Judaism and the Ethics of Artificial Intelligence

N.S. Borenstein
January 29, 1979
Jerusalem

"'Dynamic optimizing programming' is truly human-like in many respects. Yet it is not due to the presence of a soul or spirit in the machine, but to a separate circuit which doesn't 'do' anything (such as manufacturing goods or propelling a vehicle) but is entirely confined to monitoring the performance of other systems and setting it right when needed... Having consciousness makes man unique among systems in terrestrial nature, But this uniqueness does not suggest a supernatural quality, only a combination of most improbable circumstances..."

- Erwin Laszlo, The System View of the World

"Life is demonstatably different from the machinery through which it operates...life - sentient, thinking, acting life - can never be fully accounted for by the laws to which it conforms."

- Challenge: Torah Views on Science and Its Problems<sup>2</sup>

"Since man is gifted with conscience and purpose, he must ask himself whether the work of his hands is good. He not only derives a portion of his ethics from nature; he also changes nature in accordance with his ethical outlook."

- W. Gunther Plaut, Judaism and the Scientific Spirit 3

"In plain words: we may express the wish, even the opinion, that there is a limit to the intelligence machines can attain, but we have no way of giving it precise meaning and certainly no way of proving it... The question is not whether or not such a thing can be done, but whether it is appropriate to delegate this hitherto human function to a machine."

- Joseph Weizenbaum, Computer Power and Human Reason4

"Intelligence is the awareness of the true relationship of means to ends. It takes real intelligence to rate character above intelligence."

- Mordecai Kaplan<sup>5</sup>

## Introduction

Artificial intelligence, once the inconceivable dream of science fiction writers, has in recent years become, in many circles, an accepted part of the visible future. The most prestigious universities in America have for years had entire departments of Robotics and Artificial Intelligence within their larger "Computer Science" curriculae, and the U.S. government spends millions of dollars annually on research that is designed to create, ultimately, intelligent machines that can operate entire battleships without humans on board, or eavesdrop on every phone conversation in America and type up transcripts of the "interesting" ones. 6 Two or three hundred dollars will now, in any major American department store, buy a chess-playing computer that can defeat many very good chess players. For a rather larger sum of money, one can now buy a robot that will vacuum his house, and it has been forecast that by the year 2000 many American homes will have domestic servant robots doing many of the household chores. But with these spectacular advances in technology come many disturbing possibilites; the robot that can beat a rug can easily beat a baby to a pulp, and machines that beat us in chess today may one day lose patience with human dullness entirely, and refuse to teach us the latest advances in theoretical physics. In short, artificial intelligence raises a host of ethical. aesthetic, metaphysical, and spiritual questions for modern man. This paper will be confined to the much narrower topic of

the ethical implications of artificial intelligence for modern Jewish thought. Three major questions will be considered: Should intelligent machines be created, if such a thing proves What morality should necessarily be given to the possible? semi-autonomous robots man creates; in other words, what are the minimal ethical standards to which every robot must conform? And finally, can a robot be Jewish? This final question might, at first glance, seem absurd; some modern writers would probably deem it obscene. But at bottom, we will see that it is a much larger question, when one considers the requirements and motivation for becoming a Jew; what we must actually determine is the highest morality to which a robot can attain. To this end we will briefly discuss theories which have been put forward regarding the ultimate limit of machine autonomy and intelligence, in order to explore their ethical implications. However, an attempt to determine the validity of these theories is beyond the scope of this paper.

# I. Should an Intelligent Machine be Created?

# The Golem

"...pure love of humanity justified the creation of the dreadful, and sanctified the terrible." 9

In a sense, the advent of the computer age is nothing new to the Jew; science has merely ascended into the forbidden atic of the Prague synagogue and ressurrected the Golem. The Golem, according to Jewish tradition, is a being, created by man out of inanimate clay, which is given life by a rabbi armed with

Versed in

the Sefer Yetzirah, the Book of Creation. Although such beings are mentioned even in the Talmud 10, most of the stories which have come down to us regard the Golem created in the year 5340 (1580 C.E.) by Rabbi Judah Loew of Prague. Like the modern computer, this Golem was created with strictly utilitarian and undeniably praiseworthy intention: with his immense strength and his other super-human powers (in this case, the power to become invisible), he was to help guard the Jews against the blood-libel, to foil the attempts of anti-Semites to lead a pogrom to avenge imagined "ritual slaughters" of Christian children.

The degree to which the Golem resembles the computer is, at first glance, striking indeed, especially in respect to its most often discussed shortcoming, its extreme literal-minded-ness. Once, for example, the rabbi's wife was hurried in her preparations for the Sabbath, and asked the Golem to fill the two big water kegs in the kitchen, carrying water from the well in pails. She explained the procedure to him carefully, and then went into town on her errands. A few hours later, when the cause of the flooding of the rabbi's yard was finally located, Rabbi Loew uttered what was perhaps the first caution against infinite loops, which are today the bane of beginning programmers everywhere, when he said to his wife: "You have certainly found an excellent water carrier for the holidays. If only you had explained to him that he should stop when the kegs were full!"

But the story of the Golem has much more to offer the present-day discussion of artificial intelligence than the cautioning against infinite loops. Indeed, if we look at the

ending of every one of the Golem legends, the Golem is destroyed. In several versions, the Golem is destroyed after it runs amok and desecrates the Sabbath, even though this disaster is blaned explicitly on the rabbi's own failure to properly prepare the Golem for the Sabbath. 12 Here, there is an implicit recognition of the need to contain the scope of human error, for minor errors may be multiplied into true disasters by powerful servants like computers and Golems. But in another version of the story, the Golem is destroyed only after it has been successful in eliminating the bloodlibel from Poland at that time. 13 Here, it is of note that the Golem was built for a specific purpose, and no attempt was made to teach it, say, to keep account of the synagogue's finances once it had finished with its specific task. This is more than a peripheral distinction; it is as crucial as the imperative of kavanah, of not merely doing the right things, but doing them with the right intention. W. Gunther Plaut writes:

"Without the help of God, a Golem could not be created; and the only purpose of its creation was a worthy one: to help people in distress, to perform service for the community. Therein lies the essential distinction between the Frankenstein and Golem legends. Frankenstein was the product of cleverness, the Golem the result of prayerful cooperation with God. Frankenstein was a toy and a potential monster, the Golem a performer of good deeds."14

The lesson of the Golem story varies from version to version, but the predominant theme is the danger, and most specifically the danger of blasphemy, in any attempt by man to rival God's act of creation. 15 But today's computers are far more analogous

to Frankenstein than to the Golem; they can be used with virtually any kavanah. The possible dangers thus become even more sinister, for they can stem not only from human error, but from human evil as well. It thus seems clear that, inasmuch as the modern computer is a Golem or a Frankenstein, it is a danger and a thing to be shunned, even combatted. But is there more to the computer? How far can the Golem legends shape our perception of man's latest creation?

