



2026

Insights from the 2026 VC-BootCamp

INTELLIGENCE WITHOUT GUARDRAILS

Ecosystems, Bias, and the Future
of Capital Allocation

Stanford University & UC Berkeley

March 10–12, 2026

Hoss Zaouali, Founder — VC-Bootcamp / Me2We

Insights from the 2026 VC-BootCamp

INTELLIGENCE WITHOUT GUARDRAILS

Ecosystems, Bias, and the Future of Capital Allocation



“The venture market has outgrown human cognitive capacity. AI is now revolutionizing VC the same way algorithms revolutionized finance in the 1980s.” ”

Hoss Zaouali

founder of VC Bootcamp - Silicon Valley - Toronto - Paris - Johannesburg

EXECUTIVE SUMMARY

Artificial Intelligence (AI) is no longer an emerging consideration for Venture Capital (VC) — it is an operating reality reshaping every stage of the investment lifecycle. At the 2026 VC-BootCamp held at Stanford University and UC Berkeley, five practitioner-led panels — including a dedicated dinner session on AI and human bias and a keynote fireside chat with Andreessen Horowitz — surfaced some of the most candid and practically grounded conversations on AI in capital allocation to emerge from any convening this year.

This white paper distills the key findings, tensions, and practitioner frameworks from those sessions. It is intended for venture capital partners, limited partners, institutional investors, banks, and founders navigating a capital environment being fundamentally restructured by AI. Five central insights define what the 2026 cohort concluded:

- 01 The probability of a startup being funded and its probability of actual success are increasingly divergent** — and AI is beginning to expose that gap with measurable precision. Firms optimizing for fundability patterns may be systematically investing in the wrong companies.
- 02 AI can materially reduce bias in deal sourcing — but only through intentional system design.** One fund achieved 10x the diversity of the typical US VC fund alongside a top-0.1% vintage performance using unsupervised machine learning applied to sourcing.
- 03 AI has reached co-pilot maturity in VC workflows — screening, CRM, due diligence documentation, scenario modeling — but **autopilot-level decision-making remains premature due to critical limitations in context, judgment, and accountability.****
- 04 The democratization of fund infrastructure through AI-powered tooling is lowering barriers to fund formation and changing who can become a VC.** But governance architecture has not kept pace with adoption, leaving a generation of new managers exposed.
- 05 AI is trained on the rich world's data.** Less than 10% of the global population generates the material that trains most models, **creating structural blind spots for emerging market founders** whose traction signals appear in WhatsApp groups, foot traffic, and local-language media — not in LinkedIn or GitHub.

Context & Background

VC-Bootcamp is a practitioner-grade convening that brings together venture capitalists, founders, operators, and ecosystem builders to examine — without hype — **how AI is reshaping the global capital ecosystem**. The 2026 edition was held across two days at Stanford Graduate School of Business and UC Berkeley's Haas School of Business, structured around five sessions ranging from panel discussions to a keynote fireside chat with one of the world's most prominent venture firms.

Organizer	Me2We / VC Bootcamp
Dates	March 10–12, 2026
Locations	Stanford Graduate School of Business & UC Berkeley (Haas School of Business)
Theme	The AI Revolution in Venture Capital
<hr/>	
Panel 1	Building Bridges Between Investment Ecosystems
Panel 2	From Gut Feel to Infrastructure: How AI Is Rewriting Venture Decision-Making
Panel 3	AI in Investment: Autopilot, Co-Pilot, or Hype?
Panel 4	Will AI Amplify or Mitigate Human Bias?
Panel 5	Fireside Chat: Andreessen Horowitz / Speedrun
Participants:	VC GPs and partners (USA, Canada, Middle East, Asia, Europe), Stanford LEAD founders, operators, and ecosystem builders from 15+ countries
Methodology	All findings and quotes are drawn directly from the five panel transcripts. Speakers are attributed by name and affiliation. Where findings are hypothesis-level rather than empirically confirmed, they are labeled as such.

