Research Statement

Alexander Glaser

School of Public and International Affairs and
Department of Mechanical and Aerospace Engineering, Princeton University

July 2024

Over the past decade, my research has centered on technical and policy analysis of nuclear monitoring and verification, nuclear arms control and nonproliferation, and next-generation nuclear energy technologies, including both fission and fusion systems. It bridges the policy-driven agenda of SPIA's Program on Science and Global Security (SGS) and the technically-driven research of our small laboratory, established in the Department of Mechanical and Aerospace Engineering (MAE) in 2016. This work is underpinned by several technical capabilities, including state-of-the-art neutronics calculations and nuclear fuel-cycle simulation and analysis. We have used these tools to do innovative work in all areas of my technical and policy research.

Here, I offer a brief account of my research, with an emphasis on more recent work. It is divided into several broad research areas, each of which combines science, technology, and policy. I explain what motivates my work in each of these research areas, the specific contributions I have made and their importance, and how I expect my research to move forward in the future.

Area I. Nuclear Verification

Nuclear arms control is in crisis, and it is currently difficult to anticipate what future bilateral or multilateral agreements could look like and what their objectives might be. Possible frameworks could include reductions with verified warhead dismantlements, numerical or geographical limits on the total stockpiles of nuclear weapons, or approaches that avoid warhead inspections altogether. Regardless of these possible pathways, future arms-control initiatives and progress toward nuclear disarmament may be difficult and perhaps impossible unless viable technical verification approaches are available to provide policy makers confidence in such agreements. A major area of my technical and policy research aims to contribute to this effort by developing new verification approaches for confirming the authenticity of nuclear warheads and accounting for the fissile materials that make these weapons possible.

My work in this area has attracted policy attention at a high level, and I have given invited briefings at the United Nations Conference on Disarmament (CD) in Geneva and at the United Nations in New York, including to the UN Secretary-General's Advisory

Board on Disarmament Matters on "New Verification Technologies for Arms Control and Disarmament." My work has also attracted attention from researchers at national laboratories in the United States and in other countries. In 2017–2021, verification became a key issue in distinct policy debates. While the Treaty on the Prohibition of Nuclear Weapons (TPNW) was being negotiated, we provided critical input on relevant verification provisions for this new treaty [125] and presented these ideas at the United Nations [126]; later that year, the International Campaign to Abolish Nuclear Weapons (ICAN) would be awarded the Nobel Peace Prize for this effort (↗). Similarly, as a potential denuclearization of North Korea suddenly appeared as a possible path forward, we published in *Science* our assessment of key verification challenges associated with such a process while addressing North Korea's security concerns that would persist along the way [147, 162]. In 2021, I was invited to co-edit a volume on emerging nuclear verification challenges for the German Federal Foreign Office [161].

Much of my research in this area has been funded by foundation and federal-government grants. In particular, I have been the Principal Investigator on two five-year grants awarded by the U.S. Department of Energy, which established consortia bringing together research groups from several U.S. universities and national laboratories. I led one of the three thrust areas of one of the Consortia, and our team received generous funding to support this work over the years.

Zero-knowledge Nuclear Warhead Verification

How can an inspector be convinced that a nuclear weapon is real without learning anything about its design? This is a long-standing challenge in arms-control verification, and has been a central theme of research in my group. To resolve the dilemma of simultaneously certifying and authenticating electronics, I proposed a fundamentally different approach to warhead verification, one that avoids the measurement of sensitive information at the outset as a way to address concerns about its potential leakage. The core of the idea is the concept of "zero-knowledge" proofs, a class of interactive proof systems developed by mathematicians working in cryptography and intended to prove an assertion while yielding nothing beyond the validity of the assertion being proved. Among the many possible strategies to implement zero-knowledge verification, we chose to prioritize and demonstrate a simple but powerful approach that is based on differential neutron radiography measurements combined with preloadable non-electronic (bubble) detectors.

The first major publication featuring this research was the 2014 *Nature* article [88]. The core of the paper was a series of large-scale computer simulations in which I showed that small diversions of heavy metal from a representative test object can be reliably detected with this combination of technologies and concepts. The response of the expert community far exceeded any expectations. Remarkably, this research also has triggered renewed interest in "physical implementations" of zero-knowledge proofs in cryptography, which is typically constrained on the digital world (↗).

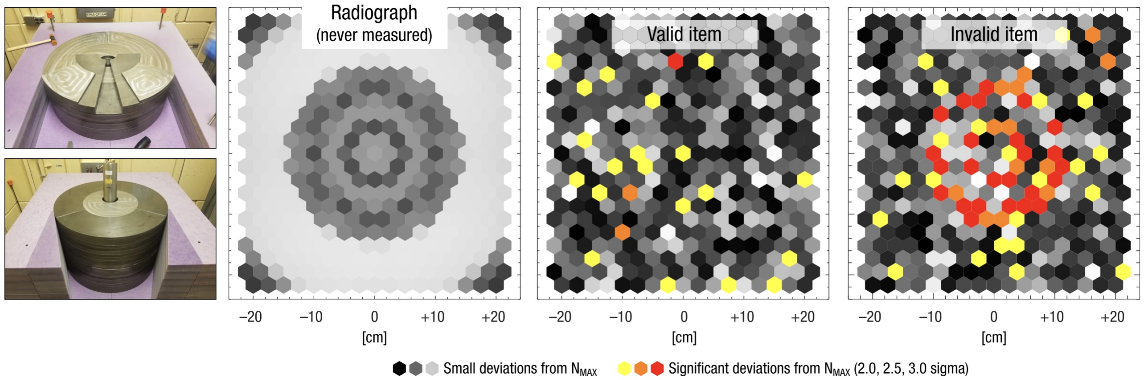

How can an inspector be convinced that a nuclear weapon is real without learning anything about its design? In 2014, we proposed a new approach following a "zero-knowledge protocol" combining neutron radiography with non-electronic preloaded detectors [88]. To demonstrate the viability of the approach for practical applications, we have since designed and built a full-scale experimental device at PPPL (Excalibur, shown on the left), which can be configured and operated in different modes depending on the inspection task [141, 163, 171, 182].

