Seven Minutes by Design

Seven minutes is not a number. It is a design discipline.

Why focused learning works best when the unit is small enough to finish, trace, and maintain

A Seven Minute Modules / Group Moovs methodology paper · Version 1.1

For L&D leaders, instructional designers, learning consultants, procurement reviewers, compliance training owners, and executives who want the evidence without an academic literature review.

Executive summary

Much workplace learning does not fail because it is too short. It fails because it is too much. A single course often tries to cover too much, so too little reaches the point of use. Objectives compete. Content piles up. Completion drops, and what little is finished is mostly forgotten within days.

Seven Minute Modules answers this with a constraint: one objective, seven minutes. The number is not magic, and it is not a claim about human attention spans. It is a design envelope. The discipline is in the limit.

The evidence does not crown seven minutes. It supports something more useful: short, single-objective, well-structured, spaced learning. The clearest single data point is that engagement with educational video drops sharply after roughly six minutes.[1] Working memory is small.[2] Retrieval and spacing, not exposure, are what make learning last.[3][4]

There is no fixed attention span. The honest reading of the research is that attention is governed by design and active engagement, not by a biological timer.[5] So we do not sell the "attention span" story. We design for the conditions that actually hold attention.

Seven minutes sits at the upper edge of the effective envelope. It is long enough to teach one thing properly, with context, an example, practice, and a check. It is short enough to stay inside the

window where people engage and finish. And it is small enough to keep current, to source, and to regenerate.

Short is not automatically good. A weak seven-minute module fails like any other. Seven minutes works only when the module is designed well: one objective, concise context, a clear explanation, an example, practice, a retrieval check, and a next action.

Seven minutes is also a unit within a larger system, not a cure for everything. Longer formats such as workshops, coaching, and group practice remain essential. Some outcomes need a pathway of many units. We say so plainly, because overclaiming would betray the trust the whole framework is built on. Seven minutes is not a number. It is a design discipline.

The problem with bloated learning

The default failure mode in workplace learning is the opposite of what people assume. It is not brevity. It is bulk.

Many workplace courses are built to be comprehensive. A course often carries many objectives at once, on the theory that more coverage means more value. The learner cannot hold all of it, so most of it does not land. This is cognitive overload by design, and adding content makes it worse, not better.

Length then collides with reality. Most people have very little time for formal learning, and industry benchmarks consistently put the formal-learning budget at a small fraction of the working week.[6] A ninety-minute course rarely fits that week without being scheduled, deferred, or interrupted. Completion figures for long-form content circulate widely, but most are poorly sourced, so we do not lean on a specific number. The direction is not seriously disputed: the longer and less focused the content, the less of it gets finished.

Even when a long course is completed, it tends not to stick. Without focused objectives and reinforcement, people forget most of what they meet within days.[7] And without room to practise and apply, little of it transfers to the job. The result is the familiar complaint that training does not change behaviour. The cause is usually that the design never gave behaviour a chance.

What the evidence actually says

The research here does not name a perfect duration. Read together, it points consistently toward short, focused, active, and spaced learning. We use careful language on purpose, and we flag industry evidence as industry evidence.

Video and engagement. A large-scale study of educational-video engagement, based on 6.9 million viewing sessions, found that engagement drops sharply after roughly six minutes, with length the single strongest predictor of how much of a video people watch.[1:1] The same study found that personal, direct delivery held attention better than polished but impersonal production. This is a finding about engagement and design, not a measurement of a fixed attention limit.

Cognitive load and working memory. Cognitive Load Theory starts from a simple fact: working memory is limited, and learning fails when demand exceeds it.[2:1] Later work narrows the limit further, finding that only a few elements can be actively processed at once.[8] The design consequence is direct. One objective, and a small number of concepts, is not a stylistic preference. It reflects how limited working memory actually is. (The familiar idea that working memory holds about seven items concerns chunks of information, not minutes of learning, and is not the basis for the seven-minute unit.)

Segmenting. Segmenting complex material into smaller units supports learning, especially when each segment has a clear purpose.[9] Short, single-objective units are segmentation applied at the level of the whole learning experience.

Retrieval practice and the testing effect. Retrieving information beats rereading it for durable memory, even though rereading feels more productive at the time.[3:1] Practice and a check are not decoration added at the end of a module. They are where the learning is made.

Spacing. Learning spread across separate sessions is retained far better than the same learning crammed into one sitting, with sessions ideally at least a day apart.[4:1] This is the case for units that can be revisited and reinforced, rather than a single long event.

