

HumanóAndroid Interaction in the Future

Neal J. Roese and Eyal Amir

HumanóAndroid Interaction in the Near and Distant Future

Neal J. Roese<sup>1</sup> and Eyal Amir<sup>2</sup>

*<sup>1</sup>Department of Psychology, University of Illinois at Urbana-Champaign, and*

*<sup>2</sup>Department of Computer Science, University of Illinois at Urbana-Champaign*

Address correspondence to Neal J. Roese, Department of Psychology, University of Illinois at Urbana-Champaign, Champaign, IL 61820; e-mail: roese@illinois.edu.

**ABSTRACT**—In this article, a psychologist and an artificial-intelligence (AI) researcher speculate on the future of social interaction between humans and androids (robots designed to look and act exactly like people). We review the trajectory of currently developing robotics technologies and assess the level of android sophistication likely to be achieved in 50 years time. On the basis of psychological research, we consider obstacles to creating an android indistinguishable from humans. Finally, we discuss the implications of humanóandroid social interaction from the standpoint of current psychological and AI research and speculate on the novel psychological issues likely to arise from such interaction. The science of psychology will face a remarkable new set of challenges in grappling with humanóandroid interaction.

How would it feel to interact face-to-face with another person, without knowing whether that person was real? The technology needed to create androids— robots designed to look and act just like human beings— is advancing rapidly. We may reach a point within the next 100 years at which androids are sufficiently sophisticated that they may be indistinguishable (at casual glance or brief conversation) from a real person. If and when we reach this threshold of indistinguishability, a new psychology to describe human—android interaction may well be needed. In this essay, a psychologist and an artificial-intelligence (AI) researcher speculate on (a) what androids might be like 50 years from now, given current technological trajectories, (b) the challenges suggested by psychological research that must be addressed in order to create realistic androids, and (c) some psychological implications of human interaction with androids that have passed the threshold of indistinguishability.

## ANDROIDS

Androids have long been a staple of science fiction. Through such fare as *Blade Runner*, *Star Trek*, and *Battlestar Galactica*, the general public has become widely familiar with the idea of social interaction with highly realistic artificial persons. Beyond these fictional incarnations, the actual technology to construct realistic androids is progressing rapidly. Frequent news stories announce the latest technical advances in reasoning, learning, appearance, mannerisms, body movements, speech synthesis, and facial expression of emotion. Today, if you walk into a room and catch a glance of an android, you might for a moment assume it to be a real person— but only for a moment. Walking closer to it, physical features such as skin and hair will immediately reveal the person to be artificial; further observation will reveal the body movement and

mannerisms to be artificial, and the simplicity of any conversational interaction will instantly reveal the limitations of current natural language AI. In the film *Blade Runner*, by contrast, the androids were fully conscious, highly intelligent individuals with distinct personalities. They not only spoke fluent conversational English, but also exhibited goal-directed behavior, visual-scene understanding, and recognition of others' intentions and beliefs, all motivated by their own personal agendas. The gap between the current state of the art and the ideal depicted in *Blade Runner* is vast, and the largest piece of this gap is AI. We elaborate on this gap by considering what androids might be like 50 years hence, given the development trajectories of key technologies.

### **50 YEARS HENCE**

What will androids be like 50 years from now? Will they be cunningly realistic or merely cartoonish? In reviewing the current state of various robotic technologies, it becomes clear that several fundamental hurdles make it unlikely that the androids of 50 years from now will be indistinguishable from human beings.

#### **Morphology**

The technology to create realistic skin, hair, and eyes poses some thorny challenges. Research on realistic-looking flesh for burn victims (Atiyeh, Gunn, & Hayek, 2005) and biological robots (Brooks, 2001) are promising, but the distance from realistic flesh for robots is still great, barring significant revolutions in material science. In contrast, computer graphics and animation has reached a stage in which very realistic humans are portrayed automatically (e.g., Blanz, Scherbaum, & Seidel, 2007). Fifty years hence, we anticipate that the distance of indistinguishability (how close you must be to an android to detect its artificial nature) will be as much as several meters.

