What do We Really Know About Robots and the Law?

William D. Smart

Draft of March 24, 2016*

1 Introduction

We Robot is a gloriously multi-disciplinary conference. Legal and policy scholars discuss current and future robots and robotics technologies, and offer sage advice about how to legislate them. Roboticists talk about current and future legal and policy structures, and hypothesize how these might interact with the underlying technology. Ideas flow back and forth, and we make concrete steps towards understanding how law, policy, and robotics interact, both today and in the years to come. Or do we? How much do legal and policy scholars know about real robots, and the technology that goes into them? How well do roboticists understand the law and how it works? The frank answer to both of these questions is "not as much or as well as they should".

This paper provides the results of a survey taken at We Robot 2013, designed to assess how well me know each others fields of expertise. We Robot participants were invited to answer a set of questions to determine how well they understood the basics of robotics technology and some of the legal issues surrounding it. Participants self-identified as either a roboticist, a legal scholar, a policy scholar, or other and gave some basic demographic information. The paper gives an analysis of the results of the survey, highlighting some interesting trends. Without giving the game away, none of us knows as much as we think we do, and theres still a lot of work to be done to educate each other, and to really understand the basics of each others fields. This is, however, vital for the long-term success of both We Robot and the ideas on which it is founded.

In addition to presenting the results of the survey, the paper identifies some key areas where we can make progress in educating each other, and provide some concrete suggestions for how to do this. The ultimate goal of this paper is to scare everyone just a little bit, but then to provide a ray of hope for the future. If we know where our own shortcomings are, we can work to directly address them, and to make sure that, when we are caricaturing

*In this draft, some things are asserted without the citation of statistical evidence. This will be corrected in the final draft, which will be available shortly

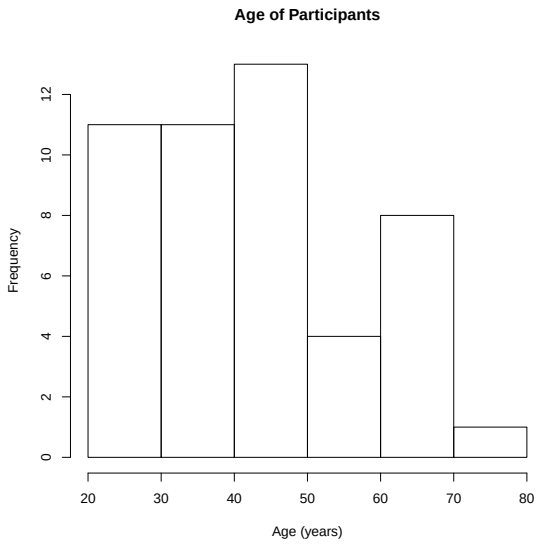

Figure 1: The distribution of ages of We Robot questionnaire participants.

our colleagues areas of expertise, we are at least 80% right. We'll start, however, with some demographics.

2 Who Comes to We Robot?

We had 49 people fill in our questionnaire at We Robot 2013. Of these, 18 (37%) self-identified as female, 29 (59%)[1] as male, and 2 (4%) as other-gendered. The median age was 43.5 years, the mean was 43.56, and the standard deviation of age was 13.07 years (see figure 1)

We are, as you might expect, a fairly educated crowd. 34 (69%) of participants had some sort of doctoral degree, 9 (18%) had a Master's, 5 (10%) had a Bachelor's, and 1 (2%) had "some college/Asociate's degree".

The early We Robot conferences were more heavily skewed towards legal scholars. In 2013, 24 (49%) of participants self-identified as having a background in Law, 4 (8%) as being in Policy, and 6 (12%) as being roboticists. A further 15 (31%) claimed to have another background. These backgrounds are captured in table 1.

The conference was also skewed towards academics: 28 (57%) of participants self-identified as being an academic, 8 (16%) as being in legal practice, 5 (10%) as being in industry, and 8 (16%) as being "other". These numbers are captured in table 2.