From the beginning, it was clear that we needed experimental data to complement our simulations, especially to examine the robustness of the method in real-world situations, where there may be systematic errors in measurements, small misalignments of equipment, or variations in environmental conditions that are difficult to anticipate with computer simulations. This became a major focus of our work since 2015. The Princeton Plasma Physics Laboratory (PPPL), which has been very supportive of this research gave us laboratory space to conduct experiments with a dedicated DT neutron generator that we acquired thanks to a grant from the U.S. Department of Energy. Over the past couple of years, as part of their PhDs, Sébastien Philippe and Mike Hepler conducted important experimental and conceptual research [121], which contributed to a deeper understanding of the approach and helped optimize the experimental setup, now called "Experiment for Calibration with Uranium" or "Excalibur" for short. Excalibur can be used in two distinct configurations: a collimated mode that produces a relatively narrow beam of 14 MeV neutrons for radiography, and a moderated mode that produces a uniform broad flux of sub-MeV neutrons to induce fission events in fissile isotopes, which are expected in a typical item offered for inspection [171].

We have now established a robust experimental workflow at PPPL and conduct measurement campaigns every other week. As part of this effort, we have also further characterized bubble detector technology, which is so critical for this use case [146, 172]. Already in 2022, we transitioned from one-dimensional (linear) measurements to more realistic two-dimensional radiographic imaging [182] and have performed additional tests to support a demonstration of the concept in the field. This work is led by MAE PhD student Jihye Jeon, who is currently in the final stages of finishing her thesis. In late 2024, we expect the delivery and ability to use enriched uranium, which will enable a fundamentally new series of experiments. We are also partnering with experts from Pacific Northwest National Laboratory to further improve the sensitivity and reproducibility of the system [163].

One impact of all this work has been to further establish Princeton as an academic center for both arms-control science and policy. Since this work is being in done in an unclassified setting, published in peer-reviewed journals, and presented at open conferences, it is attracting attention and interest in collaboration. One important outcome was our collaboration with experts from the China Academy of Engineering Physics (CAEP), which included workshops in Mianyang and Beijing to discuss current work and develop a basis for future cooperative research. The Beijing workshop brought together 11 researchers affiliated with Princeton's SGS and more than 30 participants from the Chinese host. It laid the basis for an important technical and policy collaboration between Princeton-based academic arms-control researchers and Chinese technical experts; in fact, we hosted one postdoctoral researcher from CAEP in our program for 18 months, which led to joint publications that remain quite unique today [105, 156]. Unfortunately, such collaborations are not currently possible.

Information Barriers

Standard warhead confirmation measurements envision the protection of sensitive information acquired during a radiation measurement system—and yet, after thirty-five years of R&D, no viable inspection system with information barrier has been successfully demonstrated. Fundamentally, the challenge is rooted in a lack of trust in the electronics that process the data and protect the sensitive information. In an attempt to facilitate confidence in verification systems and to resolve the certification/authentication dilemma, we have developed and demonstrated several inspection systems, exploring different approaches to resolve some of these longstanding challenges.

Several earlier research efforts have produced prototype systems, but their designs are typically not published in detail, making it difficult to enable trust in these devices. In 2016–2017, we developed the Information Barrier Experimental (IBX) built around the Red Pitaya computing platform using passive gamma spectroscopy [135]. It is the first such platform designed to help a broad research community study vulnerabilities and define the required specifications for a common, trusted inspection system. It is low-cost, simple to assemble, and enables comprehensive hardware and software authentication studies. The device follows a digital data acquisition approach, which significantly reduces the number of components between scintillator and spectrum output. The IBX was based on earlier work of our team on "multi-criteria template matching" approaches [122]. We used the device for an extensive measurement campaign at the Device Assembly Facility (DAF) on the Nevada Test Site in August 2017. Ultimately, we built three prototypes of the device, sharing one unit with Sandia National Laboratories where it was used to support their own research on information barriers.

We have also proposed and examined the use of vintage electronics from the 1970s, i.e., from a time when integrated circuits first became widely available. Establishing trust in such hardware should be much easier because old hardware is significantly less capable than modern electronics, which would make the implementation of exploitable software backdoors or hidden switches extremely difficult. Moreover, it is implausible that microprocessors (with a few thousand transistors) produced more than forty years ago would have been manufactured with suitable hardware vulnerabilities, introduced accidentally or on purpose at the time, that remain undetected to this day. The information barrier we designed and built (IBX II) uses a MOS 6502 processor, proved very reliable, and produces results comparable with measurement devices using modern electronics. In addition to journal publication of this research [153] and some related work [154], the project enjoyed significant public interest and led to a plenary talk at the 34th Chaos Communication Congress in December 2017; it was later also featured in a SEAS story (↗).

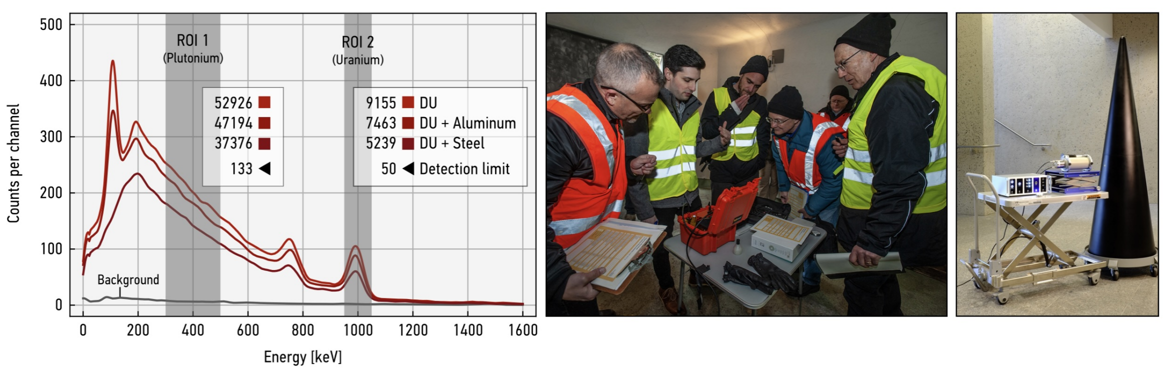

Information barriers in action. The gamma radiation spectrum on the left shows the signature of a depleted uranium target measured at PPPL to validate the inspection algorithm for a system to confirm the absence of nuclear weapons [185]. Participants during a verification exercise organized by the United Nations Institute for Disarmament Research (UNIDIR) in Switzerland [184], where our team was in charge of all radiation measurements and provided dedicated inspection equipment (center). Another information barrier, shown on the right, was developed as part of the course MAE 354/574 in 2016 [135] and later used for radiation measurements on weapon-grade plutonium and other nuclear materials at the Device Assembly Facility on the Nevada Test Site.

