SAND2021-12411C

Estimating the Value of Automation for Concentrating Solar Power Industry Operations

Laura A. McNamara[a)], Randy Brost [b)] and Daniel Small [c)]
Sandia National Laboratories, Albuquerque, New Mexico 87185
[a)]lamcnam@sandia.gov, [b)]rcbrost@sandia.gov, [c)]desmall@sandia.gov

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Abstract. This poster summarizes findings from a small, mixed-method research study examining industry perspectives on the potential for new forms of automation to invigorate the concentrating solar power (CSP) industry. In Fall 2021, the Solar Energy Technologies Office (SETO) of the United States Department of Energy (DOE) funded Sandia National Laboratories to elicit industry stakeholder perspectives on the potential role of automated systems in CSP operations. We interviewed eleven CSP professionals from five countries, using a combination of structured and open comment response modes. Respondents indicated a preference for automated systems that support heliostat manufacturing and installation, calibration, and responsiveness to shifting weather conditions. This pilot study demonstrates the importance of engaging industry stakeholders in discussions of technology research and development, to promote adoptable, useful innovation.

1. INTRODUCTION

As automation becomes increasingly pervasive across every area of society, it is easy to take good automation for granted. Well-designed automation is seamless and invisible, seemingly performing laborious, time-consuming tasks without requiring a great deal of involvement, oversight, or supervision on the part of the human beneficiary. However, as anyone who has worked on designing, developing, and refining automated systems can attest, implementing high-quality automated technologies is an extremely challenging process, often requiring years of research, development, testing, and iterative refinement before an autonomy concept is ready for deployment in real-world applications.

Because the lead time for consumer-ready automation is usually measured in years rather than weeks or months, deciding which automation concepts merit investment can be a challenging forecasting problem, particularly if one's long-term goal is provision of a marketable technology that demonstrably benefits its intended users. Oftentimes, decisions about what forms of automation are worth pursuing occur in research environments that are, to a greater or lesser extent, removed from the envisioned domain of application – i.e., the real-world operations of a factory floor or a busy kitchen. Speaking as engineering researchers who have spent our careers in a laboratory workplace, we know how easy it is for researchers to come up with technically enticing research projects that, even when implemented as high-performing technologies, fail to be adopted for use, simply because the concept did not address a salient, meaningful stakeholder problem.

This mismatch between automation research dreams and stakeholder reality provided the topic for a small mixed-method study recently performed by Sandia National Laboratories in Albuquerque, New Mexico, USA. In this case, we were interested in collecting input from concentrating solar power (CSP) industry professionals about the potential for genuinely useful automation in CSP operations. As CSP proponents work to establish a viable niche in the burgeoning renewable energy sector, there are many potential opportunities for developing automated systems that could make CSP operations more reliable, cost effective, efficient, even safer. Rather than making educated guesses about beneficial automation for CSP facilities, why not ask the people working on the front lines of CSP to prioritize automation concepts intended for their use?

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Overview

Our poster summarizes findings from a study funded by the United States Department of Energy's Solar Energy Technology Office, or SETO. The authors spoke with CSP professionals working in different aspects of the CSP industry across four countries, using a semi-structured interview protocol to elicit their preferences for automation investment specific to CSP industry needs. Several of our interviewees quite literally took time away from operating a CSP facility to speak with us. Others are working in supporting sectors that provide technology and services to CSP facilities. Each of our respondents spent roughly an hour discussing the topic of automation with us, providing thoughtful, insightful commentary on how automated systems can help mature CSP into a sustainable element in the world's renewables portfolio. We are grateful for their time, preparation, and attention.

In the following pages, we briefly discuss the trajectory of this research project, including a methodological switch we made midway through the project when we realized that the elicitation approach preferred by our sponsoring agency (SETO) was poorly suited for one-on-one interviews with schedule-constrained CSP industry professionals. We quickly devised a more efficient protocol aimed at getting our respondents to indicate automation priorities, while also encouraging them to share their experience with automation in CSP operations. We summarize their responses, highlighting the automation concepts they collectively prioritized for investment, along with concepts that our respondents suggested could be beneficial for CSP industry futures. We also provide some lessons learned for colleagues interested in pursuing CSP stakeholder elicitation projects of their own.

