Slide 1: Smart and Autonomous Systems (S&AS); NSF 18-557

Reid Simmons: Good afternoon, and thank you for joining this webinar. I am Reid Simmons; I'm a program director in the Division of Information and Intelligent Systems of CISE and the lead program director for the Smart and Autonomous Systems program. I am joined today by Jim Donlon and Jie Yang, who are also program directors in the Information and Intelligent Systems division of CISE.

Before we start, I would like to outline the schedule: The webinar will be one hour in duration; we will present for approximately 30-40 minutes, and then we'll take questions for the remainder of the hour. The question and answer session will be guided by the operator, so please follow the instructions from him or her when we have finished our part of the presentation. We have a quite a few participants today, so there may be a delay before you get an opportunity to ask a question. In addition, you may send questions to the email address listed on the bottom of each slide (aside from the title slide), and we will try to respond to them promptly.

Slide 2: The S&AS Team

Reid Simmons: The Smart and Autonomous Systems solicitation is a continuation of the CISE program initiated in 2017. This year the program is sponsored by the Division of Information and Intelligent Systems, or IIS. We will highlight other changes to the program later. IIS is represented in the S&AS program by myself and program directors Jim Donlon and Jie Yang.

Slide 3: Welcome

Reid Simmons: Let me begin by calling on Erwin Gianchandani, the Deputy Assistant Director and Acting Assistant Director in CISE, to welcome you and to say a few words about how Smart and Autonomous Systems fits into the broader picture at NSF.

Erwin Gianchandani:

Slide 4: Outline

Reid Simmons: Thank you, Erwin. The discussion today is intended to introduce you to the S&AS program, and to help you prepare and submit proposals that are consistent with the goals of the program. We will begin with an overview of why the program was created, what scientific and societal needs it fulfils, and discuss the overarching goal of the program and on what research themes the program is focusing.

We will then discuss the solicitation in more detail, introducing the two classes of projects supported by the program, their budget ranges, and information about the anticipated size and number of awards.

We then discuss the scope of the Smart and Autonomous Systems program and its relationships to other cross-cutting programs at NSF, specifically the National Robotics Initiative and the Cyber-Physical Systems program.

Finally, we will present some details that proposers should keep in mind when submitting proposals and answer questions. As mentioned, we will take questions by phone at the end, or you can submit questions to the email address listed at the bottom of each slide.

Slide 5: Autonomy: Definitions

Reid Simmons: The focus of this program is autonomy, but what is autonomy, exactly? Dictionary definitions are somewhat varied, but tend to agree that autonomy implies independence in making decisions, especially when making ethical decisions. Note that the second definition on this slide states that autonomous systems **can** make choices free of outside influence; this does not imply that autonomous systems **must** be free of outside influences, only that they have the freedom to disregard those influences. So, for instance, a truly autonomous system can take commands from a human, but also decide to disregard them if some other, more critical, issue arises.

Finally, a recently released OSTP report on the future of AI states that autonomy implies the ability to adapt behavior to different contexts. This is in contrast to automation, where machines behave precisely as specified by a person, according to a fixed set of rules. For true autonomy, it is not sufficient to know what to do, but also involves knowing why actions should be taken, reasoning about the consequences of those actions, adapting decisions to changing circumstances, and understanding the contexts in which certain actions may not be appropriate.

Slide 6: Need

Reid Simmons: Systems that interact with the physical world, such as robots and cyber-physical systems, are becoming increasingly prevalent in society, from manufacturing, to health-care, to service industries, to helping to monitor and maintain our infrastructure and natural resources. In the near future, we expect that they will be operating everywhere within our society, as well as places that are difficult, or impossible, for humans to go. The overall aim of such systems is to assist and/or augment people in their work and life but, to do so, it is often the case that the systems need to be highly autonomous. Note that this is not to say that autonomous systems should exclude human **interaction**, but rather that they should minimize the need for human **intervention**. For instance, just as we, as people, interact with one another on a regular basis, that does not contradict our beliefs that we are autonomous creatures. Likewise, our systems may interact with people, but still be highly autonomous, monitoring their own functioning, and repairing themselves or asking for assistance, when needed. They should be aware of their own capabilities and limitations, anticipate potential problems, and not act outside of their competency.

