

# Task Force on Artificial Intelligence, Emerging Technology, and Disability Benefits

## **Phase One Report**

April 2025



NATIONAL ACADEMY  
OF SOCIAL INSURANCE

# Acknowledgments

Founded in 1986, the National Academy of Social Insurance is a nonprofit, nonpartisan organization made up of the nation's leading experts on social insurance policy, practice, research, and innovation. From Social Security and Supplemental Security Income; to Medicare, Medicaid, and the Affordable Care Act; to Workers' Compensation and Unemployment Insurance; to the Supplemental Nutrition Assistance Program, Temporary Assistance for Needy Families, and income-boosting tax credits, America's social insurance ecosystem serves as the bedrock of economic protection against the risks of life — such as when we retire, lose a job, experience disability/illness, or lose a family breadwinner.

For nearly 40 years, the Academy and its powerful, diverse member network have championed the safeguarding, strengthening, and modernizing of social insurance and worked to increase public understanding of how it contributes to economic security. Reports of the National Academy of Social Insurance aim to assess policy options without advocating for any particular option. The options and perspectives discussed in this report do not necessarily reflect the views of individual members of the Task Force or of the organizations with which they are affiliated.

The National Academy of Social Insurance gratefully acknowledges the work of the individuals who serve on the Task Force on Artificial Intelligence, Emerging Technology, and Disability Benefits. Their affiliations represent their positions at the time of the Task Force's deliberations and report publication.

As with all Academy reports, this report went through peer and Board review to ensure that it meets the organization's standards for clarity, completeness, accuracy, and objectivity. Special thanks to Paul Van de Water for serving as the final Board reviewer of this work.

To learn more about the Academy's work, please contact Rebecca Vallas, Chief Executive Officer at [rvallas@nasi.org](mailto:rvallas@nasi.org)



# Task Force Members:

**Chantel Boyens**, *Urban Institute*, Co-Principal Investigator

**Jack Smalligan**, *Urban Institute*, Co-Principal Investigator

**Jen Burdick**, *Community Legal Services of Philadelphia*, Co-Chair

**David Camp**, *National Organization of Social Security Claimants' Representatives*, Co-Chair

**Henry Claypool**, *American Association of People with Disabilities and Center for Democracy and Technology*, Co-Chair

**Tracey Gronniger**, *Justice in Aging*, Co-Chair

**Ariana Aboulafia**, *Center for Democracy and Technology*

**Ly Xīnzhèn Zhǎngsūn Brown**, *National Disability Institute*

**Ashley Burnside**, *Center for Law and Social Policy*

**Fay Lomax Cook**, *Northwestern University*

**Jeff Cruz**, *American Federation of Government Employees*

**Indivar Dutta-Gupta**, *National Academy of Social Insurance*

**Hannah Groch-Begley**, *The Roosevelt Institute*

**Mia Ives-Ruble**, *Center for American Progress*

**Dom Kelly**, *New Disabled South*

**Aparna Mathur**, *Amazon*

**Graham McDonald**, *Urban Institute*

**Tom Nicholls**, *AARP*

**Shaun O'Brien**, *AFSCME*

**Aron Paquette**, *Unum*

**Rylin Rodgers**, *Microsoft*

**Steven Rollins**, *Former Associate Commissioner, Social Security Administration*

**Meilani Santillán**, *Code for America*

**Tim Shaw**, *Aspen Institute*

**Jude Soundar**, *Microsoft*

**Emerson Sprick**, *Bipartisan Policy Center*

**Shelley Washington**, *AFGE Local 1923*

**David Weaver**, *University of South Carolina*

**Darren Webb**, *Mindset Care*

**David Wittenburg**, *Westat*

**Dora Mendelson**, *National Academy of Social Insurance*, Research Assistant

**Rebecca Vallas**, *National Academy of Social Insurance*, Project Director

# Executive Summary

The National Academy of Social Insurance formed the Task Force on Artificial Intelligence (AI), Emerging Technology and Disability Benefits to assess and inform how the Social Security Administration approaches the use of AI in the Social Security Disability Insurance (SSDI) and Supplemental Security Income (SSI) programs for purposes of determining benefit eligibility. The goal of the Task Force and this accompanying report is to bridge general AI guidance to the Social Security Administration's (SSA) specific programs and challenges. SSA processes millions of disability claims a year. Eligibility for SSDI and SSI is strict, and claims undergo a complex, multi-step determination process before benefits can be paid to eligible individuals. Currently, SSA faces an unprecedented backlog of cases waiting for review, imposing long, costly wait times on applicants. SSI and SSDI are critical programs that often act as a lifeline for people with disabilities or disabled family members. They are intended to provide income support due to significant work-limiting disability, old age, and/or death of a spouse or parent.

The Academy Taskforce brings together experts on disability policy, including legal and social service practitioners who assist disability claimants in accessing benefits and services, and technology experts from the private and non-profit sectors. The Taskforce also consulted experts from within government, including the Social Security Administration's Office of the Chief Information Officer and the Chief AI Officer.<sup>1</sup> The Taskforce, its Co-Chairs and the Co-Principal Investigators conducted a series of meetings and informal interviews to gather information, learn and share insights across areas of expertise, and produce the report that follows.

In this report, we begin by describing the challenges SSA faces and how it is already using AI in the disability determination process. We then discuss the risks involved in using AI in the context of disability benefits, including types of bias that must be prevented and mitigated, followed by the potential opportunities for program improvement that AI may provide. Next, we summarize existing guidance to Federal agencies on the use and development of AI in programs and contexts that are potentially "rights-impacting."

We conclude that additional principles and guardrails are needed to guide the use of AI in the Social Security disability determination process. We identify four principles that should guide the Agency's approach to development and use of AI and act as a "north star" for future regulations, guidance and policies. These principles are:

- Do not use AI to limit or impede the rights of applicants and beneficiaries
- Prioritize human decision making in disability determinations
- Prevent, reduce and mitigate bias
- Improve fairness

In addition to these guiding principles, the Task Force has also identified a set of guardrails intended to help protect against the risks identified in this report.

The Task Force is invested in preventing potential harm and ensuring responsible use of AI, especially as new uses are introduced, and AI applications evolve. The Task Force has identified the following guardrails for the testing, deployment, and evaluation of AI:

- Establish internal and external governance processes for oversight of AI
- Make tools that are explainable, transparent, testable, accessible to people with disabilities, and subject to independent evaluation
- Strengthen procurement policies
- Conduct ongoing monitoring and evaluation of AI applications
- Protect data and personal privacy

Together, these principles and guardrails are intended to serve as a foundational framework for the development of more detailed policies and procedures governing use of AI in Social Security programs. More work is needed to ensure that the potential benefits and costs stemming from the use of AI in SSA programs are fully understood and evaluated, and that applicant and beneficiary rights are protected. As this technology rapidly evolves, we urge stakeholders across government, the private and non-profit sectors, and academia to engage in ongoing monitoring of the opportunities and challenges identified in this report and to build upon this framework.

## Introduction

Social Security Disability Insurance (SSDI) and Supplemental Security Income (SSI) benefits are a vital component of the U.S. social insurance system, offering critical economic protection across the lifespan to people unable to support themselves through work due to disability. SSDI provides disability benefits based on past work history to a diverse array of Americans, from older workers who did not expect disability to become part of their lives, to younger adults who have been navigating life with a disability since birth.

Today, a young person starting their career has a 1 in 3 chance of dying or qualifying for SSDI before reaching retirement age.<sup>2</sup> In 2024, SSDI served approximately 7.3 million Americans. SSI—which is not based on prior work—serves older adults and disabled people with very low incomes and few assets, including disabled children. SSI served approximately 7.5 million people in 2024. SSI and SSDI’s modest benefits provide essential protection for disabled workers and their families. The SSDI and SSI programs are administered by the Social Security Administration (SSA) and over half of the agency’s administrative budget is devoted to operating the programs.<sup>3</sup>

As artificial intelligence (AI) tools have become increasingly common in everyday life, government agencies—like individuals and private businesses—are exploring how AI can be used to improve service delivery and reduce costs.

For purposes of this paper, we use the National Institute of Science and Technology (NIST) definition of AI systems “as an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.”<sup>4</sup>

Other papers on AI have used a more expansive definition that includes simple computer-based algorithms. Under an expansive definition, some use of AI may seem inevitable. In addition, we recognize that there are many types of decision-supporting algorithms that can be used as part of AI systems. Some of these algorithms use AI to predict future outcomes of individuals in situations where the individual is applying for a job, consumer credit, or early release from incarceration. These are known as predictive algorithms, and some AI experts are highly critical of these AI algorithms, both raising ethical concerns and questioning their efficacy.<sup>5</sup>

The use of AI in administering government benefit programs carries a special set of potential risks and opportunities because they are potentially “rights-impacting,” meaning the use of AI can affect decisions or actions that have a significant material effect on an individual’s civil rights, privacy, or access to critical resources or services. These risks and opportunities are apparent across SSA’s programs, which touch the lives of nearly every American and are often the sole or primary source of income for individuals and families who can no longer work or are retired.