We have also pioneered possible verification approaches that would be relevant for an "absence regime." In such a scenario, one could simply accept as weapons all "items declared as weapons" by the host. Verification of such an agreement would be drastically simplified, but it could still require instruments that can confirm that certain containerized items are indeed not treaty accountable, even when the host cannot offer visual access to these items. To support such a regime, we have proposed and demonstrated a measurement approach that relies exclusively on passive gamma measurements using a low-resolution sodium-iodide detector [158, 175]. The system would be able to confirm the absence of both plutonium-based and uranium-based weapons, especially when complemented with simple neutron counters. In a verification regime based on absence measurements, no weapons should ever be part of an inspection, and safety and security concerns would therefore be dramatically reduced. We had the unique opportunity to demonstrate the approach and a dedicated prototype device during a verification exercise in March 2023 hosted by the United Nations at a former military site in Switzerland [184, 185].

Area II. Non-Intrusive Approaches to Monitoring and Verification

Intrusiveness of inspections, especially when involving sensitive facilities and items, has been a major concern with regard to nuclear verification, even when there was more support among weapon states to consider reductions or limits on their stockpiles. This dilemma raises the question of whether alternative verification approaches can be devised that could avoid onsite inspections altogether, simplify verification objectives, or use specialized equipment to perform tasks that have previously been completed by human inspectors. Indeed, a 2021 National Academy study (↗) concluded that future arms control treaties that limit the total number of weapons "would likely require access to storage areas either directly or remotely."

Remote inspections are particularly relevant in a world, where remote collaboration has become an effective and widely accepted practice. Having available non-intrusive approaches could be particularly valuable when bilateral or multilateral arms control resumes and may also help convince those nuclear weapon states that have not engaged much on verification research internationally to join and support future efforts. Over the last five years, we have proposed and partly demonstrated several complementary approaches to monitoring and verification. We have reviewed some key concepts as part of a manuscript prepared for the German Federal Foreign Office in 2021 [159].

Remote Inspections

Onsite inspections have played an important role in verifying compliance with nuclear nonproliferation and arms control agreements. Recent advances in remote and standoff monitoring may complement such inspections, which could make verification approaches more robust, less intrusive, and possibly also less expensive. We have, for example, begun to explore the idea of separating the host and the inspector, such that only the host has access to a sensitive site or item, while the inspector follows relevant activities remotely.

In 2023, we completed a major collaboration with the Max Planck Institute for Security and Privacy in Germany. The experimental setup included a challenge response system installed in a large shipping container, where a number of "treaty accountable items" were placed. While our work was inspired by arms-control applications, the system could be used to monitor any type of "valuable items" that are stored in adversary-controlled environments where regular direct access may be difficult or impossible. Building on the concept of physically unclonable functions, the system uses radio-waves to generate fingerprints of a room and its content using an array of randomly oriented mirrors to verify that nothing changes over time. The system does neither require secure communication channels nor tamper-resistant sensor hardware at the inspected site. The inspection system yields reproducible results, can detect physical changes on the order of a few millimeters, and is robust against major physical and computational (machine-learning) attacks [186].

Another possibility to implement remote inspections is via wide-area environmental monitoring or sampling using sensor networks, either deployed locally, regionally, or globally. For example, in order to determine the requirements for detecting clandestine plutonium separation (reprocessing) using the traces of the noble gas krypton-85 that such activities release, I have worked with Michael Schöppner on extensive atmospheric transport modeling to characterize the current variability of the global background of krypton-85 and assess the detectability of hypothetical clandestine reprocessing plants in various regions of the world [102, 111]. The aim was to set goals for researchers working on specific krypton-85 detector technologies and sampling strategies for mobile and fixed sensor networks and to establish criteria to benchmark and compare different methodologies. Atmospheric transport modeling has since become a major focus of ongoing research at SGS led by other researchers on our team.

Satellite Earth Observation

The development of the first imaging satellites is often credited with enabling nuclear arms control based on the assumption that more information brought transparency and thus facilitated restraint and cooperation. Challenging this conventional wisdom, Tamara Patton (STEP PhD, 2021) showed that the increasing availability of satellite imagery during the Cold War instead encouraged arms control based on limits rather than outright bans, which otherwise may have been favored by the superpowers; it also became entangled with (and ultimately supported) modernization efforts, especially in the United States, because it was felt that the introduction of new weapon systems could be "managed" and would there therefore not be a destabilizing factor.

Given how rapidly satellite Earth observation is evolving today, now also with commercial providers offering access to imagery with high resolution and high revisit rates, it is crucial to better understand the potential impact of these developments on monitoring and verification. This is relevant not only in nuclear arms control and nonproliferation but also for many other areas, including for greenhouse-gas emission assessments, for environmental monitoring, or in humanitarian contexts. Modern machine-learning techniques that can be used analyze satellite imagery at scale further enhance these trends [187]. In addition to developing and teaching a graduate course on this topic, I have also begun to examine these issues in my research along with other members of our group.

A first project funded by German Foundation for Peace Research examines the potential role of synthetic satellite imagery to facilitate systematic assessments of opportunities and challenges in this field. The use of synthetic imagery allows us to carefully control relevant parameters such as seasons, revisit times, environmental conditions, and resolution. We can then use this imagery to develop and examine concrete scenarios, in which parties are either compliant or non-compliant with an agreement. Synthetic imagery can also be used as training data for detection algorithms, which is particularly valuable when real imagery is sparse. Our work so far has been focused on developing the toolkit to accomplish this task [178, 180]. Naturally, the existence of synthetic media also raises concerns about misuse of the very same technologies, and experts have begun to warn about the possibility of satellite image spoofing for malicious purposes. Our ongoing work involves collaborators in digital data and information science to examine the plausibility of such concerns and to understand how trust in digital data is established. This is an area of ongoing and future work.

