


MATRIX

EXPERTISE MODEL

Whitepaper v1.5 — how people can build expertise faster

Don't teach the learner every answer. Build a small library of generative models that enables the learner to see the signal, understand the principle and generate a fitting action on their own.

Expertise = automation of the most generative principles – and the ability to know when they do not apply.



SEE THE SIGNAL · UNDERSTAND THE PRINCIPLE · GENERATE THE SOLUTION

Lennert Kvamme

Summary

The MATRIX Expertise Model is a learning architecture for how people can build independent, adaptive expertise faster. Its starting point is that many complex domains can be compressed into a small library of generative principles: models that are activated when the learner recognises signals in the situation, and that can produce many situation-specific actions.

The method has five learning steps in a fixed order: **principle → generative model → signal → automatic response pattern → freed cognitive capacity**. When basic actions run on their own, working memory is freed for awareness, choices and creativity.

MATRIX IN SIX SENTENCES

1. Working memory is the bottleneck.
2. Generative principles are compression.
3. Automation is installed compression.
4. Expertise is the automation of the most generative principles.
5. Freed capacity enables better navigation.
6. Adaptive expertise is knowing when the automation does not apply.

Definition: *Expertise = automation of the most generative principles – and the ability to know when they do not apply.*

Unfolded: expertise is experience compressed into signal-activated, automated generative models that free and expand cognitive capacity for navigation, even when the situation is new.

MATRIX does not claim that principles alone create expertise. Technical skill, experience, feedback, physical and cognitive capacity, motivation and the quality of practice remain decisive. MATRIX proposes a way of organising knowledge and practice so that every experience has a higher transfer value.

THE MATRIX HYPOTHESIS

Learning can be accelerated when many individual situations are compressed into generative models that (1) are learned as a principle and a mechanism, (2) are activated when the learner recognises relevant signals, (3) are automated as response patterns in game-like practice, and (4) are updated through feedback. The effect should first appear as better transfer to new variations within the same domain.

The Matrix Playbook series (Levels 5, 10, 15 and 20) is the model's first implementation, in football. This whitepaper describes the model, its relationship to research, a method for discovering principles, a development ladder, testable predictions, and how AI can be used as scaffolding without making people dependent on it.

MATRIX is positioned as the player's self-training alongside club and academy training, not instead of it. The books build understanding and perception at home. Club training and matches are where the models become automatic (chapter 4).

New in v1.5: the whitepaper now describes MATRIX's place in the whole: self-training alongside club and academy, what each book is for, and how home, club and match share the work (chapter 4). The pilot is redesigned as self-training in one club with two groups that act as each other's control, the transfer terms are made unambiguous, and it describes how the boundaries of a principle are trained. Version 1.4 added adaptive expertise to the definition, extended freed capacity with long-term working memory, based communication on external focus and analogy learning, operationalised generativity and limited the transfer claim.

Background: how MATRIX came about

MATRIX grew out of my own way of learning across domains.

When I later taught my children football, I started from principles they could understand and use to find their own solutions. I did not yet know enough about football to always choose the most generative principles. Through years as a grassroots coach, and by following player development from children's football onwards, I have revised and collected the principles in Matrix Playbook Levels 5, 10 and 15. Level 20 draws especially on my experience from martial arts and mental training.

The book series is therefore the result of both a learning method and an ongoing refinement of the models it teaches. It is also an example of the model itself: principles tested in practice, feedback from players and matches, and models that are compressed and improved over time (chapters 9 and 12).

— Lennert Kvamme

1. The problem MATRIX tries to solve

Beginners often meet a domain as a large collection of isolated rules, techniques and situations. This creates a high cognitive load and weak transfer: the learner has learned the answer to situation A but lacks a system for solving situation B.

MATRIX therefore does not first look for more answers. It looks for structures that can produce families of answers.


Level	Example
Catalogue (low-level rules)	"Pressure from the right → touch to the left." "Pressure from the left → touch to the right."
MATRIX compression	Directional Touch: first touch AWAY from pressure, into space.

Level	Example
Generated in the situation	The specific foot, angle and touch.

Such a principle need not be absolute. At a high level an expert may break, combine or manipulate the basic model. That does not disqualify the principle as a highly effective default at earlier stages of development.

2. The method: five learning steps

MATRIX is learned in five steps in a fixed order. The first three build understanding and perception; the last two turn understanding into automaticity and free capacity.



Step	What happens in the learner	Example: Directional Touch (Level 5)
1 Principle	Gains a core insight into how the domain works.	“Whoever owns the first touch owns the ball.”
2 Generative model	Understands cause and effect, and which actions open up.	A touch away from pressure buys time and space. The same logic applies wherever the pressure comes from.
3 Signal	Recognises the pattern in the situation, in time. It is the pattern recognition that activates the model. The coach may use a trigger word as an attention anchor.	The presser’s direction before receiving. Trigger word: “Touch!”
4 Automatic response pattern	The signal triggers the pattern without hesitation, through many varied repetitions.	The touch into space happens without thinking.
5 Freed capacity	Working memory is free for awareness, choices, creativity and reading feedback.	The player is turned and sees the next pass before the pressure arrives.