To enable long-term autonomy, systems will need to learn and adapt to changes. These include changes in the environment, either short or long term, changes to the system's hardware, such as failing components or hardware degradation, and changes to the system's tasks and missions, including learning new tasks without the need for explicit programming.

Finally, with the great power of autonomous systems comes great responsibility. Autonomous systems will need to act ethically, explicitly reasoning about tradeoffs between the achievement of mission goals and the social good. For instance, if a child darts into a roadway, an autonomous vehicle might need to decide between serving out of its lane and crashing versus hitting the child. Given uncertainty about what might occur – the child might stop in time, cars in the other lane might themselves swerve to avoid the accident – the decision about what is the most ethical action to take is often not easy to make. But, if we desire our systems to be truly autonomous, researchers will need to address such concerns.

Slide 7: Need

Reid Simmons: The critical need for research into autonomous systems is highlighted in several recent reports. In September 2015, an NSF-sponsored workshop took place to consider future directions in cyber-physical systems and robotics. The workshop participants highlighted **autonomy** as a potentially disruptive technology, and therefore strongly called for programs that helped advance a science of autonomy. The report presents drivers and metrics for developing and evaluating autonomous systems. This report was very influential in the formation of the Smart and Autonomous Systems program.

Just prior to the initiation of this program, OSTP released two reports: 1) on preparing for the future of AI; and 2) a strategic plan for research and development in artificial intelligence. While the reports cover all of AI and touch on a broad range of research areas, they also call out autonomy and learning as two of the critical elements for future AI systems.

Finally, in its 2016 report, the 100 year study on AI raises concerns about the need for ethical and social considerations in the development of autonomous systems. While the report mainly concerns itself with ethical issues involved in creating such systems, it does give nod to issues surrounding the ethical behavior of such systems, themselves. For instance, it notes that "Ethical questions are also involved in programming cars to act in situations in which human injury or death is inevitable, especially when there are split-second choices to be made about whom to put at risk."

Slide 8: Why True Autonomy is Hard

Reid Simmons: One reason that NSF has a special program devoted to autonomy is that true autonomy is hard to achieve. A now-familiar aphorism attributed to Donald Rumsfeld illustrates the issue – current autonomous systems mostly know how to deal with the known-knowns – things that we expect to happen; we really don't know how to deal with the unknown-unknowns – things that we don't even know that we don't know about; what this program is focusing on is how systems can deal with known unknowns – situations that we are aware may occur, but are not currently prepared to handle. The reason why this category is important is that, in general, there are uncountably many problems that can occur with non-negligible probability. Since designers of IPS cannot possibly enumerate them all, it is important that we develop approaches that can enable IPS to recognize or anticipate a situation that is outside of its area of expertise and figure out a way to act appropriately, safely, and ethically. Getting to true autonomy is further complicated by the fact that all models are approximations of the real world, and they will inevitably be wrong in some given situations. Systems need to understand the limitations of their own models, and work to improve them when possible.

Slide 9: Goal

Reid Simmons: So, hopefully, the need for autonomous systems, along with some of the concerns that they raise, have been made clear. These issues helped give rise to the Smart and Autonomous Systems program. As in the inaugural version of Smart and Autonomous Systems, the focus of this program is on fundamental research into providing **Intelligent Physical Systems** with a high degree of autonomy, especially for long-term autonomy. We define Intelligent Physical Systems (or **IPS**) in terms of several defining characteristics. Please take note that these characteristics have been modified to clarify the focus of the program. The original five characteristics are now four: **cognizant, taskable, adaptable, and ethical**. These characteristics are intended to make clear that the program focus is on systems that can robustly handle uncertain, unanticipated and dynamically changing situations through high-level cognition, self-awareness, and adaptation. We will go through each of these characteristics in turn, but I want to point out that the research themes of the program are aligned with these four characteristics, and every proposal to the S&AS program must address at least one of those themes. It is allowable for proposals to cover additional themes, providing its primary focus is on one, or more, of these defining characteristics. Also note that "Knowledge-Rich" has been removed from this list, to clarify that this should be a characteristic of any IPS research, regardless of the characteristics emphasized in that investigation. I will next describe each one of these characteristics.

