2), 254-284. <https://doi.org/10.1037/0033-2909.119.2.254>
- Mayer, R. E., & Moreno, R. (2003). Nine Ways to Reduce Cognitive Load in Multimedia Learning. *Educational Psychologist*, 38(1), 43-52. [https://doi.org/10.1207/S15326985EP3801\\_6](https://doi.org/10.1207/S15326985EP3801_6)

Shute, V. J. (2008). Focus on Formative Feedback. *Review of Educational Research*, 78(1), 153-189.  
<https://doi.org/10.3102/0034654307313795>

About this paper: written by the DuckyHelper team at MingLLM, Inc. It describes the design of DuckyHelper and cites published research. It has not been peer reviewed and reports no study of DuckyHelper. HTML version: <https://dukyhelper.com/research/papers/showing-not-telling/>