Why this order

Working memory is the bottleneck. It is too small to control everything in a match, but large enough to build the models that later run by themselves. The first three steps use working memory to build and understand the model. Step 4 installs it, so that it runs without working memory. Step 5 is the payoff: the capacity that is freed is used for navigation.

Capacity is not only freed – it is expanded. Experts build what Ericsson and Kintsch called *long-term working memory*: domain-specific structures that let them keep far more relevant information active than beginners can. Chess masters see meaningful wholes where beginners see single pieces (Chase and Simon). The models therefore free up space and also give working memory a larger vocabulary to navigate with.

Element	What it is
Generative principle	Compression: many situations gathered into one model
Automation	Installed compression: the model runs without working memory
Expertise	Automation of the most generative principles – and the ability to know when they do not apply
Navigation	What the freed and expanded capacity is used for

“Automated” and “generative” are not opposites. What becomes automatic is the recognition and the default response. The adaptation to the situation remains generative.

THE TEMPO RULE

Think slowly to build the model. Train with variation until you see the pattern quickly. Play fast when the signal is clear – and think again when it isn’t.

The last part is as important as the first three. Automaticity is right when the situation resembles what the model was built for. When the signal is unclear or something unexpected happens, the player should bring working memory back in. Knowing the difference is part of navigation, and it is what separates expertise from habit.

Boundaries are trained with contrasts. Every principle has an antidote: situations where the same signal calls for a different answer. When the player practises both the rule and the exception, they learn to see the difference instead of just repeating the answer. That is how adaptive expertise is trained, not just described.

ROUTINE AND ADAPTIVE EXPERTISE

Hatano and Inagaki distinguished routine expertise, which is fast and accurate in familiar situations, from adaptive expertise, which also copes with the unfamiliar. Automation alone produces routine expertise. The tempo rule and navigation make MATRIX adaptive: the player knows when the model applies and when it does not. That is why adaptive expertise is part of the definition.

How the principle is communicated

Explicit rules about one’s own body can break down under pressure, because the player starts to control the movement consciously (Masters). MATRIX principles are about the game and the space, not the body, and so offer little of the “move your foot like this” kind. Four measures reinforce this:

- **The principle as an image.** The memory image is an analogy, not a manual. Analogies produce learning with less explicit rule knowledge and hold up better under pressure (Liao and Masters).
- **Focus outwards.** Trigger words point to the signal and the space (“Eyes!”, “Space!”), not to the body. An external focus of attention produces better learning and performance than an internal one (Wulf).
- **The coach falls silent.** When the model sticks, the word disappears too. Then there is no rule left to think about in the match.
- **Adapted to age.** Younger children learn the principle through images, play and games. Older players can also learn through explanation and by comparing examples.

Three clarifications

1. **Learning and performance run in opposite directions.** The method is learned from principle to capacity. In the situation it runs the other way: the signal comes first, activates the model and triggers the response pattern in a fraction of a second (chapter 4).
2. **A response pattern, not a single response.** What is automated is the link from signal to model and the robust default. The foot, the angle and the tempo are still generated by the situation. If one fixed action is automated, the model becomes a rule and loses transfer.
3. **The local principle comes first; the big law is discovered.** Each mission starts with its own principle. The laws that compress many missions are often revealed last, so that the learner sees the pattern themselves. This is not a break with “principle first”, but compression learned through discovery.

THE TRIGGER WORD IS THE COACH’S TOOL

The coach uses the word as an attention anchor that helps the player direct their gaze to the signal. As the player masters the model, the word becomes their own: they say it to themselves, or no longer need it. Then the coach should be quiet. A trigger word the coach still has to shout is a sign that step 4 has not been reached.

Each step is supported by existing research (chapter 10). What is hypothetical is that the whole chain in this order produces expertise faster than other training approaches.

3. The Matrix Playbook series as a first implementation

Four books implement the model in football, from perception to identity. Each book has 14 missions that are compressed into a few laws, and each mission is built as a MATRIX loop. The books are currently published in Norwegian; titles and quotes are translated here.

Level	Book	Tagline	The player learns	Compression
5	Matrix Playbook – The Fundamentals (Fundamentene)	“See the game before it happens”	To SEE the game	14 fundamentals → 1: SEE – BEFORE YOU ACT
10	Matrix Playbook – The Hacks (Hackene)	“Play against his brain, not his shirt”	How one brain fails	14 hacks → 5 laws
15	Matrix Teamplay – Team Hacks	“See the same. Think the same. Act together.”	How eleven fail at once	14 team hacks → 5 laws
20	Matrix Identity	“From individual to part of something bigger”	Who you become inside it	14 installations → 5 engines → 1 system

The laws that compress the books

- **Level 10:** People don’t see everything. People guess more than they see. The body has limits. Timing is everything. People follow habits.
- **Level 15:** A team looks at one place at a time. A team is pulled towards the ball. The last man is the worst covered. A team has seams. A team needs five seconds to become a team again.
- **Level 20:** Focus (a player doesn’t see everything), Resilience (you don’t decide the mistake), Courage (fear is a signal, not a command), Leadership (everyone sees you), Discipline (you fall to what you have decided).