**Locomotion**

Robotic locomotion and motion control is developing rapidly. The past 20 years have seen advances, some led by the Japanese industry, leading to the reasonably realistic humanoid robots of the present day (Ishiguro, 2006; Sakagami et al., 2002). It seems likely that androids will soon be capable of human-like walking and running, if the current rate of progress continues. The emerging engineering challenges are reliability and robustness of machinery (broken-down robots blocking a city sidewalk would not be a pretty sight). Nevertheless, widespread commercial development is picking up, suggesting that reliability issues will be addressed by the economy of mass production. For other reasons (detailed below), we believe that the majority of android applications in 50 years will be in office or service settings, which would require little, if any, walking and running.

**Speech**

Speech synthesis and recognition are undergoing rapid research and development. Synthesis of speech from text now includes intonations and prosody that are very close to human (Dutoit, 1997). Speech recognition has benefited in recent years from advances in machine learning, and it is showing better precision, though still far from human-level performance (Jelinek, 1998). These facts suggest that androids will speak with clear, human-like voices and will identify spoken words with precision 50 years hence.

**Facial Emotion**

The mechanics of human facial muscle action are well understood, such that their essential design is currently emulated in android face construction. Controlling discrete facial motion so as to assemble a palette of facial emotion displays has made some

progress (Breazeal, 2002), but is attracting only limited research at present. Within 50 years, we expect that research will advance to the point where androids will be able to display a range of realistic facial expressions. Of course, the limitations on skin technology (especially with regard to stretching and wrinkling) will limit, to an extent, the believability of such displays, except from a distance of several meters.

### **Detecting Emotions From Speech**

Visual emotion detection, with the goal of modifying detecting emotions from speech, will be quite feasible. This currently is an active and successful area of linguistics research (Shih, 2005), mirroring the active inquiries into human emotion detection (e.g., Porter & ten Brinke, 2008).

### **Computer Vision**

Anyone reading the printed media and following advertisements may believe that current cameras and computer technology are capable of recognizing people, objects, and activities. Unfortunately, technology is far behind such advertisements, and there are still major hurdles. Most important, there is little understanding today of how to combine knowledge and common-sense reasoning in a system interpreting visual scenes. For example, given a picture of a human writing a letter, no general-purpose vision system today can identify that this is a picture of a human holding a pen and writing a letter. To reach such a level, vision systems need to go beyond low-hanging fruit and address the connection of knowledge and vision. Some progress is being made on this front today (Barnard et al., 2003; Hoiem, 2007), but much progress remains to be made. It is unclear if androids 50 years from now will be able to interpret what they see. If not, this will pose a significant barrier to realistic interactions with humans. However, if they do, then a

major breakthrough has occurred, and computers will have learned to reason with images. In that case, androids will be able to perform many human tasks, and many stimulating humanóandroid social interactions may revolve around the interpretation and execution of vision-based tasks. Coordinated activities, such as building a fence, painting a house, or catering a party, might become the province of joint humanóandroid teams.

### **Natural Language Processing (NLP)**

Another barrier to meaningful interactions between humans and androids is the (mis)understanding of natural language, such as English. Natural language generation (Reiter & Dale, 2000) seems today to be somewhat easier than the inverse problem of understanding. Most research in the past 15 years has focused on applications of statistical techniques to coordinate and generate natural language messages based on very large databases (e.g., documents available on the Internet together with a large number of documents) to simple NLP problems (e.g., Manning & Schutze, 1999). More recent approaches involve integration of complex structure and knowledge in answering simple questions posed by humans, with some success (Punyakanok, Roth, & Yih, 2008; Zettlemoyer & Collins, 2005). Nonetheless, it is far from possible to answer questions from text in a general-purpose manner. For example, a system of today can read the Old Testament and be unable to tell if Isaac was born after Abraham's birth. It seems likely that NLP will achieve a breakthrough in the coming 50 years that will enable such question answering. However, we believe that linguistic proficiency at a conversational level will still remain outside the reach of androids. Conversations between humans and androids will thus be limited to those that do not require deep understanding of the state of mind of the listener.