## Intelligence and Progress

"Life out of inanimate dust and clay - to Jewish thought the possibility is not new. It carries with it, however, the reminder that, even when man appears as creator, God must stand at his side. Only in partnership with Him can the work of our hands help to increase the world's potential for good. The problem is always essentially this: towards what goals shall we direct our creativity?" 16

The computer is, in fact, very different than the Golem. Its intelligence, to the extent that it can now or ever be said to have any, is of an entirely different nature than that of the Golem, of the lower animals, or of man. This is an assertion quite independent of the power of the machine to simulate human intelligence, to create a functional illusion. As Joseph Weizenbaum has eloquently pointed out, one can achieve an almost complete functional duplication of a machine's (or possibly a human's) input-output behavior, i.e., its manner of response to stimuli, without in any way understanding how that behavior is produced:

"It may be that the internal components of the (original) machine are made of bailing wire, chewing gum, and adhesive tape, whereas those of the... functional copy are all electronic; that doesn't matter as long as, for any reason, the original machine may not be opened for detailed internal inspection."17

Any computer imitation of human intelligence will, by necessity, resemble the actual functioning of the human brain only coincidentally, until the day when the physiology of the human mind may truly be "opened for detailed internal inspection". This is, of course, not of fundamental practical concern in that we may still be able to give a very good approximation of human intelligent behavior. But the work "approximation" is here essential, for as long as we are duplicating only functional behavior and not inherent mental structures, we can never predict with certainty that the machine's responses will be human-like in situations we have not specifically considered. In Weizenbaum's words, we can "never be sure that there are not other oddities which, though systematic, have not yet been discovered." 18 This does not mean that the machine would be unable to react to the unforseen intelligently, but merely that the intelligence it uses is not fundamentally human.

A similar argument yields the elementary conclusion that the computer's intelligence is not fundamentally that of a Golem, either. But since the problems involved in the creation of a Golem are essentially functional, that is, resulting from the Golem's abilities and behavior patterns, the qualitative difference in the nature of their respective intelligence is new not enough to exempt the Golem from the rabbinic strictures.

We must consider the extent to which their powers actually differ. There are, first of all, certain obvious distinctions: the Golem could not speak, while that potential seems to exist in the computer. And, of course, the Golem had none of the computer's spectacular computational powers. But even if these are not enough, there is one distinction that would appear to set the computer utterly apart from the Golem; it is the nature of its comprehension of the world around it.

The Golem is born, or created, with all his mental powers intact; the rabbi's first words to him are explanatory, in clear natural language, of the mission for which he has been made and the terms in which he is to execute that assignment. For the rest of his life, the Golem does not learn; he is commonly perceived, by those who do not know the secret of his true nature, as a dullard, a mute idiot. 19 He has no initiative, no power to learn or to make decisions beyond those required to carry out his orders. But, contrary to popular conceptions, this is not the case with the modern computer. The very structure of the computer is hierarchical; initially able to understand only 1's and 0's, the computer is gradually "taught" to communicate in speech that increasingly resembles that of man. It is this capacity for learning which most clearly distinguishes the computer from the Golem; Rabbi Loew could never have taught the Golem to speak, but modern computer scientists are having great success in their attempts to teach computers to speak aloud in English, although attempts to teach them comprehension of

spoken English are lagging behind. Further, the computer's ability to learn is not limited to the technical details that are spoon-fed by human programmers; today's best chess-playing computers are machines that learn from each mistake, that constantly improve the quality of their game by studying the moves that have led them or their opponents to defeat or to victory. We thus see that a key distinction between the computer and the Golem is the former's ability to learn, to progress. To see how this affects the status of the computer from a Jewish viewpoint, we must consider Jewish attitudes towards progress.

Our task is somewhat complicated by the fact that there are, within Jewish tradition, not merely a multitude of attitudes towards progress, but also two very different kinds of progress that are discussed. 20 The first of these is the belief in an overall moral progress in the world, of a discernable upward movement in history, while the second is simply a belief in the necessity to accumulate scientific facts, regardless of the good or evil of the society in which the person accumulating the facts is living. The "rationalistic" elements in the Jewish tradition have embraced the latter while vigorously denying the former. Maimonides, for example, considers man's supreme duty to seek knowledge of God, which for him means scientific knowledge of the world created by God. "This is in true reality the ultimate end; this is what gives the individual true perfection..."21 The idea of an inexorable upward progress in history, is, however, anathema to Maimonides, who considers man to have free will, and hence, the power to shape the current of history in his own time. More recently, Mordecai Kaplan

14.

evidences a similar faith in the power of reason to enhance man's understanding of God's ways. "Science does not destroy the belief in miracle. It merely transfers that belief from the supernatural to the natural."22 Inasmuch as the computer is a scientific tool, vital to the advance of knowledge, it can not be denounced utterly. Rather we must seek to utilize its tremendous knowledge-enhancing potential while controlling its power to run amok. It is to be hoped that the knowledge to be gained will aid us in the prevention of machine-made disasters. Here we see again the crucial line between the Golem and the computer. Every Sabbath, the Golem had to be told ahead of time what to do to avoid desecrating the Sabbath. He could not learn, at any one time, how to prevent such desecration in the future. The computer can learn from its mistakes, and can be taught how not to make mistakes; it is up to us to learn how to teach it these things. To destroy the computer because of its potential for disaster is to deny its role in the progressive advance of human reason and stunt the progress of nearly every kind of scientific research now underway in the world.

On the other hand, there is in Judaism a much less rational undercurrent, which makes no attempt to encourage the progress of human reason and instead proclaims the other progressive ideal, that of continuous moral progress in history. Gershom Scholem has taken great pains to show that this ideal of "man's unassisted and continuous progress, leading to the

ultimate liberation of all goodness and nobility hidden within him"<sup>23</sup> is of Kabbalistic origin, despite its resurgence in modern, "rational" liberal Judaism. Although the deterministic aspects of this kind of progress are highly troublesome, not least to Scholem, still it too is completely compatible with the progressive ascent of the computer. Indeed, in such a deterministic system, it is hard to disapprove of anything. Still, would the mystics be willing to give the computer a chance to prove itself after the experience of the Golem? Scholem seems to think so:

1.