Anatomy of a mission: the five steps on every page

Step	Levels 5 and 10	Level 15	Level 20
1 Principle	Memory image, part motto	LAW on every card, memory image	The engine, Identity Hack
2 Generative model	Why it works	Why it works, “Can I explain why?”	Matrix reveal
3 Signal	Trigger word, four pictures from “What do you see?”	Trigger word, three pictures	SIGNAL and TRIGGER
4 Automaticity	Mini-drill, “Mastered when”, three ticks	SEE → TRY → OWN	UNDERSTOOD → ACTIVATED → AUTOMATED
5 Freed capacity	Implied in match mode	Implied	Implied
Limits	Antidote, Matrix Defence	Most common mistake, The team’s answer	System error
Composition	Match mode, Missions	Combines with, bosses	Your Matrix, The complete player
Feedback	Matrix log	Matrix log, Team Score	Evidence file

What the books have taught the model

- **The golden rule:** “Shout the trigger word – never the explanation” and “Ask ONE question” are the scaffolding principle in practice. Over time the word becomes the player’s own, and the coach falls silent.
- **The mirror:** every principle has an antidote. Knowing both sides of a model gives clarity about its limits.
- **Read their team (Level 15):** context decides which models work today. This is the CONTEXT stage made operational.
- **The evidence file (Level 20):** identity is built from dated evidence, not from words.
- **Using it:** the power must make the whole team better, never humiliate anyone. Ethics is part of the model.
- **Mastery climate:** the player competes against themselves, chooses freely and has no deadline. Children are never ranked against each other. This keeps learning in a mastery climate rather than a results climate.

4. Place in the whole: self-training alongside club and academy

MATRIX does not replace club training. The club and the academy train the team, tactics, technique and physical capacity. MATRIX is the player’s own layer on top of this: it teaches the player to *see the game*, so that every session and every match they already have produces more learning.

With MATRIX the player does not train more football, but better. The self-training is short, 10–15 minutes, and the aim is for the principle to follow the player onto the pitch. Club training and matches are not competitors to MATRIX. They are where the models become automatic.

THE POSITION IN ONE SENTENCE

MATRIX is the player’s self-training in seeing the game – alongside the club and academy, not instead of them – so that every session and every match produces more learning.

What already exists

Type of offer	What it trains	What is missing
Technique programmes and drill libraries	Ball skills and technique	Why and when a technique should be used
Cognitive programmes (apps, computer games, VR)	Scanning, reaction and decision speed	Trained detached from a principle, and transfer to matches is uncertain
Books with advice on game understanding	Tips and rules of thumb	A learning method, mastery levels and a development ladder
Club training	Team, tactics, fitness and play	Time for each individual player's models

MATRIX fills the space between these: principles the player understands, signals they recognise, and a path from self-training into real play. This is what separates MATRIX from cognitive training programmes: the player does not train faster vision in general, but learns what to look for, and why.

The division of labour: home, club and match

Step	Where it happens	Who helps
1-2 Principle and generative model	At home: the book or the app, 10-15 minutes	The player, with a parent for the youngest
3 Signal	At home: pictures and video, or 1v1 with siblings and friends	A parent or the app says the trigger word
4 Automatic response pattern	In club training and matches	The player says the word to themselves, then no longer needs it
5 Freed capacity and feedback	The Matrix log after training and matches: "What did I see? Did it work?"	The player, with one question from an adult

Self-training alone cannot install automaticity. That requires real opponents, pressure and variation, which club training and matches provide. MATRIX therefore builds understanding and perception at home, and uses the pitch the player already has to make it automatic.

What the books do, and what they are for

Book	Purpose	Used by	Works
Level 5 – The Fundamentals	Seeing the game. The foundation everything else builds on.	The player, with a parent helping the youngest	At home, then in training and matches
Level 10 – The Hacks	Reading one opponent and how his brain can be fooled	The player	In 1v1 situations in training and matches
Level 15 – Team Hacks	Reading a whole team, and giving the team a shared language	The player alone, but with the greatest effect when coach and team share the language	At home and in the club track
Level 20 – Identity	Who the player becomes: focus, resilience, courage, leadership and discipline	The player	The whole football life, also off the pitch

Together the books have one purpose: to give the player their own library of models that they understand, recognise in the game and eventually use without thinking. Over time the player should discover the boundaries, combine the models and develop their own. The books are designed so that the player can use them without a coach, but with the pitch as the testing ground. They are a delivery channel for the method, not the method itself (Appendix B).

Three ground rules for working together

- 1. The coach decides in training and matches.** MATRIX principles are used within the team's instructions, never against them. If there is a conflict, the player follows the coach and raises the question afterwards.
- 2. Parents help at home, but do not coach from the sideline.** The golden rule applies at home too: the trigger word, never the explanation, and one question afterwards. In matches, the coach does the talking.
- 3. The trigger word moves.** At home a parent or the app says it. In training the player says it to themselves. In matches it is no longer needed.

The sign that MATRIX is working

The question parents and coaches can ask themselves is not whether the child does it right, but this:

THE MATRIX QUESTION

Is the child beginning to notice something they used to miss, understand why it might matter, try a solution and discover something new when the outcome turns out differently?

What you see in the child	The step in MATRIX
Notices something they used to miss	Signal: pattern recognition has woken up
Understands why it might matter	Generative model: the child sees cause and effect

What you see in the child	The step in MATRIX
Tries a solution	Response pattern: the model becomes action
Discovers something new when the outcome turns out differently	Feedback and boundaries: the model is updated, and the child learns where it does not apply

The question measures learning, not performance. An attempt that fails, but where the child discovers something new, counts as much as one that succeeds. The last part matters most: it is the beginning of adaptive expertise, and of the player discovering boundaries and developing their own models.