2. DEVELOPING A VIABLE ELICITATION PROTOCOL

The Solar Energy Technologies Office of the United States Department of Energy invests considerable resources in research and development to support ongoing evolution and growth of the solar energy sector, both in the United States and abroad. In fall 2020, this project's principal investigator, Randy Brost, approached the Department of Energy's (DOE) Solar Energy Technology Office (SETO) with an idea about autonomy research. With a background in robotics and solar technologies, as well as experience working in industry, Brost was aware that innovative R&D concepts do not always transition effectively to industry if they fail to address the on-the-ground needs of industry practitioners. Accordingly, Brost proposed developing a survey that would be deployed to concentrating solar power (CSP) industry professionals, asking them to evaluate candidate autonomy applications in terms of their potential value for CSP industry operations.

When responding to Brost's proposal, the SETO program office agreed that the topic was important. They also suggested that the Analytic Hierarchy Process, or AHP, would be a useful framework for the analysis Brost had proposed. Never having employed AHP methods, Brost reached out to Laura McNamara, a Sandia social scientist with expertise in knowledge elicitation, for methodological assistance. McNamara agreed to work with Brost and colleague Daniel Small to develop and pilot an AHP protocol for deployment with CSP professionals.

AHP has been used many times to elicit stakeholder preferences in energy and environmental decision-making. However, properly deploying an AHP study requires some commitment on the part of respondents, who need to understand how AHP works if they are to provide reliable inputs to the analysis. As we explain below, AHP was not a good fit for the deployment conditions we were facing, which necessitated a shift to a more user-friendly elicitation approach.

Take One: The Analytic Hierarchy Process

First developed by Thomas Saaty in the 1980s, the Analytic Hierarchy Process is one of several multi-criteria decision making (MCDM) methods used in knowledge elicitation and decision-making research [1]. Today AHP is widely used in academia, industry, and government to evaluate and prioritize investments and technologies [2]. AHP is considered a "user friendly" technique because it organizes the decision criteria into a relatively simple two or three-level hierarchy, then uses systematic pairwise comparisons to help interviewees evaluate and rank preferences. These techniques minimize cognitive load for interviewees and facilitate transparency in the elicitation process, even among stakeholders with divergent perspectives on a problem space [3, 4].

Over the past two decades, AHP methods have been applied to dozens of energy and environmental decision-making projects, including solar technology investment opportunities. For example, Aragones-Beltran et. al. implemented an AHP methodology to help a major solar power company prioritize solar thermal facility projects in its portfolio [5], while Ho et al. developed an AHP elicitation process to evaluate designs for a high-temperature heat exchanger for use in solar thermal energy facilities [6].

Initially, we assumed that developing and deploying an AHP protocol would be a straightforward process. Working with a Sandia colleague who had recently left a career in CSP industry operations, we identified a set of decision criteria that we believed would plausibly represent the factors a CSP professional would consider when evaluating the potential benefits of various candidate automation concepts. We used our professional networks to identify a few dozen potential interviewees actively working in the CSP industry. We generated a set of candidate automation investment R&D concepts for prioritization, using the decision criteria we had identified with our industry colleague. We then set up a few pilot interview sessions with researcher colleagues to determine if the protocol we had designed needed any refinements, prior to beginning subject recruitment from our list of CSP industry professionals.

The pilot interviews were enlightening, and not in an encouraging way. Our initial set of AHP decision criteria items included impact on capital expenditures, impact on operational costs, impact on capacity factor and impact on a facility's total output. When deploying this protocol, we expected the criteria would be somewhat self-explanatory, given that our industry colleague had helped us identify them from his own recent experience commissioning a CSP facility. Instead, during the pilot sessions, each of these eleven criteria required careful explanation and some discussion before our pilot interviewees were satisfied they understood why each criterion had been included. One pilot respondent pointed out that our interviewees might have other criteria they would include, which was not something we could easily accommodate with our pre-planned decision model.

This was a fair point. Ideally, when AHP is deployed across a group of domain experts, they are directly engaged in development of the decision model. This ensures that everyone is on the same page about why specific decision criteria are included and what those decision criteria mean for the broader decision goals. Negotiating consensus among domain experts is a labor-intensive process requiring multiple iterations with engaged respondents [7]. Given our respondents' busy schedules, COVID-19 working conditions, and the pilot nature of this project, we did not have the time nor resources to run a multi-iteration process.