The potential risks and opportunities inherent in the use of AI are especially evident in Social Security’s disability programs which include a complex and lengthy claims determination process. There is significant pressure to adopt more efficient practices since SSA is currently operating amidst an unprecedented backlog of people waiting for a decision on their disability claims. In addition, individuals applying for and receiving disability benefits are already facing serious challenges to their health and income security. The potential to automate this process comes with risks, including the risk that algorithmic systems trained on pattern recognition will result in biased decisions.

By virtue of living with a disability, many disabled people exist outside of the confines of what is “typical.” Data that is used to train algorithmic systems and models is often limited due to the smaller number of people with disabilities and the fact that people with disabilities that are being included in these larger datasets may not be properly identified, or “coded” as people with disabilities.



When people don't identify as disabled (due to issues in definitions, stigma, concerns over privacy, or any number of other factors), this does not mean that disabled people do not exist in datasets it means they are not properly identified. Small sample sizes in surveys and few observations in administrative data put people with disabilities at inherently greater risk of discrimination when interacting with algorithmic and AI systems more generally than those without disabilities. This makes it particularly vital to focus on mitigating the risks of AI in Social Security programs that serve people with disabilities.

Managed well, SSA should seek to prevent, reduce, and mitigate risks inherent in the use of AI and emerging technology in its programs and ameliorate weaknesses in existing operations. In an ideal world, SSA would also address the root causes of rising backlogs and wait times—including understaffing, high rates of staff turnover, and unnecessary program complexity—instead of looking to AI to fix these systemic problems.

To succeed, SSA needs to clearly and transparently identify and measure areas of risk, establish targets for improvement and undertake continual evaluation. In addition, while SSA made progress establishing internal governance structures to guide AI development (including via the appointment of a Chief AI Officer), external input is also critical. This report intends to provide an external resource and perspective for understanding both the risks and opportunities with using AI to administer SSA's disability programs.

This report focuses on rights-impacting disability benefit determinations and introduces issues that may be common to the provision of medical care, as well as determining access to many other government benefits, especially means-tested benefits. Consequently, this discussion and lessons from SSA's experience with AI likely have relevance for other agencies, including the Health and Human Services (HHS) Department's oversight of Medicare, Medicaid, Temporary Assistance to Needy Families (TANF) and other health and means-tested programs and the United States Department of Agriculture's (USDA) oversight of the Supplemental Nutrition Assistance Program (SNAP) and other nutrition programs.

It is worth noting that people with disabilities may also be disproportionately likely to use these programs, as they are more likely to be low-income and require these forms of assistance. As a result, the integration of best practices across a host of benefits programs can make all of these programs more inclusive and accessible for their beneficiaries.

Perspectives on the use of AI in general, and for the delivery of public benefits, are diverse. Some emphasize the ways AI can harm low-income Americans with examples of how AI algorithms impact low-income people in various contexts, including housing, employment, and programs like SNAP, SSI and SSDI. One paper warns "AI provides a cloak of unwarranted rationality to actions that intensify or perpetuate injustices... Revealingly, the AI systems applied to low-income people almost never improve access to benefits or other opportunities."<sup>6</sup> This perspective underscores how AI can reinforce the marginalization of low-income folks across not only government systems, but in the private sector as well. As a result, it is both possible and necessary to adopt harm mitigation for AI usage.

Another perspective cautions against the risk of a “cascade of rigidity” in how AI is used: “In guidance, it would be helpful to explicitly promote the use of judgment and discretion on the part of civil servants, and to acknowledge that no risk is effectively stasis, and stasis has its own risks that must also be considered.”<sup>7</sup> This perspective warns of dangers associated with not making efforts to positively harness the potential of AI and assumes that in many ways, the use of AI is inevitable, a view addressed below.

This report is being released at a critical moment for AI in the policy landscape. Under President Trump’s administration, AI has become a huge focus as Trump and his allies attempt to pare down the federal government as much as possible. President Trump’s nominee for Commissioner of the Social Security Administration, Frank Bisignano, has decades of experience in the AI space, and if confirmed, promises to bring that into his work as Commissioner. This report offers a framework for the development of more detailed procedures governing use of AI in Social Security programs, in hopes that any use of AI is accompanied by policies that center the rights of claimants and beneficiaries.

## **Challenges Facing the Disability Determination System**

SSA receives about 2.2 million claims for disability benefits each year. After being reviewed by SSA field offices, claims are sent to state disability determination service (DDS) offices for the initial review. In recent years chronic underfunding of SSA operations and staffing disruptions during the COVID-19 pandemic have caused a historically high increase in backlogs of claims.<sup>8</sup> In 2017, SSA had 523,000 pending initial disability claims; as of July 2024, the agency has 1.216 million pending initial claims. Wait times for an initial determination have grown from 110 days in 2017, to 231 days as of August 2024. SSA is experiencing similar large backlogs and long wait times for cases that are appealed for a reconsideration by the disability determination services (DDS) office or an administrative law judge hearing.<sup>9</sup>

SSA manages the disability determination process while having to give flexibility to key decision makers, state DDS and administrative law judges (ALJs). The Social Security Act directs SSA to conduct the initial disability review through each state DDS office. The Social Security Act directs SSA to contract with state disability determination services (DDSs) to conduct initial determinations and reconsiderations (the first level of appeal), as well as continuing disability reviews (CDRs).



SSA funds DDS operations and conducts quality reviews, and state disability examiners must follow SSA policy guidance. However, there are significant variations in practices and outcomes among state DDS offices. SSA uses ALJs to conduct hearings when claimants appeal a denial at the DDS level. Federal personnel law limits SSA's ability to review an ALJ's decision. For many years widely differential allowance rates across ALJs, with grant rates varying from 10% to 90%, raised concerns with the consistency and fairness of ALJ decisions.<sup>10</sup> SSA has made strides in improving the consistency of ALJs decisions, in part through the use of data analytics, though variation among ALJs still remains.<sup>11</sup>

The growth of electronic medical records presents both challenges and opportunities for SSA. Medical records are growing in volume, requiring longer review by disability examiners. Conversations with individuals who work on disability claims indicate that the quantity of electronic medical evidence for each disability benefit application is growing.<sup>12</sup> Medical records for one individual can exceed 1,000 pages with around 80% of the records in an unstructured format. The length and redundancy of outpatient records are also increasing with one large medical facility estimating that the volume of records increased by 60% from 2009 to 2018.<sup>13</sup> One challenge for SSA's use of these records is the extensive redundancy in the files. One study estimates that half of the evidence was redundant reports from prior patient visits.<sup>14</sup>

SSA must also process benefit claims where the applicant brings very little medical evidence, either because the applicant is uninsured or has faced interruptions in health insurance; faces barriers to health care access; lives in an area with fewer medical facilities; has experienced homelessness; or some combination of these or other factors. SSA's OIG estimates the agency spends a half a billion dollars each year collecting medical evidence.<sup>15</sup>

The challenges facing SSA are clear: historic backlogs are rooted in understaffing, including high staff turnover rates. In addition, claims reviewers are faced with an unprecedented volume of medical records and the impact is seen in the growing backlog of claims. The instinct to turn to automation as a solution to these issues is understandable, but needs to be approached with care. Additionally, it is critical to note that use of AI in any one of these areas will not solve issues rooted in the understaffing and under-resourcing of SSA.

## **SSA's Current and Future Use of AI**

SSA deploys a number of AI systems that fall into the general category of algorithmic decision support tools. While SSA is using more traditional AI, the agency is exploring the use of more recent types of AI, some of which are referred to as Large Language Models or generative AI. Below we describe some of the more prominent uses of AI within disability programs, as well as areas identified by the agency as new potential use cases identified by the agency. A more complete list of SSA's current AI investments can be found in Appendix A.

The SSA AI application that has received the most recent attention is the Intelligent Medical Language Analysis Generation (IMAGEN) program. IMAGEN is focused on medical records in the file of an applicant for disability benefits. It uses natural language processing and predictive analytics to organize and provide visibility to key medical evidence. IMAGEN is used by a portion of state disability examiners to review initial applications and reconsiderations and conduct continuing disability reviews.

SSA has been using predictive models and expert systems for several decades.<sup>16</sup> During the initial disability determination, these models are used in the Quick Disability Determination and Compassionate Allowance programs to identify applications for accelerated review.<sup>17</sup>

At the hearings level, the Insight program is an AI application distinct from IMAGEN that uses natural language processing to review and identify weaknesses and inconsistencies in draft ALJ opinions given the requirements in SSA regulations.<sup>18</sup> SSA also uses algorithms to better manage the assignment of ALJ cases. Using claim specific data from SSA's case management system, these tools group cases to permit specialization in the review of claims.<sup>19</sup> SSA is also using Insight at the Appeals Council level, as in aid in reviewing ALJ decisions to identify potential quality issues.<sup>20</sup>

SSA also has used and is exploring new uses for AI systems to improve customer services. In a meeting with this Task Force, the agency identified the following as potential future applications of AI: using generative AI to support writing determination decisions and coding; use of AI to review customer service calls to find out if the needs of customers were met; and use of AI to help agency staff quickly find and understand AI policy. These and other uses are beyond the scope of this paper.