The question is for the adult, not the child. The child is asked one question at a time, for example “What did you notice today that you didn’t see before?”, “Why do you think it matters?” or “What happened when it turned out differently than you expected?”

Two tracks

- **The main track: self-training.** The player uses the books and the app on their own, alongside regular club and academy training.
- **The club track: a shared language.** Coaches and clubs that want to use MATRIX together can add the 15-minute template from the books to their sessions and use Level 15 as the team’s shared language. The club track is an option, not a requirement.

5. The MATRIX core in performance

In the situation itself, the expert follows this chain:

**SIGNAL → STRUCTURE → PRINCIPLE → MODEL → POSSIBILITIES → CHOICE → ACTION →
FEEDBACK → UPDATE ↻**

The chain is not a conscious, slow thought process. With training, the links between signal, model and action become fast and automatic. The point is that the action is informed by structure, not retrieved as an isolated answer.

Once the model is installed, the chain can be simplified:

**SIGNAL → PATTERN RECOGNITION → PREDICTION → ACTION → OUTCOME → UPDATED
MODEL ↻**

Working memory does not need to be involved in every link. It is freed for navigation: awareness, choices and combining models in situations the player has not seen before. The outcome and the update close the loop, so the model keeps improving (see Appendix C).

The books show this chain in every mission with four pictures: *What do you see?* → *What does it mean?* → *What happens now?* → *You're there first!*

Instruction, rule and MATRIX model

Form	Wording
Instruction	“Do X.”
Rule	“When A happens, do X.”
MATRIX model	“When you detect A, recognise structure B. Because of mechanism C, the space of possibilities changes. Choose the action that fits the situation.”

6. The development ladder

The ladder follows the method's five steps. Stages 2–6 are measured per principle. Stages 7–9 are measured across the library, because they always involve several models.

Stage	Characteristic	Step in the method
1 CHAOS	The learner does not know what matters.	–
2 PRINCIPLE	The learner can state the core insight in their own words.	1
3 MODEL	The learner understands cause and effect and can suggest several actions in new situations.	2
4 SIGNAL	The learner recognises the signal in time during play, and the recognition activates the model.	3
5 AUTOMATICITY	The response pattern runs without hesitation under pressure.	4
6 FREED CAPACITY	The learner scans, communicates and chooses while the pattern runs. Context and exceptions become visible.	5
7 COMBINATION	Several models are used at once.	Across
8 MANIPULATION	The performer influences other people's signals and expectations.	Across
9 GENERATION	New solutions are built from known structures in unknown situations.	Across

Three levels of mastery per principle

For the learner, the ladder is simplified to three levels, as in Levels 15 and 20.

Level	The learner says	Typically takes	Stages
UNDERSTOOD	"I know what's happening."	Five minutes	2–3
ACTIVATED	"I've tried it on purpose."	One match	4
AUTOMATED	"I do it when it counts, without thinking."	A season	5–6

7. Model composition and unfamiliar situations

A model library is most valuable when the situation is new. The new situation may contain familiar sub-structures, even if the whole has never been seen before.

Situation	What happens
Familiar situation + familiar model	Fast default response
New situation + familiar structure	Transfer
New situation + new combination of familiar structures	Model composition and a new solution

The books train composition explicitly. Match mode shows chains in which three models work within one second: scan behind you (“Look!”), receive half-turned (“Open!”), first touch into space (“Touch!”). Matrix Missions set a goal without a recipe, and Level 15 has “Combines with” on every card.

This is a central, testable hypothesis: MATRIX training should not only improve tasks similar to the training, but also increase the ability to recognise and combine relevant structures in new problems.

8. Pattern recognition and fractal reach

MATRIX assumes that expert learning often means recognising a known structure beneath a new surface. The model distinguishes three forms:

- **Horizontal:** the same structure in different situations within the same domain.
- **Vertical:** related structure at different scales, such as one opponent, eleven opponents and one’s own inner state.
- **Cross-domain:** an abstract structure that seems to appear in several fields.

The book series is an example of vertical reach. The same engine is used at four scales:

Level	Scale	The same structure
5	The game	One place at a time: his eyes reveal where the ball is going (The Eye Reader).
10	One brain	People don’t see everything: he blinks, and you disappear (Saccadic).
15	Eleven brains	A team looks at one place at a time: eleven heads, one direction, one gap (Same Picture).

Level	Scale	The same structure
20	Your own system	A player doesn't see everything, not even inside: what did you actually look for first? (the Focus engine, The Spotlight)

Cross-domain similarity is a hypothesis, not proof. Two phenomena can look alike without sharing a causal mechanism. MATRIX therefore requires a mechanism test before an analogy is used as an explanation.

FRACTAL REACH

Across how many situations, levels or scales does a model retain its explanatory and generative power? The term is used analogically, not mathematically.

