



Design Principles for Automated Advocates

The field of Automated Advocates is relatively new. An Automated Advocate is a technology tool that both works on behalf of someone to help them achieve their goal more quickly, affordably, or easily than they otherwise would have by reducing their administrative burden through the use of automation; and uses insights from users' aggregate experience to advocate for improved functioning of the overall system for everyone.

Here is a collection of suggested design principles, developed in consultation with existing tool builders from across the civic tech landscape, to guide the growth of AAs so that they can be ethical, scalable, and effective. Because this field is still growing, there is no one perfect way to create an Automated Advocate tool, and some entrants may achieve some of these Design Principles but not all. However, all Automated Advocate builders should be able to articulate their progress relative to this list. Each principle is accompanied by a set of questions that can be used by government, users, or AA builders to benchmark the quality of their tools. Use these Design Principles to inspire development areas for new Automated Advocates and to assess existing ones.

The 9 Design Principles:

- 1.** Be better than the status quo
- 2.** Be transparent about what the tool is and how it works
- 3.** Be accessible
- 4.** Be a good data steward
- 5.** Be prepared for roadblocks with an acceptable failure mode
- 6.** Be loyal to the user's ultimate goals
- 7.** Be a responsible digital citizen
- 8.** Be open with data that could lead to systemic improvements
- 9.** Be responsive to changes in user needs

1. Be better than the status quo

An Automated Advocate is asking a user to invest time, trust, and energy in interacting with their tool. AAs should measure themselves against a user's existing options to meet their goal and be able to show how they are measurably better.

Why It Matters: Technology is not a panacea. Automated Advocates should provide better outcomes, and avoid inadvertently contributing to worse ones, for their users. Notably, “better” can be across a number of dimensions, such as affordability, discoverability, accessibility, legibility, and comprehensiveness.

Questions to ask:

- How are people solving this problem now?
- In what ways are you an improvement?
- In what ways might you risk being a step backwards?
- What data can government and community members use to benchmark your success, either now or after you are in existence for a period of time, against the current system?

2. Be transparent about what the tool is and how it works

An Automated Advocate should clearly tell a user what the tool will help them accomplish and how it works. Users should be clear on the business model, authorship, data collection, and role of automation of any AA they interact with.

Why It Matters: Users entrust Automated Advocates with sensitive information and critical requests. They need enough information to be able to make informed decisions about whether or not a given AA is the right fit for them.

Questions to ask:

- How can users find out how your tool works?
- How easy is it for your users to find:
 - your business model?
 - your data policy?
 - how recently your tool has been updated?
 - what parts of the tool are run by automation versus humans?
 - and,
 - who created the tool?

3. Be accessible

Automated Advocates exist to serve everyone, equally. They must follow current laws around web content and should follow [best practices in digital design](#) to ensure that the tool is accessible to everyone regardless of their individual abilities, device, environment, or quality of access. Automated Advocates should seek to be understood and use plain, clear language.

Why it matters: Technology should democratize service delivery, not deepen existing inequities. Accessible design allows everyone to participate and designing for the margins creates better tools for everyone. It also helps the Automated Advocates by maintaining the broadest possible community of users.

Questions to ask:

- What are you doing to make your Automated Advocate as accessible as possible to all audiences? Please consider:
 - users who depend on Web Content Accessibility Guidelines;
 - users with reduced access to technology devices;
 - users with lower literacy levels; and
 - users who do not speak English.
- If you are not currently as accessible as you could be, what mitigation efforts are you making? How and when do you plan to change that?
- How does the makeup or training of your team contribute to your ability to make a tool that is as accessible as possible?
- What else have you learned about the accessibility needs of your users and how can you use that knowledge to guide service delivery in the future?

4. Be a good data steward

Automated Advocates will often need to request data from users to help them accomplish their goal. They should clearly identify what data is needed, as well as how they will use, store, or share those data. AAs should show their respect for users' privacy and security by maintaining clear, rigorous policies and practices.

Why it matters: Automated Advocates may ask for a lot of detail from their users in order to help them achieve their goals. They can only do that if they merit the trust of their users. Additionally, AAs may collect sensitive information, and users could face serious repercussions to their data being shared inadvertently or without their consent.

Questions to ask:

- What data do you gather from your users? Whom else do you share that data with?
- How do you protect user data?
- How do you communicate your privacy and security policies?
- How often do you update those policies?

5. Be prepared for roadblocks with an acceptable failure mode

Like any technology tool, Automated Advocates may face unexpected obstacles that stop them from perfect performance or uptime. These might be failures of external dependencies, like a fax machine that stops receiving messages, or bugs inadvertently added by an AA development team. AAs need to plan for a resilient system in advance by having a fallback mode to ensure that regardless of the AAs' performance, users' applications will still be delivered, requests made, and/or data protected. To ensure process integrity, AAs should go beyond self-serve offerings; all AAs should have a way to escalate questions to a human.

Why it matters: Every tool is fallible. Admitting that up front, and planning for it, means that an AA can be used to deliver critical services without fear of users accidentally stumbling into a dead end where they begin a process and unknowingly will never reach their goal.