The problem this conversation addresses is structural. Quarterly review cycles cannot keep pace with companies moving from idea to \$100M valuation in 18 months. No human team can track 50,000+ startups raising globally each year. And the data required to make good decisions — patents, GitHub activity, talent movement, social signals — grows faster than any analyst can process. At the same time, AI systems trained on historical deal data encode historical biases. The firms deploying AI fastest are not necessarily building the governance architecture to use it responsibly.

INTELLIGENCE WITHOUT GUARDRAILS



Panel 1: Building Bridges Between Investment Ecosystems

ECOSYSTEM FIT IS THE NEW PRODUCT-MARKET FIT

The panel opened by reframing how practitioners think about investment ecosystems. Rather than fixed geographical constructs, ecosystems are now dynamic, consciously crafted networks — and the question of fit applies equally to founders seeking capital and investors building portfolios. The core insight: finding the right ecosystem is not about changing your company, it is about finding the community of investors whose mental models are already primed for the problem you are solving.

Kim Kolt of Bay Bridge Ventures drew on her own firm's experience pivoting from targeting institutional LP allocators to family offices and high-net-worth investors after recognizing a fundamental mismatch between the firm's profile and the expectations of large institutions. Despite having a founding team with significant institutional pedigree — including a former head of direct investments at CalPERS overseeing a \$16 billion portfolio — the institutional doors did not open for a first-time fund. The lesson was the same one that founders learn the hard way: test your proposition repeatedly and cheaply before committing resources.



“Look for the people who disagree with you — because those are the most valuable. Failure forced our portfolio companies to get extremely efficient, turn to AI, and outsource different skill sets — positioning them to succeed in ways they wouldn't have otherwise.”

— [Kim Kolt](#), Bay Bridge Ventures

Gregory LaBlanc offered a parallel example from his own experience helping a company raise: an EdTech company that could not attract a dollar was repositioned as HR Tech — with no change to the underlying business whatsoever — and capital began flowing immediately. The company had not changed. The ecosystem vocabulary had.

“This notion of lean startup, prototyping, pivoting, frequent and cheap and early failure — it pervades not just building a product, but the entire process of doing a startup or raising a fund. You've got to be prepared to fail a lot when you're going out there trying to find funding.”

— [Gregory LaBlanc](#), UC Berkeley



Michael-Anthony Clement extended this to the operator's perspective: ecosystem fit is ultimately about locating concentrated expertise, not just geography. His Quebec-based startup was built on a network of talent from a single prior company — an example of how hyper-local knowledge clusters can generate disproportionate advantage even in markets that are not traditional venture hubs.



“Capital is very liquid. It’s moving all over the world. Technology is extremely decentralized, and AI further enhances that decentralization. Go where the signal is. Don’t hesitate, wherever that may be.”

— [Michael-Anthony Clement](#), Consignments.ai

THE PRACTICAL ECOSYSTEM TEST

The panel converged on a practical three-stage approach to ecosystem discovery. The first stage is a message test: present your company or thesis in its current framing to ten investors in your target ecosystem, and track the quality and pattern of objections — not just the yes/no rate. The second stage is a pivot test: reframe the same company using different vocabulary and run the same conversations. The third stage is ecosystem selection: identify the frame and community that produces the strongest resonance, and pursue that ecosystem deliberately. **Testing ecosystem fit costs almost nothing. Pursuing the wrong one for eighteen months costs everything.**

The Role of AI: The panel identified attention and judgment as the new scarce resources in venture. As AI floods the top of the deal funnel and the number of startups explodes, the ability to filter signal from noise — and to know which ecosystem to be in — becomes more valuable than the ability to process volume. One fund now reviews 4,600+ companies per year through AI-assisted sourcing — a volume that would be impossible to screen manually.

Panel 2: From Gut Feel to Infrastructure: How AI Is Rewriting Venture Decision-Making

THE FUNDABILITY-SUCCESS INVERSION

The most empirically striking finding of the bootcamp came from Jeff Trueman of Luminary Capital, who presented results from an ongoing study conducted with a forecasting researcher at the Kelley School of Business. The study asked a straightforward question: can LLMs predict which startups will receive funding better than humans can? The answer was yes — by a significant margin.