Robotic Inspections

We first began to consider using robots for inspection tasks as part of a grant awarded by the U.S. Department of State in 2017. As part of this project, we developed a special neutron detector that uses sets of boron-coated-straw detectors embedded in a polyethylene cylinder [148]. This detector can be mounted on a mobile robotic platform, which can be remote controlled or, in principle, also navigate autonomously. The "inspector bot" is spectrally sensitive and can perform single-shot directional measurements, i.e., it can determine the location of a neutron source in a single measurement and without having to rotate in place. We later demonstrated the performance of the inspector bot experimentally during an extensive measurement campaign conducted at PPPL. These measurements confirmed the spectral and directional sensitivity of the system [166, 174]. For improved performance, we later upgraded the design from a 3-

detector to a 6-detector system. We currently have a total of three units available, and these can in principle coordinate their actions (for example, using explore-vs-exploit policies) to complete their detection task in a minimum amount of time.

Possible applications of the inspector bot are in nuclear safeguards to confirm the absence of undeclared activities in gas-centrifuge enrichment plants, which could include detection of hidden feed and withdrawal stations in cascade areas or monitoring of declared feed and withdrawal areas in the plant. The International Atomic Energy Agency has also been considering robotics for "repetitive measurements in areas that can be difficult to access, or with elevated radiation levels" (↗).

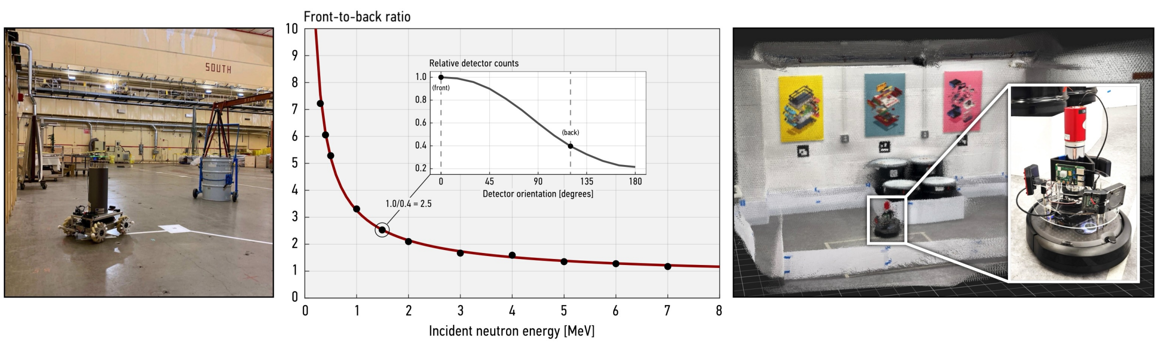

Inspector bot facing a neutron source during a measurement campaign at PPPL. The robot is equipped with three neutron detectors, which can not only determine the direction of a neutron signal but also estimate the dominant neutron energy based on the front-to-back ratio of detector counts, for example, when the robot has rotated such that one detector is facing the source directly [174]. Shown on the right is another mobile robotic platform, which can explore an area without ever creating a map and therefore maintaining the privacy of the inspected site [188]. See also a Princeton story from November 2019 (↗), a PPPL story from June 2023 (↗), and a story about MAE PhD student Eric Lepowsky from November 2023 (↗).

Similarly, robotic inspectors could be used in nuclear arms-control verification. In 2022, using state-of-the-art path-planning algorithms, we experimentally demonstrated that the inspector bot can localize concealed radiation sources, confirm the absence of sources in a room that is declared to be empty, and confirm the correctness of a declared radiation field in a storage facility [176].

Building on this work, we also explored the idea of minimally intrusive, privacy-preserving robotic inspectors, capable of performing a source verification task without requiring, nor providing, any information about the search environment. Most reported methods in robotic radiation detection either require *a priori* information about the environment or reveal site-specific information as part of the search, for example, by creating a map. Instead, rather than deciding what to remember and what to forget, the robot never learns in the first place. Based on this paradigm, we established a method for autonomously performing radiation measurements without revealing the site's layout and without learning specifics of the radiation field, thereby affording a high level of privacy for the host. To demonstrate the algorithm, we performed extensive simulations followed by a series of experiments using another mobile robot based on the *iRobot Create 3* equipped with a number of radiation detectors [188].

Other Minimally Intrusive Verification Approaches

We have examined and further developed a number of additional verification approaches that, when taken together, provide an extensive "menu" of options for future efforts to implement arms-control or nonproliferation agreements.

As part of one such effort, we developed a prototype of a "proximity tag" or "buddy tag" [155]. In a tagging regime using buddy tags, a party would declare a certain number of treaty accountable items and receive exactly one (unique and unclonable) tag for each. The monitored party would then co-locate these tags with the items. During a short-notice onsite inspection later on, the inspected party must be able to present one buddy tag for each item present in a particular area or site. It's critical that such a tag cannot be moved without being detected, which is the major technical challenge for this verification approach [112, 127, 136b]. Developed as part of a grant by the U.S. Department of State, based on an earlier collaboration with Sandia National Laboratories, the buddy-tag concept offers an elegant solution to treaties that involve numerical limits on treaty accountable items when their locations are considered sensitive.

We have also examined in some detail the possibility of verification regimes that rely on declarations only while protecting sensitive information, i.e., without revealing locations, numbers, or operational details, except during inspections. Here, the concept of "hashed declarations" has been previously recognized as a possible basis for verifying limits on the number of nuclear warheads in the arsenals. In fact, the general idea was first proposed in 1980, but the concept only became viable with the invention of modern hash functions in the mid 1990s. As part of such a framework, parties would regularly exchange hashed declarations with entries for each declared treaty accountable item. These declarations could even be made public as they do not contain any information besides the total number of declared items [156]. Prior to a short-notice inspection, relevant entries would be revealed for the inspected site providing a snapshot in time for that site only. Ultimately, this and other concepts of cryptographic escrow allow for step-by-step verification of the correctness and completeness of declarations so that the information release keeps pace with parallel diplomatic and political processes. We applied these ideas to the case of a possible North Korean denuclearization, which was widely discussed before 2020 [151].

