

WELCOME

A little map. *A fresh hand.*

The Lantern Poker Quick-start.


Lars Rocha

Twelve pages.

Eight discoveries.

A welcoming first seat.

Patterns, price,
and a little curiosity.



WELCOME

A legend for the whole book

The denominator is part of the idea.

Mark or term	Meaning
0 away	A rank in the reference set
1 away	A circular neighbor of a reference
2 away / 3+ away	Farther rank buckets; 3+ combines all larger distances
LIKELY	At least seven of thirteen ranks are near
UNLIKELY	At most six ranks are near
MC	Monte Carlo estimate
WSRF	Prediction from an exported pytoolzx WSRF forest
BB	One big blind: twenty play chips in the room
Range width	A synthetic weighting parameter, not literal range coverage
Additional chips	Return measured from the current decision; earlier contributions are spent

Ranks are numbered 2=0 through A=12. The four reference expressions are {a, b, (a+b) mod 13, |a-b|}. Distances are circular, but legal straights follow poker rules. Suits do not affect the rank map; they do affect hand strength.

All examples are play-chip studies. A practice peek changes what the human sees, while the adviser continues to use its own pocket and public information. The books explain the implementation rather than assuming every site, HUD or poker program uses the same conventions.

Our way of looking

The joy is a complicated situation becoming readable. A compact map helps you choose where to look next. The recorded experiments keep that curiosity connected to what was actually counted, sampled and learned.

FOUNDATIONS

A way to see the hand

Start with a question you can answer.

A poker hand offers more information than a player can comfortably hold at once: two private cards, a changing board, several stacks, positions, prices, and the actions that brought everyone here. The distance map gives some of that information a compact shape. Its value as a study tool begins when the shape helps you ask a precise question.

Use three passes. First, describe your actual hand and the public board. Second, describe the map: the pocket zone, the dynamic zone, and which ranks are near. Third, price the action. These passes work together. A large near band can coexist with a weak hand, while a small near band can accompany a powerful made hand.

The Lantern room puts these passes beside one another. A WSRF forest estimates showdown share; another compares betting continuations. The player can inspect a tree, request a hint, change a range assumption, and test the spot. Curiosity becomes useful when you can name exactly what changed between two experiments.

Begin gently. Pause before acting, write one sentence about the decision, and compare it with the explanation afterward. The goal of the first session is a repeatable way to think. Winning an individual pot is a pleasant outcome, but explaining the decision is the skill you can carry into the next hand.

Hand → map → people → price

Try it first

You spot a LIKELY label and a high call price. Can exploring your made hand, the people, and the eligible pot help you choose your next move?

Worked answer. Yes. Inspect the actual made hand, opposing range assumptions, active players, and eligible pot. The label alone does not answer whether the call is worth its cost.

FOUNDATIONS

Distance around the circle

The shortest route can cross the ace–two seam.

Place the thirteen ranks around a circle in their usual order. The distance between rank numbers a and b is the shorter of the clockwise and counterclockwise routes. In symbols it is $\min(|a-b|, 13-|a-b|)$. The largest possible pairwise distance on this odd circle is six.

This convention makes ace and two neighbors in the map. It does not mean that every poker straight wraps around the circle. A-K-Q-J-T is a straight, and A-2-3-4-5 is a straight; Q-K-A-2-3 is not. The circle is the geometry of the feature, while legal poker hands are evaluated by separate rules.

Distance to a reference set means the smallest distance to any reference in that set. If one reference is far away but another is adjacent, the rank is one away. Repeated references do not add extra weight. Four reference expressions can collapse into fewer distinct points.

On each atlas page, the reference points have one color, their immediate neighbors another, and the remaining ranks a third. Trace a route yourself before reading the numerical row. This makes the zero/one/two/three-plus buckets feel like a physical map instead of a memorized label.

$$d(a,b) = \min(|a-b|, 13-|a-b|)$$

Try it first

Can you travel from A to 3 in two circular steps while keeping poker straight rules as a separate guide?

Worked answer. Yes. The map distance is two: $A \rightarrow 2 \rightarrow 3$. Poker straights follow their own five-card sequences; A-2-3-4-5 is the ace-low straight. Keeping both rules clear lets you explore each kind of pattern.

FOUNDATIONS

The four reference expressions

A compact construction, with duplicates removed.

Given the two rank numbers a and b , form the reference set $\{a, b, (a+b) \bmod 13, |a-b|\}$. The first two terms anchor the map to the cards. The sum wraps around the thirteen-rank circle. The difference is the nonnegative separation before taking circular distances to the resulting set.

For $7\spadesuit 5\diamondsuit$, the rank numbers are 5 and 3. The reference numbers are 5, 3, 8, and 2, corresponding to 7, 5, T, and 4. These four locations generate all thirteen distance values. Their immediate neighborhoods overlap, so counting three ranks per reference would overcount the near band.

Pocket pairs provide especially useful examples. If $a=b$, the difference is zero and the two direct references coincide. Depending on the doubled rank modulo thirteen, the entire reference set can contain only one, two, or three distinct points. A pair can therefore have a narrow map even though its poker strength is substantial.