## AI

AI—the study of computational architectures of reasoning, learning, and knowledge—will remain the primary limitation on reaching the threshold of indistinguishability. AI continues to develop at a glacial pace, with several significant revolutions required to approach the order of complexity needed to carry on even a small conversation in natural language. In particular, connecting knowledge at the level required to process both high-level knowledge and low-level sensory input remains a challenge (Amir & Maynard-Zhang, 2004; de Salvo Braz, Amir, & Roth, 2007). Current research strives to model reasoning directed at understanding other agents (e.g., people or self-directed robots) in a real-time, general-purpose manner about beliefs of agents (Shirazi & Amir, 2008), but connecting such reasoning to the real world (through language and vision) remains a challenge (Amir & Doyle, 2002). For that reason, we believe that androids 50 years from now may have a conversational capacity slightly more complex than the automatic telephone services used by corporations today. Blade Runner-esque androids with personal agendas will remain the province of science fiction.

## Free Will

At present, some AI researchers contend that people do not want androids with free will. For example, we prefer machines to be under our control, such as recorders that record television according to our demands rather than anticipating our preferences (McCarthy, 2000). Accordingly, investigations into the free will of robots may be far off. In any case, AI technology pales next to human intelligence on many fronts, and examinations of artificial free will are considered philosophical at present.

To summarize, the android of 50 years hence will still be far from the threshold of indistinguishability. Current developments in synthetic skin, locomotion, and emotion detection may produce workable androids in fifty years. However, major breakthroughs are needed for continued progress in artificial facial emotion, NLP, AI, and free will. The android of 50 years from now will look, walk, and talk quite well, but it will have a narrow behavioral repertoire a far cry from the sentient, autonomous beings of science fiction more. What more will it take for androids to reach that threshold of indistinguishability?

### **DESIGN CHALLENGES SUGGESTED BY PSYCHOLOGICAL SCIENCE**

In this section, we take into account psychological research findings to point out some challenges that must be addressed by roboticists in order for androids to become truly realistic.

#### **Eye Gaze**

People are remarkably sensitive to eye gaze, particularly as to whether a person (or animal) is looking directly at them (Macrae, Hood, Milne, Rowe, & Mason, 2002). Moreover, eye contact reflects a complex social dance, signaling recognition of intent, intimacy, status, and power (Calder et al., 2002). Androids must therefore regulate their gaze direction with exquisite subtlety: making eye contact upon first meeting and coordinating the back-and-forth of looking toward and looking away as a function of the unfolding social interaction (Minato, Shimada, Itakura, Lee, & Ishiguro, 2005). There might be nothing quite so unnerving as being stared at by an android! Designers will no doubt recognize that the stares of android may be counterproductive (scaring customers away) as well as useful (warning customers against theft).

## Body Language

Moving your body into specific positions (e.g., cocking a head sideways, gazing at the ceiling, crossing one's legs) contributes an important signal for communication (de Paulo, 1992). Nevertheless, the information contained in such communication is of much lower fidelity than that found in symbolic language, hence the engineering challenges are much simpler than it is for natural language interaction. A reasonably-sized repertoire of several thousand distinct body movement routines would be computationally straightforward and easily married to developments in robotic motion control. However, there is little or no active current research on android body language.

## Personal Space

Face-to-face social interaction operates at a particular, socially normative interpersonal distance (‘personal space’, Evans & Howard, 1973) that reflects cultural norms learned early in life. Social interactions in which one person is positioned too near or too far from the other are awkward and stressful (Middlemist, Knowles, & Matter, 1976). Androids would thus need to regulate their body positioning with regard to personal space. Because space boundaries vary considerably across cultures (Evans, Lepore, & Allen, 2000), this design aspect must be either culture-specific or designed to be flexibly regulated on the basis of observation of the current social situation (cf. Kaplan, 2004).

## Theory of Mind

People effortlessly understand the meaning of other people's speech and actions, inferring underlying motives, emotional states, and the like. This ability is termed *theory of mind* (e.g., Carruthers & Smith, 1996; Herrmann, Call, Hernandez-Lloreda, Hare, &

Tomasello, 2007). How human brains achieve this capacity is an active area of research across several disciplines of psychology. For example, social cognition theorists (e.g., Malle, 2004) have isolated a sequence of cognitive processing steps underlying theory of mind capability (e.g., inferences about another person's category membership, intentionality, goals, dispositions), and decision theorists (e.g., Dastani, Hulstijn, & van der Torre, 2005) have modeled these computational steps. How best to integrate the fruits of psychological research with AI research remains an enormous challenge (Ishiguro, 2006).