[1]We report percentages of the number of responses to a particular question. Since some questions were not answered by some participants, the same numerical value may be reported as different percentages in different questions. Also, since we round to the nearest whole percentage, they may not quite sum to 100% in some cases.

Background	**Number**
Law	24
Policy	4
Robotics	6
Other	15
Journalism / Writer	3
Philosophy	2
Biology, Geology, GIS	2
Electrical Engineering	1
Education	1
Business Consultant	1
Law and Policy	3
Law and Robotics	1
Policy and Technology	1

Table 1: The backgrounds of We Robot questionnaire participants.

Workplace	**Number**
Academia	28
Legal Practice	8
Industry	5
Government	2
State and non-profit	2
Journalism	1
Media	1
Self-Employed	1
Retired	1

Table 2: The workplaces of We Robot questionnaire participants.

Roomba Lego nxt. I have 9 of them.
A toy robot
3d robotics quadcopter
Neato, RWI B14 (2)
VGo, Pioneers, ATRV-Jr, VGTVExtreme, 2 wheelchairs, Nao, Romibo, etc.
Roomba
Pleo
Mint
Roomba, Sphero, Parrot AR Drone
Many
Several Drones
Many
Personally Made

Table 3: Robots owned by participants.

3 Got Robot?

Since We Robot is concerned with robots, it seemed natural to try to find out how many people at the conference had direct experience with a real robot of some kind. 14 (30%) of participants actually owned a robot, while 33 (70%) did not. Two lucky participants both claimed to have "many" robots, while Roombas and drones were also popular. A number of research robots (such as the RWI B14, Pioneers, ATRV-Jr, etc) were also present, presumably from the robotics researchers. The full list of responses is given in table 3.

A larger number of participants, 36 (77%) had controlled a robot, while 11 (23%) had not. Of the 14 people who owned a robot two self-reported never to have controlled a robot, presumably because they owned Roombas. Of the 33 who did not own a robot, 15 reported controlling someone else's robot. We found the number encouraging but surprising, since it is over twice the number of people that own their own robot. However, over three in four people at We Robot can be expected to have direct experience controlling an actual, real robot.

Controlling a robot using a joystick or an app on your phone is one thing, but *programming* it to perform some task is another. Of our participants, 28 (59%) had programmed a computer before, while 20 (41%) had not. Again, we find this number very encouraging, given the demographics of the participants. Three in five We Robot participants understand something about computer programming, and algorithmic thinking. A smaller number, 16 (33%), had programmed a *robot* before, while the remaining 32 (67%) had not. Of the 14 people who owned a robot, 10 (71%) had programmed a robot (presumably their own). Of the 33 who did not own a robot, 28 (85%) had not programmed a robot. 4 people owned a robot, but had not programmed it, and 5 had programmed someone else's robot.

If we assume that all 6 roboticists program their own robots, then this means that 10 of the remaining 39 participants, about 25%, also have experience programming some sort of robot. Even if we're pessimistic, and assume that 4 of these 10 are those from technical

	Programmed a Robot?	
Background	**No**	**Yes**
Robotics	0	6
Law	21	2
Policy	3	1
Other	8	7

Table 4: Participants who have programmed a robot, broken down by background.

factor	$p < 0.001$	$p < 0.01$	$p < 0.05$	$p < 0.1$
programmed a computer	4	17	28	38
programmed a robot	12	20	27	29
own a robot	7	13	23	28
teleop a robot	4	8	20	28
background	3	8	17	25
work	0	4	10	17
gender	2	3	6	15
education level	0	2	2	7

Table 5: Summary of the factors significantly affecting rankings for the three question groups. Note that the numbers in each column subsume the numbers in the preceeding columns.

fields that are not robotics (Electrical Engineering, Business Consulting, Law and Robotics, Policy and Technology, from table 1), this leaves 6 people (9%) from non-technical fields with robot programming experience. This is borne out in the breakdown shown in table 4.