Forgetting. The forgetting curve is real and has been replicated under modern conditions: retention drops steeply and early when nothing reinforces it.[7:1] Short units that can be spaced and revisited are a direct response to it.

Transfer of training. What reaches the job is shaped by more than the content. Across studies, a supportive work environment and manager support predict transfer alongside the learner's own

motivation, and weigh at least as heavily as the quality or format of the material.[10] This is why we treat a module as one part of a larger system, not as the whole system.

Motivation and workplace context. Learning has to fit the day it lives in. The time budget is real, and attention in a workplace is fragile and easily interrupted. Designing as if people have long, quiet stretches is designing for conditions that do not exist.

The myth we refuse to sell. You will often hear that people have an attention span of a few minutes, or that it has collapsed to seconds. As stated, this is not supported. Reviews of the primary evidence find no fixed attention limit, and conclude that the most consistent driver of attention is teaching and design, not a biological clock.[5:1] The most detailed study found that attention lapses begin early and recur, but drop sharply whenever the session turns active, and that the benefit carries forward.[11] The lesson is not that attention dies at a set minute. It is that attention is sustained by short, segmented, active design. That is what we build for, and it is why we do not rest the case on "attention span."

It is true that short-form media has changed what many people expect from digital content. People are used to faster pacing, clearer hooks, and less tolerance for irrelevant material. Some recent studies also suggest that rapid switching between short videos can increase boredom rather than relieve it,[12] and that interruptions from short-form video feeds can impair the memory needed to resume a task.[13] But that is not the same as proving a fixed or collapsing attention span. Entertainment feeds and learning experiences ask different things of the brain. The lesson for Seven Minute Modules is not to imitate the feed. It is to respect attention, remove waste, create active engagement, and make every minute earn its place.

A note on microlearning research. The microlearning literature often studies interventions between one and fifteen minutes. Seven minutes sits inside that range, but the number is not derived by averaging those studies. It is a practical design choice: long enough for a complete learning arc, short enough to stay focused, and small enough to govern and maintain. When a topic needs fifteen minutes, our answer is usually not one longer module. It is two focused modules.

Why seven minutes

Given that the evidence supports a short, focused, active, spaced envelope rather than a single number, why commit to seven minutes at all?

Because a stated limit is the most powerful tool there is for forcing quality. Left open, content expands to fill the time available, and the objectives multiply. A fixed envelope forces the designer to

cut everything that does not serve one objective. The constraint is the discipline.

Seven minutes is chosen for what it makes possible.

It is completion-friendly. It fits a coffee break, a gap between meetings, or a short window in a crowded week, without needing to be scheduled.

It is long enough for one meaningful objective. Seven minutes leaves room to set context, explain the core idea, show an example, let the learner practise, and check understanding. That full arc is what turns a viewing into a capability.

It is short enough for focused attention. Not because attention expires at seven minutes, but because a short, single-objective unit stays inside the window where engagement holds, and avoids the drift that sets in during long, passive stretches.

It is small enough to keep current. A short, single-objective unit is far easier to source, review, update, and regenerate than a long, mixed one. This is the link to assurance, and we return to it shortly.

And it is practical for real work. It runs on a desktop or a phone, in a quiet office or on a noisy floor, for a fluent reader or a second-language one. The method stays constant. The execution adapts.

We design the unit for a seven-minute envelope, and we redesign rather than stretch when a topic will not fit. Seven is the upper edge of the envelope, chosen on purpose, not a law of nature.

Why not shorter

If shorter content is more engaging, why not push to sixty seconds, two minutes, or three?

Because very short units are too compressed for a complete learning arc. You can state a fact in sixty seconds. You cannot, in that time, set context, show a worked example, let the learner practise, give feedback, and prompt a next action. Those are the steps that produce retention and transfer, and the evidence on retrieval and application is clear that they need room to happen.[3:2][10:1]

So sub-three-minute units have a real and valuable role, as reminders, nudges, reference cards, and quick refreshers. They are excellent at refreshing something already learned. They are not enough to teach one thing to the point where a person can use it. Compression past a certain point does not trim fat. It removes the parts that make learning stick.

This is the difference between a flashcard and a lesson. Both are useful. They are not the same unit, and the seven-minute module is a lesson.

Why not longer

In the other direction, why not fifteen, twenty, or thirty minutes?