Slide 10: Cognizant

Reid Simmons: The first characteristic, and the first research theme, of the program is that IPS must be **cognizant** of their own capabilities and limitations. The issue here is that, in order to operate in complex domains with minimal human supervision or intervention, the systems must recognize when they have gotten into trouble and adapt their behaviors accordingly. More prudently, IPS should anticipate when they *possibly* may be heading for trouble and act to forestall any impending problems. They need to do act in this self-aware fashion, despite the fact that there may be significant uncertainty about what is happening, due to unexpected events, uncertain perception, and inaccurate modeling of the situation.

For IPS that interact with humans, they should behave transparently in a way that people find predictable and legible, to avoid confusion about what the system is trying to achieve. In situations where the system must ask people for assistance, it is important for the system to provide suitable explanations for what it has been doing, and why, in ways that are readily understandable by people. Note that the system should anticipate any potential problems far enough in advance that they might request assistance in time for a person to react in time. Such explanatory capabilities can also be used for introspection, where the system uses meta-reasoning to understand how its own behavior has resulted in certain outcomes, to avoid bad outcomes, or reinforce good outcomes in the future.

As an illustrative example, consider oceanic exploration using unmanned robotic vehicles. Unmanned marine systems that are cognizant will rely upon on-vehicle intelligence to enable longer duration deployments of unmanned underwater and surface vehicles and reduce the need for human intervention and control. A robust cognizant capability allows unmanned marine vehicles to adapt in situ to changing environmental conditions, mission objectives, and operator specifications, and facilitates failure recovery and introspection to identify and recover from faults and anomalies during long-duration operation.

Slide 11: Taskable

Reid Simmons: IPS should be versatile, that is, they should have the ability to achieve a wide range of tasks, in diverse environments, without the need for extensive programming. The system should be capable of accepting high-level, goal-oriented objectives, and planning to achieve those objectives in a context-dependent way, while still taking longer-term objectives into account. **Taskable** IPS should be able to take different objectives into account, such as time, energy, safety, or privacy, when deciding how to behave. To aid in achieving their high-level goals, IPS may need to reason about their own sensing and communication behaviors, as well as taking actions that help them better achieve their own situation awareness.

People use many modalities for describing goals, including natural language, gestures, and graphically, such as by drawing maps. Similarly, we expect IPS to be able to accept instructions using similar modalities and, if necessary, engage in multi-modal dialog with humans to clarify any ambiguity in the goals. In addition, IPS should be easily interruptible, postponing a current task to carry out another, more important, task, and then returning, when feasible, to complete the original task (or to indicate when further progress is not possible). Ideally, IPS should reason about these issues – for instance, if achieving a higher priority task may preclude finishing the current lower priority task, or if the lower priority task is nearly done, then the system might decide to complete the lower priority task, if it determines that by doing so it will not jeopardize the higher priority task.

As an example, consider the development of a household robot that is capable of independently performing long-term tasks. Such a capability requires a system that is sufficiently autonomous that it can be issued natural, high level commands such as, "Keep the room tidy," and from the basis of that high-level goal, process the complexity and uncertainty of the task environment, form objectives that consider complex interacting goals, respond appropriately to interaction with multiple people and many objects, and choose actions that are rational with regard to those considerations. Taskability enables such a system to act sensibly over long periods without the need for constant supervision.