9. How generative principles are discovered

MATRIX Discovery is the process from experiences to compressed models:

1. Map the domain: recurring tasks, environments and problems.
2. Find the expert signals: what do skilled performers notice that beginners miss?
3. Compare contrasts: similar situations with different outcomes. What changed?
4. Remove the surface: ignore names, techniques and cosmetic differences.
5. Find the invariant: which relationship holds across the examples?
6. Find the mechanism: why does that relationship change the situation?
7. Formulate the principle: "When [signal/structure], because [mechanism], [space of possibilities] changes."
8. Test generativity and transfer: does the model give good answers in new variations?
9. Find the limits: when is the model insufficient or subordinate to another?
10. Compress again: are several principles expressions of a deeper structure?

**EXAMPLES → CONTRAST → PATTERN → INVARIANT → MECHANISM → SIGNAL →
GENERATIVE MODEL**

The Discovery protocol is for the author and the subject-matter expert. The books give the player selected core models to explore in play. The aim is for the player, over time, to discover boundaries, combine models and develop their own (stages 7–9 of the development ladder).

The cases do not disappear. Experts do not only have principles, but large libraries of familiar situations they recognise (Klein). Compression therefore does not mean discarding the examples. Each principle keeps its best examples and contrasts, because those are what the learner needs to build the model, and what makes recognition fast.

10. Relationship to existing research

MATRIX is a new synthesis and design hypothesis, not an established scientific theory. Each learning step connects to existing research:

Step or element	Field of research	What the field says
Principle and generative model	Schema induction and analogical transfer (Gick and Holyoak)	Comparing several examples supports abstraction of a general schema, and the quality of the schema is related to later transfer.
Signal	Perceptual-cognitive expertise in sport (Mann, Williams, Ward and Janelle)	A meta-analysis of 42 studies shows that experts pick up perceptual cues more accurately and faster than non-experts, and use more efficient visual search patterns.
Game-like practice	Representative learning design (Pinder, Davids and colleagues)	Training should preserve the coupling between information and action found in the performance environment.
Automatic response pattern	Skill acquisition (Fitts and Posner)	Practice makes performance progressively more automatic.
Freed capacity	Cognitive load theory (Sweller)	Working memory is a bottleneck in learning.
Varied practice	Contextual interference (Shea and Morgan)	Varied and mixed practice lowers performance during training but improves learning and transfer.
Communicating the principle	External focus and analogy learning (Wulf; Liao and Masters; Masters)	An external focus produces better learning than an internal one. Analogies and implicit learning hold up better under pressure than explicit rules.
The expert's working memory	Long-term working memory and chunking (Ericsson and Kintsch; Chase and Simon)	Experts expand working memory with domain-specific structures and see meaningful wholes.
Adaptive expertise	Routine and adaptive expertise (Hatano and Inagaki)	Expertise that copes with the unfamiliar requires understanding why, not just fast execution.
Guided discovery	Instructional research (Kirschner, Sweller and Clark)	Minimally guided instruction works poorly for novices. Guidance should be faded gradually.
Feedback and repetition	Deliberate practice (the Ericsson tradition)	Structured training and feedback are central, but practice alone does not explain all individual differences (Hambrick; Macnamara and Maitra).

MATRIX's distinctive contribution is the combination: a fixed learning order, signal-based model activation, explicit compression, mastery per principle, model composition, fractal reach and a ladder where basic principles can later be refined or manipulated.

MATRIX may also bridge two research camps. Cognitive psychology speaks of internal models and schemas. Ecological dynamics speaks of the direct coupling of perception and action, and of possibilities for action (affordances) in the environment. MATRIX's *space of possibilities* is close to the latter (Gibson). And the trigger word as an attention anchor is in practice what ecological researchers call the *education of attention*: teaching the player to direct their gaze at the information that matters.

An important caveat about transfer. Far transfer, from one domain to a completely different one, has proved hard to demonstrate (Barnett and Ceci; Sala and Gobet). MATRIX's transfer claims therefore apply first to near and structural transfer within the same domain. Far transfer is a long-term hypothesis, not a promise.

IMPORTANT LIMITATION

This whitepaper does not claim that MATRIX has already been shown to accelerate expertise. It proposes a testable architecture and predictions that can be investigated empirically.

11. How to measure the quality of a MATRIX principle

Measure	Question
Signal clarity	How easily can the learner detect the information that activates the model?
Generativity	How many situation-specific actions can the model produce?
Transfer	How well does the model work in new situations with the same structure?
Compression	How many low-level rules or cases does the model replace?
Mechanism clarity	Can we explain why the model works?
Automatability	Can the response pattern be practised until it runs without hesitation?
Independence	Does the model reduce the need for the coach's specific solution?
Fractal reach	Does the model have useful structure at several scales?
Boundary clarity	Do we know when the model must be refined or overridden? Does it have an antidote?

How the key measures can be assessed

Measure	Operationalisation
Generativity	Number of different, correct actions the player proposes or performs across five new variations with the same structure, rated independently by two coaches
Transfer	Share of new variations solved, distinguishing near transfer (new surface, same situation) from structural transfer (new situation, same structure)
Automatability	Hesitation from signal to action in play under pressure, measured on video
Boundary clarity	Share of exception situations in which the player sees that the model does not apply

PROVISIONAL HEURISTIC

MATRIX VALUE $\approx$ Signal clarity $\times$ Generativity $\times$ Transfer $\times$ Compression $\times$ Mechanism clarity.
This is a conceptual rule of thumb, not a validated formula.