### **PSYCHOLOGY OF HUMAN-ANDROID INTERACTION**

Let us imagine that the engineering challenges described above have been addressed, and highly realistic androids are a common feature of everyday life. This achievement is at least 100 years away, but what would it be like? In this final section, we consider some psychological implications of living in a world populated by sentient robots (cf. Libin & Libin, 2004).

#### **Mindless Acceptance**

Many social interactions are so scripted that careful attention is unnecessary: reflexively "mindless" social interactions suffice (Langer, Blank & Chanowitz, 1978). Interactions at coffee counters, playgrounds, and toll booths are good examples. If the threshold of indistinguishability passes, people would quickly become blasé about interacting with androids. In many casual interactions in which the human observer may not immediately know whether the other person is real or artificial, that human observer would probably not care nor give it a second thought.

#### **Sentience Ambiguity**

In more important social interactions, such as arguing with an airline agent over a canceled flight, it begins to matter whether the interaction partner is real or artificial. In an exchange involving a complaint, one may need to follow different social scripts depending on whether the other person is real or artificial. In this case, the individual benefits from making a rapid inference: Is that person real? To the extent that android technology is sophisticated, this may not be obvious, at least on the basis of brief observation. The momentary feeling of not knowing *sentience ambiguity* might be one of disquieting unease. Generally speaking, people find uncertainty of any sort to be unpleasant and hence are motivated to reduce it by forming explanations and reaching conclusions (e.g., Kruglanski, 1989). *Sentience ambiguity* would probably elicit a quick behavioral response, but what might that be?

### **Turing Anger**

One straightforward means of relieving *sentience ambiguity* is anger. Even if AI progresses to the point of permitting small talk conversations, the simplest way to reveal the status of an interaction partner may be an angry verbal attack. The Turing test refers to a verbal interaction designed to unveil an unseen speaker as human or artificial (Turing, 1950), yet nearly all subsequent depictions of the Turing test centered on gymnastics of semantic memory (‘‘Who won the World Series in 1992?’’). A Turing test built on angry attack would achieve vastly swifter results, as would nearly any highly arousing emotional content, the reason being that the AI needed to grasp the complexity of human emotion is perhaps the highest of all AI hurdles. If ‘‘Turing anger’’ turns out to be an effective way to confront *sentience ambiguity*, then a culture filled with advanced androids might well be an angry one.

## **Caricatures**

Do we, as a society, want androids to be truly realistic? An alternative is that people may prefer androids that are obvious deviations from humanity, such as caricatures or cartoon characters. There is a widely discussed phenomenon among roboticists called "the uncanny valley" in which a highly realistic robot or puppet that comes very near to human form, but still contains subtle flaws or discrepancies, creates an eerie feeling among human observers (Ishiguro, 2006). Highly realistic but not-quite-perfect androids can come across as unsettlingly creepy (MacDorman, 2006). To avoid creepiness, androids that are deliberate caricatures might be preferred. If technology can transcend the threshold of indistinguishability, it would certainly solve the problem of the uncanny valley, yet many people may still desire caricatures. Many tastes may have their way, with some preferring cartoon robots (like C-3PO from the *Star Wars* films) and others preferring highly realistic artificial people.

## **New Outgroup, Old Stereotyping**

There is a deeply human tendency to look upon outside groups (*outgroups*) with disdain (Fiske, 1998). Basic mechanics of categorization (dividing the world into us vs. them) along with motives of self-protection (avoiding danger from strangers) and self-enhancement (feeling good by putting others down) together contribute to the formation of negative stereotypes about outgroups (Plous, 2003). Such stereotypes hamper social interaction and result in suboptimal outcomes (e.g., hiring discrimination). If androids constitute a social group, then human beings will have stereotypes about androids, the exact same mechanisms of stereotype formation, activation, and alteration that underlie perceptions of human groups will apply to perceptions of androids. Although androids

will not (unless so designed) feel angered by discrimination, stereotypes may nevertheless hamper interaction to the extent that stereotyped beliefs are false and rigidly held. Advertising the benefits of a new model android will be tougher than advertising for new cars or computers if the general public views androids through the lens of social stereotyping. Might there come a time when humans and androids might interact so successfully that some humans count androids as members of their ingroup? Given research into how computer-mediated friendships can form across group boundaries (Postmes, Spears, & Lea, 1998), there is no reason to rule out the formation of human-android friendships.