The reference set is symmetric in its two inputs: exchanging the reference cards does not change it. Suits also do not change the map. Those invariances help you check an implementation. They do not imply that swapping a hole card with a public board card preserves the meaning of the hand.

$$R(a,b) = \{a, b, (a+b) \bmod 13, |a-b|\}$$

Try it first

For $2\clubsuit 2\diamondsuit$, do all four reference expressions meet at one point: rank 2?

Worked answer. Yes. Both rank numbers are zero. Every expression gives zero, so there is one distinct reference point: the rank 2. Its near ranks are A, 2, and 3.

FOUNDATIONS

What LIKELY names

Coverage of the rank circle.

A rank is near when its shortest distance to the reference set is zero or one. The whole reference pair is labeled LIKELY when more than half of the thirteen ranks are near: seven or more. Six or fewer gives UNLIKELY. This book follows the executable LRO convention documented in the source receipt.

The label refers to the size of a rank band. It is not a statement that the player is favored to win, that a particular card is due, or that a bet is profitable. Those questions have different denominators and need different information. Treating all of them as one question makes a compact map less useful.

The word UNLIKELY deserves just as much attention as LIKELY. It identifies a narrower near band, so more rank locations lie farther from the references. An unusual-looking map can help you notice a change that a familiar map hides. In the training data, both labels receive deliberate coverage so the model has examples from each.

A useful habit is to say the full sentence once: “This pair is LIKELY because nine of thirteen ranks are near.” That sentence preserves the definition. Later, you can shorten it to the label while still remembering what was counted.

$$\text{LIKELY} \Leftrightarrow \text{near-rank count} \geq 7$$

Try it first

Your map has six near ranks and your hand is a full house. Can you celebrate that made strength and explore the board, opponents, and price together?

Worked answer. Yes. The map is UNLIKELY by coverage, and the full house brings its own made strength. Explore that strength together with the board, opponents, price, and available actions.

FOUNDATIONS

Remove the cards you know

A rank count is not an unseen-card count.

A full deck contains four cards of each rank. Once cards are known, those multiplicities change. A nine-rank near band starts with thirty-six physical cards, but some may be in your pocket or on the board. Counting unseen cards requires removing every card known to the observer exactly once.

For the pocket $7\spadesuit 5\clubsuit$, the near band contains nine ranks. Both pocket cards lie in that band, leaving thirty-four near cards among fifty unseen cards before the flop. Under a uniform unconditioned deck, the next physical card belongs to the near band with probability $34/50$, or 68%. Nine thirteenths and thirty-four fiftieths answer related but different questions.

On later streets, update the count from the observer's pocket and public board. Do not remove an opponent's actual private cards merely because a practice peek is visible. The adviser's information boundary stays the same in peek mode. A separate omniscient exercise would need to say explicitly that it assumes those cards are known.

Card removal also changes blockers. Holding a particular ace removes one ace from possible opposing pockets. Holding a heart removes a heart from possible flush combinations. The same act of removal can alter both a distance bucket and a poker range, through different relationships.

Uniform next-card share = cards in bucket ÷ unseen cards

Try it first

With 15 bucket cards among 47 unseen cards, is the uniform next-card bucket share $15/47$, about 31.91%?

Worked answer. Yes. $15/47$, about 31.91%. This is a bucket event. It becomes a winning-out calculation only if every counted card has been checked against the relevant opposing hands.

FOUNDATIONS

Where the board sits

Alignment is a third question.

Pocket coverage asks how much of the rank circle lies near the pocket references. Dynamic coverage asks the same question around the newest reference pair. Board alignment asks whether a majority of the revealed board ranks lie near the pocket references. These can disagree without any contradiction.

Count board cards, including repeated ranks, when measuring alignment. On a paired board, two cards of the paired rank are two observations of that rank. If both are near, both contribute to the majority. This makes alignment a description of the visible board, rather than a second calculation of the number of distinct near ranks.

A board-majority label uses a strict majority in this implementation. A two-near, two-far turn board does not have a near majority. Before the flop the board is undealt, so asking for its alignment has its observations still ahead. An empty board should not be read as evidence of disagreement.

Try saying three independent sentences: “My pocket covers nine ranks. The dynamic pair covers five. Two of the three flop cards are near my pocket.” Each sentence describes a different count. Keeping them separate makes it much easier to understand a tree split or to reproduce an example by hand.

$$\text{Alignment} = \text{near board cards} \div \text{revealed board cards}$$

Try it first

A turn board has two near and two far cards. Is its near alignment exactly half, with a strict majority beginning above half?

Worked answer. Yes. Two of four is exactly half, and the rule requires more than half. The pocket and dynamic coverage labels must be calculated separately.

FOUNDATIONS

The price of a call

Use the pot you can actually win.

At a decision with betting held fixed, a call costing C into an eligible current pot P breaks even at a showdown share of $C/(P+C)$. The current pot includes chips already committed by the other players. Your call adds C to the final amount being contested.