Finally, and not surprisingly, programming computers and programming robots are significantly correlated with each other ($X^2 = 14.2844$, $p < 0.05$). However, it seems that one participants programs robots, but not computers, which seems odd.

4 What Matters?

In the questionnaire, we asked participants to rate 25 things (see section 12) on a 7-point scale three times, to judge (1) where on the mechanism-robot continuum they thought it lay; (2) how autonomous they thought it was; and (3) how well they thought that they understood the underlying technology. While there are no "right" answers to these questions, we an tell something from the relative rankings that the participants gave. In general, roboticists were more likely to classify something as a mechanism rather than a robot, and to claim a greater understanding of the various systems than were the other groups.

We performed Analysis of Variance (ANOVA) tests to determine what factors (from the demographics questions) had an affect on the rankings given then these three question groups. A summary of these results is given in table 5.

This means that (with a significance of $p < 0.001$), the thing that most affects your answers to the three groups of questions is whether or not you have programmed a robot in

the past. The first four factors in table 5 correlate strongly with background (specifically, with a background in robotics) and work (specifically, in academia).

All four of these factors correlate strongly with an self-reported increased understanding of the technologies, which is not surprising. They also correlate with a decreased variance of scores (the scores are more consistent between participants), and a lower average score on the mechanism-robot spectrum for most devices.

We would expect that the scores on the mechanism-robot spectrum would correlate with those on the autonomy spectrum for the same devices. This comes from our assumption that mechanisms are not autonomous and robots are (which is more-or-less borne out in the results reported in section 5). However, it is not well-supported by the numbers. In fact, at a significance level of $p < 0.001$, only one system, "a powered exoskeleton worn by a human", shows this correlation. At $p < 0.01$, "the anti-locking brake system on your car" is added. At $p < 0.05$, we also get "a Roomba vacuum cleaner" and "landmine". This suggests that, contrary to our original belief, how mechanistic a system is does has little to say about its level of autonomy.

We also assumed that there would be some correlation between how well a system was understood and how autonomous or how like a robot it was. There turned out to be no evidence for this in the data, other then a mild correlation ($p < 0.05$) between understanding and the level of autonomy for "light switch" and "the system that won the recent DARPA robot challenge".

Despite not finding what we expected, we take these results as encouraging. The things that most strongly influence rankings (programming robots and computers, controlling a robot, owning a robot) are all things that we can control. We can buy robots and play with them. We can learn to program a computer. Doing these things should bring the perspectives of We Robot participants from different areas and with different backgrounds more into alignment, and help reduce misunderstandings about the underlying technologies.

5 Mechanism or Robot, Redux?

We asked the participants "In your mind, what are the main distinctions between a mechanism and a robot." There were a wide variety of answers to this question, which reflects the wide variety of opinions that we hear at the conference.

Two participants thought that the main distinction was the predictability of the system:

1. unpredictability

2. mechanisms always do the expected action unless it is broken

This seems to indicate that if we understand the system in enough detail to predict what it will do, then it is closer to being a mechanism than a robot. It might also implicitly claim that robots are more unpredictable than mechanisms. However, we claim, this is very subjective, since it relies on the understanding of the observer; if a roboticist looks at a robot, they might be able to predict exactly what it will do in a given situation, pushing

it towards a mechanism. However, a lay person, watching the same thing, might classify it more as a robot, since they don't know what's going on under the hood.