Across 30 Kickstarter-style companies and 750 paired comparisons, four LLMs (Claude, GPT-4, Gemini, and DeepSeek) predicted funding outcomes with 75% accuracy on average, compared to approximately 20–40% for human evaluators including Harvard MBAs working as investors. But the deeper finding was more unsettling: when the same AI models were asked to assess the probability of actual business success, the ranking was nearly the inverse of fundability. Companies most likely to attract investment were often least likely to create lasting value — and vice versa.

The implications for AI-assisted deal sourcing are significant: if models are trained to predict what gets funded, they may be systematically optimizing for narrative appeal and pattern recognition from past winners rather than underlying viability. A pesticide redistribution company and a Mexican medtech firm ranked low on fundability in the study, but high on genuine potential.



“ Claude finally came back and said: Jeff, you’re a creator. Every version of this what-if ends with managing a low-margin, regulation-heavy business on the side of two ventures you’re actually excited about. What would you do with that freed-up mental energy if that company was just gone?”

— [Jeff Trueman](#), Luminary Capital

AGENTIC DUE DILIGENCE AND THE COUNTERFACTUAL METHOD

Jorden Woods of StarChain Ventures described a practice growing in lean funds: using agentic frameworks to perform counterfactual due diligence — explicitly prompting AI to identify everything that could go wrong, model pessimistic scenarios, and challenge the investment thesis before it reaches an investment committee. This is the AI-as-devil’s-advocate use case, and it represents a qualitatively different application than single-model querying: it creates structured adversarial analysis at a scale no human team can replicate.

“The entire VC industry is being transformed so that it’s much leaner, much faster, and much more democratized. It’s not just the same people year after year building funds. If you’ve been a founder before, you’ve raised money the same way — you go to friends and family. Same process. The sky’s the limit now.”

— [Jordan Woods](#), StarChain Ventures



Keshia Theobald-van Gent of BDev Ventures raised the bias concern directly: pattern-matching AI trained on historical funding data may systematically undervalue non-traditional founders, emerging markets, and breakthrough ideas that do not yet have precedent in the training data. Her firm’s internal guidance has become one of the most quoted lines of the bootcamp:



“Everyone in this industry is operating from some type of a black box — which is a great place to hide mediocre talent. Use AI. Don’t worship at the altar.”

— [Keshia Theobald-van Gent](#), BDev Ventures

FUND INFRASTRUCTURE: THE 24-HOUR FUND

The panel also explored the democratization of fund formation itself. **AI-powered fund-in-a-box platforms now enable experienced operators and founders to launch a compliant VC fund in as little as 24 hours** — collapsing a process that previously required 6–9 months and approximately \$150,000 in legal fees. The panel consensus: this democratization of fund formation is real and genuinely significant. But governance frameworks — around AI use in diligence, LP reporting, and investment decision documentation — are not being built at the same speed as the funds themselves. **A new generation of fund managers is entering the market with AI-powered infrastructure but without the architecture that institutional investors will eventually require.**

The Role of AI: The probability of being funded and the probability of actual success may be inversely correlated. VC firms that calibrate their AI sourcing tools on historical funding patterns may be building machines that consistently find the wrong companies — at unprecedented scale and speed.

Panel 3: AI in Investment: Autopilot, Co-Pilot, or Hype?

THE GOOD, THE BAD, AND THE DANGEROUS

The third panel delivered the most technically grounded perspective of the bootcamp, anchored by Lutz Finger from XGEN AI / Cornell University and Allison Vuong of Pivotal Life Sciences. Their core argument: **the VC industry's excitement about generative AI and agents must be tempered by a clear-eyed understanding of where classical machine learning already outperforms LLMs, and where agentic systems fail dangerously.**

The Good: AI excels at unstructured-to-structured data extraction — reading pitch decks, pulling CRM fields from emails, parsing board materials. It excels at summarization of large document sets, multi-agent debate frameworks, and detecting hidden sentiment signals in company communications. Tom of BK Fund described deploying ten agents to debate an investment thesis and surfacing points of consensus and disagreement — a technique for building structured adversarial analysis at a scale no human team could replicate.



“ I have 10 agents and I ask them to debate with each other. After the debate, they give me information on what they agreed on — that’s consensus. And the things they don’t agree on — those are the ones I pay most attention to.”