12. Research and validation programme

The most important test is not whether users like the model, but whether it produces better transfer and more independent problem solving.

1. Choose a bounded domain and 5–10 candidate principles. First pilot: six fundamentals from Level 5.
2. Create two self-training programmes with equal time, each with its own set of principles. Both groups continue regular club training, and each group is the control for the other group's principles.
3. Measure learning in trained situations.
4. Measure near transfer: a new surface, the same situation.
5. Measure structural transfer: a new situation, the same structure. Far transfer, to other domains, is not measured in the first pilot.
6. Measure automaticity: hesitation from signal to action in play under pressure.
7. Measure freed capacity: can the learner scan or communicate while the response pattern is executed (dual task)?
8. Measure independence: how often does the learner need specific instruction?
9. Measure model quality: can the learner explain the signal, mechanism and limits?
10. Follow development over time and test whether the library improves through feedback.

A simple qualitative measure. The four parts of the MATRIX question (chapter 4) are also used in the pilot. The Matrix logs and short conversations with the players are coded by which parts the

player shows: notices, understands why, tries, discovers something new. This allows progress to be followed week by week, including for players who are not yet succeeding in matches.

Pre-registration

Before the pilot starts, the hypotheses, measures and what counts as the method *not* working are registered. A poor result then cannot be explained away afterwards by saying the principles “were not generative enough”. Control principles are chosen with a different structure from the trained ones, so that learning does not spill over. Two coaches rate the video independently of each other; and the agreement between them is reported.

Predictions

- The MATRIX group identifies relevant signals faster in new variations.
- Self-training at home produces progress on the trained principles in matches too, where nobody says the trigger word.
- The MATRIX group shows better near and structural transfer when the surface changes but the structure remains. Far transfer to other domains is not expected to appear in the first pilot.
- Players who learn the principle as an image and with an external focus maintain performance under pressure better than players who learn explicit body rules.
- The effect is largest for principles with high compression and clear signals.
- Automated response patterns produce measurably better performance on dual tasks.
- Too many caveats early reduce the benefit of compression. Nuances should come developmentally.
- At advanced levels, combining models becomes more important than single models.

13. Roadmap

The next step is to test the model empirically.

1. A pilot in one club for 8 weeks as self-training alongside regular club training. The players are split into two groups that each train their own set of three fundamentals from Level 5, and are tested on both sets before and after. Each group is then the control for the other, and differences in how hard the principles are even out. Hypotheses and measures are pre-registered.
2. Three profiles in very different domains, to test whether the model compresses rather than merely summarises.
3. An evaluation set with “tip versus generative model” examples.
4. Expert review of the principles, labelled VALIDATED, SYNTHESIS or HYPOTHESIS.
5. Further development of the book series based on the pilot.
6. MATRIX Expertise Model v2.0 based on transfer data and expert critique.
7. Then: the self-training app, the club track for coaches, domain profiles and an agent library.

14. Conclusion

MATRIX IN ONE SENTENCE

Expertise = automation of the most generative principles – and the ability to know when they do not apply. MATRIX is a method for finding these principles and installing them faster, so that capacity is freed for navigation.

The most ambitious hypothesis is not only that MATRIX can make learning faster, but that a good model library makes the learner better equipped to solve unfamiliar problems by recognising known structure beneath a new surface. This is the model’s most important promise – and its most important empirical test. The test begins within the domain, where transfer is most likely and easiest to measure.

The book series shows that the model can be made concrete for children and young people. It also shows the model’s ethics: the power to read and trick others only means something if it makes the whole team better.

SEE THE SIGNAL · UNDERSTAND THE PRINCIPLE · GENERATE THE SOLUTION

Appendix A: Definitions

Term	Working definition
Principle	A core insight into how the domain works, stated in one sentence.
Generative model	A compact model of cause and effect that produces competent responses in many situations with the same structure.
Signal	Observable information in the situation. The signal activates the model through the player's pattern recognition: the player recognises a pattern they have seen before.
Trigger word	A very short word the coach uses as an attention anchor for the player. It is the coach's tool, not part of the skill: once the player masters the model, the word is their own, and the coach need not say anything.
Response pattern	The automated link from signal to model and a robust default response. The details are generated by the situation.
Freed capacity	Working memory that is freed when low-level tasks run automatically, and that can be used for navigation.
Structure	The relationship or pattern beneath surface differences.
Mechanism	Why the structure affects the space of possibilities or the outcome.
Space of possibilities	The actions that become available, or more or less appropriate, in the situation.
Transfer	Learning that affects performance in new or changed problems.
Compression	Several low-level rules being replaced by one deeper model.
Fractal reach	Whether a related structure recurs at several scales or in several contexts. Used analogically, not as a mathematical definition.
Model composition	Combining several models when no single model is sufficient.
Installed compression	Automation: a generative model practised until it runs without working memory.
Navigation	What freed and expanded capacity is used for: awareness, choices, creativity and combining models in situations one has not seen before.

Term	Working definition
Routine expertise	Fast and accurate performance in familiar situations.
Adaptive expertise	Expertise that also copes with new situations, because the performer knows when the automation applies and when it does not.
Long-term working memory	Domain-specific structures that let the expert keep far more relevant information active than working memory alone can.
External focus	Attention directed at the environment and the effect of the action, not at one's own body movements.
Self-training	The player's own training alongside club and academy. In MATRIX: 10–15 minutes of principle, model and signal, taken into training and matches.