## **Risks in Using AI in the Social Security Disability Determination Process**

Using AI in rights-impacting disability benefit determinations introduces many risks. Below we describe two of the largest sources of risk: the potential to perpetuate and exacerbate bias in the eligibility determination process and the potential for disparate treatment of applicants to result in inequitable outcomes for applicants. Other areas of risk include a lack of transparency and explainability in eligibility determinations, a weakened role for human judgement, and increased burden on beneficiaries. Many of these risks are already present in the existing disability determination process but may be exacerbated by use of AI without proper guardrails.

# Potential bias in rights-impacting AI applications

AI technology is evolving quickly. In this paper, we focus on the risks associated with deploying AI to assist in rights-impacting disability determinations. In addition, many risk factors are interconnected with the need to identify and mitigate bias whether in existing processes or AI augmented processes.

Bias in disability benefit determinations can manifest itself in many ways. Bias can exist based on race, ethnicity, sex, sexual orientation, age, education, type of disability or medical condition, among other factors. Existing areas of bias, such as in access to medical services, can be compounded when medical evidence is needed to allow a claim. Bias can be especially difficult to detect where a particular factor, such as race or sex, is heavily correlated with the presence of a particular disability. Bias can and often is multifactorial. For example, an analysis of Black and Hispanic adults found that those with disabilities were also more likely to also lack access to healthcare and to forego needed care. As a result, it is likely that the combination of racism and ableism contributes to disparities in health and health outcomes for racial and ethnic minorities with disabilities.<sup>21</sup>

Evaluators have long identified the potential for AI-enabled and algorithmic tools to perpetuate and reinforce many types of bias, in various settings. For example, the integration of AI tools in hiring processes has been shown to expose potential employees to bias on the basis of race or disability.<sup>22</sup> Also, clinical algorithms used in healthcare settings have also been shown to be biased.<sup>23</sup> One study of the use of machine learning technology to schedule medical appointments found “that state-of-the-art scheduling systems cause the Black patients in our data set to wait about 30% longer than non-Black patients.”<sup>24</sup>

In the case of evaluating applications for disability benefits, the range of potential areas of bias is extensive. AI tools are informed by and learn from large data sets but these data sets can reflect underlying biases in the data, such as if the sample used for the data set does not appropriately represent the entire population.

For example, the Census Bureau’s American Community Survey has been criticized for underestimating disability incidence and certain types of disabilities experienced.<sup>25</sup> As the authors of one paper noted “Algorithms are built by humans. They are trained on data generated by humans. Humans discriminate, and so the algorithms they construct can discriminate as well.”<sup>26</sup> Understanding existing levels of discrimination—both implicit and explicit—based on race, gender, class, and ability is critical to understanding how AI might perpetuate bias. In a report focused on bias in AI, NIST described several risks associated with AI:

Systemic and implicit biases such as racism and other forms of discrimination can inadvertently manifest in AI through the data used in training, as well as through the institutional policies and practices underlying how AI is commissioned, developed, deployed, and used. Statistical/algorithmic and human cognitive and perceptual biases enter the engineering and modeling processes themselves, and an inability to properly validate model performance leaves these biases exposed during deployment.... These biases collide with the cognitive biases of the individuals interacting with the AI systems as users, experts in the loop, or other decision makers.<sup>27</sup>

NIST guidelines instruct users of AI tools to seek to manage bias in acknowledgement of the constant risk that new biases can emerge, and old biases can re-emerge. NIST and the Coalition for Health AI identify three types of bias:

- Systemic bias “can be present in AI datasets; in the organizational norms, practices, and processes across the AI lifecycle; and throughout the broader society that uses AI systems.”<sup>28</sup>
- “Computational and statistical biases may occur in datasets used to train AI systems and may also be present in the resulting algorithmic processes. Such bias often stems from systematic errors due to nonrepresentative samples.”<sup>29</sup>
- “Human-cognitive biases, as defined by NIST, are those that relate to how an individual or group perceives AI system information to make a decision or fill in missing information, or how humans think about the purposes and functions of an AI system.”<sup>30</sup>

Algorithmic bias presents a special risk for people with disabilities given the reality of structurally ingrained ableism. This does not mean, however, that AI needs to be avoided-- it just needs to be properly managed. As researcher Ian Moura puts it, “strategies from other areas of research can serve as possible means of responding to and preventing algorithmic bias towards disabled people. In particular, participatory methods, the use of inclusive design, and efforts to improve disability data must be part of an overarching effort to ensure that disabled people are not further marginalized by algorithmic tools.”<sup>31</sup>

In its AI risk framework, NIST cautions that “systems in which harmful biases are mitigated are not necessarily fair. For example, systems in which predictions are somewhat balanced across demographic groups may still be inaccessible to individuals with disabilities or affected by the digital divide or may exacerbate existing disparities or systemic biases.”<sup>32</sup> We consider this point further when we discuss the potential for disparate treatment of applicants below.

## Specific risks of AI decision-making processes that lack meaningful human review

Some AI applications are designed to produce program eligibility decisions with little or no human involvement. The Center for Democracy and Technology (CDT) labels some of these applications as “algorithm-driven decision-making.” In a report focused on how these applications impact benefit determinations for people with disabilities, CDT focuses especially on algorithmic systems that have been integrated into state Medicaid applications used to allocate home and community-based services (HCBS).<sup>33</sup>

CDT identifies a variety of risks with rules-based AI decision making tools, including potential violation of due process rights, the Administrative Procedure Act and the Americans with Disabilities Act. When agencies use applications, especially when designed by contractors, there is a risk that proprietary claims prevent individuals from obtaining the information needed to understand and perhaps challenge a decision on their eligibility. As we discuss later in this report, we recommend that SSA and other agencies focus on AI tools where humans are still the ultimate decision makers.

## Automation bias or the risk of undermining the thoroughness of a disability review

Another type of bias is automation bias, which occurs when an adjudicator has an uncritical acceptance of an automated recommendation. Automation bias is, by and large, an unconscious sort of bias that leads most people to defer to decisions made by computerized or algorithmic systems. Automation bias is a good example of a risk that may not be a serious concern at the initial deployment of an AI tool when users are testing and evaluating the technology. Over time, automation bias may undermine the vigilance with which staff exercise their own professional judgment regarding the recommendations from the AI tool. Given the backlogs of disability cases and pressures on DDS staff discussed earlier, automatically deferring to the output from an AI application could be tempting for overworked staff.

The performance metrics around an AI tool should not discourage adjudicators from deeply weighing the harder to quantify factors in an application for disability benefits, particularly in circumstances where the evidence is sparse or somewhat inconsistent with the claimant's experience.

This Task Force recognizes the importance of the “human in the loop” framework, which advocates for the prioritization of the human role in algorithmic decision-making at various levels of AI programming and implementation. Increased use of advanced AI tools presents the risk of delegating important decisions to technology with only a token human role. In some programs, algorithm-driven decision-making has fully automated the decision-making process, removing any human review.<sup>34</sup> Even if agencies take clear steps to have a person make a final benefit determination, there is the potential for human oversight to erode over time as AI tools become more robust.

## Disparate and inequitable impact across claimants

SSA already uses algorithms to identify disability claims that have a high probability of meeting the criteria for an allowance. As is discussed in Appendix A, the Quick Disability Determinations (QDD) predictive model is used to expedite allowances for certain claimants with conditions that are easier to resolve.<sup>35</sup> SSA contracted with Abt Associates to investigate the potential for bias in the QDD model. Abt’s report is not definitive in terms of the existence of bias, but it does note that QDD cases are predominantly cancer related, and they found that certain groups, including Black and Hispanic people, were underrepresented.<sup>36</sup>

QDD cases represent a relatively small portion of overall cases. As we discuss below, we see an opportunity for SSA to successfully develop additional AI applications to review medical evidence and identify cases with a high probability for an allowance, functionality IMAGEN is starting to provide. Should SSA pursue this strategy, the agency also needs to evaluate how to process and prioritize cases that do not as readily benefit from this technology. SSA may not be able to develop new AI tools across the full range of types of disabilities. It is valuable for SSA to identify areas where AI can provide an effective solution, as well as areas where AI tools may not be a good fit. For example, medical experts have cautioned that fewer AI applications have been developed in the area of mental health because of the complexities of a mental health disability allowance.<sup>37</sup>

Another area of concern is claimants with little medical evidence, either because they live in an underserved community or lack health insurance. Staff time that is saved due to automation should be devoted to people with underdeveloped claims. Additionally, AI systems should be trained to identify those underdeveloped claims and suggest the type of information that would help complete the claim.

SSA needs to take care when privileging easily escalated cases over those that are more complex or time-consuming. To the extent that AI tools enable certain claims to be more quickly adjudicated, SSA should reallocate that time to reviewing claims that might be more challenging or complex, like certain behavioral health claims.



