Machine Learning/Artificial Intelligence Strategy

Developing a Trust-Focused Strategy for Technology Adoption

Leslie A. McFarlin
Senior UX Architect, MF Business Unit, EDP Valuestream

Introduction

Machine Learning and Artificial Intelligence (ML/AI) are becoming integral parts of technology strategies in nearly every industry.

Success of ML/AI depends upon gaining users' trust:

> Do users trust the creator of the technology?
>
> Do users trust the technology itself?

Bottom line: CA needs to establish itself as a **trusted creato**r of ML/AI-based products.

Introduction

In order to do this, a disciplined approach is needed that begins with understanding:

What is the perception of CA?

What is the general perception of ML/AI?

Narrative strategy can help CA address any negative perceptions of ourselves and ML/AI while supporting positive perceptions, a tactic necessary for gaining acceptance among the key audience.

Introduction

Trusting ML/AI products also depends upon meeting users' expectations:

> If users have low expectations, how can we create products that pleasantly surprise them?
>
> If users have high expectations, how can we meet them?

User Experience (UX) Strategy with a robust research program can foster evidence-based interaction and visual design standards throughout CA, as well as monitor the success of ML/AI products that go GA.

Introduction

Trustworthy products also support users in their use environment:

> Can users understand how ML/AI functions in the product?
>
> Do users know why the ML/AI product arrived at its output?

Content Strategy can link clear and concise in-application guidance with external help sources such as DocOps and CA's help desk so users never feel unsupported during ML/AI product use.

Introduction

Collaboration is key to gaining acceptance of ML/AI technology, and creating trustworthy ML/AI applications.

Gaining technology acceptance requires User Experience Researchers to be aware of the narrative strategy and what is driving it.

Creating trustworthy ML/AI applications requires collaboration between UX teams, information engineers, ML/AI engineers, and data scientists.

Why Trust

Gain Users by Connecting with Them

Why Trust?

What Is Trust?

Cultivating Trust

Trust in ML/AI

Developing for Trust

Why Trust?

Trust is a prerequisite for technologies to be successful regardless of their domain, intended user group, or use environment.

Trusted technologies get used to their fullest extent, and can even result in novel uses.

Trusted technologies make users more efficient and effective at performing tasks.

Trust also leads to risks among users that must be anticipated and managed via user research.

What Is Trust?

Belief that the relationship between someone (trustor) and another entity (trustee) will provide the trustor with benefits.

Indicates the willingness of a trustor to give control to a trustee over some activity and its outcome.

Technology-based example:

> Security administrators rely on Compliance Event Manager to alert them to all administrator account signon violations.

Cultivating Trust

Regardless of the trustee or the environment surrounding the trustor-trustee relationship, trust develops in a 3-step process:

Trustor forms **trusting beliefs**, the belief that a relationship with the trustee will be beneficial.

Trusting beliefs progress to **trusting intentions**, in which the trustor intends to rely on the trustee for a specific activity/outcome.

Trusting actions are the culmination of this process, in which the trustee has executed the action upon which the trustor is depending.

Cultivating Trust

Execution of a trusting action only marks the start of trusting relationships.

> Trust evolves with time, either becoming stronger or weaker to the point that the relationship is dissolved.

Early stages of trusting relationships rest upon **deterrent-based trust**, in which the trustor evaluates the relationship with a trustee based upon cost.

> Uses relationship calculus: If the cost to maintain the relationship outweighs the benefit, the relationship will end.

Cultivating Trust

Deterrence-based trust progresses with time and familiarity into **knowledge-based trust.**

For technology, knowledge of the technology and its creator fosters knowledge-based trust.

Knowledge-based trust offers the trustee more room for mistakes and underperformance as the trustor has enough experience with them to make excuses for such failures.

In comparison to deterrence-based trust, it is harder to erode.

Trust in ML/AI

Trust in ML/AI products has to be built on multiple levels:

In the algorithm driving the technology.

In the output of the algorithm.

Both of these aspects depend upon the quality of the user experience, as well as the content within the application and supporting it.

Trust in ML/AI

Trusted technology empowers users, enabling them to work on their own without relying on others for assistance.

