Crucial Questions for Applying Al in Startups

Barney Pell, Ph.D.

Founder: Powerset, Moon Express, LocoMobi Associate Founder: Singularity University Fellow: Creative Disruption Lab Advisor: CrowdFlower, Uber/Otto, LinkedIn, ...

barneypell@gmail.com @barneyp

My background

- I'm a long-time technologist, entrepreneur, investor and advisor
- Specialty and passion for applications of Artificial Intelligence (AI)
- SRI
- NASA
- Powerset
- Microsoft (Bing, Cortana)
- Singularity University
- Creative Disruption Lab
- Participated in hundreds of AI application efforts!

Challenges for Startups Applying Al

- Al technology is advancing rapidly
- But: Applying AI successfully is hard, with many issues leading to failure
- Art of the application is harder than the technology, and less understood
- Needs of the application often drive (new) requirements on technology
- Opportunity:
 - Ask the key questions early on can dramatically increase chances of success.
 - May also lead to different design and approach

Al Goldilocks Zone



Al Goldilocks Zone



Al Goldilocks Zone



Goldilocks Zone Questions

An AI Start-up should offer a solution to real customer needs that actually works and is much better than alternative approaches.

- Who are the target customers? Which is our best market segment?
- Which of their problems or needs are the most important ones we are planning to address?
- How well do our capabilities align with the core requirements?
- What are we bringing that lets us deliver a superior solution?
- Why is the problem too hard for the competitors and alternatives, but not too hard for us?

Customers

It is crucial to understand the net economic value of our offering to targeted customer segments

- What are the primary intended benefits of our system?
- How big is the potential market? Is one segment big enough?
- How well do we understand their needs?
- Have we identified the key use cases?
- How well do we understand our potential value?
- What are the core requirements to be selected and successful?
- How well do we understand the full costs to the customer to receive our value? Is it really worth it for them?

Capabilities

Al Applications rely on data, knowledge, team, and technology. Do we really have what it takes to succeed?

- What is our solution?
- What capabilities do we bring to the AI problem? (Data, knowledge, tech, team)
- How well do our capabilities align with the core requirements?
- What are we bringing that lets us deliver a superior solution?
- How will we get any data or knowledge we need?
- What's the optimal combination of data and knowledge for your application?

Alternatives & Competition

A startup's AI solution needs to be superior to existing competition and to obvious alternatives. What's the optimal mix of man and machine?

- What are customers doing today with respect to the needs we are addressing?
- How much better will we be?
- What can the customers do better even if they don't work with us?
- If we are offering a fully-automated solution, why is this better than human-in-the-loop?
- If we are supporting human-in-theloop, why can't we fully automate?
- Why can't other AI systems or simpler automation work well enough?

Human and Machine Division

- Questions
 - What's the optimum division of labour between humans and machines for this application?
 - What requirements does this division place on humans, and on machines?
 - What's the extreme version of automation that could be possible?
 - What are the most natural insertion points of the automation?
- Examples
 - Data extraction
 - X.Al
 - CrowdFlower

CrowdFlower

- Crowd-sourcing digital work platform
- Started as a human-centric work platform
- System initially routes all the labeling work to the humans, and the humans attach a label to the unlabeled data
- Now have added active machine learning capability
 - System can use this human-labeled data on its own as training data, potentially supported by data scientists. In this way, the system starts to learn its own models to predict the labels, and to attach confidence to its predictions
- When it gets a very high confidence prediction, it can just give the answer and not send to the humans at all
- Over time, the system can start automating a larger percentage of the labeling work and increasingly focus the human labeling effort on the more difficult to classify items
- Also an example of using a team of multiple different learning systems and multiple humans, each doing what they are good at, for an overall efficient, robust and adaptable system

Conclusion

- Ask the right questions early on will help in business success across the board
- Searching for the answers may transform various aspects of the business
- Prioritize R&D based on what really matters
- Minimize risk of product failure
- Accelerate time to key business milestones
- Increase chance of funding, as investors will ask many of these questions too