Because length quietly reintroduces every problem the constraint was meant to solve.

It invites content stuffing. Once there is more time, a second objective slips in, then a third, and the concept count climbs past what working memory can hold. Overload returns.[2:2][8:1]

It weakens completion. Longer content crosses the engagement window and starts to need scheduling, which the realistic time budget does not allow.[1:2][6:1] The unit that does not fit the day is the unit that does not get finished.

It lowers repeatability. A thirty-minute course is hard to revisit and hard to space, so it works against the reinforcement that retention depends on.[4:2]

And it is harder to maintain. A longer unit mixes more sources and more claims, which makes it harder to keep current, harder to trace, and harder to map to a specific policy. When the world changes, a large mixed course is a slow, uncertain thing to put right.

The discipline dissolves as the minutes grow. The envelope holds it in place.

For reference, the envelope looks like this.

Duration	Best used for	Why it is not the 7MM unit
1 to 3 minutes	Reminders, nudges, reference, quick refreshers	Too compressed for context, practice, feedback, and application
About 5 minutes	A single simple point	Often too tight for a full arc with practice and a real check
About 7 minutes	**One meaningful objective, taught to the point of use**	**This is the unit: one objective, an example, practice, a check, a next action**
10 to 15 minutes	A small cluster of related points	Crosses the engagement window; invites a second objective; harder to space and maintain
20 to 30 minutes	A full topic, a workshop segment	Content stuffing, weaker completion, lower repeatability, harder to keep current

The 7MM learning unit

The seven-minute envelope is only half of the method. The other half is what goes inside it. Every unit follows the same internal shape, and each step earns its place against the evidence.

- **One objective.** The single thing the learner should be able to do afterwards. This keeps demand inside working-memory limits and gives a clear sense of completion.
- **Focused context.** A short framing of why this matters and what the learner already knows, so the new idea has somewhere to attach.
- **Core explanation.** The substance, delivered in a few clear chunks rather than one dense block.
- **Example or scenario.** A concrete instance that turns an abstract point into something recognisable.
- **Practice.** An active step: a choice, a decision, a sort, a short application. This is where the work moves from watching to doing.
- **Retrieval check.** A brief moment of recall that strengthens memory, the testing effect applied in miniature.
- **Next action.** A small, clear step that turns the objective toward real behaviour.

> **The seven-minute design equation**
>
> one objective
>
> - focused context
> - core explanation
> - example
> - practice
> - retrieval check
> - next action
>
> **= one seven-minute learning unit**

This shape is format-agnostic. An interactive module expresses it as timed segments. A video expresses it as beats. A podcast or a storybook expresses it in its own structure. The same governed knowledge can become any of these. The objective is chosen first. The format follows. The knowledge stays the same, and only the delivery changes.

A short unit built without this shape is just short content, and short content is not the claim. The claim is short content designed well.

Where longer learning still belongs

Seven Minute Modules is not an argument against workshops, classroom training, coaching, roleplay, group discussion, or longer learning journeys. Those formats are how Group Moovs has built behaviour change for years, and they remain essential when the goal is practice, reflection, social learning, feedback, culture change, or behavioural rehearsal.

The point is different. When the task is to transfer a specific piece of knowledge or build one clear capability, the unit should be small enough to finish, trace, review, and maintain. Longer programmes can and often should be built from several Seven Minute Modules, combined with human interaction around them.

Around the module sits the human layer: assignments, coaching conversations, group discussions, demonstrations, roleplay, manager check-ins, peer feedback, and workplace practice. These are not exceptions to the method. They are how the method becomes behaviour.

A three-hour classroom session is not the enemy. A three-hour content dump is. The problem is not duration itself. The problem is unstructured duration that hides too many objectives in one block. A well-designed three-hour session is a sequence: several seven-minute knowledge units, with practice, discussion, reflection, application, and commitment around them. Seen this way, longer learning is not a contradiction of the method. It is a different layer that the method makes stronger.

A short note on the Assurance Grade

Within Seven Minute Modules, the Assurance Grade is the visible trust label attached to a learning unit. It shows whether the module is sourced, reviewed, current, traceable, and owned. It is not a decorative badge. It is a view of the governance behind the module.

The seven-minute unit is not only a learning choice. It is what makes that trust operable at scale.