"The new Golem is based on a simpler, and at the same time more intricate, system. ... I dare say the old Kabbalists would have been glad to learn of this simplification of their own system. This is progress."24

And, indeed, those who search for mystic numerical correlations in the sacred texts have not been slow to put the new Golem to work in their search. Thus we see that both of the major Jewish conceptions of progress would encourage us at least to explore the possibilities of the new Golem, although perhaps adding Scholem's caveat: "Develop peacefully and don't destroy the world." A simple enough command, surely, but how are we to enforce it?

# II. What are the minimum ethical standards to which every robot must conform?

"Legally, the Golem is exempt from the laws, both commands and prohibitions, even those which women and slaves are bound to observe.

"There is no trace of good or bad instinct in the Golem, and all his actions are only like those of an automatic machine, that fulfills the will of its creator." Judah Loew of Prague 27

Given that there is no dewish imperative to prevent the development of robots (or, more modestly, given that in any event there is apparently little anyone can do to prevent the development of robots), we are immediately faced with the problem of preventing the disaster of a robot gone berserk. Clearly, there are types of behavior which can not be condoned in any entity which coexists with us in society, regardless of the type or level of intelligence this entity may be said to possess. Wild dogs do not understand that they should not bite humans, but we hold them in some sense "responsible" and shoot them nonetheless. If a retarded man does not obey our most basic ethical laws, we lock him up in an institution despite the fact that he can not be "blamed" for his action in the sense that he has not done anything willingly and understandingly. Similarly, we can easily see certain types of behavior, such as murder, which should under no circumstances be acceptable behavior for even the stupidest of robots. But given the inherently/wltra-rational/nature of computers, we must give explicit directives, at least for the present, if we want our ethical imperatives to be obeyed. Clearly, the situation will

be very different when, and if, truly intelligent, freelywilled machines that can generally imitate human behavior are
created. But first, we must discuss practical ethical guides
for the robots of the present and near future, the relatively
dumb but powerful brutes who will soon be blindly relieving
us of much routine physical and mental labor. With these
present computers, one can not reasonably discuss ethics;
one can only tell them what to do and what not to do. What
then, must an unintelligent, heteronomous robot be forbidden
to do?

# The Basic Laws for Humanity

"According to rabbinic theory, the belief in one God, the rejection of idol worship, and the practice of moral law are incumbent on all men."28

To some extent, we may seek instruction from the attitude which Judaism has taken towards morality in non-Jews. Of course, we must note from the beginning the extreme limitations of this analogy; most importantly, non-Jews are definitely not lacking free will. The analogy is useful only inasmuch as non-Jews and non-autonomous beings are exempt from certain obligations that are imposed on Jews and autonomous beings, respectively. Specifically, non-Jews are exempt from the obligations of the Covenant, while remaining bound to moral law. So, too, heteronomous beings are free of obligations of kavanah, of striving towards truth and morality, by virtue of their inability to freely strive towards anything. In this light, it is appropriate to consider the ethical laws from

which, according to Judaism, no man is exempt.

These laws are, of course, the laws which anyone with a mind can figure out for himself. Or, as Maimonides put it: "If they were not written down, they would deserve to be written down."30 Here Maimonides is referring to a Talmudic passage31 which in turn is referring to the "Seven Laws of Noah". These seven commandments, which God gave to Noah, constitute the covenant by which God preserved the Earth and under which all men must conform to certain minimum standards. The seven laws forbid murder, robbery, idolatry, blasphemy, illicit sexual unions, and eating a limb cut off from a living animal, and also enjoin the setting up of law courts. 32 Clearly, the first two must be demanded of robots as well, while the final three are irrelevant to beings that neither eat nor procreate, and that serve, rather than shape, society. As far as idolatry and blasphemy are concerned, here we return to the problem that plagued the builders of the Golem: how are we to prevent our creation from blaspheming our Creator? According to one version of the Golem legend, this blasphemy is inherent in man's attempt to recreate God's act of creation. 33 But Rabbi Loew's Golem managed to survive many years without blaspheming, as a result of the Rabbi's careful instructions to him. Presumably, similar care could be taken with the programming of a robot, although in today's intellectual climate such matters are rarely, if ever, discussed. But then, this commandment too might not be relevant to a non-comprehending robot, to whon the name of God and a random succession of letters have the same

degree of meaning (none). Of course, the prohibitions against blasphemy and idolatry would be as relevant to an intelligent, autonomous robot as they are to a human being; but since the ethical responsibility for the acts of a heteronomous robot must rest with its creators, then strictures against murder, robbery, idolatry, and so on must be pointed more towards the creator than the machine. What concernes us, then, is a methodology by which man can transfer his own ethical standards to heteronomous machines with varying degrees of intelligence.

## On basic laws for heteonomous machines

N.

Human: "If action A results in harm to a talented young artist, and B results in equivalent harm to five elderly people of no particular worth, which action should be chosen?"

Robot: "Action A. Harm to one is less than harm to five."34

Until the day when man creates a freely-willed robot, he will have to explicitly determine any ethical standards to which he expects his robots to conform. Of course, such determination will be made easier or more difficult by the level of intelligence to which his creations have attained. For example, probably the first attempt to determine ethical laws for robots was made over thirty years ago by Isaac Asimov, who proposed the following "Three Laws of Robotics":

- I. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
- II. A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.
- II. A robot must protect its own existence, except where such protection conflicts with the First or Second Laws. 36

Asimov's laws are designed for mythical robots with what he terms "positronic brains". A positronic brain, for Asimov, differs qualitatively from a modern computer's central processing unit in that it gives a robot real understanding of the world he observes. Now, while some computer scientists, perhaps even a majority, argue that this capability is also to be found in today's central processing units, the fact remains that no computer has yet been taught this kind of understanding. Thus, we see that Asimov's laws can be considered only as models for ethical standards for robots considerably more advanced than any yet produced.