SSA's experience with deploying new technologies and internet applications gives an indication of the potential disparate impacts on individuals as well as the potential benefits to those who are able to access the technology. When SSA deployed the SSDI internet applications called iClaims in 2008, research found the new technology privileged those with internet access. One study found that counties with internet connectivity experienced an on average a 1.6 percent increase in SSDI applications, and a 2.8 percent increase in appeals after the reform compared to predominately rural counties without broadband internet. They estimate that the increase in applications due to iClaim explained 15 percent of the overall increase in applications between 2008 and 2011.<sup>38</sup> Another study focused on areas with different income levels and internet availability and found similar results.<sup>39</sup> These results are encouraging for those with access to technology, but shows how the use of AI in certain cases can perpetuate existing inequities. While access to broadband has certainly increased since 2008, an area of current concern is access to devices.

AI applications that have disparate impact across individuals can have different implications depending upon the context. For example, if the AI application is helping select individuals for a fixed number of jobs in a company or selecting who should receive a hard to obtain medical treatment, it is unacceptable for an AI application to have a disparate impact in the selection due to bias.<sup>40</sup> In the case of a disability eligibility determination by SSA, an application that helps accelerate a decision for one claimant does not necessarily adversely affect another claimant. Eligibility for SSDI and SSI disability benefits are legal entitlements for everyone who meets the eligibility criteria. Equity considerations are introduced when claimants with hard to adjudicate cases are allowed to languish for many months or years. The longstanding underfunding of SSA and the historically high backlog of disability claims already makes this a serious equity issue.<sup>41</sup> To not exacerbate existing inequities, care and oversight is needed in how SSA selects which populations to target for its AI technology research and development.

## Other risks:

- **Lack of transparency or explainability.** The most advanced forms of AI, built from large language models and often referred to as generative AI, frequently function as a “black box” with even the experts who built the applications unable to explain why the tool produced its specific results. A lack of explainability presents obvious and serious problems if a tool operates in the domain of government benefit determinations. This same issue applies to much simpler automation because of the use of private vendors and trade secrets claims. In this case, even if the technology itself would be possible to understand, the public would not be able to access information about how the AI is making a decision.

- **Increased burden on beneficiaries.** A poorly tested AI application could put unnecessary burden on beneficiaries and exacerbate erroneous overpayment determination. SSA already uses pattern recognition software in its fraud detection efforts to spot irregularities. These tools can be useful unless they produce too many false positives that unnecessarily burden claimants, and worse, result in an erroneous improper payment. With enhanced AI tools, SSA has the ability to identify more potential cases of improper payments at lower cost to the agency. One example of this risk is SSA’s experience using private databases of property ownership to identify potential excess assets among SSI beneficiaries, an effort that resulted in many false positives.<sup>42</sup> Quality controls need to be in place so these efforts do not unnecessarily shift burden to beneficiaries to resolve erroneous issues. Further concerns are raised about redress, even if applicants are able to challenge an automated determination. Appeals processes can be time-consuming and inaccessible. Appeals are the only way erroneous decisions get flagged, and not everyone is able to successfully appeal.
- **Data privacy.** The most advanced AI programs require large amounts of data for training. Protecting data privacy while also ensuring the quality of data AI programs use is a widely recognized challenge. The International Social Security Association identified data privacy and quality as one of the greatest challenges facing Social Security agencies across the world.<sup>43</sup>

## Bias in the Existing Social Security Disability System

A discussion of the risks of bias in AI tools would not be complete without acknowledging that bias already exists in the Social Security disability system, including both decades of program history and human history more broadly. The legislative history of the Social Security programs contains statutory provisions that may represent bias. Large portions of the workforce were initially excluded from the program with workers in those sectors disproportionately poor and people of color, legislation providing access to certain benefits, such as access to SSI ABLE savings accounts, are arbitrarily limited by the age of disability onset.

SSA has not collected data on race and ethnicity for many decades, creating a serious obstacle to assessing bias in these areas.<sup>44</sup> Similarly, many obstacles exist to collecting baseline data on types of disabilities.<sup>45</sup>

As discussed earlier, the Social Security Act directs SSA to use state disability determination service (DDS) offices to make initial disability determinations. The Social Security Advisory Board (SSAB) has documented wide variations in practices across DDS offices, such as in the frequency of procuring consultative exams.<sup>46</sup> Commissioner O'Malley's SecurityStat performance reports includes a chart that shows the enormous variation in processing times for initial disability determinations.<sup>47</sup> Research has found that most of the geographic variation in SSDI and SSI participation rates can be explained by underlying disability prevalence rates and variation in socio-economic characteristics.<sup>48</sup> SSA's reviews of DDS allowances generally find high accuracy rates across DDS. However, the SSA review is based on the file the DDS adjudicator has assembled and it is not known if variations in practices across states because of bias influence how a determination file is assembled.

Weaknesses in current Social Security disability policies, processes, and funding can be perpetuated with new AI applications. In its AI risk framework, NIST states "Decisions that go into the design, development, deployment, evaluation, and use of AI systems reflect systemic and human cognitive biases. AI actors bring their cognitive biases, both individual and group, into the process. Biases can stem from end-user decision-making tasks and be introduced across the AI lifecycle via human assumptions, expectations, and decisions during design and modeling tasks."<sup>49</sup>

Moving forward, SSA should mitigate the bias that exists in its systems by conducting further research and working with communities when creating AI applications that could be rights-impacting. As previously mentioned, AI systems heavily rely on existing data sets. This poses concern since SSA has not collected data on race, ethnicity, and disability for decades. Further research and data collection is required to ensure that the AI systems SSA is using or will use in the future are properly informed.

## **Opportunities to Improve the Disability Determination Process Using AI Tools**

Given the range of potential risks associated with the use of AI, why should SSA or other similar agencies explore the use of AI? Experts from NIH and elsewhere note that SSA "Adjudicators manually review hundreds of evidence pages per case to determine eligibility based on financial, medical, and functional criteria. Natural Language Processing (NLP) technology is uniquely suited to support this adjudication work."<sup>50</sup> Below we discuss some of the ways in which an effective use of AI could improve SSA's operations and service to the public. These areas are not intended to be exhaustive and some of the opportunities at this stage are at best hypothetical. Additionally, implementation of AI systems would make SSA more efficient only in conjunction with proper staffing and resources.

# Improving the quality and timeliness of disability benefit decisions

SSA has the potential to use AI to improve the quality and speed of decisions. SSA's current determination process forces many applicants to repeatedly appeal a denial before ultimately being approved for benefits. At the initial DDS determination, we estimate that 48 percent of initial denials that are appealed are ultimately allowed.<sup>51</sup> In addition, many other applicants who might eventually be allowed drop out of the process.

SSA should focus a substantial portion of its AI development efforts on AI capacity to organize and interpret the evidence in a benefit claim and identify claims with a high probability of being eventually allowed. In some cases, SSA's AI tools could identify claims that would under prior processes be denied but if examined more closely-- or supplemented with further medical evidence-- could be allowed during the initial review. In other countries, AI tools are being evaluated to see if they can improve the targeting of social supports for at risk families.<sup>52</sup> In this manner, AI tools could go much further than existing programs like Quick Disability Determinations.

Improving the quality and timeliness of initial determinations may increase the cost of the initial DDS review. However, identifying claims that will ultimately be allowed can lower appeal rates and reduce the cost of later adjudications, especially at the costly administrative law judge (ALJ) hearing level.

SSA should make one of the goals of its AI efforts to reduce bias in the disability determination process. SSA should not be satisfied with ensuring new AI tools do not perpetuate or exacerbate existing bias. SSA should go farther and set a goal of using AI tools to reduce bias. Looking at AI applications across a wide range of use cases, some experts identify an opportunity to reduce bias by analyzing the AI tools: "Algorithms by their nature require a far greater level of specificity than is usually involved with human decision making, which in some sense is the ultimate "black box." With the right legal and regulatory systems in place, algorithms can serve as something akin to a Geiger counter that makes it easier to detect—and hence prevent— discrimination."<sup>53</sup>

## Improving productivity

Many discussions of the potential of AI focus on improved worker productivity. Earlier we described the challenges SSA faces with historically large backlogs in disability claims. While we are intrigued by the productivity potential for AI in this area, we conclude it is premature to assume AI tools will increase worker productivity. Projections show that even with an increase of productivity in general, the backlog will grow substantially unless agency funding and staff are increased.<sup>54</sup> Prominent labor economists have reached conflicting decisions with some, including David Autor, expressing optimism<sup>55</sup> while others, including Daron Acemoglu, estimating the likely amount of productivity gain will be very modest.<sup>56</sup>

We are uncertain how much AI can improve the efficiency of SSA’s operations, especially if SSA prioritizes improving the quality and fairness of disability determinations as core objectives of its AI initiatives. In the near-term investments in AI development should not remove the urgency of providing SSA essential administrative funding to address very serious backlogs.