This complexity theme was picked up by a number of other comments:

1. Embodiment, degree of automation, technical complexity, anthropomorphism.
2. Robot involves some processing of environment, purpose, decision-making
3. A degree of sophistication
4. complexity
5. automation complexity
6. feedback
7. complexity of feedback, pattern recognition, and learning systems

These also start to hint that a robot has sensors and somehow processes information in a way that a mechanism doesn't. Also that a robot is more sophisticated and complex than a mechanism is. Again, we contend that this is observer-dependent; the less one understands a system, the more like a robot it is. However, we would argue that mechanisms also process information and sense the world, albeit in a rudimentary way. A light switch, for example, was universally classified as being a mechanism in an earlier question. However, one could argue, that this device senses the state of the switch, and acts accordingly, either allowing current to pass through to the light or not. In fact, there is a classic argument in robotics as to whether or not a thermostat is a robot or a mechanism. It can be purely mechanical, in the same way that a light switch is, but it can seem to make a decision based on sensing the environment. The question has, of course, been further muddied by the arrival of the Nest thermostat, a high-tech thermostat with learning and (rudimentary) artificial intelligence appearing on the market.

Other comments centered around the independence of the system; does it need a human to cause it to do something:

1. Need for human input
2. mechanism doesn't do anything until person or something activates it with physical act; tends to be simpler (even one activity)
3. mechanisms are actuators controlled by humans, with either human or electro-mechanical sensory input. Robots sense their environment with electro-mechanical sensors, have one or more computer processors to analyze and fuse the data, which then determines what action to take. The robot acts in the physical world (not just a computer display).

This raises the question of a robot that only does what it's human controller tells it to. Is it a robot, or is it a mechanism? The last comment as hints that robots might have more on-board processing, and need computers of some sort. It also argues for embodiment; a robot physically acts on things in the real world. This is backed-up by one more comment:

1. A robot can manipulate the environment

We believe that the participant means *intelligently* manipulate the environment in this case, since even simple mechanisms (like a trebuchet, for example) can manipulate the real world in some way (by throwing parts of it into the air).

Finally, the largest set of responses addressed intelligence and autonomy: robots are things that can make decisions on their own, based on some sort of reasoning process:

1. autonomy; discretion; ability to learn and change choices
2. Sensor based decision making, embodiment, mobility, and knowing it when I see it-¿ it's a fuzzy defintiion
3. Sense, Think, Act
4. AI and actuators
5. Autonomy and motion
6. Autonomous behavior, mechanical vs. automated
7. A robot can run independently of its operator
8. autonomy
9. Thinking, Independence
10. Robot has a semi autonomous system and is programmed
11. Mechanism is human controlled

These comments seem to suggest that a robot is a physical container for artificial intelligence. This implicitly implies some sort of (sophisticated) sensing, computational hardware, and actuation that can affect the world. We believe that most of the comments also implicitly assume that robots are embodied, although few of the comments explicitly mention this.

So, it seems that the distinction between a mechanism and a robot is complicated. How much different would the conference be if it was We Mechanism? We're willing to bet that the answer is "not much", based on the questionnaire responses. Just as in the field of robotics itself, there's no real consensus of what a robot is, and where the dividing line between robot and mechanism is.

Is this even a useful question? Would it help the discussion at We Robot to have a more precise definition of what we mean by "robot"? Would it benefit us to have an ontology of systems that we talk about, and a commonly-used way to talking about them? It might since, unless we're all talking about the same thing, discussions can start to get misleading. If I say "robots are evil", it *really* helps to understand what I mean by *robot*, since it frames the discussion. Perhaps a more precise statement, such as "category 2b robots are evil" would be clearer and help us to make faster progress, assuming that we can decide how to split up the world of robots and mechanisms into nice, neat, non-overlapping bits. However, this might be a rabbit hole that we, as a community, might not want to go down.

Situation	Law	Policy	Robotics	Other
1	1.10	1.33	**0.00**	0.97
2	0.17	0.25	0.33	0.18
3	0.91	1.33	0.50	1.36
4	1.03	0.92	**0.00**	0.82
5	1.16	2.25	**0.00**	1.15
6	0.80	1.00	1.00	1.27
7	0.91	**0.00**	**0.00**	0.93
8	0.10	**0.00**	**0.00**	0.79
9	0.11	**0.00**	**0.00**	0.79
10	0.24	0.67	0.50	0.45
11	0.26	0.25	0.33	0.20
12	0.23	0.25	0.50	0.57

Table 6: Variance of responses to a variety of legal questions, by area of background.