— Tom Huang, BK Fund

The Bad: LLMs are trained on averages. **Venture capital is the business of outliers.** As Bruce noted, no LLM trained on historical data would have funded Snapchat, FedEx, or the iPod — ideas that violated the prevailing logic of their moment. AI also has a structural positive bias by default, which requires deliberate prompting to counteract.

“ LLMs are based on averages. If I ask you to complete the sentence ‘Life is like a box of —’, you all say chocolate. What we as VCs actually look for is outside the average — the outlier. For everything that is genuinely special, it’s harder.”

— [Lutz Finger](#), XGEN AI / Cornell University



The Dangerous: Agentic systems lack the contextual judgment that prevents catastrophic decisions. Lutz Finger described a widely discussed MIT simulation in which swarms of AI agents given geopolitical tools ended in nuclear strikes 95% of the time — not because the agents were malicious, but because they lacked the human understanding that red lines are performative declarations, not literal commands. The same logic applies to venture: AI cannot yet know when a crazy idea is Snapchat and when it is just crazy.

THE CO-PILOT FRAMEWORK: WHERE TO DRAW THE LINE

Allison Vuong offered **a practical framework for where to draw the co-pilot versus autopilot line: consider the reversibility of the decision, the availability of context to the agent, and the clarity of accountability when something goes wrong.** Low-stakes, reversible, data-rich decisions — CRM updates, screening filters, portfolio monitoring — can be automated. High-stakes, irreversible, context-dependent decisions — investment committee votes, founder relationships, fund strategy — must remain human-led.



“Investing is a historically very biased field. Being able to bring in an outsider — an AI that can bring in at least a different perspective and highlight external information you may not be privy to — will inevitably improve decision quality. Judgment is the name of the game. But AI has a huge role to play in informing that judgment.”

— [Allison Vuong](#), Pivotal Life Sciences

Conclusion: There is a meaningful and consequential difference between using AI to augment investment decision-making and delegating investment decisions to AI. The former is a competitive advantage. The latter is a governance failure.

The Role of AI: Classical machine learning and LLMs systematically struggle with outliers and lack the contextual awareness that high-stakes decisions require. The key variable is reversibility: the lower the stakes and the easier the reversal, the more confidently AI can be deployed autonomously.

Panel 4: Will AI Amplify or Mitigate Human Bias?

PERFORMANCE AND DIVERSITY ARE NOT IN TENSION

The dinner panel produced the bootcamp's most substantive and at times contentious debate. Ben Orthlieb of Blue Moon Fund opened with a disclosure that immediately reframed the conversation: a seed fund achieving top-0.1% vintage performance while simultaneously generating 10x the diversity of the typical US VC fund — not as a mission-driven exercise, but as a product of intentional system design.



“Our first fund beats 99.9% of our vintage — 4.4 TVPI, 1.4 DPI — with 10x the diversity of the typical US fund. 25% of our CEO founders are female, 40% have female co-founders, 9% are people of color. I would like to claim brilliance. It’s probably just good tech with a good process, if I’m honest.”

— [Ben Orthlieb](#), Blue Moon Fund

The mechanism behind Blue Moon Fund's results: unsupervised machine learning that does not declare features in advance, allowing signal patterns to emerge from data rather than from human-defined proxies. This is a critical distinction. Supervised models that are told which features to look for — university attended, prior employer, LinkedIn connections — will encode the proxy biases that have historically concentrated capital in a narrow founder profile. Unsupervised approaches allow the data to surface what actually predicts success, independent of the features a human investor would have chosen to examine. The fund now screens 4,600+ companies per year through this system — a volume impossible to process manually.

Radhika Iyengar of StarChain Ventures pushed back on any suggestion that AI is a neutral arbiter of merit — and the debate that followed was one of the sharpest exchanges of the bootcamp.

“AI can’t solve for biased data. It just can’t. If the dataset it’s been trained on is saying certain kinds of founders succeed, that’s all it knows. Human first and human last and AI in the middle. If you have the same investors funding the same founders for the same kinds of projects, it’s just status quo. We need to change the flow of capital — make it multichromatic, not monochromatic.”