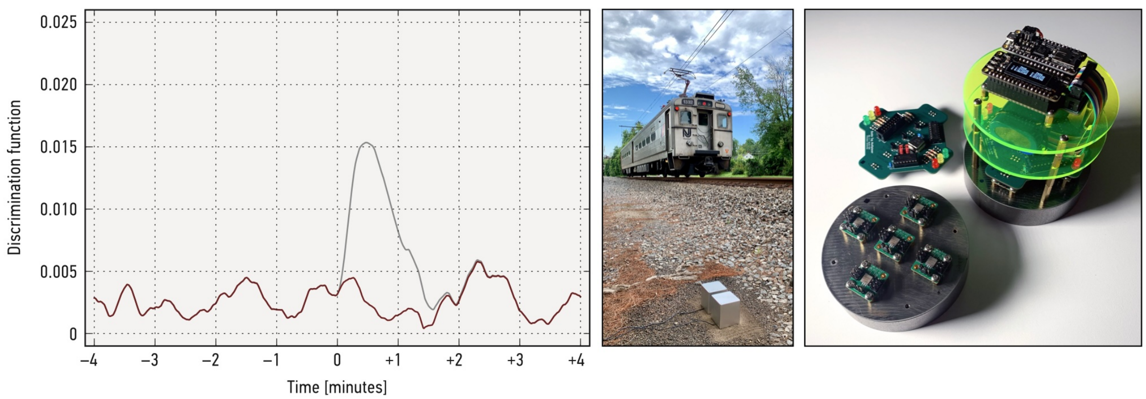

Even while the Dinky is passing by, the buddy tag can clearly distinguish accelerations due to environmental noise from (simulated) stealthy displacements of the tag. An extremely small acceleration of 20 μg has been added post measurement to one of the axes to simulate a stealthy displacement attempt while the tag is experiencing strong vibrations [155]. The device can clearly distinguish these vibrations (red) from the displacement (gray). During field tests near train tracks in Princeton, NJ, the buddy tags were emplaced in aluminum enclosures to reduce environmental effects (center). The final prototype of the buddy tag is shown on the right. Five accelerometers are mounted on a steel ground plate for redundancy and improved performance. The buddy-tag concept offers an elegant solution to treaties that involve numerical limits on treaty accountable items when their locations are considered sensitive. The project evolved from an earlier collaboration with Sandia National Laboratories [112, 127, 136b]. See also the SEAS story from October 2020 (↗).

Finally, we have collaborated with researchers at the United Nations Institute for Disarmament Research (UNIDIR) to further develop the concept of deferred verification [152]. Parties divide up their territory into "open" and "closed" segments, and inspections are only allowed in the open segment while the contours of the closed segment shrink over time, as a country consolidates and reduces its stockpiles of materials and items. The open segment would be under IAEA safeguards to ensure peaceful use of all nuclear materials. Deferred verification might require concepts of absence verification (as described in the previous section). Partly thanks to our efforts in this area, confirming the absence of nuclear weapons at an inspected site has emerged as a "hot topic" in verification research. The approach might also require robust definitions of what is not considered a nuclear weapon, and I have highlighted related challenges of this task [183].

Virtual Inspections

We have also introduced virtual reality (VR) technologies to the area of nuclear inspections, beginning in 2016/2017 with a grant on "Virtual Reality for Nuclear Material Security and Arms Control: Engaging the Public & International Government Partners.” On the research side, our team developed several virtual environments to provide a new pathway to support experts and governments in developing a shared, hands-on understanding of the challenges involved in nuclear security and verification. Here, VR offers a means to create and interact in flexible environments, cooperatively explore new concepts and approaches, and lay a basis for live exercises and new policy initiatives. (↗). In particular, we were able to use our expertise in radiation transport modeling to implement real-time (gamma) radiation fields in VR using a hybrid approach that combines pre-computed radiation signatures and detector response functions (based on extensive Monte Carlo simulations) with deterministic methods to handle shielding and attenuation effects allowing the movements of sources, detectors, and shielding materials in the VR experience [128, 149]. Collaborating with a team of researchers who work on the “science of magic” (↗), we also implemented elements interactivity, curveballs, and gameplay to test the robustness of verification approaches when unexpected events occur [169]. Virtual reality offers an ideal platform to explore these aspects thanks to the ability to perfectly reproduce all elements of a given scenario.

Our work was presented at safeguards symposia organized by the International Atomic Energy Agency, inspired similar efforts elsewhere and, in particular, prompted the International Partnership on Nuclear Disarmament Verification (IPNDV) to include VR exercises into its ongoing work.

Area III. Fissile Materials, Nuclear Energy, and Nuclear Nonproliferation

At the outset of the nuclear age, leading scientists recognized that the new technologies for making fissile materials developed by the United States as part of its World War II nuclear weapons program might have possible peaceful applications. They also realized such materials and technologies in civilian use could pose the enduring risk of being diverted to weapons purposes should a government wish to do so and argued for the need for a system of international controls. A major area of my research combines technical and policy analysis to understand and reduce the risks from today’s large global fissile material stockpiles and from the nuclear technologies now associated with civilian nuclear energy programs around the world.

In this section, I describe my research on verifying past production of nuclear-weapon materials, ending the use of highly enriched uranium as a fuel for research reactors, reducing the proliferation risks from the gas centrifuge technology used to make the enriched uranium that today fuels most of the world’s nuclear reactors, and from plutonium production at civilian research reactors. I have approached these critical issues with a perspective that I have helped develop as part of the SPIA’s Program on Science and Global Security. According to this view, the only reliable way to reduce and end the threat from fissile materials is to verifiably end production and use of these materials, and to reduce existing stockpiles as transparently and irreversibly as possible.

The key technical analysis and policy arguments that support this fissile material perspective view on arms control and nonproliferation are laid out in *Unmaking the Bomb,* the book I co-authored with my SGS colleagues, published by MIT Press [92].