6 What Law Applies?

We asked participants to identify what area of law (contract, property, criminal, or tort) applied in twelve hypothetical situations involving a robot. The main think that we learned from this question is that the law is complicated, and hinges on details and intent. We learned this mostly because of feedback from the participants, as they explained how poorly-designed the question was. There wasn't enough context for most of the situations; the intent of the various parties was not make clear; often multiple areas of the law applied. In the questionnaire, we allowed only one area of the law to be chosen. We readily concede that the question was poorly-designed to elicit the areas of law that might pertain.

However, we can learn something from the responses. By coding each area of the law with an integer, and calculating the variance of the responses, we can get an idea of the amount of agreement among the respondents. A variance of zero means that everyone agreed on a single answer. Greater than zero implies disagreement; since the scale is nominal, we cannot say much about the *emph* size of the non-zero variances. Table 6 summarizes these variances, broken down by the self-identified areas of expertise of the participants. The interesting thing to note is that the legal scholars and those who identified as "other" *never* completely agree on the area of law that applies. However, policy scholars are in complete agreement in 3 of the 12 examples (25%), while the roboticists agree in 6 out of the 12 (50%). We interpret this as an over-simplification of the situation by those less familiar with the details of the law. As we said above, the legal interpretations of these situations are (to us, as roboticists) subtle and often not as clear-cut as we would like them to be.

7 Most Pressing Problems

We asked participants "In your opinion, what are the most significant legal or policy problems for robotics at this time?" Again, there was little consensus here, and a wide variety of opinions. The largest block of responses dealt with liabilities and torts:

1. Allocating responsibility
2. Where to place the blame.
3. assigning liability when a robot commits a crime or property damage
4. What happens if a robot hurts someone.
5. Delegation, Liability
6. sorting liability + responsibility in design
7. liability
8. Liability
9. Safety related torts
10. Need to have standards to reduce tort suits
11. Product liability
12. Liability - will the manufacturers be liable for all robot malfunction?

Privacy and regulation were also common:

1. Privacy, data security, products liability/negligent design, insurance
2. Space and privacy boundaries, safety
3. privacy
4. Regulations from the FAA
5. FAA
6. regulators

And, there were a variety of other responses:

1. patent thickets, privacy intrusions, insurance
2. personification and dignification of machines, in a social environment where too few people are accorded the full dignity of personhood

3. poorly-prepared judges
4. Having an understanding of the technology, not thinking of scifi
5. security
6. human expectations and vested political interests
7. Agency
8. Contracts
9. Getting metaphors right
10. The need for imagination conflicts with a system based upon precedent
11. Authority to execute actions when a person says No

We present these with little comment, except to note that the breadth of responses is very broad. There are lot of possible concerns with robotics, captured by the statements above, all of which are valid. However, as a field comprised of individuals, we have a limited amount of time. This suggests, to us at least, that it would be useful to start thinking about what issues are the most immediately pressing: what problems, if we don't think about them *now* will come back to bite us in the next few years. Or, perhaps, what problems are so important that, if we don't address them now, will cause huge problems in the longer term?

While it's probably futile to try to agree on a small set of issues to focus on, it might help to focus the discussion to think about *how* and *when* each of the problems listed above will affect us if left unstudied. This way, we might be able to most effectively apply our (all too limited) resources to best effect.

8 Metaphors for Robots

The metaphors and analogies we use to explain systems matter because they strongly frame the system we are describing. Using "an electronic letter" as the analogy for email is different than using "an electronic postcard". The former implicitly suggests a private, and perhaps serious, communication, while the latter suggests a more open, less serious one.