Questions to ask:

- What critical infrastructure do you rely on to function?
- What are your plans to make sure users do not hit a dead end if any one of those pieces of critical infrastructure no longer work? (This includes outside systems, like government benefits portals that sometimes go down.)
- How do you communicate any potential dead ends to users? These include moments when the technology is down as well as moments when an issue means a user request cannot be processed through the usual channels.
- How can a user escalate a challenge they are facing with your AA to a human support team?
- How will you track whether or not your users have successfully completed their goals?
- How will you ensure an elegant exit that does not put users at risk if your AA shuts down in the future?

6. Be loyal to the user's ultimate goals

Users come to Automated Advocates with a clear end goal in mind: usually to apply for (and receive!) a service or benefit or to defend themselves from a threat. When designing a tool, builders should include the minimum number of steps that allow a user to accomplish that goal. Even better, AAs should focus on supporting not just on a user's immediate need (e.g., to understand and finish a form quickly) but also their ultimate desire (e.g., to win a case or receive recurring benefits). Automated Advocates should consider their relationship with a user to be as sacred as a traditional fiduciary duty; by definition, they are an advocate for the success of their users over competing priorities.

Why it matters: Like lawyers or investors, Automated Advocates provide advice to and take action on behalf of users. It is critical that those actions are as unconflicted as possible. Automated Advocates must, as a field, differentiate themselves by their ethics and avoid shadow practices that appear on the surface to help users but ultimately impede them from reaching their ultimate goal.

Questions to ask:

- What do you see as your responsibility to your users? What industry ethical guidelines, if any, do you subscribe to?
- What are the ultimate goals of most of your users?
- How are you helping them to reach those goals, relative to the status quo?
- Where might your interests and those of your users conflict? How do you plan to make that conflict clear to your users? How do you plan to navigate that conflict with your behavior?
- How are you aligning your business incentives with your users' success in achieving their interim or ultimate goals?

7. Be a responsible digital citizen

Automated Advocates often function within a broader technological ecosystem in which they are drawing upon open source tools, existing APIs, and industry-specific integrations. To the extent possible, Automated Advocates should seek to be responsible digital citizens that leave the digital commons as good or better than they found it. This means being careful not to inadvertently overload and crash interdependent systems, supporting technical standards and data sharing, and contributing back to shared technical knowledge bases where possible.

Why it matters: Related to Principle 6, sometimes the quickest way to achieve a user's goal serves your individual user but crashes the system for everyone else. Automated Advocates should think about these tradeoffs up front and seek to minimize disruption to other users in need of services wherever possible. Automated Advocates should also recognize that working in the open often indirectly creates a better outcome for all potential users in the future, helping them to avoid vendor lock-in and making it easier for other builders to suggest improvements or build in collaboration with what has been created.

Questions to ask:

- What is your relationship to existing data sources and technical tools, and are there ways in which your Automated Advocate meaningfully impacts the ability of any of those other tools to function?
- If one of those tools or data sources is government-managed, is there anything government entities could do to make that integration smoother for both parties?
- In what ways are you contributing back to the knowledge of the commons with your technology (e.g., use of open source code, use of shared technical or data standards, or publication of learnings)?

8. Be open with data that could lead to systemic improvements

What differentiates an Automated Advocate from existing technology tools like standalone robo-lawyers, document-assemblers, and bots is their dedication to doing more than just assisting individual users. Automated Advocates commit to using their experience with individual users to learn about the functioning of a system as a whole and then to advocate for that system to work better. This means: illuminating and pushing back against roadblocks at both an implementation and a policy level, deciding strategically which data points to gather over time, being straightforward and bold in calling out observed challenges as they discover them, and trying to be a help to anyone who is trying to help their users. This also means being responsible and deliberate in how they gather and share user data to make sure users' intent and privacy are respected.

Why it matters: Automated Advocates have a unique advantage: they interact with enough individual users to gather both qualitative and quantitative data about their experiences. They can escalate individual concerns about a system's function to the level of class action and demonstrate patterns (or lack thereof) where there may only have been hunches. These data, combined with a receptive systems-level partner, can help make the case for meaningful improvements at specific points in both policy and operations. For example, if an AA shows that the majority of its users experience delays at a certain point in the process of getting a particular benefit, it can advocate for better staffing or reduced overhead to ensure that the users can achieve their ultimate objective.

Questions to ask:

- What data do you track to learn about your users' journey through a system?
What external-to-you roadblocks have you already been able to uncover
- through this process?
What information can or should be shared with users, the public, and/or government agencies?
- How will you ensure that you share any information safely and securely, in a way that does not put your individual users at risk?
How regularly can these audiences expect data from you?

9. Be responsive to changes in user needs

Automated Advocates exist at the intersection of available technology, government policy, user needs, user expectations, and market conditions. As these ingredients evolve, so, too, should the approach of an Automated Advocate. AAs should model the best behaviors of the civic technology movement: start small, stay centered on user needs, iterate, and improve over time.

Why it matters: Automated Advocates should use the best available technologies to meet users where they are at, and users' needs and abilities change over time. Successful Automated Advocates will continue to grow alongside their users.

Questions to ask:

- How will you evolve your Automated Advocate over time in response to specific user needs?
- Give an example of a learning you had about your users, your technology, or your policy landscape and how it led to a change in your tool. What are some user needs you are not currently satisfying but may try to address in the future?
- How do you plan to improve and iterate over time?
Answer these questions from the [Good Service Scale](#) about your Automated Advocate:
 - users who depend on Web Content Accessibility Guidelines;
 - users with reduced access to technology devices;
 - users with lower literacy levels; and
 - users who do not speak English.