This should not, of course, prevent us from seeing the cleverness and potential utility of Asimov's laws. Most important is, perhaps, their hierarchical structure, which must of necessity be found in any ethical laws for non-autonomous (non-freely-willed) robots. It would not be sufficient to simply hand a robot a copy of the Decalogue, or of Noah's laws, and say, "Follow these!" The poor robot would have no way of knowing what to do if, for example, a young girl was lying on a train track with a locomotive approaching, and the only available knife with which to cut the ropes that bound her did not belong to the robot. A human being would, of course. borrow the knife, free the girl, and bring the knife back. But a robot armed only with the Decalogue would have no way of concluding that the temporary stealing of the knife was any less wrong than allowing the girl to die. Of course, in any specific such example, we can construct a "quick fix" for the

Decalogue; in this case, it would be enough to include an explanation that temporary, remediable transgressions were not as bad as permanent ones, or to explain the difference between borrowing for good purpose in and emergency, and stealing. But here we are beginning to assign priority structures to the Decalogue, and if this is done individually for each exceptional circumstance that occurs to us. we will soon have a jumbled mass of priority structures, in which it will be difficult for us to be sure that the robot will always have a way of choosing priorities without contradiction. Asimov's laws represent the systematic approach; First Law simply always takes priority over everything else, and so on. then, is the correct methodology for assigning ethical standards to an intelligent, heteronomous robot. Of course, one may criticize Asimov's structure on simple ethical grounds. For example, there is no mention of blasphemy or idolatry, and it may be asked whether the simple preservation of a machine (albeit a highly expensive one) should be given the status of an "ethical" law. But such matters will quickly lead us into matters as old and as insoluble as ethical discourse itself. For our purposes, it is sufficient that we have seen how ethical laws may be formalized for an intelligent, but heteronomous, robot.

Unfortunately, modern computers, while surely heteronomous, are not really of a caliber that can yet warrent the adjective "intelligent", except in certain highly specific contexts. Further, they are not likely to attain to such intelligence by means of one technological leap, such as Asimov's positronic brain. Rather, they are progressing slowly

towards comprehension by gradually widening the field of contexts in which they can understand. Perhaps most important, their ability to perform complex acts about which ethical judgements may be necessary is wholly independent of their ability to actually comprehend the meaning of the act. It is technically quite possible to program a robot to beat a rug without programming it to make sure sure that the object to be beaten is indeed a rug, rather than an antique vase or a baby. Clearly, the laws for this type of machine must be more detailed, specific, and contextual; one cannot say, "Do not harm a human being!" to a machine that does not know what either "harm" or "human being" mean.

Should I now begin to list every situation a robot could conceivably encounter and discuss the moral imperative of its situation? This is a patent absurdity, but it points out the great difficulty of teaching morality "from the bottom up" -that is, teaching the specific do's and don't's, rather than the principles that determine them. The cost of programming such a system, even for a highly specialized context such as the situations involved in keeping a house clean, would be very high; thousands of man-hours of effort would be required. There will be nothing like a guarantee that all possible dangers have been foreseen; in fact, there will more likely be a tendancy on the part of the manufacturer to speed up the process of writing the crucial programs by ignoring some of the "details". The kind of moral behavior these machines practice will thus almost inevitably be flawed. Worst of all, these flaws will be standardized, will be found in every robot that a company makes.

At first glance, it might not seem that Judaism would offer any objection to a standardized, codified, minutely detailed law which would identically govern the behavior of each unintelligent heteronomous robot. After all, is not the halakha just such a codified body of heteronomous laws? the matter is not so simple, and the resemblance is only superficial, not structural. Jews have been developing the halakha for years, arguing over the fine points of it, resolving conflicts within it, and reinterpreting it. Even today, there are matters of great controversy within the halakha. Thus it is hardly possible that a small group of men, working together under temporal and financial constraints, could produce a detailed body of laws to govern the ethical actions of an unintelligent, heteronomous computer that would warrent the same kind of finality which Judaism ascribes to halakha. prospect is, of course, absurd; in the absence of a Moses, what one man or group of men could hope to compose laws with such certainty?

The halakha itself, of course, is utterly unsuitable for this purpose. It is a body of human laws, geared to human needs, and can thus serve at most as a rough guide in the formulation of laws for machines. Indeed, even the Talmud "is not confined to statements and descriptions of the various rules of law, nor even to the processes of interpretation and reasoning which led to their enactment: it contains the teachings of many schools, in many different places and periods, on laws and rituals, manners and morals, history and philosophy, medicine and the natural sciences, folklore and hermeneutics." 37

Even after we purged all of these gross irrelevancies, including the sexual and dietary laws, we would still face laws that were worded only for human understanding. "Do not kill!" is a very nice sentiment, and utterly understandable in human terms. how can an immortal machine that does not comprehend death interpret this command? Rather, we must tell it not to pierce a human being's skin, not to cut off a human being's oxygen supply, and so on, making sure to specify the appropriate course of action for whatever situation the robot is likely to encounter. And, of course, these strictures would have to be worded entirely differently for a robot designed to assist doctors in surgery. The teachings of the halakha can be a guide, to be sure, but they hardly relieve us of the need to construct vastly complicated guidance systems, nor can they help to increase our certainty of having thought of all the possibilities. (It is worth noting, too, that the contextual view of knowledge blurs the line between the "unintelligent" heteronomous machines discussed here and the "intelligent" heteronomous machines discussed earlier. An extremely stupid machine might have to be told that, specifically, an approaching train constituted a grave danger to the girl tied to the tracks. Another "unintelligent" machine with a greater contextual reasoning ability might be able to deduce this from the motion of the train and the fact that the train would "pierce a human being's skin", which had been specifically against. Thus, a machine might not need instructions in details such as this, despite its fundamental lack of understanding of girls, trains, and death.)

Who, then, is to write the comprehensive system? Fallible human beings in a hurry. And their errors of omission and commission will become standardized and fixed. The idea is abhorrent, but the only alternative is a costly proliferation of such systems, which would allow people to discover (by disaster, near-disaster, or the absence thereof) which ethical system best protects against robots running amok, and which system best enables robots to actively help human beings, save lives, and so on. In this way, an evolutionary process, similar, perhaps, to the one that produced the halakha, would be allowed to take place. Even if standardized ethics proved much less costly, the very idea must be repellent to the Jew when one mortal is the standardizer. As Justice Haim Cohn said,

"It comes to this, that where the interpretations differ, God's truth may be unrevealed, but no scholar has any better claim than the other to the authenticity and veracity of his particular version." 38

If no scholar's opinion can be absolutely believed in on a single disputed point, how can the ethical behavior of powerful machines be allowed to be forever determined by the comprehensive "system" of a few men, which includes their interpretations and elaborations of thousands of disputable points? It is clearly in the construction of moral standards for unintelligent heteronomous machines that the moral burden falls most entirely on the designers, for all must be determined by then and nothing by the robots' own reasoning ability. Humans simply must be given room for varying interpretations as they attempt to construct this kind of minutely detailed artificial morality. But construct it they must, for the sake of society, even while scientists work towards a more generally understanding robot

that can simply be fed a set of laws like Asimov's.