## Identifying key gaps in the evidence associated with a claim and providing a better explanation of a disability determination to a claimant

AI tools could help highlight gaps in the medical evidence that, if addressed, could result in a benefit claim being approved earlier in the disability determination process. In parallel with that, AI tools may enable disability adjudicators to better explain the rationale for denying a claim. The medical profession is already using AI tools to summarize medical diagnoses for patients.<sup>57</sup> If an AI tool was able to improve the explanation of denials—such as by highlighting missing medical documentation—it may better equip claimants and their advocates better decide whether or how to appeal a decision.

## Reducing burden on claimants

SSA can use AI tools to make the benefit application and review process less burdensome on claimants. This goal can begin with improved interfaces to collect claimant information, especially medical evidence. SSA has already made progress in this area with its Health Information Technology processes.

## Improving the allocation of reviews of beneficiary’s status

AI tools could help improve the use, quality, and timing of medical and work CDRs and SSI redeterminations. Experts analyzing SSA’s annual report to Congress estimate that 71.6 percent of initial CDR cessations of disabled worker benefits in FY 2015 that were appealed were overturned at reconsideration.<sup>58</sup> In addition, by analyzing medical evidence and reviewing earnings reports on a real time basis, an AI tool could help improve the timing and efficiency of these reviews by SSA and DDS staff. The tool could also provide beneficiaries with more customized information, such as about work incentive rules. AI tools could also provide SSA with a much better capacity to help beneficiaries navigate their participation in the SSDI and SSI programs.

# Executive Orders on AI

The Trump Administration issued the first Executive Order on artificial intelligence on February 19, 2019. Executive Order 13859, *Maintaining American Leadership in Artificial Intelligence*, focused primarily on supporting AI research and development in the United States.<sup>59</sup> It set forth requirements for executive branch agencies, including promoting a responsible approach to AI that would improve outcomes for workers and users of new AI technologies.

The Biden Administration published four documents providing guidance to inform how agencies use AI. Although these documents were repealed immediately by the Trump Administration, together they provided a helpful framework for safer AI implementation. They included an overarching executive order to guide national policy,<sup>60</sup> two more detailed OMB memorandums to federal agencies on use of AI<sup>61</sup> and responsible acquisition,<sup>62</sup> and a related executive order that addresses the nexus between advancing equity and developing, acquiring and using AI and other emerging technology.<sup>63</sup>

The OMB AI memorandum provides useful guidance on agency AI governance, including designating a Chief AI Officer (CAIO) and establishment of an agency AI Governance Board. While the memo largely focuses on internal agency governance it does require agencies to publish AI Use Case Inventories on the agency's website. The inventories for SSA and other agencies can be found [here](#). The OMB memo provides a strong foundation for an AI governance structure. Later in this paper we will discuss further agency governance in the context of rights-impacting AI applications. The memo defines rights-impacting AI as AI “whose output serves as a principal basis for a decision or action concerning a specific individual or entity that has a legal, material, binding, or similarly significant effect on that individual's or entity's:

1. Civil rights, civil liberties, or privacy, including but not limited to freedom of speech, voting, human autonomy, and protections from discrimination, excessive punishment, and unlawful surveillance;
2. Equal opportunities, including equitable access to education, housing, insurance, credit, employment, and other programs where civil rights and equal opportunity protections apply; or
3. Access to or the ability to apply for critical government resources or services, including healthcare, financial services, public housing, social services, transportation, and essential goods and services.”



In addition to the White House and OMB guidance the National Institutes of Standards and Technology (NIST) issued an AI risk management framework that we will use to inform the findings in this report. For rights-impacting AI, the OMB memorandum identified a set of minimum risk management practices, including:

- Complete an AI impact assessment showing “The intended purpose for the AI and its expected benefit, supported by specific metrics or qualitative analysis. Metrics should be quantifiable measures of positive outcomes for the agency’s mission...and it should demonstrate that AI is better suited to accomplish the relevant task as compared to alternative strategies.”
- “The potential risks of using AI, as well as what, if any, additional mitigation measures, beyond these minimum practices, the agency will take to help reduce these risks.... Agencies should be especially attentive to the potential risks to underserved communities...”
- “The quality and appropriateness of the relevant data. Agencies must assess the quality of the data used in the AI’s design, development, training, testing, and operation and its fitness to the AI’s intended purpose....At a minimum, agencies must document: a. the data collection and preparation process, which must also include the provenance of any data used to train, fine-tune, or operate the AI; b. the quality and representativeness of the data for its intended purpose...”

The Trump Administration has most recently issued the executive order, “Removing Barriers to American Leadership in Artificial Intelligence.”<sup>64</sup> This order repeals the Biden Administration executive orders and underscores the Trump Administration’s continued strong interest in the AI sector. It is unclear what additional guidance may be forthcoming from the current Administration on AI and its use in rights-impacting programs, such as those administered by SSA. In the following sections we discuss principles and guardrails to guide development and use of AI in SSI and SSDI in ways that emphasize caution and sensitivity in order to prioritize fairness.

## Principles and Guardrails

In the discussion below we build on that framework to describe in the specific context of SSA’s use of rights-impacting AI in the disability determination process. As an independent, external Task Force we understand that not only is the future use of AI in SSA inevitable, but it is already happening. We offer principles and guardrails below in hopes they help the Administration build a strong foundation for implementing future AI applications.

The principles and guardrails described below are not exhaustive. They are focused on issues of particular importance for rights-impacting AI applications. The expectation in this focus is that SSA will have already established that an AI application is a good value proposition for the agency and has demonstrated its useability and safety.

# Principles

Below we identify principles that can help guide SSA's development of AI tools. These principles have been set in hopes to improve fairness and center the rights of beneficiaries as SSA moves towards new AI tools. Importantly, these principles and guardrails also seek to increase trust among the public and beneficiaries around use of AI tools within SSA.

## Do not use AI to limit or impede the rights of applicants and beneficiaries

To prevent and protect against harm, uses of AI or other emerging technologies should be limited to applications that will support and protect the rights of applicants and beneficiaries. This includes not using AI or other emerging technology as the ultimate authority upon which to deny benefits to an individual. This requires a commitment by SSA and others who monitor the use of AI in Social Security and SSI programs to continually evaluate how AI is being deployed and used within the agency. Adhering to this principle will entail challenges SSA should anticipate.

For example, AI applications that help a disability examiner consider an applicant's claim may only identify evidence indicating an allowance. However, the absence of evidence indicating an allowance could lead a disability examiner to infer that a claim should be denied. It is not known how many claims are denied because of the absence of positive evidence. Using AI to assist in continuing disability reviews (CDRs) presents another set of challenges. For a CDR, SSA must demonstrate medical improvement for a beneficiary to lose eligibility, which provides a potential limit on how the AI tool is used.

## Prioritize human decision making in disability determinations

AI tools will at best augment human decision making when rights-impacting benefit decisions are made. While certain routine administrative processes could be automated, benefit decisions cannot be fully automated without jeopardizing the rights of a claimant. And though this is a key principle, it is not without limitations. A NIST report cautions to not overestimate how much a human oversight of an AI application can mitigate bias and erroneous decisions: "The fundamental assumption of such structures is that a human overseer, simply by virtue of being human, will be able to provide adequate governance for systems. The reality however is that without significant procedural and cultural support, optimistic expectations about how humans are able to serve in this administrative capacity are not borne out in practice."<sup>65</sup>

## Prevent, reduce and mitigate bias

While identifying bias in AI applications is difficult, SSA can learn from existing research — such in the health care system and directly from affected populations — on baseline levels of bias and the potential for AI tools to perpetuate bias.<sup>66</sup> SSA should affirm a goal that the use of AI tools will reduce the extent of bias in the determination process. As we discussed earlier, researchers are developing AI tools that are specifically designed to identify existing types of bias.

## Improve fairness and equity

SSA's use of AI should continually seek to improve overall fairness and equity. There are certain biases and inequities in the existing systems and processes used by SSA in administering disability programs. Introducing the use of AI within these existing systems will alter the types of biases and the extent to which they impact individuals and groups. In some cases, AI could be used to improve outcomes for particular groups, such as in the case of compassionate allowances. As discussed under risks, this disparate treatment of applicants can be justified in certain circumstances given the nature of an applicant's disability. However, SSA must affirm that improvements that impact select groups also contribute to improved outcomes for all. This pre-supposes, at a minimum, maintenance of resources for staffing. For example, if expediting review and approval of claims for certain conditions improves outcomes for both affected individuals and speeds up overall processing time for other cases, SSA may conclude that it increases equity in the administration of the program overall. Or it may choose to deploy such changes as part of a package that improves fairness and equity overall.

## Guardrails

The guardrails we propose would provide protections against the risks discussed earlier with the goal of preventing potential harms and ensuring responsible use of AI as new uses are introduced and to guide ongoing monitoring as AI applications evolve.

## Establish internal and external governance processes for oversight of AI

For new AI initiatives, in addition to existing review from committees and vetting with the disability advocacy community, SSA needs a governance process for procuring and implementing AI that involves: beneficiaries, benefits examiners and their respective representatives, including disability advocacy organizations and labor organizations. Ensuring accountability and trust in developing rights-impacting AI also requires an external governance oversight structure. We describe several governance models that could be used, in keeping in mind constraints such as the Federal Advisory Committee Act.