Verifying Past Production of Nuclear-weapon Materials

There currently are about 12,000 nuclear weapons in the arsenals of the nine nuclear weapon states—but global fissile material stocks are estimated to be sufficient for about 200,000 weapons. The range reflects uncertainties of the order of 20–40 percent in the fissile material production and stockpiles for some weapon states. My research in this area is aimed at developing the science and technology needed to reduce these uncertainties in understanding the production histories and estimating existing stockpiles of fissile materials. My research is also aimed at developing the technical capability to independently verify national fissile material declarations. To this end, we have developed new computational tools to support this type of research and have been making them available to the community. Here, perhaps the most important development has been the open source depletion code ONIX [157], which expands the capabilities of another open source code developed at MIT and formed the basis for Julien de Troullioud de Lanversin's MAE PhD thesis in 2019.

I have developed detailed computer models of the types of reactors used for plutonium production by weapon states and helped develop detailed histories of uranium enrichment programs to generate new estimates of national fissile material holdings [52, 57, 170]. These estimates have been a key element of fissile material estimates published by the International Panel on Fissile Materials (IPFM), which are widely quoted by government and non-government experts. The computer models that I have been developing for understanding plutonium production in weapon states have also found use in the field of nuclear forensics, which combines methods of traditional forensics, radiochemistry, analytical chemistry, material science, isotope geochemistry, and nuclear physics to characterize isotopic signatures of nuclear materials. Earlier neutronics calculations formed part of an influential joint study on the state of the art of nuclear forensics by the American Physical Society and the American Association for the Advancement of Science [28]. Building on this work, I have also explored the question of how nuclear forensics can be used for nuclear arms-control verification [95].

Another area in which I am working to combine nuclear arms-control verification and nuclear forensics is the emerging field of "nuclear archaeology," which seeks to develop methods to analyze the isotopics of trace impurities in structural materials or in waste materials at former fissile material production sites. The idea was originally proposed in the 1990s for one particular type of plutonium production reactor. In a series of papers, I have systematically expanded the horizon of nuclear archaeology. I have shown that isotopic signatures of weapon-grade uranium and plutonium can be associated with particular facility types and modes of operation [33, 45]. We also have shown how nuclear archaeology can be used in heavy water reactors [59], how to distinguish past tritium from plutonium production [150], and how it can be applied to enrichment plants that have been used to make the bulk of weapon-grade uranium [79].

In a more recent effort, we have proposed the idea of document-based nuclear archaeology, which could complement standard nuclear archaeology techniques based on forensic analysis of physical samples. To validate and demonstrate some of the relevant concepts, we were fortunate to work with the staff of a Norwegian research reactor that had recently been shut down. We collected, analyzed, and started to preserve the reactor's operating records, which exist on both analog and digital media [168, 173]. In developing guidelines for best practices that conform to existing standards for long-term digital preservation and curation, we hope this project can help lay the basis for future nuclear archaeology efforts to support nuclear arms control and disarmament.

Minimizing Proliferation Risks of the Nuclear Fuel Cycle

Over the years, my research has addressed a broad range of issues relevant for efforts to strengthen the proliferation resistance of the nuclear fuel cycle. Ending the use of highly enriched uranium (HEU) has been such a focus of my work. HEU is one of the two fissile materials that can be used for making nuclear weapons, the other being plutonium. My research in this area began as part of my doctoral work, and it has been one where I have made a significant impact. Since coming to Princeton, I made several important new contributions to the technical and policy debates aimed facilitating and accelerating reactor conversion efforts world-wide [6, 14, 15, 19, 110]. My work in this area led to invitations to serve as an expert adviser to the International Atomic Energy Agency to help shape its positions and priorities on this matter. In 2014–2016, I was a committee member of the National Academy of Sciences' study on the *"Current Status of and Progress toward Eliminating Highly Enriched Uranium Use in Fuel for Civilian Research and Test Reactors,"* which issued a number of bold recommendations on how to accelerate the global phaseout of HEU fuel from the civilian fuel cycle. Since then, there have been a number of new and unforeseen developments that could further increase the salience of the issue in coming years. In particular, Russia's war against Ukraine has effectively broken existing global supply chains for nuclear fuel. Russia has been a major supplier of research and power reactor fuels for many countries, and it's currently unclear how suppliers and operators might adjust to this new reality. At the same time, the United States is advocating for the use of high-assay low-enriched uranium (HALEU) in small modular reactors being developed. This is unprecedented and poses possible proliferation challenges that are poorly understood.

Another major research effort in this area has been on the proliferation risks of centrifuge enrichment technology, which gained prominence with the discovery of Iran's undeclared uranium enrichment program. In a series of articles, my work has analyzed the technical characteristics of gas centrifuges and how a state might use this technology in a crash program to make highly enriched uranium suitable for nuclear weapons—a prospect commonly described as a breakout [21, 32, 33, 34]. I used my technical understanding of gas centrifuges to work with SGS colleagues to develop new options, shared with U.S and Iranian officials, on how to extend Iran's breakout time and limit the proliferation risk from its enrichment program [86, 90]. My work on understanding and reducing the proliferation risks from gas centrifuge uranium enrichment technology and from plutonium production at civilian research reactors came together and found an important policy focus in the negotiations over the future of Iran's nuclear energy program. I contributed technical and policy analysis that informed negotiators and decision-makers in the United States and in Iran and helped lay the technical basis for the July 2015 agreement between Iran and the six world powers. This agreement is not currently implemented, with little hope that it can be saved.

A third contribution that I have made to the international effort to limit the long term risks from civilian nuclear programs has not yet been adopted. This is based on my research exploring the benefits of moving from national control to multinational ownership and control of uranium enrichment facilities [39, 49]. This work was undertaken first as a study for the International Commission on Nuclear Non-proliferation and Disarmament, established by the governments of Australia and Japan. In 2015, I co-authored an article for *Science* applying this idea to Iran [98]. The article suggested converting Iran's national enrichment plant into a multinational one, possibly including as partners some of Iran's neighbors and one or more of the six countries that reached a comprehensive nuclear agreement with Iran in July 2015. With the currently considered plan to share U.S. nuclear technologies with Saudi Arabia, which could eventually also include enrichment technology, these ideas may be more relevant than ever.