Empowering users with ML/AI applications requires creating stable mental models in order for them to perform all necessary functions without being tethered to coworkers or help documentation.

> User-centered design will enable users to form **functional mental models**, while clear and concise content will help them form **structural mental models**.

Trust in ML/AI

Trust, though good, also encourages development of **overreliance**, the tendency for a human operator to reduce their oversight on the ML/AI's quality of output.

Overreliance, while always an issue, becomes a dangerous threat to ML/AI application performance as data sets start to shift.

> Data sets that resemble training data should perform well; however, as dissimilarity between training data and real world data increases, performance will vary.

Product teams must also **design for user engagement** within ML/AI applications.

Developing for Trust

Creating ML/AI applications that inspire trust requires partnership across multiple groups within the product development organization.

UX teams must partner with ML/AI engineers and data scientists so that an optimal user experience can be designed based upon the algorithms used and intent of the ML/AI application.

Within UX teams, researchers and designers must work closely to ensure timely sharing of information.

Cross team collaboration with UX and content writers can ensure the development of content that adequately supports the user during product use.

Narrative Strategy

Assessing, Improving, and Maintaining User Perceptions

Defining Our Narrative

Setting Goals and Outcomes

Where Do We Start?

Defining Our Narrative Strategy

Narrative strategies are stories intended to produce specific outcomes.

For CA, a solid narrative strategy will require understanding two things:

> What is the perception of CA, and why?
>
> What is the general perception of machine learning, and why?

Based upon these understandings, a strategic roadmap complete with action plans and measurable milestones can be developed.

Setting Goals and Outcomes

At a high level, a narrative strategy should seek to counteract any negative perceptions of CA and support the positive ones.

Similarly, negative perceptions of ML/AI need to be counteracted and positive ones supported.

Awareness of individuals and groups who caution against the use of ML/AI is necessary as CA should be seen as a **responsible creator of ML/AI applications**.

Marketing needs to take the lead crafting a narrative strategy with specific goals and measurable outcomes.

Where Do We Start?

Begin a collaboration between UX researchers and marketing to find and review existing data on perceptions of CA.

Connect UX researchers with sales and pre-sales to learn about any discussions around ML/AI applications they may have had with customers.

Where more information on both topics are needed, craft research plans to gather data on perceptions of CA and ML/AI.

UX Strategy

Leveraging Research to Design Smarter

Research As a Foundation

Standards

Drafting a Process

Where Do We Start?

Research As a Foundation

Utilizing research at all levels of product development helps inform product roadmaps, ensures that product teams do meaningful work, and leads to products that benefit users.

Research provides the understanding of users' reality that product teams need to create ML/AI applications for users to rely upon as they perform their work functions.

Although research in this context refers to UX research, ideally a product's PM, PO, and UX team will work together to gather insights and share them with the development teams.

Research As a Foundation

As an exploratory tool, research can provide insight into the user's work functions and work environment, as well as the issues in the user's industry that impact their organization's processes.

> Understanding the user's ecosystem enables product teams to design task flows that minimize friction with how users perform their work tasks.

As a design tool, research can help product teams iterate on designs until they create one that satisfies user needs and business goals.

Research As a Foundation

Content writers can also use research to assess how well content supports application use and learnability, as well as how useful and understandable it is.

Once ML/AI applications go live, research can still provide useful data on product use, such as strengths and areas for improvement and growth.

As products go live, research must also look beyond assessing product usability, and assess the evolution of the user + ML/AI application relationship to examine the development of biases and their impact.

Standards

Given the uniqueness of ML/AI applications, standards defined and refined through user research should be created for teams to leverage.

> Individual business units may need unique standards, and should be encouraged to conduct the research they need to create them.

Interaction design standards and visual design standards should be based around the knowledge of how to create trust, combining to create a **trust-centered design system.**

Building a Trust-Centered Design System

Creating a trust-focused design system requires setting standards that address the 3 P's of ML/AI application use:

Performance, getting the system to help users successfully complete work tasks

Process, user perceptions of the ML algorithm's fit to the work tasks

Purpose, user perceptions of what the ML/AI application's creators intended it to do

Each of the above items relates to an important facet of product design that builds trust.