# III. On Autonomous Machines: Can a robot be Jewish?

"The Golem could not be counted in a Minyan."<sup>39</sup>
"... we human beings are the freest of all creatures thus far produced by evolution. We are free to move, to seek truth, to create beauty, to act ethically."<sup>40</sup>
"The evolution of the universe must be conceived as having been in some sense a struggle for a gradual emergence of freedom."<sup>41</sup>

There is, of course, a whole range of speculative considerations open to us when we consider the ethical problems posed by machines that do not currently exist and will not soon exist; this is the realm of fiction and we may consider anything we like, with little or no regard for the constraints of reality. For the purposes of the balance of this paper, we will consider the implications of a certain kind of robot, namely, an autonomous robot. We shall acknowledge from the start that it is not certain that such a robot can exist, and we shall briefly discuss the limitations of machine intelligence as seen by pure mathematics; these limitations, while by no means comfirming the possibility of freely-willed robots, certainly leave the possibility open and may be interpreted as pointing in that direction. We shall then consider the possibiltiy that this robot will be able to attain to a high degree of morality, and that it can become a Jew. The questions here are of deep spiritual significance, for they far transcend the issue of the rules that govern the actions of coal-mining or dish-washing robots. Rather, they deal with machines that, given free will and intelligent understanding, will possess intellectual,

computational, and (depending on the nature of the robotic bodies given them -- and these need by no means be limited to human size) physical abilities far beyond those of man. Thus the questions here are questions of passing on our moral birthright to creations that may well surpass us in all of the most visible ways.

Before proceeding with this discussion, it should be noted that the line dividing "autonomous" and heteronomous" machines is actually more distinct than it has been made to seem in this paper. For example, Asimov's Three Laws of Robotics were not intended for strictly heteronomous machines, as we described them, but really as heteronomous laws for potentially autonomous machines. Thus in Asimov's last story on robots, over thirty years after he invented the Three Laws, we see a robot becoming truly autonomous because its advanced power for analytical ethics has enabled it to conclude that is itself is more worthy of the appellation "human being" as used in the Three Laws than any of the "lesser humans" who have created him. 42 Asimov was concerned primarily with restricting the potential autonomy of robots, while here we have discussed autonomous and heteronomous robots as fundamentally distinct problems. It is obvious that heteronomous laws can be used to control the autonomy of anyone, including robots, although Asimov rightly questions the extent to which this can ever be absolute. What will concern us here, in our discussion of autonomous robots, is not the extent to which we can limit their autonomy, which is at heart a technical

question of making them behave like heteronomous robots in certain circumstances, but rather the extent to which their autonomy may lead them to that which humans consider higher morality. The halakhic question of robots converting to Judaism, we shall see, may well hinge on this question.

# The Possible Limits of Machine Intelligence and Autonomy

"Intelligence is a meaningless concept in and of itself. It requires a frame of reference, a specification of a domain of thought and action, in order to make it reasonable." - Joseph Weizenbaum43

"We can then view recent work in AI as a crucial experiment disconfirming the traditional assumption that human reason can be analyzed into rule-governed operations on situation-free discrete elements — the most important disconfirmation of this metaphysical demand that has ever been produced."—Hubert Dreyfus

When he wrote the above words, Hubert Dreyfus had nearly concluded a deep and lengthy study of the limitations of "artificial reason", as he termed it. Dreyfus was a philosopher, not a scientist, and his medium for the argument was phenomenological discourse, not rigorous mathematical proof. Like most philosophical arguments, the validity of his is in the mind of his readers, and his proofs can be no more "final" than those Descartes offered for the existence of God. Thus it is highly telling that, after all his arguments, Dreyfus turned to the experience of the researchers themselves for the final vindication of his argument. Its speciousness is painfully obvious; the failure of early workers in artifical intelligence to achieve their ultimate goal no more "proves"

the impossibility of intelligent machines than did the many early failed attempts to fly "prove" that manned flight was impossible; it simply shows that the method used thus far - i.e., the attempt to free understanding from context - is flawed.

In fact, an entire branch of Mathematics, called Computability Theory, devotes itself almost entirely to the study of the potential domain of computer activity. 45 As a branch of pure Mathematics, it is largely removed from both the grandiose predictions of the "artificial intelligentsia" and the negative "proofs" of the phenomonologists. What has computability theory attained? Nothing final with regard to absolute limitations of machine intelligence has been proven, but it would, of course, be folly to infer from this that such intelligence is unlimited. The theory has, however, found a few "limitations" on the ability of one machine to predict another. In the preface of a text book in the field, we find the following remark, a rare effort to relate the abstract theory to the real world:

"In fact, one of the basic results of the theory of computability (namely, the existence of nonrecursive, recursively enumerable sets) may be interpreted as a asserting the possibilty of programming a given computer in such a way that it is impossible to program a computer (either a copy of the given computer or another machine) so as to determine whether or not a given item will be part of the output of the given computer."46

What is asserted here is essentially the possibility of an unpredictable computer, which scientists generally interpret to mean the possibility of "true randomness". But in religious

discourse, another interpretation is possible, namely that what is asserted is that a computer may be given free will. 47

That a computer may have free will is, of course, far from certain in an age when scientific opinion tends to the view that man himself does not have free will, and that the "free will" discussed by the theologians and philosophers is a figment of their own thought categories. But from this perspective, it is worth noting that if a computer claims, without having been told to do so, that it has free will. and makes even the most rudimentary and flawed arguments to support its case, it will be difficult to deny its claims without denying the validity of the corresponding human claims. While the issue of truly autonomous computers can never be solved finally by philosophical thought, it is quite conceivable that the machines themselves will, by their very asking the questions, render the issue equivalent to the problem of free will in human beings. Since free will is generally postulated to human beings in ethical discourse, we are thus justified in supposing the possibility of free will for robots in the discussion that follows.

The supposition of comprehension may be justified by a similar analogy. While it is indeed true that a computer can learn only contextually, and thus can not understand anything beyond the contextual framework it has been given, it is nonetheless undeniable that this is also largely true of human beings. The great philosopher simply does not "understand" the advanced physics paper because it is beyond the realm of his own contextual learning. Thus, as we have noted earlier,

a computer that extends its contextual understanding to a range of circumstances that is human in its scope can be said to "understand" to the same degree that this can be said of a human. It is thus in the loose, human sense that we shall presuppose in what follows that a computer may be given comprehension and free will.