Assessing Next-Generation Nuclear Reactors

Another line of my research in this area is exploring the potential and limits of emerging nuclear energy technologies, many still in the R&D stage, which could be game changers in shaping the nuclear future. This includes next-generation nuclear fission systems that seek to overcome key problems with current nuclear reactors as well as possible nuclear fusion energy systems. A large-scale expansion of nuclear power will be needed if it is to play a significant role in reducing greenhouse-gas emissions to address the threat of climate change [47]. My research aims to inform the technical assessments and policy debates that set the requirements for these next-generation nuclear fission and fusion technologies, with a particular emphasis on understanding and establishing possible criteria for the impacts of these technologies on nonproliferation efforts.

Existing nuclear power plants are seen as having four significant challenges that limit the possible growth of nuclear power as part of a global effort to limit greenhouse-gas emissions and climate change: cost, safety, waste disposal, and nuclear proliferation. My research has contributed to understanding these aspects through computer modeling to illuminate the impact of fuel-cycle choices, uranium fuel requirements, and plutonium production rates, all of which play an important role in assessing the possible proliferation consequences and, ultimately, the viability of a proposed reactor design as a sustainable energy technology. Over the years, we have examined and assessed a wide variety of next-generation reactor technologies, including molten salt reactors [71, 93] and small modular reactors [75, 99]. An important part of my work with my SGS colleagues has also included an effort to inform the development of viable regulatory and licensing processes for SMRs [77], which I briefed the U.S. Nuclear Regulatory Commission.

My research on nuclear energy systems has also included nuclear fusion, which has long held the potential of solving many of humanity's long-term energy needs. The technical challenges involved mean commercial deployment is still over the horizon, and may only be achieved in the second half of this century. Still, using an integrated assessment model, we have found that the present value of the fusion option should be significantly larger than the estimated cost of a comprehensive R&D plan to develop the technology [108]. The coming years therefore offer an opportunity to shape the technology and, in particular, minimize its potential proliferation risks. My research on fusion energy seeks to inform the technical and policy debate with analyses that illuminate the possible proliferation impacts of candidate fusion energy systems and aims to provide a basis for building in safeguards-by-design for this technology.

Like traditional nuclear (fission) systems, nuclear fusion that is based on the deuterium-tritium reaction involves intense neutron fluxes. This means a fusion reactor could in principle be used to make weapon-usable fissile material. Working with Robert J. Goldston, I published the definitive analysis of the proliferation risks of (magnetic confinement) nuclear fusion technology [67]. In response to this article, which is still widely cited today, the International Atomic Energy Agency decided to host a consultancy meeting in June 2013 to discuss relevant aspects for IAEA safeguards [70]. More recently, I addressed these issues at a "Fusion Energy and Nonproliferation" workshop held at the Andlinger Center, hosted by PPPL, and sponsored by the U.S. Department of Energy in January 2023. With the currently emerging interest in nuclear fusion, involving for the first time many startups and large amounts of government funding, I expect this area to be an important area of my work in coming years.

Future Directions

My plans for future research include work that is already funded, some of it through 2026, as well as ideas for new work that I have proposed in grant applications and some new research problems that I am now only beginning to explore. These plans are sketched out below and include research in several of areas of interest.

Nuclear Verification. We have built and set up at PPPL, a robust, compact, and versatile neutron radiography system (Excalibur). Expanding the scale and scope of the experimental work on nuclear verification remain an important focus of my research over the next few years. Later in 2024, PPPL will receive test cubes of high-assay low-enriched uranium that is 15%-enriched in the isotope uranium-235. Since most of our experiments are aimed to detect the presence and concentration of this fissile isotope, these materials will enable us to conduct fundamentally new experiments.

We have also begun to explore the idea of "real-time" zero-knowledge radiography and published some possible implementations of the concept [181]. The idea is to capture light from a scintillator plate (exposed to the neutron beam) using an extremely sensitive digital camera. To support this idea, PPPL recently acquired a deeply-cooled single-photon camera that should be able to meet the requirements for this application. The challenge will be to capture an image in a way that protects sensitive information, which could be realized using a complement of the inspected item, possibly in conjunction with a beam splitter/combiner. If successful, this approach would overcome some of the main limitations of the current implementation using bubble detectors, namely the low efficiency and low spatial resolution of these detectors.

Remote (Virtual) and Robotic Inspections. Onsite inspections play an important role in nuclear monitoring and verification. They are difficult to substitute with other verification approaches, but they are expensive and often considered intrusive. One major theme of my planned research seeks to expand on the idea of "physically separating" the host and the inspector as a hybrid approach that preserves some aspects of onsite inspections while resolving concerns about costs and intrusiveness; in particular, we want to consider situations where only the host enters a facility subject to inspection while the inspector follows the relevant activities remotely. Specifically, as part of these future efforts, we will examine and develop solutions to enable interactivity, privacy, and strategies to confirm that activities are live and taking place at the agreed location. We are in the process of exploring different technologies and approaches based on this fundamental idea, including event-based vision cameras and inspection environments that have both virtual and physical (real) elements.

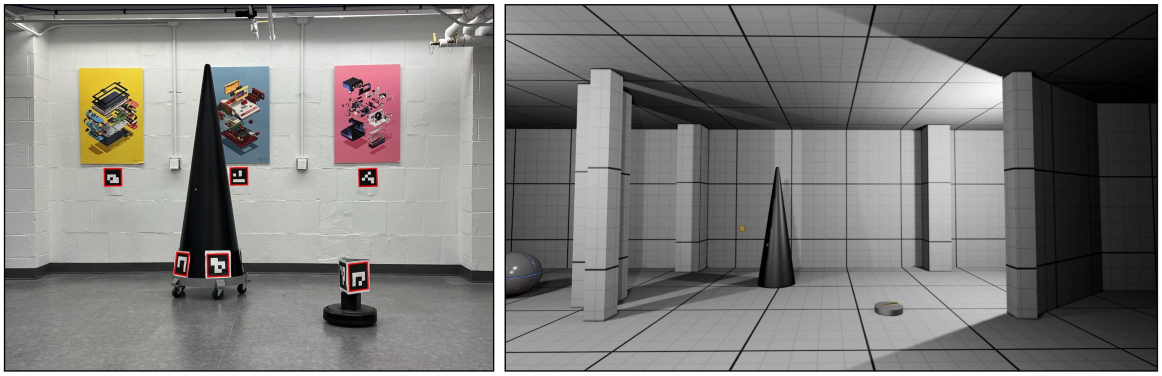

A mockup warhead in the physical and in the virtual world. Somewhat related to the concept of a digital twin, during a virtual inspection, only the host party would enter the inspected site while the inspector participates remotely. The challenge is to enable meaningful interventions by the inspector and to ensure the remote party that, whatever it is they observe, is happening live and at the agreed location [177].