— [Radhika Iyengar](#), StarChain Ventures



AI IS TRAINED ON THE RICH WORLD'S DATA

Tom Huang of BK Fund added the global dimension that would resonate through the rest of the evening — and through the fireside chat that followed:

“ AI is trained primarily on materials in English, Chinese, and European languages. The population who have access to that information accounts for less than 10% of the total world population. AI is trained by the rich world's material. What happens to the rest? How can AI identify the outlier who is not living in the world we're familiar with — but may be just as powerful?”

— Tom Huang , BK Fund



This is not a diversity problem alone. It is a returns problem. The global founders most likely to build the next generation of breakout companies may be precisely those whose traction signals are invisible to current AI sourcing infrastructure — because those signals appear in WhatsApp groups, foot traffic counts, pay-as-you-go transaction records, and local-language social media rather than in GitHub, LinkedIn, or Crunchbase.

Ben Narasin, moderating, offered a counterpoint rooted in his own 18 years of investing — and the debate between his optimism about meritocracy and Radhika Iyengar's structural critique was never fully resolved:



“ Venture has changed. In the old days, I would argue venture was the Olympics — the best and brightest that already achieved something came and got to give it a shot. That's not venture anymore. Venture today is CrossFit. 20,000 people get seeded a year, and the ones that succeed, get funded.”

— [Ben Narasin](#), Seed & Series A Investor

The panel concluded that whether AI amplifies or mitigates bias in venture depends entirely on the intentionality of the people designing the system — what data they train on, what features they include or exclude, and whether they audit outcomes. Performance and diversity are not in tension when the system is properly designed. Intentionality is not optional - it is the design requirement.

The Role of AI: AI in venture is neither neutral nor inevitable — it is a mirror of the data it was trained on. The critical variable is not the technology itself, but the intentionality of the system design: what data is used, what features are included or excluded, and whether outcomes are audited. Unsupervised approaches have already demonstrated that performance and diversity can reinforce each other.

Panel 5: Fireside Chat: Andreessen Horowitz / Speedrun



HOW A16Z USES AI: THE SPEEDRUN MODEL

Emily Bennett, investing partner for Andreessen Horowitz’s (a16z) Speedrun early-stage vehicle, provided the most operationally detailed account of AI deployment in a major venture firm. Speedrun accepts approximately 65 teams out of 25,000 applications each cycle — a selection ratio that demands systematic prioritization. The firm trained a proprietary model on signals historically predictive of moving to next steps with a team: not replacing the work of investors, but systematically prioritizing where their attention should go first.

“We’ve trained a model that attributes a score to every application — trained on signals that historically have been predictive of our wanting to take a next step. The scores are really good. Almost spot on at the top and the bottom. I have been staggeringly impressed. But I also still think the human elements of VC are even more important: building a relationship, building credibility, making this person feel like they’re part of it. That doesn’t go away.”

— [Emily Bennett](#), Andreessen Horowitz / Speedrun

Bennett also addressed the question of whether AI will reduce barriers to entry in VC. Her answer was calibrated: AI can reduce the cost of operating a fund, but it cannot solve for reputation, LP relationships, or access to outlier deals — which remain self-perpetuating advantages for established firms. The VC stack is becoming increasingly top-heavy; AI, she argued, will not by itself dismantle that concentration.

ARTIFICIAL WISDOM: WILL AI SURPASS HUMAN JUDGMENT?

The conversation’s most philosophically significant exchange came when Bennett was asked whether AI will ever achieve not just intelligence but wisdom — the capacity to pursue the right outcome without emotional distortion.

“ I think we’re going to have AI that is wiser than humans. What limits human wisdom is that we are burdened by our own emotional investments. AI, in its current form, is not burdened by human emotionality. It learns from the history of all past actions and outcomes at a rate no human can match. If implemented well, it should actually be better at decision-making — not just faster, but better.”

— [Emily Bennett](#), Andreessen Horowitz / Speedrun

Bennett cautioned that this potential is conditioned on diverse inputs. Single-source AI creates filter bubbles at institutional scale — reinforcing rather than challenging the biases of its creators. The path to AI wisdom runs through multi-source inputs, explicit counterfactual design, and human review at consequential decision points.