Appendix B: Learning architecture and AI as scaffolding

MATRIX is first and foremost a model of how people learn. The book, the coach and the app are delivery channels, and none of them covers all five steps on its own.

Channel	Strongest at step	Role
The book	1–3	The player's self-training: principle, model and signal in pictures, missions and trigger words
The parent at home	3	The trigger word and one question, never the explanation
The coach, club training and the match	3–5	Game-like practice under pressure, observation and feedback. This is where the models become automatic.
App and AI	2 and feedback	Variations, transfer tests, reflection, learner profile and session plans

The scaffolding is faded. First the principle and the model are shown. Then come questions, and the coach uses trigger words as attention anchors. Finally the word is the player's own, and the coach need not say anything. The direction is fixed: support down, model capacity up, independence up.

Loop	Sequence	Who does the thinking
Answer machine	Problem → ask AI → get answer	The AI
MATRIX	Problem → see the signal → activate the model → generate possibilities → choose → act → learn	The learner, with support as scaffolding

NORTH STAR

MATRIX should use AI to make expertise more accessible, without making people more dependent on AI. *AI should not replace the learner's thinking. It should help build it.*

The MATRIX Expertise Agent is a proposed AI agent that runs Discovery on new domains. Standard command: "Run MATRIX on [domain]. My goal is [goal/level]." The agent delivers a domain map, expert signals, 5–10 generative principles, a signal-to-model map, model interactions, compression, learning order, practice design, a feedback loop and a MATRIX card.

To avoid neat but false “laws”, the agent labels every principle:

Status	Meaning
VALIDATED	Well supported by established domain knowledge or research.
SYNTHESIS	A MATRIX compression of known knowledge. The wording is the agent's.
HYPOTHESIS	A promising model that must be tested before being treated as robust.

AI proposes and compresses, but it is not the authority. The coach or domain expert decides the status.

END POINT

The end point of MATRIX is not a smarter AI user. It is a more competent human being.

Appendix C: Predictive processing and generative AI

Two fields describe learning in a way close to MATRIX: predictive brain theory and generative AI. Both rest on the idea that a system learns by predicting, and improves when the prediction misses. This appendix uses them as explanatory language and a source of hypotheses, not as proof of MATRIX.

C.1 Predictive brain theory

Predictive processing is an influential framework in cognitive neuroscience. In brief:

- **The brain predicts.** It uses an internal, generative model of the world to predict what the senses will encounter (Rao and Ballard; Clark).
- **The error is the learning signal.** What is passed on is largely *prediction error*: the difference between expected and actual sensory input. The model is updated to reduce such errors over time (Friston).
- **Attention is weighting.** The brain weights prediction errors by how reliable they are (precision). Attention can be understood as turning up the weight on particular signals.
- **Action is prediction too.** In *active inference*, the organism acts to make the world match what it expects, rather than only updating the model.
- **The model is hierarchical.** Higher levels predict slow, abstract patterns; lower levels predict fast, concrete details.

The framework is widespread but debated. It is a proposal for a unifying theory, not a settled fact.

C.2 Generative AI architecture

Modern generative AI learns in a related way:

- **Language models are trained to predict the next word.** The prediction error is used to adjust the model. To do this well, the model must build compressed internal representations of the patterns in the data.
- **The attention mechanism** weights which parts of the context are relevant for the next prediction (Vaswani and colleagues).
- **World models** in reinforcement learning learn a compact model of the environment and use it to predict and plan actions (Ha and Schmidhuber).

C.3 Three systems, the same logic

Element	The brain (predictive processing)	Generative AI	MATRIX training
Model	Hierarchical generative model	Neural network with learned representations	Principle and generative model (steps 1–2)
Learning signal	Prediction error	Error in predicting the next word or state	Feedback: “What did I see? Did it work?”
Attention	Precision weighting	Attention	The trigger word as attention anchor (step 3)
Recognition	Sensory input matching a prediction	Pattern resembling training data	Signal recognised as a pattern (step 3)
Compression	Priors that summarise experience	Compressed representations	A few laws instead of many rules
Action	Active inference	Agents that plan and act	Automatic response pattern (step 4)
Data	Few, personal, embodied experiences	Huge amounts of passive data	Game-like, varied repetitions

C.4 What this may explain in MATRIX

MATRIX	Possible explanation in predictive processing	Status
Principle and model before signal	A good model gives precise predictions, so the signal can be recognised quickly.	SYNTHESIS
The signal is activated by pattern recognition	Recognition is sensory input fitting a learned prediction.	SYNTHESIS
The trigger word as anchor	The coach increases the weight on the right signal until the player does it themselves. Once the model is learned, the anchor is not needed.	HYPOTHESIS
Automatic response pattern	When predictions are accurate, there are fewer surprises that require conscious processing.	SYNTHESIS
Freed capacity	Fewer low-level prediction errors may make room for higher levels: awareness and choice.	HYPOTHESIS
Level 10: "The brain guesses more than it sees"	A feint works because it exploits the opponent's predictions.	SYNTHESIS
Fractal levels	Hierarchical models on different time scales resemble the ladder from game to team to one's own system.	HYPOTHESIS
The tempo rule: think again when the signal is unclear	Small prediction errors are handled automatically. Large or unexpected errors require conscious processing, and working memory is brought back in.	SYNTHESIS

C.5 Important differences

- **Data:** A child learns a model from a few experiences with their own body. A language model needs vast amounts of text. MATRIX assumes that *well-chosen* examples and contrasts can replace quantity.
- **Body and action:** The brain learns in a loop of action and perception. Most AI models learn passively.
- **Analogy is not mechanism:** That two systems both "predict" does not mean they do so in the same way. By MATRIX's own rule, cross-domain similarity is a hypothesis that requires a mechanism test (chapter 7).