### CODA

We stand at the forefront of a new machine-initiated transformation of labor and employment. In the 1950s, the advent of electronic computers sent armies of human computers, who for many decades had toiled over paper spreadsheets to hand-compile never-ending columns of financial figures, to the unemployment line. In the 1980s through today, assembly line workers met a similar fate when simple industrial robots replaced them. The vast majority of these workers found new employment in the service sector. In the coming decades, service-sector jobs may experience yet another such convulsion, with human workers replaced by androids. The social upheavals wrought by this third machine transformation may well dwarf those of the previous two. We see the threshold of indistinguishability—the moment at which technology can create an android that is indistinguishable from human beings— as more than 100 years away from current technology (maybe much further), with roadblocks centering most pivotally on the material science underlying artificial skin and the computational challenges of computer

vision and natural language AI. Nevertheless, in 50 years we suspect androids of substantial sophistication to populate our world and participate in everyday social interactions. The science of psychology will face a remarkable new set of challenges in grappling with humanóandroid interaction.

## REFERENCES

- Amir, E., & Doyle, P. (2002, July). *Adventure games: A challenge for cognitive robotics*. Paper presented at the AAAI02 Workshop on Cognitive Robotics, Edmonton, Alberta, Canada.
- Amir, E., & Maynard-Zhang, P. (2004). Logic-based subsumption architecture. *Artificial Intelligence, 153*, 167ó237.
- Atiyeh, B.S., Gunn, S.W., & Hayek, S.N. (2005). State of the art in burn treatment. *World Journal of Surgery, 29*, 131ó148.
- Barnard, K., Duygulu, P., de Freitas, N., Forsyth, D., Blei, D., & Jordan, M. (2003). Matching words and pictures. *Journal of Machine Learning Research, 3*, 1107ó1135.
- Blanz, V., Scherbaum, K., & Seidel, H.-P. (2007, October). *Fitting a morphable model to 3D scans of faces*. Paper presented at the IEEE International Conference on Computer Vision, Rio de Janeiro, Brazil. Retrieved from <http://www.mpi-inf.mpg.de/~scherbaum/publications/FittingAMorphableModelTo3DScans.pdf>
- Breazeal, C. (2002). *Designing sociable robots*. Cambridge, MA: MIT Press.
- Brooks, R. (2001). The relationship between matter and life. *Nature, 409*, 409ó411.

- Calder, A.J., Lawrence, A.D., Keane, J., Scott, S.K., Owen, A.M., Christoffe, I., & Young, A.W. (2002). Reading the mind from eye gaze. *Neuropsychologia*, 40, 1129-1138.
- Carruthers, P., & Smith, P. K. (Eds.). (1996). *Theories of theories of mind*. New York: Cambridge University Press.
- Dastani, M., Hulstijn, J., & van der Torre, L. (2005). How to decide what to do? *European Journal of Operational Research*, 160, 762-784.
- de Paulo, B.M. (1992). Nonverbal behavior and self-presentation. *Psychological Bulletin*, 111, 203-243.
- de Salvo Braz, R., Amir, E., & Roth, D. (2007). Lifted first-order probabilistic inference. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence* (pp. 1319-1325). San Francisco, CA: Morgan Kaufmann.
- Dutoit, T. (1997). *An introduction to text-to-speech synthesis*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Evans, G.W., & Howard, R.B. (1973). Personal space. *Psychological Bulletin*, 80, 334-344.
- Evans, G.W., Lepore, S.J., & Allen, K.M. (2000). Cross-cultural differences in tolerance for crowding: Fact or fiction? *Journal of Personality and Social Psychology*, 79, 2046-210.
- Fiske, S.T. (1998). Stereotyping, prejudice and discrimination. In D.T. Gilbert, S.T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (pp. 357-641). Boston: McGraw-Hill.