With the support of the University, we recently acquired a SPOT robot, which enabled and played a key role in a new course we developed ("Robots in Human Ecology"); this robot also offers the opportunity for new research. In particular, we want to examine inspection approaches where the robot either performs inspection tasks autonomously or enables the remote participation of a human inspector. Leveraging these capabilities, we also seek to further explore the concept of robotic inspections in special environments, including at PPPL where a new experimental facility will be becoming available.

With regard to robotic inspections, PPPL is interested in pursuing research on radiation hardened sensors and autonomous monitoring of fusion energy systems. Here again, SPOT could enable some innovative new research as it has been shown to be extremely radiation resilient and could be equipped with a variety of different sensors to perform inspection tasks that would otherwise be difficult or impossible. PPPL offers an ideal environment to solve some of these challenges.

Satellite Imagery, Remote Monitoring, and Treaty Compliance. Launched with seed-funding from the German Foundation for Peace Research, this project seeks to provide a first systematic assessment of emerging monitoring and verification capabilities enabled by new satellite constellations with very short revisit times, primarily leveraging synthetic satellite imagery to do so. We now have the toolkit and datasets available to start work on first assessments. These will be relevant for monitoring compliance with arms control and nonproliferation agreements, but they are equally relevant in many other contexts, including for greenhouse-gas emission assessments, for environmental monitoring, or in humanitarian contexts. For example, some research groups have tried to estimate CO_2 emissions from fossil power plants using machine-learning algorithms, but these efforts have had mixed results due to insufficient training data and the lack of ground truth. I believe our technique could fundamentally transform the potential of some of these possible use cases.

Our team has also recently also obtained exclusive access to high-resolution thermal satellite imagery. The resolution of this thermal imagery is about 10–30 times higher than imagery delivered by satellites that are currently in use. While the first satellite ceased to function in December 2023, having acquired fewer than five hundred images, the next launches are currently planned for 2025.

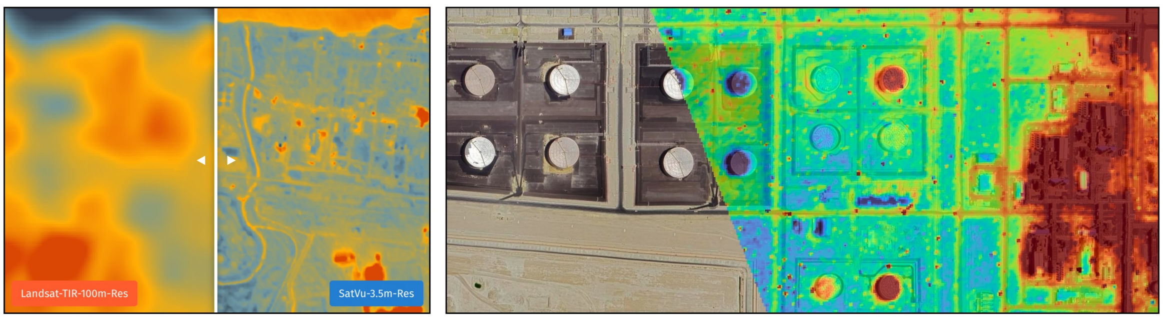

Exploring the potential of high-resolution thermal satellite imagery. Shown on the left is a comparison between standard and next-generation thermal imagery; shown on the right is a small area of an oil refinery in Kuwait. In the thermal imagery (layered on top of an optical image), filled and empty tanks can be clearly distinguished due to their temperature difference. The thermal imagery is taken at night. *Imagery courtesy of satellitevu.com and Google Earth.*

Thermal imagery at this resolution offers fundamentally new opportunities for remote monitoring as energy-intensive applications (including, for example, bitcoin-mining operations) become visible such that their energy use could be verified remotely.

Optical satellite imagery has been used for decades to support proliferation detection, nuclear safeguards and arms-control verification. Optical imagery has some important limitations, however. In particular, undeclared activities can be conducted during cloudy weather or at night-time. There have also been some instances where states have pursued visual deceptions or concealment using false structures to evade detection via satellite monitoring. Even if visually detected, it can still be difficult to confirm the purpose or the operational status of a facility. Thermal satellite imagery offers an opportunity to overcome some of these constraints, but the resolution of older satellite sensors has been inadequate for detailed analyses. We will examine the unique potential of high-resolution thermal imagery that is now becoming available.

Safeguarding Nuclear Fusion Reactors. There is renewed interest in using advanced nuclear technologies to accelerate a transition to a low-carbon energy future.

In addition to small modular reactors and "microreactors" (both based on nuclear fission), there now also exist multiple nuclear fusion startups, and the U.S. Department of Energy recently released a decadal fusion energy strategy (↗). Major technology companies are considering some of these technologies to power their data centers with large and growing energy demands. This may have important implications for the front-end and the back-end of nuclear fuel cycle, including the possible use of high-assay low-enriched uranium in fission reactors and the use of tritium in kilogram quantities in the case of fusion systems. I am planning to provide technical and policy analysis for relevant assessments leveraging the tools we have developed in our group.

The overall goal of this research will be to inform the debate about how nuclear power would fit into a low-carbon decentralized energy system that privileges flexibility, energy efficiency, and small-scale solutions rather than the large-scale capital and fuel intensive centralized power plants that define today's energy system in most countries. In particular, I plan to continue the analysis of emerging nuclear reactor technologies to inform the debate on the future of nuclear power with a particular emphasis on understanding the attendant proliferation risks.

❧