Secondly, AHP is more user friendly than more mathematically complex decision analysis methods. However, AHP does require some understanding of how the process works. Compared to other multicriteria decision analysis methods, AHP is straightforward, but it does require some training for participants to become accustomed to the scoring schemes used to assign weights to the decision criteria and to the options under evaluation.

The pilot process was a wake-up call. Running an AHP interview would obviously require some training and socialization with the interviewees, so they would understand the AHP structure and response modes prior to evaluating the decision criteria we were putting before them. Based on two pilot sessions, we estimated that a full AHP interview using our envisioned protocol would take a minimum of 90 minutes of interviewee time, likely spread over two, perhaps three interview sessions. We were quite sure that this time investment would make it difficult, if not impossible, for many of our extremely busy industry colleagues to participate, as multi-session interviews are exhausting even when one is not busily trying to run a major solar energy production facility. Moreover, we were concerned that respondents would be so focused on the AHP methodology that they would not be able to engage the topic at hand – automation futures for the CSP industry – as fully as we had intended.

Take Two: Borrowing from Stated Preference Methods

Given these findings, McNamara suggested the team set aside the AHP approach in favor of a less burdensome elicitation methodology. Social scientists working in consumer research often face significant constraints on the amount of time and effort that respondents are willing to invest in market interviews. At the same time, product development teams need insight into what features consumers value in a product, and some way of estimating how much they might be willing to pay for those features.

Stated preference methods are well-established in economics and market research to elicit consumer preferences for a product's possible features [8]. As the moniker implies, stated preference methods ask respondents to review a set of options and indicate which ones they value and by how much (in contrast to revealed preference methods, which require access to consumer behavioral data). One simple example of a stated preference survey is an option menu: respondents are given a list of possible product features, along with a menu of prices for each feature and a hypothetical budget they can use to "purchase" the features they care about. Because stated preference studies simulate common decision-making behaviors, they are easy for respondents to understand and can be completed in a few minutes with minimal interviewer guidance. Importantly, experiments comparing stated preference and AHP approaches indicate that the two generate comparable analytic outputs [9, 10].

A Stated Preference Protocol for CSP Automation

We built the revised protocol around an investment allocation exercise derived from stated preference techniques. In doing so, we set aside the decision criteria we had generated in the AHP process and focused on the automation concepts that we had identified in discussions with our colleagues at the DOE national laboratories, several of whom had extensive experience interacting with CSP operators and/or working in CSP facilities. With a few revisions, we were able to identify twelve concepts we would present as a menu of research and development investment options to our respondents. We also gave our respondents a hypothetical venture capital budget of 100M US dollars, which they could allocate as they saw fit among a portfolio of automation concepts. Figure 1, below, is the response form we used with our interviewees, with the stated preference instructions in the cyan block to the right, and the list of automation concepts to the left.

To ensure we were getting adequate buy-in from our respondents, we decided to include prompts and questions aimed at encouraging suggestions about automation opportunities in the CSP sector, based on their experience working in CSP technology, services, or operations. After all, one of the main goals of this project was eliciting input from CSP domain experts. We saw the investment allocation exercise as a means of opening a discussion about automation with our respondents, rather than an end *sui generis*.

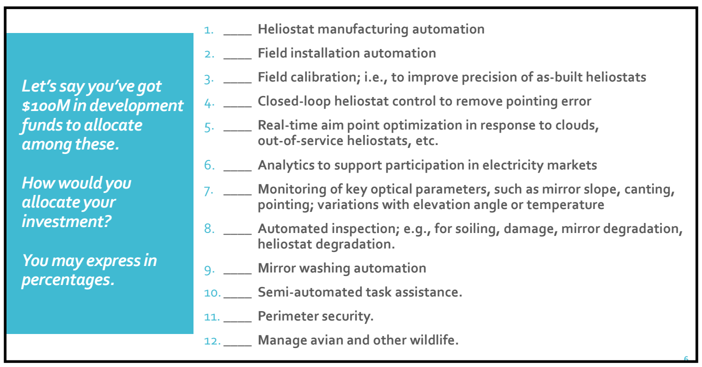

Figure 1: List of automation concepts for investment allocation exercise.