C.6 Testable hypotheses

1. Players who learn model and signal should anticipate faster in new variations of the same situation than players who learn fixed rules.

2. The trigger word should have the greatest effect early, and be removable without loss once the player has reached AUTOMATED.
3. Players who know the model behind a feint (Matrix Defence: “Don’t guess”) should be fooled by it less often.
4. Players at the AUTOMATED level should perform better on dual tasks than players at ACTIVATED.

LIMITATION

Predictive processing and generative AI give MATRIX a precise language for *why* the method may work. They do not prove that it works. The links above are labelled SYNTHESIS or HYPOTHESIS and must be tested.

References

- Barnett, S. M., & Ceci, S. J. (2002). When and where do we apply what we learn? A taxonomy for far transfer. *Psychological Bulletin*, 128(4), 612–637. <https://doi.org/10.1037/0033-2909.128.4.612>
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55–81.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204. <https://doi.org/10.1017/S0140525X12000477>
- Dicks, M., Davids, K., & Button, C. (2009). Representative task designs for the study of perception and action in sport. *International Journal of Sport Psychology*, 40(4), 506–524.
- Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, 47, 273–305. <https://doi.org/10.1146/annurev.psych.47.1.273>
- Ericsson, K. A. (2008). Deliberate practice and acquisition of expert performance: A general overview. *Academic Emergency Medicine*, 15(11), 988–994. <https://doi.org/10.1111/j.1553-2712.2008.00227.x>
- Ericsson, K. A. (2010). Expertise. *WIREs Cognitive Science*, 1. <https://doi.org/10.1002/wcs.47>
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211–245.
- Fitts, P. M., & Posner, M. I. (1967). *Human performance*. Brooks/Cole.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15(1), 1–38. [https://doi.org/10.1016/0010-0285\(83\)90002-6](https://doi.org/10.1016/0010-0285(83)90002-6)
- Hatano, G., & Inagaki, K. (1986). Two courses of expertise. In H. Stevenson, H. Azuma & K. Hakuta (Eds.), *Child development and education in Japan* (pp. 262–272). Freeman.
- Ha, D., & Schmidhuber, J. (2018). World models. *arXiv:1803.10122*.
- Hambrick, D. Z., et al. (2014). Deliberate practice: Is that all it takes to become an expert? *Intelligence*, 45, 34–45. <https://doi.org/10.1016/j.intell.2013.04.001>
- Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work. *Educational Psychologist*, 41(2), 75–86.
- Klein, G. (1998). *Sources of power: How people make decisions*. MIT Press.
- Liao, C.-M., & Masters, R. S. W. (2001). Analogy learning: A means to implicit motor learning. *Journal of Sports Sciences*, 19(5), 307–319.
- Macnamara, B. N., & Maitra, M. (2020). Is the deliberate practice view defensible? A review of evidence and discussion of issues. *Frontiers in Psychology*, 11, 1134. <https://doi.org/10.3389/fpsyg.2020.01134>
- Mann, D. T. Y., Williams, A. M., Ward, P., & Janelle, C. M. (2007). Perceptual-cognitive expertise in sport: A meta-analysis. *Journal of Sport & Exercise Psychology*, 29(4), 457–478. <https://doi.org/10.1123/jsep.29.4.457>
- Masters, R. S. W. (1992). Knowledge, knerves and know-how: The role of explicit versus implicit knowledge in the breakdown of a complex motor skill under pressure. *British Journal of Psychology*, 83(3), 343–358.
- Pinder, R. A., Davids, K., Renshaw, I., & Araújo, D. (2011). Representative learning design and functionality of research and practice in sport. *Journal of Sport & Exercise Psychology*, 33(1), 146–155. <https://doi.org/10.1123/jsep.33.1.146>
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>

- Sala, G., & Gobet, F. (2017). Does far transfer exist? Negative evidence from chess, music, and working memory training. *Current Directions in Psychological Science*, 26(6), 515–520.
- Shea, J. B., & Morgan, R. L. (1979). Contextual interference effects on the acquisition, retention, and transfer of a motor skill. *Journal of Experimental Psychology: Human Learning and Memory*, 5(2), 179–187.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30. arXiv:1706.03762.
- Wulf, G. (2013). Attentional focus and motor learning: A review of 15 years. *International Review of Sport and Exercise Psychology*, 6(1), 77–104.
- Primary sources for chapter 3:** Matrix Playbook – Level 5, Matrix Playbook – Level 10 (2nd edition), Matrix Teamplay – Level 15 and Matrix Identity – Level 20, all by Lennert Kvamme (2026), published in Norwegian.

© 2026 Lennert Kvamme · MATRIX Expertise Model, whitepaper v1.5 · Concept paper.