A short, single-objective unit, drawn from a known set of sources, is far easier to govern than a long, mixed one. It is easier to source, because there are fewer claims and each can be traced. It is easier to review, because a person can actually check seven minutes of content properly. It is easier to update, because a change touches one small unit rather than a sprawling course. It is easier to regenerate, easier to audit, and easier to map to a specific policy or regulation. And it is easier to keep current, because when a source changes, the registry can point to the exact unit that depends on it.

This is the quiet advantage of granularity. The smaller and more focused the unit, the more honestly you can stand behind it. The seven-minute constraint and the Assurance Grade are two sides of the same design decision: produce learning you can trace, review, and keep true. Trusted, traceable learning, seven minutes at a time.

The honest limits

Seven minutes does not solve everything, and saying otherwise would undermine the trust this method exists to protect.

Some topics are too large for one unit, and belong in a pathway of several, sequenced and spaced. Some skills cannot be learned in a single sitting at all, and need deliberate practice over time. Some outcomes depend on things no module controls: coaching, manager involvement, and a workplace that supports the new behaviour. The transfer evidence is explicit that this context often matters more than the content itself.[10:2]

A module can prepare, focus, reinforce, and evidence a learning moment. It cannot replace every human moment. For many outcomes, the strongest design is a seven-minute unit surrounded by practice, conversation, feedback, and application.

So Seven Minute Modules is a unit within a larger learning system, not a replacement for one. It is the atomic piece that pathways are built from, the piece that can be combined, spaced, reinforced, and kept current. Short is a discipline, not a panacea. Treating it as a panacea is exactly the overclaim we refuse to make.

Conclusion

The instinct to teach more by adding more is the thing that makes most learning fail. Seven Minute Modules takes the opposite path. It fixes the unit at one objective and a seven-minute envelope, not because the research crowns that number, but because a hard limit forces the clarity, the focus, and the discipline that good learning needs.

The evidence does not promise a magic duration, and we do not claim one. It supports short, single-objective, active, spaced learning, designed well and kept current, and surrounded by the human

practice that turns knowledge into behaviour. Seven minutes is the practical envelope that makes all of that possible, and small enough that you can stand behind every unit you produce.

Seven minutes is not a number. It is a design discipline.

Notes

Editor's Notes

Internal. Not part of the published paper.

Major choices

- Title and subtitle now carry the paper. The fixed conclusion line became the title; the explanatory line is the subtitle. The earlier "design discipline, not a gimmick" subtitle was dropped because "gimmick" read oddly.
- Positioning spine is v3; the evidence depth is recovered from v1 and v2, which v3 compressed. This paper is where that depth belongs.
- The single claim is held throughout: seven minutes is a design envelope, not a magic number, and not a fixed attention span. The attention-span myth is met early, and the case is deliberately not rested on it.
- New section, "Where longer learning still belongs," placed before "The honest limits." This is the commercial and philosophical fix: it states plainly that 7MM is the atomic knowledge-transfer unit, not a replacement for the workshops, coaching, and human practice that are Group Moovs' heritage. The keeper line is "A three-hour classroom session is not the enemy. A three-hour content dump is."
- The human layer (assignments, coaching, demonstrations, roleplay, manager check-ins, peer feedback) is named explicitly, so the paper is not read as content-unit-only. A matching line was added to "The honest limits."
- The Assurance Grade is now defined in one short paragraph before it is used, so the paper stands alone and does not assume the reader saw v3.
- A single boxed model, the seven-minute design equation, anchors "The 7MM learning unit." Its labels match the section's bullet labels exactly, for consistency. It is the only model in the paper.
- The duration-envelope table (1 to 3 / 5 / 7 / 10 to 15 / 20 to 30 minutes) is the only table; it is high-clarity for instructional designers and procurement reviewers.

Edits applied in this version

- Propagation pass, 6 June 2026 (v1.1): the Bersin figure moved from body text to the evidence register per the Appendix B rule. No other content changes.
- Softened, per agreement: "Much workplace learning" (not "Most"); "covers too much, so too little reaches the point of use"; "rarely fits that week without being scheduled, deferred, or interrupted"; "one objective, seven minutes"; "We design the unit for a seven-minute envelope, and we redesign rather than stretch"; "often carries many objectives"; "Many workplace courses are built to be comprehensive"; "the studies vary widely."
- Guo claim no longer says "largest study." It now reads "a large-scale study... based on 6.9 million viewing sessions," which is the verified part. Do not let "largest" return without verification.
- "The capacity of the apparatus" replaced with "It reflects how limited working memory actually is." Other strained phrases were de-jargoned in the same pass.
- The vague 2024 meta-analysis footnote was removed. The microlearning point now uses the range-and-design framing, which does not need the meta-analysis as a load-bearing citation.
- "Above / below" navigation language removed throughout, so the paper reads as standalone prose.