- Designate an AI Ethics Officer at SSA, similar to what many major corporations have done in recent years. This could be done as a standalone role at SSA and/or could be implemented in conjunction with establishing a similar governance role at OMB.
- Have the National Academies of Sciences (NAS) convene an advisory group with members appointed by rights-impacting agencies, including SSA, HHS and USDA, and by NIST. Such a group could, on an ongoing basis, address both agency specific and cross-cutting issues relating to rights-impacting AI applications.
- Have an independent outside group evaluate agency AI applications. Such an independent group could have similarities to the National Academy of Social Insurance's AI Task Force, which could provide transitional oversight function until the Federal government establishes an oversight function, or could serve as an ongoing resource.

The governance structure must acknowledge and give transparency to the substantial role that contractors can play in developing AI and other emerging technology.

A strong governance structure is much more than having processes for user engagement and research, as important as those features are. Governance must also include proper staff training when it comes to using AI tools. Without significant procedural and cultural support, optimistic expectations about how humans are able to serve in this administrative capacity are not borne out in practice. Staff must be confident in their ability to be the human in the loop. This will also lead to increased trust among the public for SSA's process.

## Make tools that are explainable, transparent, testable and subject to independent evaluation

SSA should clearly explain to the public how it is using AI in the review of an applicant's benefit application. SSA has already conducted forums, provided a list of AI applications consistent with an OMB directive and responded to Congressional inquiries. In communicating with benefit applicants, SSA can explain when AI is being used and describe how it is being used on its website, including the disability-specific datasets it is using to train generative AI tools.

As we discussed earlier, some AI tools have the risk of not being explainable or transparent. In some cases, the companies that create the tools and the agencies or groups that procure and use them may not be transparent about their impact towards those affected. In the CHAI criteria "testability refers to the extent to which an AI algorithm's performance can be verified as satisfactory in terms of meeting all standards for trustworthy AI, including ... robustness, safety, bias mitigation, fairness, and equity in both development and evaluation." Confidence that these standards are being met with independent evaluation and validation is a further aspect of this guardrail.

## Strengthen procurement policies

CDT notes that many federal agencies are currently making decisions about how to “procure AI-driven technology from third parties.”<sup>67</sup> Additionally, CDT cautions that limitations within existing federal procurement processes threaten to further impede the responsible procurement of AI. These include difficulties around understanding and evaluating bias, incorporating human oversight and intervention, and defining and implementing a process for redress.

While agency contractors should be observing all the principles and guardrails we discuss, one of the most important is transparency. For rights-impacting AI applications, the basis for the support an AI application provides Federal staff should be clear.<sup>68</sup>

## Conduct ongoing monitoring and evaluation of AI applications

Many AI tools fail to achieve their objectives.<sup>69</sup> SSA and all government agencies need to carefully audit and independently evaluate the performance of AI tools and for whether they avoid the risks discussed earlier. In addition, while AI evolves rapidly, the new technology can take many years to be fully incorporated into program operations. As it is incorporated, it can gradually shape how program staff perform their duties. SSA needs to establish an ongoing evaluation strategy capable of assessing continual systems change. As part of this, SSA staff should evaluate if the algorithmic systems and AI tools are accomplishing what they are designed to do and if they pose significant risk of bias. This strategy should address both known risks and new issues, and establish metrics for evaluation. The strategy should include simplifying the process for appeals and looking for other feedback loops to identify erroneous determinations. SSA should monitor the frequency of algorithmic bias among human adjudicators and confirm that SSA and DDS staff consistently exercise independent judgement. SSA should also ensure that claims that are not as efficiently handled by AI systems are still processed expeditiously.

## Protect data privacy

Disability benefit decisions require access to a large amount of sensitive and private information. At the same time, machine learning and generative AI tools require access to substantial relevant data to learn and become more useful. Before using sensitive health records to train AI tools, safeguards must be in place to protect privacy and keep data secure. SSA needs to provide clarity on how medical records are shared and protected. Furthermore, they should integrate principles of data minimization and purpose limitation—to collect only the data necessary for the tool to function and use it solely for its intended purpose—into their adoption of any AI or algorithmic tools. Further, in an ideal world, agencies would also delete any data that is not needed anymore after an agreed-upon period.

# Conclusion

SSDI and SSI benefits have been a vital piece of the U.S. social insurance ecosystem for more than half a century. These programs are key to the economic security and wellbeing of millions of American workers and families—mitigating the effects of poverty and hardship for more than 7 million Americans with disabilities and older adults, including nearly 1 million disabled children.

This report outlines key risks and opportunities confronting the Social Security Administration as it considers the use of AI and emerging technologies in administration of these vital disability programs. Any implementation of AI needs to be approached with the utmost care in recognition of the potential to change or impact Americans' access to benefits. It provides principles and guardrails to guide future use of AI and emerging technologies in Social Security programs and serve as a foundation for future development of more detailed policies and procedures. Looking ahead, more work is also needed to ensure that the potential benefits and costs stemming from the use of AI in SSA programs are fully understood and properly evaluated and that first and foremost, applicant and beneficiary rights are protected.

As the use of AI technology rapidly evolves—under the current Administration and beyond—we urge stakeholders across government, the private and non-profit sectors, and academia to engage in ongoing monitoring of the opportunities and challenges identified in this report and to build upon the framework presented here.



# Endnotes

1. This Task Force first convened in June 2024 following public announcements that the Social Security Administration was increasingly utilizing artificial intelligence and other emerging technology in its disability determinations process.
2. “Chart Book: Social Security Disability Insurance,” (Center on Budget and Policy Priorities, August 6, 2024), <https://www.cbpp.org/research/social-security/social-security-disability-insurance-0>.
3. Chantel Boyens and Jack Smalligan, “Social Security Customer Service Challenges: Causes, Impacts, and Solutions,” AARP, July 23, 2024.
4. National Institute of Standards and Technology, the U.S. Department of Commerce, “Artificial Intelligence Risk Management Framework,” January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
5. Angelina Wang et al., “Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms That Optimize Predictive Accuracy,” *ACM Journal on Responsible Computing* 1, no. 1 (March 20, 2024). <https://dl.acm.org/doi/10.1145/3636509>.
6. Kevin De Liban, “Inescapable AI: The Ways AI Decides How Low-income People Work, Live, Learn, and Survive” (TechTonic Justice, November 2024), 15.
7. Jennifer Pahlka, “AI Meets the Cascade of Rigidity,” *Digitalist Papers*, <https://www.digitalistpapers.com/essays/ai-meets-the-cascade-of-rigidity>.
8. During the height of the COVID-19 Pandemic, SSA continued to roll out the implementation of a new disability determination system, DCPS. For more information: <https://oig.ssa.gov/assets/uploads/a-01-20-50963.pdf>
9. Chantel Boyens and Jack Smalligan, “Social Security Customer Service Challenges: Causes, Impacts, and Solutions” (The Urban Institute, July 30, 2024), <https://www.urban.org/research/publication/social-security-customer-service-challenges-causes-impacts-and-solutions>.
10. U. S. Government Accountability Office, “Social Security Disability: Additional Measures and Evaluation Needed to Enhance Accuracy and Consistency of Hearings Decisions,” December 7, 2017, <https://www.gao.gov/products/gao-18-37>.
11. Jeffrey Lubbers and Gerald Ray, “A Government Success Story: How Data Analysis by the Social Security Appeals Council (with a Push from the Administrative Conference of the United States) Is Transforming Social Security Disability Adjudication,” *The George Washington Law Review* 83, no. 4/5 (January 1, 2015), [https://digitalcommons.wcl.american.edu/facsch\\_lawrev/1092](https://digitalcommons.wcl.american.edu/facsch_lawrev/1092).
12. Kevin Agatstein and Melissa Crocker, “More Patient Data? Be Careful What You Wish for...AI’s Role in Making Clinical Data Exchange Useful,” *Population Health Management* 27, no. 4 (August 1, 2024), <https://doi.org/10.1089/pop.2024.0049>.
13. Ibid.
14. Jackson Steinkamp, Jacob J. Kantrowitz, and Subha Airan-Javia, “Prevalence and Sources of Duplicate Information in the Electronic Medical Record,” *JAMA Network Open* 5, no. 9 (September 26, 2022): <https://doi.org/10.1001/jamanetworkopen.2022.33348>.