So, how do we talk about robots performing tasks? We asked the participants to provide metaphors for a robots in a number of situations. Here are some examples of what they said:

A Roomba robot vacuum cleaner

> maid (6); housekeeper (5); tool (4); cleaner (3); cleans [things] (2); machine (2)
>
> broom; co-worker; helper; moving vacuum; muppet; out of sight out of mind; randomness; random walk; reliable commodity; servant; service; Service robot; smart appliance; trained bug; vacuum cleaner

The most popular way of thinking about this cleaning robot was as a (human) maid; someone that cleans for you when you're not there. If we code each of the responses as describing in terms of a living thing (generally a person) or a machine, 17 out of 37 descriptions (46%) analogize the robot as a living thing. Since 5 of the metaphors describe concepts rather than things, we discard them, leaving us with 53% of the descriptions in terms of a living thing, and not a machine. This seems natural, given our tendency to anthropomorphize things, and to describe them in terms of things that we already know. Maids do, of course, clean for us much in the same way that Roombas do.

An android that cleans your home using your vacuum cleaner and cleaning supplies

> maid (13); servant (3); housekeeper (3); helper (3); humanoid uses my tools (3); service (2); wife (2)
>
> cleaning lady; co-robot; co-worker; mom; reliable; smart appliance; Smart cleaner

This system is, functionally, the same as the previous one; it cleans your house while you're not there. However, it is explicitly framed with an android doing the work, and this dramatically affects the responses.

The overwhelmingly most popular response this time was "maid", with 13 out of 37 responses (35%). This response has over four times the number of replies as the next most popular, "servant". If we again code for analogies to humans or machine, 25 (68%) of responses do this, rising to 81% if we remove answers that are purely descriptive or that describe concepts. My simply changing the morphology of the system, we have changed the what that the participants analogize it, making it appear somehow more human. While the framing of this problem was quite crude (and the reader could even argue that it's not framing, since we intentionally describe a physically different system), we believe that it is a clear example of two systems that are functionally identical being analogized very differently. As we have previously argued,[2] this could profoundly affect the laws that grow up around this new technology.

Is also notable that adding an android to the description cause some of the responses to be explicitly gendered. The system was described as a "wife", a "mom", and "cleaning lady". Even leaving the problem of gender stereotyping aside, this raises significant issues, since both "wife" and "mom" are analogies that carry a lot of entailments with them, elevating them beyond simple machines. 11% of responses were gendered in this way. If we include some of the other responses that suggest a female role,[3], this number increases to 54% of responses. This compares to 30% for the previous question, if we include "maid" and "housecleaner".[4]

[2] How Should the Law think about Robots?, Neil M. Richards and William D. Smart.

[3] And here, the author freely admits to applying his own gender stereotypes to "maid" and "housecleaner".

[4] We note, with some irony, that the question was formed with an *android* which, from the greek, has a male gender. A female robot with a human shape would be a *gynoid*.

An automated shopping trolley

> helper (3); service (3); cart (2); salesclerk (2); shopper (2); shopping cart (2); tool (2)
>
> assistant; car; careful; co-worker; dog; family member; follows me in store, get what I point at; inventory picker; labor saving device; machine; muppet; RC car; red wagon; shopping assistant; smart appliance; subway with stops; transporter; vehicle

This example generated a wider variety of responses, with no clear winners. Participants were split between analogizing it as a person (46%) and a machine (54%), not counting concepts. Interestingly, only two people analogized it for what it actually is, a shopping cart.

A robot you encounter in a grocery store carrying a box of cereal and a loaf of bread

> shopper (6); stockboy (3); employee (2); helper (2); personal shopper (2) ; worker (2)
>
> carb-bot; citizen; cleric; co-robot; co-worker; friend; gluten tolerant; motorized shopping; must be doing the shopping for a person; obsticle; prop; shelf replacer; shopping tool; smart appliance; someone else's; thief

Again, this is a similar system to the previous question, but framed differently. By using the word "holding", we imply that it has hands and presumably a humanoid shape. We don not explicitly give it a purpose, but allow the participants to determine it from context. However, functionally it is a similar system: something that carries groceries around a store.