Slide 12: Adaptable

Reid Simmons: Adaptable improves upon the characteristic that was called "reflective" in the previous solicitation. By adaptable, we mean the ability to improve behavior over time, including acquiring, modifying, and transforming skills by augmenting the knowledge on how to perform tasks. IPS should also autonomously improve the models they use to perceive, plan, and act. These improvements can come from different sources, including their own past experiences; observations of how other IPS or humans perform; or explicit instructions received on how to perform tasks. IPS should learn to perform new types of tasks without the need for reprogramming. As a result, IPS are expected to continually adjust their behaviors and respond to the contexts in which they are operating, even when those contexts were not initially modeled.

This definition of adaptable is meant to improve upon the previous solicitation's notion of reflective systems, to emphasize the program's focus on high-level cognition, self-awareness, and adaptation--and to distinguish such capabilities from mere onboard diagnosis and repair currently found in many types of cyber-physical systems.

An example of an adaptable system is a team of mobile robots that integrates into the social fabric of a community of people in a public space by learning and adapting to the social norms of that place over time. This system seeks to acquire knowledge about both the environment and the people within that environment over a long period of continuous operation and interaction. The platforms modify and transform their skills by augmenting their knowledge on how to perform tasks such as guiding visitors around the environment. They improve their behavior over time as they offer and receive help, and decide when to interact with whom.

Slide 13: Ethical

Reid Simmons: Ethics is one aspect of intelligent systems that has only recently begun receiving much attention by the autonomous systems community. In general, determining whether an action is **ethical** involves reasoning about the consequences of actions with respect to achieving the goals of the system and the consequences with respect to the overall social good. The social good can be defined both in terms of societal norms and legal rules. Sometimes the task and social objectives are in sync, and the system can simultaneously achieve its goals and promote the social welfare. Sometimes, however, the objectives are in conflict, and the system must reason about the tradeoffs involved, taking risk, costs, and benefits into account in deciding what to do in order to maintain the social good.

As an example of ethical reasoning, consider an aerial vehicle that is tasked with searching for a fugitive who is suspected of hiding out in a residential neighborhood. The vehicle is tasked with scouting the area and sending back photos to law enforcement, who will determine if the fugitive can be identified. As the vehicle flies around the neighborhood, it should be aware of the privacy concerns of the citizens, and try to avoid taking photos of innocent residents, when at all possible. On the other hand, the privacy concerns should not impede the system from achieving its goal of helping to catch the fugitive.

An example of research toward this goal investigates representations and algorithms that enable autonomous systems to recognize and observe human normative expectations. The objective is to enable IPS to maximize the norms they can obey at any moment in time. This might enable systems interacting with humans to make decisions that are transparent and self-explanatory.

Slide 14: Knowledge-Rich

Reid Simmons: IPS need significant amounts of knowledge to achieve the other characteristics. This knowledge comes in many forms, and Intelligent Physical Systems should be able to utilize these various forms of knowledge in whatever ways are appropriate to the task at hand.

Note that we no longer include **Knowledge-Rich** as a distinct theme in Smart and Autonomous Systems. We consider knowledge to be a critical component of *every* IPS. That is to say, each of the four research themes in this program requires IPS to be knowledge-rich, employing a variety of representation and reasoning mechanisms, such as semantic, probabilistic, commonsense, and meta-reasoning. Some example types of knowledge that are often important include both quantitative, or numeric, reasoning and qualitative, or symbolic, reasoning. In some situations, IPS may have sufficiently detailed knowledge of the world that quantitative reasoning is warranted, while in other situations, it may be appropriate to use high-level semantic knowledge to decide what actions to take. Mixing quantitative and qualitative reasoning is important, as systems may often need to plan at multiple levels of abstraction. For dealing with the physical world, probabilistic reasoning and planning may be beneficial, especially when combined with other forms of reasoning, such as commonsense reasoning. In short, there is no one type of knowledge that is best for all situations, and knowledge-rich IPS should utilize multiple models, at multiple levels of abstraction.