Claims to verify before publication

- The two short-form-media studies (notes 12 and 13) were verified during drafting: Tam & Inzlicht 2024 in *Journal of Experimental Psychology: General* (seven experiments, 1,223 participants), and Chiossi et al. 2023 at CHI '23 (N=60). The prospective-memory study is small; it is framed as suggestive, not conclusive, and is paired with the explicit statement that this is not proof of a shorter attention span.
- The twenty-four-minutes-a-week figure (Bersin) is flagged in text as an industry benchmark. Keep that flag.
- No completion percentage is stated as fact. The old marketing figures remain banned.
- The working-memory "about seven items" point (Miller) appears only to dismiss it as the basis for the seven-minute unit. Keep that framing.
- Cognitive-load pass applied: long sentences split, abstract-noun stacks removed, consultant phrasing avoided, each section closed on a short plain line, no em-dashes.
- Verify all citation details against the master bibliography once it exists. Sources here are the validated set used across v1 to v3 and the research dossiers, plus the two newly verified short-form-media studies.

1. Guo, P. J., Kim, J., & Rubin, R. (2014). How video production affects student engagement: an empirical study of MOOC videos. *Proceedings of the First ACM Conference on Learning @ Scale*, 41–50. Based on 6.9 million video sessions; engagement dropped sharply after roughly six minutes, and video length was the strongest predictor of engagement. ↩ ↩ ↩
2. Sweller, J. (1988). Cognitive load during problem solving: effects on learning. *Cognitive Science*, 12(2), 257–285. ↩ ↩ ↩
3. Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255. ↩ ↩ ↩
4. Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: a review and quantitative synthesis. *Psychological Bulletin*, 132(3), 354–380. ↩ ↩ ↩
5. Wilson, K., & Korn, J. H. (2007). Attention during lectures: beyond ten minutes. *Teaching of Psychology*, 34(2), 85–89; Bradbury, N. A. (2016). Attention span during lectures: 8 seconds, 10 minutes, or more? *Advances in Physiology Education*, 40(4), 509–513. ↩ ↩
6. Bersin by Deloitte (2014). *Meet the Modern Learner.* A widely cited industry benchmark, not a peer-reviewed result. The specific figure is held in the 7MM evidence register (009) and re-verified before client use. ↩ ↩
7. Ebbinghaus, H. (1885). *Über das Gedächtnis*; replicated in Murre, J. M. J., & Dros, J. (2015). Replication and analysis of Ebbinghaus' forgetting curve. *PLOS ONE*, 10(7), e0120644. ↩ ↩
8. Paas, F., & Sweller, J. (2012). An evolutionary upgrade of cognitive load theory. *Educational Psychology Review*, 24(1), 27–45. ↩ ↩
9. Mayer, R. E., & Pilegard, C. (2014). Principles for managing essential processing in multimedia learning: segmenting, pre-training, and modality principles. In R. E. Mayer (Ed.), *The Cambridge Handbook of Multimedia Learning* (2nd ed.). ↩
10. Blume, B. D., Ford, J. K., Baldwin, T. T., & Huang, J. L. (2010). Transfer of training: a meta-analytic review. *Journal of Management*, 36(4), 1065–1105; Baldwin, T. T., & Ford, J. K. (1988). Transfer of training: a review and directions for future research. *Personnel Psychology*, 41(1), 63–105. ↩ ↩ ↩
11. Bunce, D. M., Flens, E. A., & Neiles, K. Y. (2010). How long can students pay attention in class? A study of student attention decline using clickers. *Journal of Chemical Education*, 87(12), 1438–1443. ↩
12. Tam, K. Y. Y., & Inzlicht, M. (2024). Fast-forward to boredom: how switching behavior on digital media makes people more bored. *Journal of Experimental Psychology: General.* Seven experiments, 1,223 participants; switching between short videos increased boredom rather than relieving it. ↩
13. Chiossi, F., Haliburton, L., Ou, C., Butz, A., & Schmidt, A. (2023). Short-form videos degrade our capacity to retain intentions: effect of context switching on prospective memory. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems.* A small between-subjects

experiment (N=60); interruption by short-form video feeds significantly impaired prospective memory, while two other feed formats did not. ↩