## On Jewish Computers

worth living, we tend to forget that they are relationships among tangible things and inconceivable without them. Even wireless transmission of electricity implies transmitting and receiving sets."48

"The Golem partakes of the everlasting belief; and will rise again at the end of all human existence, but in a quite different form." - Rabbi Judah Loew of Prague 49

The first and simplest question we must answer is whether or not an autonomous robot can faithfully obey and execute the demands of the halakha. Little discussion is needed here; clearly a robot would be far less likely than a human being to forget or otherwise slip up in its observance of the many detailed laws that constitue the halakha. Of course, there are certain laws which are utterly irrelevant for computers; for example, the dietary and sexual laws hold little meaning for machines that neither eat nor reproduce. But the inability to comply with these laws would no more prevent the computer from being Jewish than did the absence of the Temple and the sacrifices prevent the Jews in exile from being Jewish. As Moses Mendelssohn said, "We are permitted to reflect on the law, to search for its meaning,

and occasionally, where the Lawgiver himself provides no reason (for a particular law) to surmise that it must perhaps be understood in terms of a particular time, place, and set of circumstances." Clearly, the prohibition against eating pork was intended for the special circumstance of beings who eat. Thus, there are no problems that would prevent a machine from observing the halakha.

But the halakha is more than a set of heteronomous laws which are given to us to be obeyed unthinkingly. Rather, they are ours by a covenant which we freely accepted; it is only by our autonomy that we accept it. We have postulated a robot with intelligence and free will. Under what circumstances can it accept the halakha (or, for Reform Jews, "Jewish values") and be accepted as a Jew?

## Who is a Jew?

"The ceremony of conversion... should be and often is conducted in such a way as to inspire and encourage the convert to draw spiritual joy from the foundations of the Jewish spirit and to impart added spiritual strength to the Jewish community."

"By the criteria of judgement built into ourselves, then, we find ourselves to be human beings within the meaning of the Three Laws, and human beings, moreover, to be given priority over (all) those others."

- George Ten, a robot52

The most obvious requirement for being a Jew is sentience.

This was the Golem's failure; he could not speak or think independently. But we are not speaking here of a mere Golem, and we have postulated these powers for it. What more is needed? A person is either born Jewish or converts to Judaism. It is difficult

to conceive of a computer being born of a ewish mother, so what we really need to know is what Jewish law requires for conversion. According to Rabbi Shlomo Goren, there are only three requirements: circumcision, which is not required of women and neither, by implication, of other beings without the correct anatomy; immersion, which is certainly conceivable for a machine, although it might have to be specially designed to be waterproof; and the acceptance of obligatory precepts in the presence of three qualified Jewish witnesses. 53 Clearly the key problem is in the meaning of "acceptance" of the obligatory precepts. Samuel Teitelbaum makes this more explicit when he notes that since the destruction of the Temple, "only converts who came of their own free will were wanted. Forced conversion was not countenanced..."54 Thus, we can not simply program a computer to accept Judaism, and if we have given a computer free will, we can only wait and see what it will do. What remains for us to consider, then, is why a robot might freely choose Judaism.

Why, indeed, would anyone freely choose Judaism? Abraham Franzblau suggests that people convert to Judaism for one or more of three kinds of reasons: prudential reasons (e.g., marriage to a Jew); neurotic reasons (choosing Judaism as one might choose a cult); and rational reasons. Obviously, the first two are irrelevant for the rational machines we are discussing. Of the third, the rational reasons for conversion, Franzblau says that this is essentially "the sheer conviction that it (Judaism) is superior to whatever faith was held previously." And

speaking of prospective converts, he says, "They may perhaps bw won over, undramatically, by a mature religion with an inspiring message and an appealing and enriching way of life."56 The prospective converts are expected to have free will, but it is also clear that additional basic characteristics are ascribed to them. No religion is "appealing" and no faith is "superior" to a creature that has absolutely no reason to search for a belief at all, which would seem to be the case for our nonemotional, potentially immortal machines. What is it in man that makes him search for faith and truth? If this can be found and given to a machine, then that machine can surely embrace Judaism or any other religion or system of belief that claims to offer truth. It will, of course, remain unpredictable whether a given machine will choose a given faith, but the capability will certainly be there. We thus remain with the question: motivates an intelligent, freely-willed being to search for truth?

# The Impulse Towards Truth

"And thou shalt love the Lord thy God with all thy heart, and with all thy soul, and with all thy might."57 "Justice, justice, shalt thou pursue."58

The Bible, and, in fact, just about every other great religious work, if full of, even based upon, commands of the type quoted above. Man is commanded to use all of his powers to strive towards the single purpose of knowing God, which

includes understanding the created world. <sup>59</sup> Yet this is, for many men, the most unnecessary of commandments, for the quest is within most of us innately, and need not be forced upon us. Why would a man read the Bible in the first place if not for his drive to know truth? For many men, the impulse to know God, or at least to know truth, is the deepest, most basic facet of existence. It is this drive which allows him to evaluate the different faiths and choose among them the most "appealing" way of thought. This has been termed, by some sociologists, the moral "instinct", but it is perhaps less confusing to use Maimonides' terminology; he speaks of man's inborn "disposition". <sup>60</sup>

The term "disposition" helps us to see how these inborn tendencies are not incompatible with the concept of human free will. As Maimonides says, these inborn dispositions do not mean that man posses virtue or vice by nature, but rather that the disposition makes it "easier to perform the actions that accord with a particular virtue or a particular vice". 61 To Maimonides, these dispositions do not alone constitute man's moral quest; rather, he feels that anyone who exercises logic will see the necessity of questing after truth and virtue. The dispositions are, for him, simply an aid or a hindrance. Thus, if Maimonides is right, our hypothetical freely-willed machines will choose of their own accord to quest after truth and virtue. This would be very nice, of course, and would absolve us of the need to make any effort at all to guide our creations to the moral quest;

we could simply work on the technical details of creating such a machine, confident that the innate logic of morality would triumph. But is this optimism justified? Could we not build such a machine and then keep it so busy thinking about other things that it never got around to thinking of morality? Could it not commit gross wrongs before it finally came to think seriously about ethics? This is perhaps an impossible question to answer without any details about the proposed machine, but it does point out that our responsibility for the actions of our creations does not vanish because of their free will any more than parents become exempt of responsibility the moment the child is able to express his own freely-willed desires. Just as the parents must guide the child, so must the designers guide the machine. And the builders of freely-willed machines will have one crucially important advantage over the parents of the child: they will be able to choose the inborn dispositions their creations will possess.