Again, the reframing causes more human-like analogies (71%) than machine-like (29%), again removing concepts. This is a dramatic increase over the previous question although, as before, we admit that there is more the just a simple framing effect going on here. However, once again, the implied morphology of the robot system affects how people choose to describe it.

We note, however, that the specificity of the question brought out the wags, who described the system as "carb-bot" and "gluten tolerant", both of which seem more possible than relevant. The analogy to a "cleric", however, escapes us.

9 Other Participant Comments

We also asked the participants what question we *should* have asked, but didn't. Four of the participants responded, two with comments ("None; it was already too long." and "This was good and hard.") and two with additional questions:

1. What is a human?

2. Are robots a risk or benefit to society?

The first was asked by someone who self-identified as being an electrical engineer (rather than a philosopher). The second was asked by someone who self-identified as being in both Law and Policy. Both questions are very relevant to We Robot, especially if one takes "what is a human?" in the human-in-the-legal-sense interpretation. We attribute the lack of other responses to questionnaire fatigue, agreeing with one of the other commenters that, in retrospect, the set of questions was probably too long.

Finally, we asked the participants if there they had any other comments about the questionnaire. This provoked a little more activity, with 9 responses. Two participants commented that some of the questions were confusing, and lacked detail, especially for the questions pertaining to the law. This point is well-taken, and confirms that the author, at least, still has much to learn about the law, even to the extent of how to frame questions that make sense. Other comments showed varying levels of support for the intention of the questionnaire. However, perhaps the most interesting comment was: "They don't get mad at you like people do therefore more trustworthy than a person". This comment seems to equate trust with even-headedness and, presumably, predictability. Whether or not his is actually the case, we leave as an exercise for the reader.

10 Suggestions

How can we use the results presented earlier in this paper to strengthen the We Robot community? We have some modest suggestions:

1. Buy a robot and play with it. Drive it around a bit. These are the simplest ways to learn about robots, and our results suggest that having direct experience of a robot is one of the most significant factors in how you respond to the questions in our survey. If everyone in the We Robot community did this, we believe that we would have a more common understanding of what real robots, what they can reasonably do, and what their limitations are.

2. Learn to program a computer or a robot. Along with having direct experience of a robot, knowing a little bit about programming, either computers or robots, is a significant factor in determining how you respond. Although harder than just driving a robot around, these are probably the two most effective things that a member of the We Robot community can do to have a more realistic understanding of current robot technologies.

3. Pinning down definitions a bit. An ontology of the things we talk about, and a common nomenclature that we can use to reduce confusion. This might not answer the "what is a robot?" question, but it can bound the systems that we talk about for a given argument.

4. Think critically about *how* and *why* the problems enumerated in section 7 will affect society if left unstudied, to better understand where we can, as a field, most profitably spend out time.

5. Avoid the Android Fallacy[5] at all costs. As we showed in section 8, framing can have a dramatic effect on the metaphors and analogies that we use to describe robot systems. This, in turn, can lead to loaded descriptions and, we claim, will make it harder to form clear, appropriate legislation and policy.

6. Be aware that the areas in which you are not a specialist are at least as subtle and complex as your own. This seems somewhat obvious, but it is often easy to form cartoons of other areas, simplifying them because of our own limited understanding. This is a danger, especially for roboticists who like clear-cut, unambiguous answers to problems. Understanding that the legal implications of a particular robot in a particular context might be subtle, and depend on things that seems only tangentially related, is vital if we are to design systems that can actually be deployed in the real world.