Finally, as mentioned earlier, IPS should be capable of extending their knowledge base autonomously, through learning. The types of knowledge learned may similarly take many forms, from simple learning of parameters for models described as equations, to learning new semantic concepts, to learning semantic maps and other characteristics of the environment, to learning new models of how the system and the environment work, to learning new behaviors and strategies for achieving tasks. All this newly learned knowledge should be smoothly integrated into the existing knowledge base, enabling the IPS to become more capable, and more self-aware.

I would now like to turn the presentation over to my colleague, Jim Donlon, who will talk about specifics of the current solicitation and submitting proposals.

Slide 15: Project Classes

Jim Donlon: For this program, there are two project classes.

The **Foundational** class encapsulates research projects that advance fundamental knowledge pertaining to IPS. These projects must focus on at least one of the research themes described earlier. Although it is not necessary to utilize a physical testbed for the project evaluation in this class of project, the project must clearly demonstrate direct relevance to some IPS.

The **Integrative** class encapsulates research projects that focus on the integration of at least two or more components of IPS into increasingly smart and autonomous systems. In contrast to the Foundational class, Integrative projects must include a plan for evaluation on a physical testbed. The evaluation must be rigorous, corresponding to accepted scientific methodology, including well-defined hypotheses, controlled experiments, specific evaluation metrics, and statistical analyses.

Slide 16: What Proposals are Good Fits for S&AS

Jim Donlon: All NSF programs promote established or anticipated research thrusts. The Smart and Autonomous Systems program is no different.

Proposals from the research community, addressing the core foci of the program as outlined in the solicitation will be considered a “Good Fit.” These proposals will primarily address how to enable autonomous systems to robustly handle uncertain, unanticipated and dynamically changing situations through high-level cognition, self-awareness, and adaptation.

As stated in the program solicitation, the proposed projects must address at least one of the four research themes and integrative projects must include an evaluation plan relevant to a well-described IPS.

Proposers interested in exploring other research themes are strongly encouraged to contact the Smart and Autonomous Systems cognizant program directors.

Slide 17: What Proposals are Not Good Fits for S&AS

Jim Donlon: The focus of the Smart and Autonomous Systems program is on research focused on high-level cognition, self-awareness, and adaptation. Proposals that are primarily focusing on the "physical system" aspect of IPS, and underemphasizing advances in the capabilities that make them "intelligent" by way of the program's research themes, will not be a “Good Fit” for the program. Similarly, proposals that focus primarily on robotics themes of cooperative robots, human-robot interaction, self-interaction, collaboration, and augmentation are also likely to not be a "Good Fit" for the program. Purely software agents, with a lack of emphasis on embodiment, are not in scope for the program. Finally, a focus on system components including hardware design and development, low-level control, formal verification and validation, and security does not sufficiently address the program's focus on autonomy. NSF has several other programs that cater to such research activities.

Clarity on the scope of the Smart and Autonomous Systems program may be better understood by drawing a contrast to closely related existing NSF programs.

Slide 18: Relationships to Other Programs

Jim Donlon: The Smart and Autonomous Systems program is complementary to the National Robotics Initiative program, NRI 2.0 for short; and the Cyber-Physical Systems program, CPS for short. While all of the three programs address research on physical systems, it is important to note that research in this program focuses on advances in high-level cognition, self-awareness, and adaptation in these systems.

In contrast, the NRI 2.0 program is primarily focused on research in collaborative robots, human-robot interaction, self-interaction, and augmentation.

The CPS program is primarily focused on research in unmanned and robotic systems, collaborative control, mixed initiative systems, design and real-time systems verification, and trustworthiness.

The Smart and Autonomous Systems program differs from these other programs in its emphasis on achieving autonomy. This program particularly welcomes research toward systems that are self-aware, exhibit adaptability, utilize high-level planning and reasoning, and behave ethically.

Slide 19: Eligibility Requirements

Jim Donlon: Universities and colleges that are accredited and have a campus in the United States are eligible to submit proposals to the Smart and Autonomous Systems competition.