- Herrmann, E., Call, J., Hernandez-Lloreda, M.V., Hare, B., & Tomasello, M. (2007). Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, *317*, 1360-1366.
- Hoiem, D. (2007). *Seeing the world behind the image: Spatial layout for 3D scene understanding*. Doctoral dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.
- Ishiguro, H. (2006). Android science: Conscious and subconscious recognition. *Connection Science*, *18*, 319-332.
- Jelinek, F. (1998). *Statistical methods for speech recognition*. Cambridge, MA: MIT Press.
- Kaplan, F. (2004). Who is afraid of the humanoid? Investigating cultural differences in the acceptance of robots. *International Journal of Humanoid Robotics*, *1*, 1-16.
- Kruglanski, A.W. (1989). *Lay epistemics and human knowledge: Cognitive and motivational bases*. New York: Plenum.
- Langer, E., Blank, A., & Chanowitz, B. (1978). The mindlessness of ostensibly thoughtful action: The role of placebic information in interpersonal interaction. *Journal of Personality and Social Psychology*, *36*, 635-642.
- Libin, A.V., & Libin, E.V. (2004). Person-robot interactions from the robotics point of view: The robotic psychology and robototherapy approach. *Proceedings of the Institute of Electrical and Electronics Engineers*, *92*, 1789-1803.
- MacDorman, K.F. (2006, July). *Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley*. Paper presented at ICCS/CogSci 2006. Retrieved from

<http://www.macdorman.com/kfm/writings/pubs/MacDorman2006SubjectiveRatings.pdf>

- Macrae, C.N., Hood, B.M., Milne, A.B., Rowe, A.C., & Mason, M. (2002). Are you looking at me? Eye gaze and person perception. *Psychological Science, 13*, 460-464.
- Malle, B.F. (2004). *How the mind explains behavior: Folk explanations, meaning, and social interaction*. Cambridge, MA: MIT Press.
- Manning, C.D., & Schütze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.
- McCarthy, J. (2000). Free will even for robots. *Journal of Experimental and Theoretical Artificial Intelligence, 12*, 341-352.
- Middlemist, R.D., Knowles, E.S., & Matter, C.F. (1976). Personal space invasion in the lavatory: Suggestive evidence for arousal. *Journal of Personality and Social Psychology, 5*, 541-546.
- Minato, T., Shimada, M., Itakura, S., Lee, K., & Ishiguro, H. (2005). Does gaze reveal the human likeness of an android? In *Proceedings of the 4th International Conference on Development and Learning* (pp. 106-111). Washington, DC: IEEE Computer Society.
- Plous, S. (Ed.). (2003). *Understanding prejudice and discrimination*. New York: McGraw-Hill.
- Porter, S., & ten Brinke, L. (2008). Reading between the lies: Identifying concealed and falsified emotions in universal face expressions. *Psychological Science, 19*, 508-514.

- Postmes, T., Spears, R., & Lea, M. (1998). Breaching or building social boundaries? SIDE-effects of computer mediated communication. *Communication Research*, 25, 689-715.
- Punyakanok, V., Roth, D., & Yih, W. (2008). The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics*, 6, 1630.
- Reiter, E., & Dale, R. (2000). *Building natural language generation systems*. Cambridge, United Kingdom: Cambridge University Press.
- Sakagami, Y., Watanabe, R., Aoyama, C., Matsunaga, S., Higaki, N., & Fujimura, K. (2002). The intelligent ASIMO: System overview and integration. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 2478-2483). Washington, DC: IEEE Computer Society.
- Shih, C. (2005). Understanding phonology by phonetic implementation. In *Proceedings of Interspeech: 2005* (pp. 2469-2472). Bonn, Germany: International Speech Communication Association.
- Shirazi, A., & Amir, E. (2008). Factored models for probabilistic modal logic. In *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence* (pp. 541-547). Menlo Park, CA: AAAI Press.
- Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.
- Zettlemoyer, L.S., & Collins, M. (2005). Learning to map sentences to logical form: Structured classification with probabilistic categorical grammars. In D.M. Chickering, F. Bacchus, & T.S. Jaakkola (Eds.), *Proceedings of the Twenty First Conference on Uncertainty in Artificial Intelligence* (pp. 658-666). Arlington, VA: AUAI Press.