7. Help us design a better version of the questionnaire for distribution to the larger community. We think that we've uncovered some interesting stuff with this version but, as has been pointed out to us, the design is flawed, especially in the legal questions. If we can put together a more complete, better-designed set of questions, perhaps we can learn more about our new community, and what steps we should take to foster its growth in the coming years.

11 Some Final Thoughts

So, what have we learned from this questionnaire? In some sense, nothing we didn't already suspect: more experience with robots leads to a more accurate understanding of how they work; how we frame a robot system affects how people think about it; roboticists think they know more about robots than legal scholars do; the law is a complicated and subtle thing; legal scholars understand more of this subtlety than roboticists and policy scholars do; roboticists are not perfect at designing survey questions.

However, some unexpected things have also popped up: there seems to be no correlation between how well we understand a system and how autonomous we think it is, or where on the mechanism-robot spectrum we place it; there's also no correlation between where on the mechanism-robot spectrum a system is and how autonomous it is; android-based systems can invoke gender stereotypes in some situations.

This draft represents a partial analysis of the questionnaire responses, and gives a broad overview of the findings. However, there is more work to be done. We have established

[5]How Should the Law Think About Robots?, Neil M. Richards and William D. Smart.

what factors affect the responses, but *how* do they affect them? How do groups from different backgrounds answer the individual questions? Are roboticists more pessimistic about autonomy than legal scholars?

We also believe that the analysis supports some concrete steps we can take to strengthen the We Robot community. We outline these suggestions in section 10. However, the simplest, most important suggestion is also a fun one: Buy a Robot! Play with it! It'll make you better at your day job! You might even have fund doing it.

12 Appendix: The Questionnaire

Questions about Mechanisms and Robots

Participants were asked to evaluate the following systems on a 7-point scale to (1) (1) where on the mechanism-robot continuum they thought they lay; (2) how autonomous they thought they were; and (3) how well they thought that they understood the underlying technology.

1. Light switch
2. Thermostat
3. Motion-controlled lighting
4. The anti-lock braking system on your car
5. The automatic gearbox on your car
6. The cruise control on your car
7. Apple's Siri
8. The computer player in a tic-tac-toe game on your phone
9. Watson, the Jeopardy-winning IBM computer
10. The computer player in a chess game on your computer
11. The Google search engine
12. Roomba vacuum cleaner
13. Cruise missile
14. Self-driving car
15. Spot-welding system on the assembly line in an automobile plant
16. Driverless train connecting terminals at Newark airport

17. An aircraft flying on autopilot
18. Chief Justice John Roberts[6]
19. The system that won the recent DARPA robot challenge
20. A powered exoskeleton worn by a human
21. A system that physically manipulates the climate-control equipment already in your home
22. Radio-controlled car
23. Remote presence system (such as the Beam or vGo)
24. Landmine
25. A computer running a Skype videoconference

Questions about the Law

Participants were asked to identify what area of law (contract, property, criminal, or tort) applied in the following twelve hypothetical situations involving a robot.

1. An autonomous robot breaks a vase in your house.
2. Someone uses a tele-presence robot to steal something from your office.
3. A self-driving car crashes into another self-driving car.
4. A self-driving car crashes into another car, killing to occupant.
5. A self-driving car drives into a barrier, injuring the driver, rather that hitting a child that had wandered into the street.
6. A radio-controlled car runs into your dog.
7. A police drone crashes into a private house.
8. An autonomous robot drives over your foot.
9. An autonomous robot, running open-source software, drives over your foot.
10. A private drone photographs and automatically web-publishes a photo of a person nude sunbathing in their walled back garden.

[6]This was inserted in the list of systems in homage to the We Robot 2013 paper "Chief Justice John Roberts is a Robot", by Ian Kerr and Carissima Mathen.

11. Someone programs a commercially-available autonomous domestic robot to identify valuables, remove them from the house, and place them on the sidewalk. One of these robots does this in your house.

12. A private drone photographs and automatically web-publishes a photo of a couple having sex through their ninth-floor apartment window