The final interview protocol consisted of four parts. We opened the sessions with introductions among the interviewers and participants, a review of the project's motivating problem, and a brief explanation of the investment allocation approach we would use in the discussion. In the second part of the interview, Brost explained each of the 12 automation investment concepts to the interviewees, to ensure that they understood what technologies were implied by each category. In the third part of the interview, we asked respondents to identify automation concepts that we had *not* included in the list, and to give us their perspective on the potential need for each of the items we had identified. Finally, we gave the interviewees their investment "funds" and asked them to indicate how they would allocate their monies across the automation concepts just discussed, including the items they had added to the list.

Again, it is important to point out that this approach was not strictly a stated preference exercise, because we allowed respondents to add their own concepts to our initial list. We also told respondents they were free to allocate some portion of the hypothetical investment funds to those concepts. This decision makes aggregation of respondent preferences quantitively iffy, as not everyone ended up working from the same list of automation ideas. However, given the motivating context of this project, we decided it was more important to give industry professionals a direct say in the automation opportunities we would report to the SETO program office, rather than adhere to a protocol built strictly on concepts we had generated ourselves.

3. DATA COLLECTION AND ANALYSIS

Participant Pool. We received approval from Sandia's Human Studies Board to recruit study participants in July 2021. We immediately began sending recruitment emails to potential participants, drawn from a list of 25 CSP domain experts that Brost and Small had generated in networking discussions with colleagues. Eight individuals in the pool failed to respond to three separate outreach emails sent at five-day intervals. Four declined participation. We received affirmative responses from 13 of the individuals we contacted. Two of these did not respond to further scheduling requests. In the end, eleven of our 25 pool members completed the interview, for a 44% response rate.

Participant Characteristics. Participants hailed from Germany, South Africa, Spain and the United States. All 11 interviewees were male and all had a minimum of a Master's degree, primarily in engineering related fields. Of the 11 participants, six had recently or were currently working in an operational capacity at an active CSP facility. The remaining five reported current employment in CSP technology and service companies. One of these respondents had recently retired, but like his colleagues, he reported ongoing consulting and/or partnership relationships with active CSP companies and operational facilities.

Interview Logistics. We completed the interviews in three weeks in August 2021. Before the interview, we sent each participant a PDF copy of the interview slide deck, in case technology failed us and we ended up falling back on telephone interviews. To our surprise, most of the interviewees treated the slide deck as read-ahead material, rather than a backup for our planned virtual meetings, and several were ready with a list of additional technologies they wanted to add to our preliminary list. All participants e-signed a statement of informed consent, as required by US government human subjects research regulations.

McNamara and Brost were present at all interviews. McNamara served as an interview facilitator, with Brost (the project's solar and automation expert) engaging the domain expert in technical discussions. We used Microsoft Teams or Zoom to run the interviews according to the interviewee's preference for a virtual meeting platform. During the interview, we used PowerPoint slides to guide the interviewee through each step of the discussion, keeping the slide notes section visible so the interviewee could observe McNamara taking notes during the discussion. All interviews were recorded and transcribed, and interviewees were provided a copy of the slides and interview notes if they indicated they wanted a copy for their records. Once we had finished the last interview, McNamara compiled the data using Excel spreadsheets for the investment allocation tallies and used AtlasTI, a qualitative analysis program, to code key comments from the interview materials.

4. SUMMARY OF FINDINGS

In this section, we briefly discuss two sets of findings from the interviews. First, we review the allocations and comments associated with the twelve items that Sandia provided in its mock portfolio of investment options. Because of space constraints, we will only discuss the top five items in the respondent list. We then present the automation concepts that our respondents volunteered in the interviews. Please note that our respondent quotes are not attributed to individuals to preserve privacy and confidentiality. All respondents who are cited had an opportunity to review their interview transcript prior to analysis. Respondent interview notes are quoted without major editing, which is standard when citing spoken comments.

Responses to Sandia's Automation Concepts

Investment in Sandia-provided items. Table 1, below, presents the respondent allocations to the twelve items that Sandia provided in the automation opportunities list. The table provides the number of respondents to decide to allocate some portion of their "funds" to this item, the total number of dollars allocated to the item, and the mean, minimum and maximum allocations made by the respondents who chose to "invest" in this item.