15. Office of the Inspector General, “The Social Security Administration’s Expansion of Health Information Technology to Obtain and Analyze Medical Records for Disability Claims” (Social Security Agency, January 3, 2022), <https://www.oversight.gov/sites/default/files/oig-reports/SSA/01-18-50342.pdf>.
16. Ed Leary, “Expert Systems at the Social Security Administration,” *Journal of Policy Analysis and Management* 8, no. 2 (1989), <https://doi.org/10.2307/3323378>.
17. See: “Administrative Review Process for Adjudicating Initial Disability Claims,” *Federal Register*, March 31, 2006, <https://www.federalregister.gov/documents/2006/03/31/06-3011/administrative-review-process-for-adjudicating-initial-disability-claims>. (To be codified at 20 C.F.R. §§ 404, 405, 416, 422).
18. Office of the Inspector General, “The Social Security Administration’s Use of Insight Software to Identify Potential Anomalies in Hearing Decisions” (Social Security Agency, April 19, 2019), <https://oig-files.ssa.gov/audits/full/A-12-18-50353.pdf>.
19. David Freeman Engstrom et al., “Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies,” NYU School of Law, Public Law Research Paper, no. 20–54 (February 1, 2020), <https://doi.org/10.2139/ssrn.3551505>.
20. Kurt Glaze et al., “Artificial Intelligence for Adjudication: The Social Security Administration and AI Governance,” Oxford University Press, *Handbook on AI Governance* (Forthcoming), August 18, 2021, <https://doi.org/10.2139/ssrn.3935950>.
21. Dorsey Holliman, Brooke, Michelle Stransky, Nathalie Dieujuste, and Megan Morris. “Disability Doesn’t Discriminate: Health Inequities at the Intersection of Race and Disability.” *Frontiers in Rehabilitation Sciences* 4 (July 6, 2023): 1075775. <https://doi.org/10.3389/fresc.2023.1075775>.
22. Ziad Obermeyer et al., “Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations,” *Science* 366, no. 6464 (October 25, 2019), <https://doi.org/10.1126/science.aax2342>. See also: Michael P. Cary et al., “Empowering Nurses to Champion Health Equity & BE FAIR: Bias Elimination for Fair and Responsible AI in Healthcare,” *Journal of Nursing Scholarship: An Official Publication of Sigma Theta Tau International Honor Society of Nursing*, July 29, 2024, <https://doi.org/10.1111/jnu.13007>.
23. Jiageng Wu et al., “Clinical Text Datasets for Medical Artificial Intelligence and Large Language Models — A Systematic Review,” *NEJM AI* 1, no. 6 (May 23, 2024), <https://doi.org/10.1056/AIra2400012>. See also: Stephanie S. Gervasi et al., “The Potential for Bias in Machine Learning And Opportunities For Health Insurers To Address It,” *Health Affairs* 41, no. 2 (February 2022), <https://doi.org/10.1377/hlthaff.2021.01287>.
24. Michele Samorani et al., “Overbooked and Overlooked: Machine Learning and Racial Bias in Medical Appointment Scheduling,” *SSRN Electronic Journal*, 2019, <https://doi.org/10.2139/ssrn.3467047>.
25. Amanda Hermans, Sarah Morriss, and Susan J. Popkin. “An Opportunity for the Census Bureau to More Accurately Estimate the Disabled Population in the US,” February 12, 2024. <https://www.urban.org/research/publication/opportunity-census-bureau-more-accurately-estimate-disabled-population-us>.

26. Jon Kleinberg et al., “Algorithms as Discrimination Detectors,” *Proceedings of the National Academy of Sciences* 117, no. 48 (December 2020): 30096–100, <https://doi.org/10.1073/pnas.1912790117>.
27. The National Institute of Standards & Technology, “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” March 15, 2022, <https://www.nist.gov/publications/towards-standard-identifying-and-managing-bias-artificial-intelligence>.
28. Ibid.
29. Ibid.
30. Coalition for Health AI, “Blueprint for Trustworthy AI Implementation Guidance and Assurance for Healthcare” (Duke University, April 4, 2023), [https://www.coalitionforhealthai.org/papers/blueprint-for-trustworthy-ai\\_V1.0.pdf](https://www.coalitionforhealthai.org/papers/blueprint-for-trustworthy-ai_V1.0.pdf).
31. Ian Moura, “Encoding Normative Ethics: On Algorithmic Bias and Disability,” February 2023, <https://doi.org/10.5210/fm.v28i1.12905>.
32. The National Institute of Standards & Technology, “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” March 2022, <https://www.nist.gov/publications/towards-standard-identifying-and-managing-bias-artificial-intelligence>.
33. Lydia X. Z. Brown et al., “Challenging the Use of Algorithm-Driven Decision-Making in Benefits Determinations Affecting People with Disabilities” (Center for Democracy and Technology, October 21, 2020), <https://cdt.org/wp-content/uploads/2020/10/2020-10-21-Challenging-the-Use-of-Algorithm-driven-Decision-making-in-Benefits-Determinations-Affecting-People-with-Disabilities.pdf>; Ariana Aboulafia and Henry Claypool, “Building a Disability-Inclusive AI Ecosystem: A Cross-Disability, Cross-Systems Analysis Of Best Practices,” (Center for Democracy and Technology, March 2025), <https://cdt.org/wp-content/uploads/2025/03/2025-03-11-CDT-Building-A-Disability-Inclusive-AI-Ecosystem-report-final.pdf>
34. Ibid.
35. “Program Operations Manual System (POMS),” Social Security Administration, <https://secure.ssa.gov/poms.nsf/lnx/0423535010>.
36. Sarah Prenovitz et al., “An Investigation on Bias in the Quick Disability Determinations Model” (Abt Associates, August 8, 2023), <https://www.ssa.gov/disabilityresearch/documents/QDD%20Model%20-%20Evaluation,%20Analysis%20and%20Options%20Final%20Report.pdf>.
37. Ayah Zirikly et al., “Information Extraction Framework for Disability Determination Using a Mental Functioning Use-Case,” *JMIR Medical Informatics* 10, no. 3 (March 18, 2022), <https://doi.org/10.2196/32245>.
38. Andrew Foote, Michel Grosz, and Stephanie Rennane, “The Effect of Lower Transaction Costs on Social Security Disability Insurance Application Rates and Participation,” *Journal of Policy Analysis and Management* 38, no. 1 (January 2019): 99–123, <https://doi.org/10.1002/pam.22095>.

39. David Powell and George Zuo, "The Impact of Affordable and Accessible Broadband on SSDI and SSI Participation," Working Paper (Ann Arbor, MI.: University of Michigan Retirement and Disability Research Center (MRDRC), September 2023), <https://mrdrc.isr.umich.edu/publications/papers/pdf/wp467.pdf>.
40. "Civil Rights Principles for Hiring Assessment Technologies," The Leadership Conference on Civil and Human Rights, July 2020, <https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/>.
41. Chantel Boyens and Jack Smalligan, "Social Security Customer Service Challenges: Causes, Impacts, and Solutions," Urban Institute, July 30, 2024, <https://www.urban.org/research/publication/social-security-customer-service-challenges-causes-impacts-and-solutions>.
42. Sarah Mancini, Kate Lang, and Chi Chi Wu, "Mismatched and Mistaken: How the Use of an Inaccurate Private Database Results in SSI Recipients Unjustly Losing Benefits" (National Consumer Law Center, April 2021), <https://www.nclc.org/wp-content/uploads/2022/08/RptMismatchedFINAL041421.pdf>.
43. Moinul Zaber, Oxana Casu, and Ernesto Brodersohn, "Artificial Intelligence in Social Security Organizations" (Geneva: International Social Security Association and the United Nations University, 2024), <https://unu.edu/sites/default/files/2024-06/2-AI%20in%20SecSoc%202024.pdf>.
44. Patricia P. Martin, "Why Researchers Now Rely on Surveys for Race Data on OASDI and SSI Programs: A Comparison of Four Major Surveys," Research and Statistics Notes, January 2016, <https://www.ssa.gov/policy/docs/rsnotes/rsn2016-01.html>.
45. Ariana Aboulafia, Miranda Bogen, and Bonnielin Swenor, "To Reduce Disability Bias in Technology, Start With Disability Data" (Center for Democracy and Technology, July 2024), <https://cdt.org/wp-content/uploads/2024/07/2024-07-23-Data-Disability-report-final.pdf>.
46. DDSs do not have equal access to providers to perform consultative exams. Some factors that drive this discrepancy include the low pay rate for medical providers who perform consultative exams.
47. "Disability Determination Processing Time," Social Security Administration, October 4, 2024, <https://www.ssa.gov/securitystat/disability-processing-time>.
48. John Gettens, Pei-Pei Lei, and Alexis D. Henry, "Accounting for Geographic Variation in Social Security Disability Program Participation," Social Security Bulletin 78, no. 2 (2018), <https://www.ssa.gov/policy/docs/ssb/v78n2/v78n2p29.html>.; Kalman Rupp, "Factors Affecting Initial Disability Allowance Rates for the Disability Insurance and Supplemental Security Income Programs: The Role of the Demographic and Diagnostic Composition of Applicants and Local Labor Market Conditions," Social Security Bulletin 72, no. 4 (2012), <https://www.ssa.gov/policy/docs/ssb/v72n4/v72n4p11.html>.