Moreover, non-profit organizations and non-academic organizations, which include research laboratories, museums, observatories, professional societies, and similar organizations within in the United States that are associated with educational and research activities are also eligible to submit proposals to the Smart and Autonomous Systems competition.

Do note that at most two Smart and Autonomous Systems submissions are allowed in a given year for any PI, co-PI, or Senior Personnel. Again, the constraint on Senior Personnel is to be especially noted.

It also is important to note that personnel associated with industry or for-profit research laboratories are not eligible to be PIs, but can be sub awardees, providing that their budget is only a fraction of the total award amount. Also important to note is that Federally Funded Research and Development Centers can be sub awardees if they provide unique resources unavailable elsewhere.

Slide 20: Proposal Submission

Jim Donlon: The Smart and Autonomous Systems program has set forth a deadline of July 31, 2018 for proposal submissions. This deadline is for both classes of the proposals. We strongly encourage the community to carefully read the Smart and Autonomous Systems solicitation and the Proposal and Award Policies Procedures Guide for a better understanding of the program and proposal preparation.

Do note that the Smart and Autonomous Systems program mandates yearly project representation at the annual PI meeting to be held in Washington DC. Consequently, travel to the yearly PI meeting must be accounted for in the submitted project budget.

Slide 21: Award Information

Jim Donlon: We anticipate that awards made to proposals received this July will be made in Fiscal Year 2019. The anticipated funding for the Smart and Autonomous Systems program is approximately $12M.

We expect to fund approximately 10 to 15 projects in the Foundational class with a budget range of $300K to $600K, and approximately 5 to 10 projects in the Integrative class with a budget range of $500K to $1M.

You may note that the minimum and maximum budget ranges for the project classes have been modified. As before, the Smart and Autonomous Systems program purposely has overlapping budget ranges to encourage PIs to submit to the appropriate project class based primarily on their research objectives and less constrained by disjoint budget ranges.

Slide 22: Supplementary & Single Copy Documents

Jim Donlon: The Smart and Autonomous Systems program will require proposers to detail a data management plan.

It is the responsibility of the Lead PI to submit a list of all project personnel and where applicable partner institutions. The list should include PIs and where applicable co-PIs, senior personnel, postdoctoral researchers, consultants, collaborators, sub awardees, and project advisory committee members.

A project with more than one PI must include a collaboration plan. The length of the plan must correspond to the complexity of the project.

If a postdoctoral researcher is involved in the project, then a postdoctoral mentoring plan is mandatory to be included.

It is also mandatory to include a single copy document that lists all of the collaborators of the proposal's PIs, co-PIs, and senior personnel. Please note that the NSF's Proposal & Award Policies & Procedures Guide provides a template and detailed instructions that you must adhere to in submitting this document.

Slide 23: Review Criteria

Jim Donlon: Proposals submitted to the National Science Foundation are evaluated for the Intellectual Merit and Broader Impact.

In addition to the Foundation's review criteria, the Smart and Autonomous Systems program will evaluate each submission on how the proposed research addresses the goal of achieving IPS that exhibit long-term autonomy in the face of uncertain, unanticipated, and dynamically changing situations, while requiring minimal human intervention. Specifically, proposals must address research that will enable knowledge-rich IPS to be more cognizant, taskable, adaptive, or ethical.

For Integrative Projects, the Smart and Autonomous System program will also evaluate each submission on the innovation in system integration. We will also judge the rigor of the required evaluation—that it is performed on physical testbeds, and corresponds to accepted scientific methodology, including well-defined hypotheses, controlled experiments, specific evaluation metrics, and statistical analyses.

Slide 24: Thanks!

Jim Donlon: On behalf of the National Science Foundation, we would like to thank all of you for your time and would also like to thank you for your interest in programs at NSF.

This presentation will be made available in the coming days on the Smart and Autonomous Systems program website.

If you have any questions pertaining to the Smart and Autonomous program, please follow the directions of the operator to call in questions now. Operator, we are ready to take phone-in questions.

Once again, we thank you all for your participation.