Table 1: Respondent "Investments" in Sandia Automation Concepts

RESPONSE ITEM *(from the original Sandia list, not including items suggested by respondents)*	Respondents allocating funds to item	Total virtual dollars allocated (M$)	Mean virtual dollar allocation (M$)	Minimum virtual dollar allocation (M$)	Maximum virtual dollar allocation (M$)
1) Automated monitoring of optical parameters	8	177	22	2	100
2) Heliostat manufacturing automation	5	145	29	10	50
3) Field installation automation	7	108	15	5	25
4) Automated closed-loop heliostat control	7	104	14	2	40
5) Automated real-time aim point optimization	8	99	12	2	30
6) Automated inspection of heliostat field	7	87	12	2	30
7) Field calibration automation	4	67	16	10	20

RESPONSE ITEM *(from the original Sandia list, not including items suggested by respondents)*	Respondents allocating funds to item	Total virtual dollars allocated (M$)	Mean virtual dollar allocation (M$)	Minimum virtual dollar allocation (M$)	Maximum virtual dollar allocation (M$)
8) Mirror washing automation	6	65	10	5	30
9) Semi-automated task assistance	4	42	10	2	25
10) Analytics for electricity market participation	4	35	8	5	15
11) Automation to manage avian/other wildlife	4	14	3	2	5
12) Automation to monitor perimeter security	2	7	3	2	5

Eight of the eleven respondents told us that systems that provide *automated, real-time monitoring and measurement of heliostat optical parameters* (#1), such as mirror slope and canting, could play an important role in optimizing CSP facility performance. In fact, one of the respondents allocated his entire $100M portfolio to this item, saying, "If the heliostats can't hit their optical requirements, you have spillage losses… you've got to hit that number to establish a viable consumer industry, and you can't get to that if you're losing 10% or more of your photons. There's no compensation for [spillage], nothing can make up the difference." The remaining seven respondents who allocated money to this item agreed that the systems to automatically monitor and measure pointing, mirror slope, canting and other parameters merited investment, although their allocations were not as generous, $2M to $25M among those respondents.

The next two items that received a high number of votes and/or monetary allocations were associated with facility construction: namely, *automated systems for manufacturing high performance, highly reliable heliostats*; along with *automated systems to install heliostats in the solar field*. Two of the seven respondents who voted for this item each allocated $50M to manufacturing automation. Importantly, both of these respondents were based in the United States, where labor costs are relatively high and automation could help reduce capital expenditure costs. As one of these respondents said, "We need to be able to make [heliostats] cheaply. Capital costs are always an issue. Making and manufacturing components inexpensively should be our number one priority." This same respondent expressed similar concerns about field installation processes, saying, "Putting a hundred thousand heliostats in the ground, you should be able to remove touch labor, it's something that lends itself to automation."

In contrast, participants responding from *outside* the United States were far less enthusiastic about systems that would reduce human labor requirements in the manufacturing and installation process. Pointing to their experience working in countries with high unemployment rates and low labor costs, these participants were decidedly unenthusiastic about automation that reduces human labor requirements. Not surprisingly, these individuals expressed a preference for item #9, *semi-automated task assistance*, even when they did not invest money in this concept directly. This is not an automated system per se, but a model for human-automation integration in which technology helps people do better work, more safely and efficiently, without removing humans from the loop. Respondents with experience working in countries with high unemployment rates emphasized the need for systems that grow labor force participation, because doing so builds goodwill toward the CSP industry, and expands economic opportunity.

The fourth item that won a clear majority of supporters was *automated closed loop heliostat control*, which describes the potential for automated systems that measure and adjust heliostat positioning to achieve more precise aiming. One highly experienced CSP professional felt that this technology could play an important role in reducing operational costs by increasing the lifetime of heliostat components. This individual invested $40M of his portfolio in this class of automation. Another CSP domain expert who allocated $10M to this concept commented, "I really care about closed-loop tracking, because that could help us reduce slope error and improve heliostat performance. I would invest in this."