49. The National Institute of Standards & Technology, “Artificial Intelligence Risk Management Framework,” January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
50. Development of Natural Language Processing Tools to Support Determination of Federal Disability Benefits in the U.S. 2020. Bart Desmet, Julia Porcino, Ayah Zirikly, Denis Newman-Griffis, Guy Divita, and Elizabeth Rasch <https://eprints.whiterose.ac.uk/196497/1/2020.lt4gov-1.1.pdf>
51. Authors estimate using Table 3.23 in FY 2024 SSA Congressional Justification: “Justification of Estimates for Appropriations Committees Fiscal Year 2024” (Social Security Administration, March 2023), <https://www.ssa.gov/budget/assets/materials/2024/FY24-JEAC.pdf>.
52. Dario Sansone and Anna Zhu, “Using Machine Learning to Create an Early Warning System for Welfare Recipients,” *Oxford Bulletin of Economics and Statistics* 85, no. 5 (October 2023): 959–92, <https://doi.org/10.1111/obes.12550>.
53. Kleinberg et al., “Algorithms as Discrimination Detectors.”
54. Stephen C. Goss and Michael L. Stephens, “Social Security Disability Claims Pending Determination: Past and Projected” (Social Security Administration, May 2024).
55. David Autor, “Applying AI to Rebuild Middle Class Jobs” (National Bureau of Economic Research, February 2024), <https://doi.org/10.3386/w32140>.
56. Daron Acemoglu, “The Simple Macroeconomics of AI,” Massachusetts Institute of Technology, May 12, 2024, <https://economics.mit.edu/sites/default/files/202405/The%20Simple%20Macroeconomics%20of%20AI.pdf>.
57. Junaid Bajwa et al., “Artificial Intelligence in Healthcare: Transforming the Practice of Medicine,” *Future Healthcare Journal* 8, no. 2 (July 2021): e188–94, <https://doi.org/10.7861/fhj.2021-0095>.
58. **See:** <https://nosscr.org/wp-content/uploads/2020/01/NOSSCR-comments-continuing-disability-reviews-FINAL.pdf>
59. “Maintaining American Leadership in Artificial Intelligence (February 19, 2019), <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>
60. “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” (October 30, 2023), <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

61. Office of Management and Budget, Memorandum for the Heads of Executive Departments and Agencies on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (March 28, 2024). <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>
62. **See:** <https://www.whitehouse.gov/wp-content/uploads/2024/10/M-24-18-AI-Acquisition-Memorandum.pdf>
63. “Executive Order on Further Advancing Racial Equity and Support for Underserved Communities Through The Federal Government,” (February 16, 2023), <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/02/16/executive-order-on-further-advancing-racial-equity-and-support-for-underserved-communities-through-the-federal-government/>.
64. “Removing Barriers to American Leadership in Artificial Intelligence,” January 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>.
65. The National Institute of Standards and Technology, “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence” March 2022, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>.
66. Michael P. Cary et al., “Mitigating Racial and Ethnic Bias And Advancing Health Equity In Clinical Algorithms: A Scoping Review,” *Health Affairs* 42, no. 10 (October 2023): 1359–68, <https://doi.org/10.1377/hlthaff.2023.00553>.
67. Hannah Quay-de la Vallee, Ridhi Shetty, and Elizabeth Laird, “The Federal Government’s Power of the Purse: Enacting Procurement Policies and Practices to Support Responsible AI Use,” Center for Democracy and Technology, April 29, 2024, <https://cdt.org/insights/report-the-federal-governments-power-of-the-purse-enacting-procurement-policies-and-practices-to-support-responsible-ai-use/>.
68. <https://www.whitehouse.gov/wp-content/uploads/2024/10/M-24-18-AI-Acquisition-Memorandum.pdf>
69. Inioluwa Deborah Raji et al., “The Fallacy of AI Functionality,” in 2022 ACM Conference on Fairness, Accountability, and Transparency (ACM Conference on Fairness, Accountability, and Transparency, Seoul Republic of Korea: ACM, 2022), 959–72, <https://doi.org/10.1145/3531146.3533158>.
70. <https://www.ssa.gov/ai/SSA-AI-Inventory%202023.csv>
71. Ibid.
72. <https://www.ssa.gov/ai/policy/SSA%20M-24-10%20Compliance%20Plan.pdf>
73. Ibid.

# Appendix A: How SSA Uses AI and Mitigates Risk

In August 2024, Senator Wyden, then Chair of the Senate Finance Committee, requested information on how SSA—along with other agencies including HHS, DHS, and CBP—is using AI systems. SSA’s letter in response, which can be accessed [here](#), thoroughly examined SSA’s AI policies and systems.

## AI in the Disability Determination and Appeals Processes

When asked what AI systems SSA uses in the Disability Determination process and each system’s success metrics, SSA noted that they “also use the following systems in decision support for disability determinations or appeals.” SSA provides additional context on its AI systems in an AI inventory that the Office of Management and Budget required agencies to publish. The definitions listed below are sourced from [the AI Inventory](#). (see hyperlink).

- **Insight:** Insight is decision support software used by hearings and appeals-level Disability Program adjudicators to help maximize the quality, speed, and consistency of their decision making. Insight analyzes the free text of disability decisions and other case data to offer adjudicators real-time alerts on potential quality issues and case-specific reference information within a web application. It also offers adjudicators a series of interactive tools to help streamline their work. Adjudicators can leverage these features to speed their work and fix issues before the case moves forward (e.g. to another reviewing employee or to the claimant). Insight features are powered by several natural language processing and artificial intelligence packages and techniques.
- **IMAGEN:** IMAGEN is an IT Modernization Disability Analytics & Disability Decision Support (ADDS) Product that will provide new tools and services to visualize, search and more easily identify relevant clinical content in medical records. These tools and services will improve the efficiency and consistency of disability determinations and decisions and provide a foundation for machine-based decisional guidance. IMAGEN will transform text to data and enable disability adjudicators to leverage various machine learning technologies like Natural Language Processing (NLP) and predictive analytics and will support other high-priority agency initiatives such as fraud prevention and detection.



- **QDD:** The Quick Disability Determinations (QDD) process uses a computer-based predictive model to screen initial applications to identify cases where a favorable disability determination is highly likely and medical evidence is readily available. The Agency bases the QDD models predictive scores on historical data from application forms completed by millions of applicants. By identifying QDD cases early in the process, the Social Security Administration can prioritize this workload and expedite case processing. The Agency routinely refines the QDD model to reflect the characteristics of the recent applicant population and optimize its ability to identify strong candidates for expedited processing.
- **PATH Model:** This model uses machine learning to identify cases likely to receive an allowance at the hearing level and refer them to administrative law judges or senior adjudicators for prioritized review.
- **CDR Model:** This model uses machine learning techniques to identify disability cases with the greatest likelihood of medical improvement and flag them for a continuing disability review.”<sup>70</sup>

The SSA letter to Senator Wyden also outlines policy guidance offered to SSA employees who use these systems. For example, in terms of “Insight and IMAGEN – SSA instructs users that both Insight and IMAGEN outputs and feedback provide a starting point for the user’s disability determination analysis. The user holds the ultimate responsibility to ensure the accuracy of any AI output or feedback. Further, the user is trained to evaluate AI output and feedback to ensure policy compliance at every step of the disability determination process.”

## AI and Program Integrity

SSA also uses AI to prevent overpayments, detect fraud, and overall promote program integrity. According to SSA, the majority of their AI usage is applied to these areas. These AI systems include:

- “Anomalous iClaim Predictive Model – A machine learning model that identifies high-risk iClaims. These claims are then sent to Operations for further review before additional action is taken to adjudicate the claims.
- Pre-Effectuation Review / Targeted Denial Review Models – These review models use machine learning to identify cases with greatest likelihood of disability eligibility determination error and refer them for quality review checks.
- CDR Model – This model uses machine learning techniques to identify disability cases with the greatest likelihood of medical improvement and flags them for a continuing disability review.
- SSI Redetermination Model – This model uses machine learning to identify Supplemental Security Income (SSI) cases with highest expected overpayments due to changes in financial eligibility and flag them for technician review.

- Mobile Wage Reporting – This application uses AI to extract text/data from scanned images/documents representing pay stubs or payroll information to enable faster processing.”<sup>71</sup>

## Risk Mitigation

In terms of leadership, SSA created a “Chief Artificial Intelligence Officer who leads and is accountable for our AI efforts, and an AI Senior Executive Council that provides strategic leadership and oversight.” (1) These positions were created in compliance with OMB M-24-10 on “Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence” and EO 14110, “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.” SSA also created a Responsible AI (RAI) Core Team “to serve as the primary review body for all AI-related matters” within the agency.<sup>72</sup>

SSA also has an internal risk management framework called the SSA RAI Implementation Framework. This framework “required artifacts for all use-cases that encompass the information needed to support AI governance evidence, including risk management:

- “The Impact Assessment - provides guidance to teams on how to explore responsible AI challenges and how to minimize the impact of any potential unfairness.
- The Risk Assessment Matrix - captures the specific risks associated with the AI use case.
- The Model Card - documents the AI model usage, provenance, ethics considerations, and evaluation of performance.

In addition, for those use cases identified as Rights-Impacting the agency requires a bias identification and mitigation report which captures methods and results of the bias mitigation efforts.”<sup>73</sup>

Since these actions were taken in response to Executive Orders and guidance issued during the Biden Administration, it is unclear which will be continued. The agency has recently undergone a large-scale reorganization as well, which could further impact SSA’s approach to risk mitigation.