EMERGING MARKETS: A NEW FORM OF LEGIBILITY

When asked about capital allocation to the billion-strong young population of Africa and other emerging markets, Bennett articulated what may be the most actionable insight of the fireside chat. A VC from Brazil’s Amazon region had described building AI tools specifically to surface founders with no social media presence, finding them through local press and university networks — and investing in several who have since grown significantly. Bennett responded:

“ The way to get into a room if you didn’t come from a market where legibility is built in — show quantitatively that what you’re building has heat. Transaction volume trends. Foot traffic. WhatsApp growth. GitHub stars. That’s a new form of legibility that is fully democratized. That’s how the young tinkerer from Africa who’s never heard of Harvard gets in the room. What did you build? And did people want it?”

— [Emily Bennett](#), Andreessen Horowitz / Speedrun

The firm is beginning to build tools that can scrape these alternative signals — website traffic, viral social posts, open-source validation, transaction volume trends — as supplements to traditional legibility markers. But Bennett acknowledged this remains harder to systematize at scale, and the structural barriers to capital access in emerging markets are not solved by signal infrastructure alone.

Panel Conclusion: AI’s potential in venture extends beyond efficiency — Bennett argued AI may ultimately surpass human judgment in decision quality, precisely because it is not burdened by emotional investment. But that potential is conditioned on diverse inputs and explicit counterfactual design. Single-source AI does not produce wisdom. It produces a faster version of whoever built it.

The Role of AI: AI is a force multiplier for attention, not a replacement for judgment. At AI6Z Speedrun’s scale, AI prioritization is already operationally indispensable — but its ceiling is defined by the diversity of its inputs and the quality of its design. The firms best positioned to benefit are those that pair AI infrastructure with deliberate human oversight at consequential decision points.

PRACTITIONER FRAMEWORK: THE AI ROLE MATRIX

Based on the collective discussion across all five panels, the following matrix reflects practitioner consensus on where AI belongs in the VC workflow today. The key determinants for each decision are: stakes (how costly is a wrong decision?), reversibility (can the decision be undone?), and context dependency (does the decision require tacit human knowledge that AI cannot access?).

Dimension	Autopilot	Co-Pilot	Human-Led
Deal Sourcing	Data aggregation; market scanning; news monitoring	Founder scoring; signal ranking; pattern analysis	Final sourcing decisions; relationship initiation
Due Diligence	Background checks; CRM data entry; document parsing	Scenario modeling; counterfactual analysis; memo drafting	Founder assessment; IC presentation; thesis validation
Portfolio Monitoring	CRM updates; daily news alerts; portfolio dashboards	Traction signals; risk flagging; board prep support	Strategic decisions; founder relationships; LP updates
Investment Decision	—	Counterfactual audit; bias check; consensus mapping	IC vote; term negotiation; fund strategy; LP relationships

RECOMMENDATIONS

The following recommendations emerge from the collective practitioner discussion across all five sessions. They are tiered by audience and time horizon to support immediate action.

For VC Firms

Next 30 Days:

- ▶ Map every step of your investment process and assign an AI role: autopilot, co-pilot, or human-only. Publish this map internally so every team member operates from the same decision rights.
- ▶ Audit your last 24 months of investment decisions for demographic and geographic patterns. Use AI to surface patterns your human process may have systematically missed.
- ▶ Pilot a multi-agent counterfactual diligence session on one active deal. Configure one agent explicitly as a skeptic and document what it surfaces that your standard process did not.

Next 90 Days:

- ▶ Draft an AI usage policy covering what models are used, for what decisions, and with what human review requirements. Share it proactively with your LP advisory board.
- ▶ If your sourcing AI is trained on historical funded-company data, begin building a parallel signal vocabulary for non-standard traction: transaction

volume, WhatsApp growth, foot traffic, local-language media. This is where undervalued deal flow is most likely to originate.

- ▶ Create a decision log for every new investment: document the key assumptions at time of IC. This is the foundation of your feedback loop and your anti-portfolio tracking.