A related item, *real-time aim point optimization*, received support from 8 of the 11 respondents, generating $99M in virtual investment from these individuals. This concept refers to systems that monitor ambient weather conditions and adjust heliostat positioning to optimize photon capture; i.e., in response to clouds moving over the solar field. Several respondents expressed skepticism that this was possible. One CSP facility engineer said, "It would be nifty if that could be done, it would be great to have." In contrast, another CSP domain expert commented, "It would be nice to develop, but there are ways of dealing with cloud transients now and I am not sure industry would rate that really highly at this point." It may be worth exploring this and other concepts in greater detail with industry professionals to determine if systems that support real-time monitoring and responsiveness to transient conditions merit significant investment.

Respondent Automation Concepts

Interestingly, although we offered respondents the opportunity to volunteer automation ideas for inclusion in their mock investment portfolio, only six of the respondents took us up on the offer, and most of their suggestions were also oriented to the solar field. This is perhaps because our Sandia list focused on concepts related to the solar field. However, during the interviews, about half the respondents pointed out that Sandia's list dealt solely with solar field technologies and processes, rather than systems and processes in the remainder of the plant. As one respondent observed, "The conventional part of the plant is important too… the turbine, pumps, valves and so on. My experience is that, in the CSP industry, we forget all of these conventional components, but in the day-to-day they create a lot of work for maintenance and is costly. That is also important." This individual suggested that automated systems that enable operational teams to monitor, anticipate and respond to changes in the status of "conventional" systems would enhance the performance and efficiency of these facilities. Below we list the suggestions that our respondents volunteered, along with a brief description of what the concept entails and why it is important, using the respondents' own words.

Table 2: Automation Concepts Volunteered by Respondents

CONCEPT	**RESPONDENT ALLOCATION?**	**DESCRIPTION** (FROM RESPONDENT INTERVIEW NOTES)
Automated supervisory control of flow control valves	$15M	If the inlet valves to the solar field loops are automated, it offers a whole new world of optimization opportunities. Valves can be adjusted to cater for transient weather conditions, inhomogeneous solar field cleanliness/incorrect calibration… and thereby ensure that the maximum thermal energy is harvested from the solar field.
Greater use of machine learning across the facility	$20M	Our biggest challenge is low solar field efficiency, which translates into major production/revenue losses, especially during winter months. There are many factors that can cause the reduced solar field efficiency, but it is hard to identify these factors and quantify their individual contribution… I think greater use of machine learning technologies might create a means of estimating, for example, the cleanliness of each solar field loop, or contributions from things like calibration, flow, and other factors.
Real time condition monitoring across the facility	$15M	If the plant starts building up historical database of behavior of equipment onsite, not just in solar field, but power block, then creating a normalized trend of what's normal and has certain bounds for when things or alarms like that, can be hooked up to your maintenance plan onsite, kicks out automate inspection or maintenance tasks to be done on certain systems. In my experience, I think, if we had something like that onsite [it] would have helped a lot with regards to unnecessary downtime.
Improvements in heliostat reliability	$30M	After being involved in CSP facility operations for six years, I would like to design a heliostat that needs less maintenance activity… the maintenance and availability of heliostats is critical… this is a place for improvement.
Improvements in measuring/monitoring receiver temperature and windspeed	$20M	It's difficult to when you're trying to measure the state of your equipment, [you] place as close to source as possible. On and near the receiver, things don't survive very well, especially instrumentation, even thermocouples rated for these temperatures. We wish there were better ways to measure the actual temperature of the receiver surface, the outer surface of the tube [and] back surface, also the windspeed at the receiver. Instruments don't survive up there, so that's another area I can think of.
Drones applied to mirror characterization	–	There is work being done in this area and I am happy to see this technology being used to monitor the solar field.
Wireless communications and control in CSP facilities	–	You are dealing with a large number of heliostats in extreme weather conditions, and there are wiring failures, so wireless communications with heliostats, ability to control wirelessly, even distribute power to heliostats.
Advanced simulations to monitor, test, provide situation awareness of plant components, subsystems and systems	$20M	Perhaps a real time model could operate in the background alongside the operators, with inputs from the plant, to help detect deviations, monitor the operation of the actual plant, and warn provide warning. Simulations of plant operations could be useful for training operators.