For a call of thirty into ninety, the threshold is twenty-five percent. At a true showdown share of thirty percent, the simplified expected result from the decision is $0.30 \times 90 - 0.70 \times 30$, or plus six chips. That calculation treats chips you previously put in as already spent.

The word eligible matters. A short stack cannot win every side pot on the table. Using the total displayed pot in place of the portion you can contest can make a call look cheaper than it is. Multiway ties also belong in the share estimate: a three-way split contributes one third, not one half or one.

Earlier streets add future betting and strategic choices. A hand may not realize its raw showdown equity if pressure forces it to fold later. Conversely, later folds by opponents can add value. The call threshold is a clean reference point, and the learned continuation model attempts to account for more of the decision tree under its simulated opponents.

$$\text{Break-even call share} = C / (P + C)$$

Try it first

With a river call of 40 into an eligible pot of 160 and all betting ending after the call, does a 20% share break even in a rake-free game?

Worked answer. Yes. $40/(160+40)=20\%$, assuming betting complete and rake set to zero. If the pot is not fully eligible to you, first replace 160 with the relevant contestable amount.

FOUNDATIONS

The price of a bluff

The fold assumption does the work.

For a pure bluff that loses whenever called, expected return is $F \times P - (1-F) \times B$, where F is the probability everyone relevant folds, P is the pot you win immediately, and B is the amount you risk. The break-even fold rate is $B/(P+B)$.

A half-pot bluff needs one third folds. A pot-sized bluff needs one half. A two-and-a-half-pot bluff needs 2.5/3.5, or about 71.43%. Increasing the size can change how opponents respond, but the response must improve enough to pay for the larger risk.

This is why an assumed fold percentage deserves scrutiny. Betting twenty-five into ten with a seventy-five-percent fold rate gives plus 1.25 chips in the simplified model. At sixty-five percent folds, the same bet gives minus 2.25. The cards and size did not change; the behavioral assumption did.

A semi-bluff retains some chance to win when called and therefore needs a more complete calculation. A multiway bluff must consider the whole field, not one opponent's fold statistic. Use the simple equation as a transparent baseline, then identify which additional branches matter in the actual spot.

$$\text{Pure-bluff threshold} = B / (P + B)$$

Try it first

When a pure bluff risks 60 to win 40, does a 60% fold rate reach the simplified break-even point?

Worked answer. Yes. $60/(40+60)=60\%$. That is the simplified break-even point. At exactly that rate the expected result is zero before any additional costs.

TAKE YOUR SEAT

A hand worth exploring

Your pace. Your possibilities.

Choose the Lantern Lounge for continuous heads-up, six-seat, or nine-seat play with familiar AI company. Start with 2,000 play chips, choose your pace, and let your next hand arrive automatically. Your exact hand and stack save on this device.

Can a little curiosity guide the next hand?

Yes. Begin with your two cards and the call price. Open Coach to see the estimated share and action comparisons. Explore the pocket and dynamic maps, then return to the table when you feel ready.

Can you explore at your own pace?

Yes. The table holds your place while you visit a panel. A practice peek adds a learning opportunity for you while the advisers keep using their own cards and public information. Soft sounds, HUDs, quick play, and a WSRF partner are options you can choose.

Can your table story continue another day?

Yes. Your Lounge match saves between visits. Return to the same table size and resume. Save a notebook in Options to carry that story to another device. The original study room also keeps its own matches.

Can you take the discovery further?

Yes. The 180-page Field Guide expands this first course with worked decisions, zone journeys, and experiments. The 300-page WSRF Zone Bliss adds every rank pair, sixty affirmative challenges, and the complete training story.

A small invitation

Pick one hand you enjoyed. Name a pattern you noticed, a price you understood, and one possibility you would like to explore next.

TRAINING AND REFERENCE

Sources and further reading

Trace the ideas, the implementation, and the measurements.

Source	Role
Lars Rocha's local poker guides	Distance ideas and the project's starting conventions
Executable LRO compute_de_zone	Rank-zone threshold and 169-pair convention check
pytoolzx / WSRF source	Actual regression fitting and deterministic seed generation
Lantern table and simulator	Rules, sampling, public views and action continuations
Training JSON receipts	Parameters, hashes, sample sizes and measured results
Zone-atlas JSON	Every numerical study fixture in the long book

For general forest background: Leo Breiman, *Random Forests*, Machine Learning 45 (2001), 5–32. doi:10.1023/A:1010933404324. The WSRF implementation and this room's routing choices should be read from their own source; the citation does not imply that every design here is the original paper's algorithm.

For interval interpretation: NIST/SEMATECH e-Handbook, Confidence Limits for the Mean. The worked experiments in this book use their own simulated data and label approximate intervals explicitly.

For a different approach to multiplayer poker: Noam Brown and Tuomas Sandholm, *Superhuman AI for multiplayer poker*, Science (2019). doi:10.1126/science.aay2400. This room is not presented as a reproduction of Pluribus or as a professional-player benchmark.

The shortest path to a number

Use the page's fixture identifier to find its exact pocket, board, seed namespace and estimates in zone-atlas.json. Use the model hash to connect a displayed prediction with the evaluated export.