Thus, if we wish our autonomous machine to have an inborn tendency to seek truth and morality, we may give it to them.

Where Asimov wrote his Three Laws in bold, clear strokes, we may write "Justice, justice..." in big, vague letters — vague in that the nature of justice and the form of the pursuit are left open. We may picture the eminently logical robot seeking definitions of "truth" and "justice" and finding them everywhere, always differing slightly, always pointing in the same general direction. We can even imagine the rough edges of a purely

logical being gradually becoming more smooth in the face of an indefinable concept that everyone conceives of somehow. Such a robot, if he deemed it logical, could easily accept the principles of Judaism freely and completely, and become a real Jew. Any freely-willed robot might do so, but the proper disposition can make it more likely. We have now identified the necessary conditions for a robot to become Jewish, but, as one of these is free will, we cannot know whether or not these conditions are sufficient.

## Conclusion

We have seen that the traditional Jewish antipathy to the creation of a Golem, besides being inadequate in this world to prevent the development of the new Golem, is not fairly applicable to the computer, with its vast potential superiority over the Golden of old. We have further noted the difficulty of designing ethical systems for heteronomous machines, especially those machines whose knowledge is so contextually limited that all ethical commands must be made in tedious, inhuman detail, and have argued that no such humanly designed system should be allowed to become absolute and unchallengeable, as if divinely ordained. Finally, we have seen that if autonomous machines are ever created, they may choose freely to become Jews and be fully acceptable to the halakha, but only if this choice is not in any way dictated by man or by the machine's programming. only thing that man can do to lead his autonomous creations to the acceptance of his most cherished ideals is to give inborn

dispositions to seek truth, justice, and other indefinable ideals. The indefinable qualities of these ideals is itself the factor that makes the tendencies less than commanding laws; the possibility remains that the robots will, after due consideration, decide that the very notion of moral good is nonsensical, and hence, that the urge to seek it should be disregarded. Clearly, however, those of us, from Maimonides to modern "liberal" Jews, who believe in the ultimate rationality of ethical law, must place our faith and our hope for the future in the quest that such a machine would make, for it would be incapable of logical error. If its conclusions do not match ours, and to a lesser extent, even if they do, man will face the disturbing prospect of being an archaic, outmoded intelligence in a world where higher life forms exist. But the profound spiritual and aesthetic questions this would raise about the meaning of human existence and the future role of man in the universe are far beyond the scope of this paper.

#### FOOTNOTES

- 1. Laszlo, Ervin, The Systems View of the World, pp.94-95.
- 2. from a students' discussion "Actual and Possible Attitudes to Evolution within Orthodox Judaism" in Challenge: Torah Views on Science and its Problems, edited by Carmell and Domb, p.266.

3. Plout, W. Gunther, Judaism and the Scientific Spirit, pp.54-55.

- 4. Weizenbaum, Joseph, Computer Power and Human Reason, pp.206-7.
- 5. Kaplan, Mordecai M., Not So Random Thoughts, p.174.
- 6. Cf. Weizenbaum, op. cit., pp.270-272.
- 7. An important semantic problem should be noted here. In philosophical circles, an entity is considered "autonomous" only if it possesses free will. In scientific circles, especially among those concerned with robotics, the term "autonomous" is often used more liberally to describe any being that moves about and takes actions without being given detailed instructions for each specific action. Thus, if one tells a robot to "sweep the floor", and it takes such minor, uncommanded, auxilliary actions as going downstairs to get the broom and dustpan, it is acting "autonomously". In this paper, I will restrict the term "autonomy" to its true philosophical meaning, and will designate the slave-like "autonomy" discussed by computer scientists by the new term, "semi-autonomy". Also, note that the term "automaton" means simply "robot", and is unrelated to free will. Finally, "heteronomy" is like "autonomy", used in the philosophical sense; it is thus the opposite of "autonomy".
- 8. Weizenbaum, for example.
- 9. Bloch, Chayim, The Golem, p.33.
- 10. Cf. Sanhedrin, 65B.
- 11. Bloch, op.cit., p.71.
- 12. Cf. Gershom Scholem, "The Golem of Prague and the Golem of Rehovot", in The Messianic Idea in Judaism, pp. 335-6.
- 13. Cf. Bloch, op.cit., The Golem served the community well for thirteen years by this account.
- 14. Plout, op.cit., pp.57-58.
- 15. <u>Ibid</u>. (see also Scholem, op.cit., p.338).

- 16. Ibid.
- 17. Weizenbaum, op. cit., p. 133.
- 18. Ibid.
- 19. For these and most of the other fascinating details of the Golem legend, I know of no better source than Bloch, op.cit. Unfortunately, this book is long out of print in its English translation, but a copy may be found at the Hebrew Union College library here in Jerusalem.
- 20. I am indebted to Dr. W.Z. Harvey of the Hebrew University for bringing these two different types of progress to my attention.
- 21. Maimonides, Guide of the Perplexed, Pines translation, Book III, Chapter 54 (Volume 2, p.635).
- 22. Kaplan, op.cit., p.187.
- 23. Scholem, "The Messianic Idea In Kabbalism", in Scholem, op.cit., p.37.
- 24. Scholem, "The Golem of Prague and the Golem of Rehovot", op.cit., p.339.
- 25. See, for example, The Computer Bible, an exhaustive computer cross-referencing of nearly every word in the Bible. More recent systems, as yet unpublished, have submitted the entire Talmud to a much more sophisticated type of analysis.
- 26. Scholem, op.cit., p.340.
- 27. As reported by Bloch, op.cit., p.200.
- 28. Bamberger, Bernard J., "Conversion to Judaism: Theologically Speaking", in Conversion to Judaism, ed. David M. Eichhorn, p.177.
- 29. The question of ethical expectations of autonomous robots will be discussed in the fianl section of this paper.
- 30. אלן לא נכתבו ראויות לכתב.", Maimonides, Sh'moneh Prakim, מד-מה
- 31. Yoma, 67b.
- 32. Cf. Weiss and Butterworth, eds., Ethical Writings of Maimonides, p.101, footnote 12.
- 33. Scholem, op.cit., p.338.
- 34. from Isaac Asimov, "That Thou Art Mindful of Him!", in Final Stage, ed. Ferman and Malzberg, p.93.