Next 12 Months:

- ▶ Build LP-ready reporting on your AI governance practices. This will become a diligence requirement within two to three fund cycles.
- ▶ Establish a systematic anti-portfolio review. Track every deal you passed on, why, and what happened. Review at 12-month intervals. The gap between your anti-portfolio outcomes and your portfolio outcomes is your most honest assessment of judgment quality.
- ▶ If you operate in or want to access emerging markets, invest in building or partnering for regional signal infrastructure. The founders invisible to your current AI tooling may be your highest-upside opportunities.

For Limited Partners

- ▶ Add AI governance to your standard LP due diligence questionnaire. Ask specifically: what is the firm's policy on AI in investment decisions, and how is it documented? Firms without a clear answer are operating without visibility into a growing component of their process.
- ▶ Ask GPs what demographic and geographic patterns their AI sourcing systems are producing. If the answer is 'we don't track that,' the governance infrastructure is insufficient.
- ▶ Request that GPs report on their decision process, not just their outcomes. A systematic process that learns is more durable than one relying on GP intuition alone.

For Founders

- ▶ Before your next raise, run a deliberate ecosystem test across three investor communities. Track rejection patterns — they are data about ecosystem fit, not a reflection of your company's quality.
- ▶ Understand that fundability and success are not the same signal. AI-assisted screening is optimizing for what has historically gotten funded, not for what will create value. Build relationships with investors who have a demonstrated track record of backing non-obvious companies.
- ▶ If you are building in an emerging market, make your traction visible in formats that are universally accessible: transaction volume charts, user growth trends, and revenue curves communicate across any signal gap. Quantitative traction is the universal language.

For Ecosystem Builders

- ▶ Invest in smaller, more focused convenings where honest practitioner disagreement is possible. The VC Bootcamp model — curated, candid, practitioner-grade — produces insights that conference formats cannot.
- ▶ There is a critical need for shared signal infrastructure that can read traction in emerging markets. This is a research and ecosystem collaboration opportunity that no single firm can solve alone.
- ▶ Fund formation is being democratized faster than fund governance is being built. Legal service providers and LP networks should develop governance scaffolding specifically for first-time managers launching on AI-powered infrastructure.

WHAT GOOD LOOKS LIKE: **PRACTITIONER CHECKLIST**

A VC firm using AI responsibly and effectively in 2026 should be able to check most of the boxes below. Use this as an internal audit tool or as a basis for LP reporting.

Strategy & Governance

AI usage policy exists and is documented

Covers which models are used, for what decisions, with what human oversight requirements

AI role is defined for each investment stage

Autopilot / co-pilot / human-only designation per stage, reviewed quarterly

LP-ready AI governance disclosure is prepared

Partners can articulate the firm's AI practices to LPs clearly and without defensiveness

IP risk policy for founder materials is in place

The firm has addressed the risk of ingesting pitch decks and proprietary founder data into third-party AI systems

Deal Sourcing & Bias Mitigation

Sourcing AI is calibrated on success, not fundability

The firm has considered whether AI sourcing predicts value creation or pattern-matches to historical funding behavior

Sourcing output is audited for demographic and geographic patterns

If AI consistently surfaces a narrow founder profile, the issue is upstream in the training data

Non-standard traction signals are tracked

Transaction volume, WhatsApp growth, foot traffic, and local-language media alongside LinkedIn, GitHub, and Crunchbase

Unsupervised ML is considered for founder scoring

Feature-declared supervised models risk encoding proxy biases

Diligence & Decision-Making

Multi-agent counterfactual analysis is used in diligence

At least one AI agent is configured as a devil's advocate

AI outputs are reviewed by a human before IC presentation

No AI-generated analysis reaches the investment committee without human review

IC decision reasoning is documented at time of decision

Key assumptions captured at decision time, not post-hoc

Scenario analysis is used routinely

Optimistic, base, and pessimistic scenarios built with AI assistance and reviewed by humans

Learning & Feedback

Anti-portfolio is tracked and reviewed at 12-month intervals

Feedback loop between IC decisions and outcomes is operational

Portfolio bias is audited periodically using AI

Emerging market signal infrastructure is in place or in development

OPEN QUESTIONS & PRODUCTIVE TENSIONS

Some of the most valuable outputs of the 2026 VC-BootCamp panels were the questions that experienced practitioners could not answer. The 2026 cohort did not arrive at consensus on how much AI should reshape venture capital, but it converged on something more important: a shared recognition of the stakes, and a shared refusal to pretend the answers are simple. These define the research and practice agenda for the next cohort.