CONCEPT	RESPONDENT ALLOCATION?	DESCRIPTION (FROM RESPONDENT INTERVIEW NOTES)
Predictive control assistance for balance of plant	$10M	Coming from a background in chemical engineering, I often think about how we can use automation in CSP facilities. You've got the tower and a power block and a lot of other process. All of these have to talk to each other. Relevant data could be recorded by sensors instantaneously, or previous data, historical data; how do we interface all these together to have an automated facility that's also predictive? Historically you might know a day is really cloudy but you've got high demand, so how do you decide to store more during the day, deliver at night?
Multi-tower coordination across multiple facilities	$10M	Can we use automation to have multiple towers working together, using all kinds of different data in order to have some kind of predictive ability for conditions or demand? But robustness is necessary, and with machine learning or a neural network, you might not know where it will fail. So those are considerations to keep in mind.
Automated maintenance	$5M	Replacing a broken mirror. That's one example. Relubricating a gear box.
Maintenance on automated systems ("Meta-maintenance")	$5M	Automated systems need to be maintained, so can we have systems that help with that?
Cybersecurity	$1M	Automation to detect and alert to cyber security problems.

5. CLOSING COMMENTS

In this paper, we have summarized findings from a small, mixed-method knowledge elicitation study with CSP industry experts, aimed at eliciting their thoughts about what forms of automation might be most beneficial across the CSP facility lifecycle. Although we initially sought to use the Analytic Hierarchy Process, we redirected toward a simpler elicitation framework based on stated preference methods to develop a "respondent friendly" elicitation process. Our approach paid off in two ways: first, every interview except one (with a very thoughtful, conscientious respondent) was completed in less than an hour. Second, our revised approach made space for respondents to offer their own ideas about automation needs, drawing on their experience working in or alongside these facilities. When asked to provide feedback on the process, all 11 interviewees commented that the investment allocation exercise made sense and was easy to execute. Several emailed afterwards to say they had enjoyed the discussion. As interviewers, we similarly enjoyed learning about the experience of CSP operations from individuals working on the front lines of this industry and are grateful for the time they invested in this process.

REFERENCES

[1] N. J. Cooke. 1994. "Varieties of Knowledge Elicitation Techniques." *International Journal of Human-Computer Studies* 41(6): 801-849.

[2] S. Sipahi and M. Timor. 2010. "The Analytic Hierarchy Process and Analytic Network Process: an Overview of Applications." *Management Decision* 48(5-6): 775-808.

[3] ~~R. W~~T.L. Saaty. 1987. "The Analytic Hierarchy Process: What It Is and How It Is Used." *Mathematical Modelling,* 9(3): 161-176.

[4] T. L. Saaty. 2001. "The Seven Pillars of the Analytic Hierarchy Process" in Koksalan, S. and Ziont, M., eds. *Multiple Criteria Decision Making in the New Millennium*, Lecture Notes in Economics and Mathematical Systems, Volume 507, Springer-Verlag.

[5] P. Aragones-Beltran, F. Chaparro-Gonzalez, J. P. Pastor-Ferrando, and A. Pla-Rubio. 2014. "An AHP (Analytic Hierarchy Process)/ANP (Analytic Network Process)-based Multi-Criteria Decision Approach For The Selection of Solar-Thermal Power Plant Investment Projects." *Energy,* 66: 222-238.

[6] C. K. Ho, M. Carlson, K. J. Albrecht, Z. Ma, S. Jeter, and C. M. Nguye. 2019. "Evaluation of Alternative Designs For A High Temperature Particle-to-sCO2 Heat Exchanger." *Journal of Solar Energy Engineering,* vol. 141 (2).

[7] M. Meyer and J. M. Booker. 2001. *Eliciting and Analyzing Expert Judgment: A Practical Guide*. ASA-SIAM Series on Statistics and Applied Mathematics. Philadelphia, PA: SIAM Press.

[8] E. P. Kroes and R. J. Sheldon. 1988. "Stated Preference Methods: an Introduction." *Journal of Transport Economics and Policy* 22(1): 11-25.

[9] Z. Kallas, F. Lambarraa, and J. M. Gil. 2011. "A Stated Preference Analysis Comparing the Analytical Hierarchy Process versus Choice Experiments," *Food Quality and Preference* 22(2): 181-192.

[10] M. E. Ben-Akiva, D. McFadden, and K. Train. 2019. *Foundations of Stated Preference Elicitation: Consumer Behavior and Choice-Based Conjoint Analysis*. Norwell, MA: Now Publishers.