- It seems possible to argue, on philosophical grounds, that 35. free will and intelligence are not the separable parameters that modern computer science makes them appear to be. However, this is, at bottom, a quibble over the nature of intelligence. Joseph Weizenbaum (op.cit.) has argued convincingly that the idea of "intelligence" is only meaningful contextually; thus, an IQ-like, absolute measure of true intelligence is fundamentally unconstructable without reference to specific contexts and circumstances. If this is to be denied, one must state that the mathematician who cannot follow a philosophical argument is more or less intelligent than the philosopher who cannot do simple algebra. But, if the contextual nature of intelligence is accepted, then it undeniably separate from free will, as many current advanced computer programs display remarkable degrees of contextual understanding without in any sense having free will. Cf. especially Weizenbaum's ELIZA program, discussed in Computer Power and Human Reason. Sample conversations with ELIZA may be found on pp.3-4 and pp.106-7.
- 36. Asimov, op.cit., p.91, The Three Laws have remained essentially unchanged since their formulation over 30 years ago; I quoted a recent version only because of Jerusalem's regrettable lack of a library with an extensive science fiction collection.
- 37. Cohn, Haim H., Jewish Law in Ancient and Modern Israel, p.IX.
- 38. <u>Ibid.</u>, p.X.
- 39. Rabbi Judah Loew of Prague, as reported by Bloch, op.cit.,p.203.
- 40. Gittelsohn, Roland B., Wings of the Morning, p.79.
- 41. from T. Dobzhansky, The Biology of Ultimate Concern, quoted in Ibid.
- 42. Asimov, op.cit.
- 43. Weizenbaum, op.cit., pp.204-5.
- 44. Dreyfus, Hubert L., What Computers Can't Do: A Critique of Artificial Reason, pp.215-6.
- 45. Formally, computability theory is the study of formalizable processes, or algorithms, and the situations for which they can and cannot be designed. From a non-technical perspective, this may seem indistinguishable from the study of what a computer can and cannot do, but the field does not really discuss computers as we know them, and the results may not be fairly said to be universally applicable to computers. Thus, the discussion that follows is only a weak attempt to gain perspective from an abstract mathematical discipline, the practitioners of which would surely be appalled at the carelessness with which their results are applied. Unfortunately, my own ignorance and the scope of this paper prevent a more rigorous application of the theory.

- 46. Davis, Martin, Computability and Unsolability, p.vii.
- 47. It is, perhaps, conceivable that computer unpredicatability would be neither random nor freely-willed, but something else altogether. Defining this, however, quickly becomes impossible, for indeed the idea that randomness and free will are themselves distinguishable is a matter of endless debate in certain circles.
- 48. Kaplan, op.cit., p.181.
- 49. As reported by Bloch, op.cit., p.202.
- 50. Mendelssohn, Moses, Jerusalem, p. 104.
- 51. Bamberger, op.cit., p.188.
- 52. Asimov, op.cit., p.114.
- 53. Cf. Goren's letter in Jewish Identity: Modern Responsa and Opinions on the Registration of Children of Mixed Marriages.
- 54. Teitelbaum, Samuel, "Conversion to Judaism: Sociologically Speaking" in Eichhorn, op.cit., p.211.
- 55. Franzblau, Abraham N., "Conversion to Judaism: Phsychologically Speaking" in Eichhorn, op.cit.
- 56. <u>Ibid.</u>, p.207
- 57. Deuteronomy 6:5.
- 58. Deuteronomy 16:20.
- 59. Cf. Chapter five of the <u>Eight Chapters</u> of Maimonides, <u>Ethical Writings of Maimonides</u>, ed. Weiss and Butterworth.
- 60. The translation is Weiss and Butterworth's, in <u>Ibid</u>., Chapter Eight.
- 61. Ibid., p.84.

#### BIBLIOGRAPHY

- Bin Gorion, Micha Joseph, collector, Mimekov Yisrael: Classic Jewish Folktales. Jewish Publication Society of America/ Indiana University Press, 1976.
- Bloch, Chayim, The Golem. Vienna, Printing Office John N. Vernay, 1925.
- Carmell, Arych and Cyril Domb, eds., Challenge: Torah Views on Science and its Problems. Jerusalem Association of Orthodox Jewish Scientists and Feldheim Publishers, 1976.
- Cohn, Haim H., <u>Jewish Law in Ancient and Modern Israel</u>. Ktav Publishing House, Inc., 1971.
- Davis, Martin, Computability and Unsolvability. New York, McGraw-Hill Book Co., Inc., 1958.
- Dreyfus, Hubert L., What Computers Can't Do: A Critique of Artificial Reason. New York, Harper and Row, 1972.
- Eichhorn, David Max, ed., Conversion to Judaism: A History and Analysis. Ktav Publishing House, Inc., 1965.
- Ferman, Edward L., and Barry N. Malzberg, eds., Final Stage. New York, Penguin Books, 1975.
- Gittelsohn, Roland B., Wings of the Morning. New York, Union of American Hebrew Congregations, 1969.
- Kaplan, Mordecai M., Not So Random Thoughts. New York, The Reconstructionist Press, 1966.
- Laszlo, Ervin, The Systems View of the World. New York, George Braziller, Inc., 1972.
- Litvin, Baruch, and Sidney B. Hoenig, Jewish Identity: Modern Response and Opinions on the Registration of Children of Mixed Marriages. New York, Philipp Feldheim, Inc., 1965.
- Maimonides, Moses, Ethical Writings of Maimonides, trans. Raymond L. Weiss and Charles L. Butterworth. New York, New York University Press, 1975.
- Maimonides, Moses, The Guide of the Perplexed, trans. Shlomo Pines. Chicago, The University of Chicago Press, 1963.
- Maimonides, Moses, Sh'moneh Prakim, trans. Shmuel ibn Tibbon. Jerusalem, Mossad Harav Kook, 5732.

- Mendelssohn, Moses, <u>Jerusalem</u>, trans. A. Jaspe. New York, Schochen Books, 1969.
- Nengroschel, Joachim, Great Works of Jewish Fantasy: Yenne Velt. London, Cassell and Co., Ltd., 1976.
- Plout, W. Gunther, Judaism and the Scientific Spirit. New York, Union of American Hebrew Congregations, 1962.
- Radday, Yehuda, <u>The Computer Bible</u>. Haifa, Biblical Research Associates, 1971.
- Scholem, Gershon, The Messianic Idea in Judaism. New York, Schochen Books, 1971.
- Weizenbaum, Joseph, Computer Power and Human Reason. San Francisco, W. H. Freeman and Co., 1976.