Building a Trust-Centered Design System

Performance
4 Facets

Competence
Information Accuracy
Reliability over Time
Responsibility

Process
4 Facets

Dependability
Understandability
Control
Predictability

Purpose
3 Facets

Motives
Benevolence
Faith

Drafting a Process

At the heart of ML/AI interactions rest a kernel of uncertainty that must be understood and assessed for threat to building technological trust.

> Prioritize all uncertainties by their potential threat to technological trust.

Threats should be related to one of the facets of the 3 P's in order for UX designers to create a solution.

Product teams can then specify the functional requirements to make the trust-centered solution a reality.

Where Do We Start?

Exploring general awareness of machine learning and artificial intelligence among CA products' end users to gauge their understanding.

Surveying existing research on biases surrounding artificial intelligence, including the threat it poses to task outcomes.

Content Strategy

Supporting Users Inside and Outside Applications

Talking to Users

Developing Explanations

Support Outside the Application

Where Do We Start?

Talking to Users

ML/AI applications must communicate in ways that users will accept, and that encourage them to continue using the application

Tone of communication (e.g., friendly vs. professional) influences users' perceptions of a product.

> Example: Use an empowering tone when you want to help users understand why they may have received unwanted results so that they feel the product is helping them.

Matching tone to user's expectations for an interaction encourages them to use the product, and discourages reading between the lines of a message.

Developing Explanations

Explanations in ML/AI applications should empower users by supporting and providing depth to the functional mental models they develop from application use.

Two major components of ML/AI explanations are soundness and completeness.

> Soundness refers to how well each feature of a ML/AI application is explained (faithfulness).
>
> Completeness refers to how thoroughly a ML/AI application is explained (breadth and depth).

Developing Explanations

When crafting explanations, too much information is as dangerous as too little information.

Overcomplication, too much information, impedes structural mental model creation by negatively impacting a user's ability to process and understand information.

Oversimplification, too little information, affects both soundness and completeness and leads to the formation of incorrect or unstable mental models.

Support Outside the Application

External help documentation tends to be more detailed, making it ideal for adding depth to users' functional mental models.

> Detailed documentation also runs the risk of overcomplication and oversimplification, so steps should be taken to evaluate content and rewrite it as needed.

UX designers and content specialists should work together to identify logical points for accessing external help documentation.

Where Do We Start?

Existing ML/AI applications should evaluate content for appropriateness of tone with respect to tasks, and oversimplification/overcomplication of internal and external help content.

Identify content best practices that may already exist for ML/AI products.

Linking UX, Engineering, & Data Scientists

Collaboration as a Way to Increase ML/AI Quality

Form Influences How Users Function

Realizing Internal Dependencies

Institutionalize ML/AI Excellence

Form Influences How Users Function

Machine Learning relies upon algorithms, the form of which impacts how input is evaluated, how predictions are made, and how the application learns.

Tasks influence the type of algorithm selected for machine learning.

Complexity of an algorithm influences how quickly it works: Increasing complexity requires increasing the time to complete all steps.

UX teams should understand the anatomy of the tasks ML algorithms will perform, as well as the strengths and limitations of the algorithms selected by data scientists.

Realizing Internal Dependencies

Creating trustworthy ML/AI applications requires acknowledging dependencies between the various groups involved in product development.

> UX and engineering rely upon data scientists for their expertise in algorithm selection and design.
>
> Data scientists, engineering, and UX rely upon users for their expertise in the tasks that algorithms aim toperform.

As CA looks to become a trusted developer of ML/AI products, they need to develop a process that encourages interactions between UX, engineering, and data scientists.

Institutionalize ML/AI Excellence

Provide a space physically or digitally for those involved in the development of ML/AI products to exchange information and share their discoveries.

Publish standards for ML/AI application development, the strategies supporting them, and the research and design activities that influence them.

Enforce global standards, but also allow product lines the flexibility to create unique standards as needed.

Thank you!

For any questions, please reach out to the creator of this strategy:

Leslie A. McFarlin, Senior UX Architect

Mainframe Business Unit, EDP Valuestream