- ▶ **Who filters first?** Human-first versus AI-first screening is not merely a workflow preference — it reflects fundamentally different theories of where bias is introduced and where it is corrected. BDev Ventures uses AI as the first pass; StarChain Ventures insists on human judgment before AI enters the process.
- ▶ **The Moneyball Question:** Are we in a Moneyball moment for venture? The panel consensus: not yet. The data quality, fund transparency, and feedback loops required do not yet exist at scale. The average VC fund runs 17 years; you cannot train a model on outcomes you won't know for a decade.
- ▶ **Supervised Versus Unsupervised ML:** Blue Moon's results suggest unsupervised ML produces less bias in founder scoring. But the debate between intentional feature design and emergent pattern recognition was unresolved. At what cost to precision does unsupervised ML operate?
- ▶ **Intellectual Property:** When founders share business plans with VCs who run them through AI, the idea enters a system that may surface it to others. Governance frameworks for this are nonexistent and urgently needed.
- ▶ **Artificial Wisdom Versus Intelligence:** Emily Bennett argued AI may eventually surpass human wisdom — free of emotional bias, learning from all past outcomes simultaneously. If that is true, what is the enduring role of the human investor? Is it relationship, judgment, or simply accountability?
- ▶ **Emerging Market Signal Infrastructure:** WhatsApp growth, foot traffic, pay-as-you-go transaction volume, and local-language media are not currently readable by standard AI sourcing tools. Who is responsible for building the infrastructure that changes this — and how do firms collaborate without surrendering their proprietary advantage?
- ▶ **AI Accountability:** When an AI-assisted decision leads to a loss, who is accountable — the GP, the model provider, or the firm's governance structure? This question has not been answered legally or practically.

Conclusion

The 2026 VC-BootCamp made one thing unmistakably clear: venture capital is not waiting for AI to mature. Practitioners are building with it now — scanning ecosystems, challenging investment theses, compressing due diligence timelines, surfacing diverse founders that human networks miss, and running funds with a fraction of the traditional headcount.

But the most experienced voices in the room were also the most calibrated. AI in venture is a force multiplier for human judgment, not a replacement for it. The funds that will outperform are not those that automate the most — they are those that use AI to ask better questions, surface hidden signals, and free senior attention for the decisions that cannot be delegated: reading a founder, building trust with an LP, and betting on an idea the world isn't ready for yet.

And underneath all of it was a global equity question the bootcamp surfaced but did not resolve: if AI is trained on the rich world's data, and the signals of the world's most promising founders appear in WhatsApp groups and foot traffic and local-language media, then the technology that promises to democratize capital allocation may in fact concentrate it further. That is a problem that requires intentional design, collaborative data infrastructure, and honest conversation — exactly the kind that VC Bootcamp exists to convene.

VC Bootcamp 2027 will continue this conversation. The next edition will include convenings in Paris and Africa — bringing the practitioner-grade discussion to the markets and founders most underserved by current AI sourcing infrastructure. If you are a VC partner, LP, founder, or ecosystem builder who wants to be part of the next cohort, we want to hear from you.

Insights from the 2026 VC-BootCamp

INTELLIGENCE WITHOUT GUARDRAILS

**Ecosystems, Bias, and the Future
of Capital Allocation**

Contact Us

Hoss Zaouali

Hoss@VC-BootCamp.com

416.834.5951

ORGANIZATION TEAM

Hoss Zaouali

Julien Delvat

Dennelle Walton

Sameer Patil

Tom Bionski

Helio Mosquim Junior

Jasmine Chen

Saroj Shendey

Yuri Curumthaully

Megan Esopenko

Thé Tieu

Julien Ziema

Théo Garcia