University of São Paulo

Polytechnic School of Engineering

Department of Electronic Systems

Handouts of slides of class T6 – Artificial Intelligence and Cognitivism – Part I

Date: May 1, 2020 Author: João E. Kogler Jr.

During the COVID-19 pandemic, with remote classrooms, I am providing this text to compensate the lack of face-to-face contact that we are experiencing. I expect that you read this material and prepare for the remote discussion session, on May 5, via teleconference. Enjoy it and keep yourself safe and healthy.

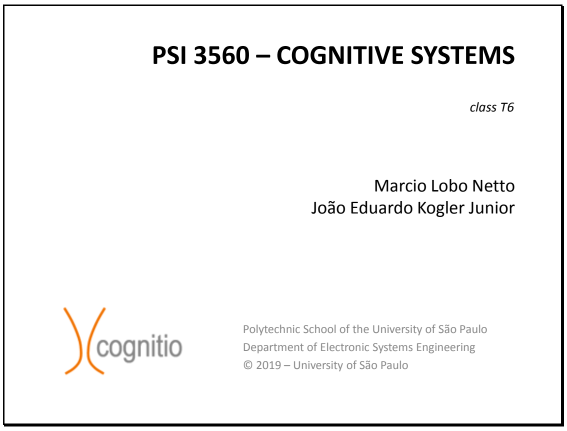

The exercise corresponding to this class T6 that is due to May 5 will be only the following below.

You are not asked to present anything else for T6. Just this exercise. So, do it carefully.

Exercise:

Write in a blank page the following and send it in pdf format to my email (kogler@lsi.usp.br) and submit it in the homework entry of Moodle for PSI3560:

- Name, date.
- I want to build a: _*put the theme of your Project here (a cognitive application)*_.

Just this. You will have to tell more about your project on the following week, for the T7 homework. But, for this T6 homework you just have to tell your project's theme (which, of course, you are allowed to ask to change before T7, in case you have a good explainable reason).

If you don't have good ideas, you can pick one from this list (which I created in less than five minutes, so, try to be creative, please!), or can adapt any of them, changing as you want. Nevertheless, keep in mind: whatever be your choice, for T7 you will be asked to justify why your choice is a cognitive application. Notice, however that once you committed to a choice, you will have to study and understand the application, in order to conceive its details.

List of ideas:

- Surveillance drone
- Autonomously guided taxi cab
- Chatter robot for…
- Cognitive supervisor for intelligent home
- Warehouse security cognitive robot watchdog
- Cognitive super toy for kids on the age…
- Early teenager personal assistant
- Homecare medical interface cognitive attendant
- Stock exchange investments broker

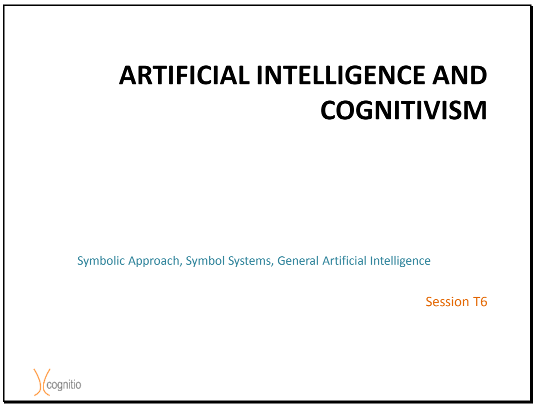

As we have already said before, the main target of this course is the study of the nature of cognitive systems and explore this knowledge for building cognitive applications. Let us emphasize again that an application is said a cognitive one if it depends on cognitive processes to be viable. Therefore, when you would have to justify that your application is cognitive, you must answer this question: “does my application require some cognitive process(es) ?”. Nevertheless, answering it positively although is sufficient, is not a necessary requirement. Eventually you can employ cognitive processes in your application even if it is not required under strict terms, but as a way to achieve better results. In this case the application can be turned into a cognitive one, given that it effectively uses some of the benefits provided by the cognitive processes involved. Then, recalling what we have already studied before, we can make a checklist to verify the requirement stated above:

- Does the application build knowledge that it extracts from gathered information?
- Does this knowledge improves the autonomy of the application or of its user or its own?
- Is there any process involved in the application that requires transforming knowledge?
- Are there decisions made by the application relying upon learned features?
- Does the application have to make predictions to complete partial information gathered from the operational environment, based on adaptive dynamic models?
- Are there computational-evolutive components in the application’s architecture?
- Does the application employ in its dynamics models that must be adapted on the fly, guided by invariant features detected on data?
- Does the application have to control actively (acting upon the environment) the process of information gathering?

This is a short list and the questions involved are very general. You may add to it other questions, maybe of more specific character. If any of these questions is answered positively, then your application is a strong candidate of being benefited by some kind of cognitive process.

Slide 4

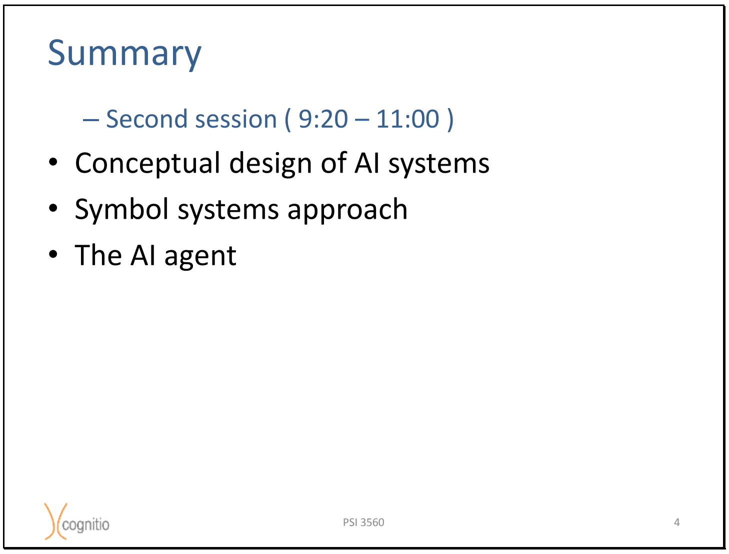

In this T6 class, we will examine the conceptual aspects of artificial intelligence (AI). As a methodological approach, AI constitutes a way to solve the challenges posed to the design of cognitive applications at the executive step of the project. It brings several techniques and algorithms that practically provide the required functionalities present in the application. However, our main interest in this course is in the conceptual step of the application's design: this is the step in which you explain why your application requires a cognitive character, and ow this contributes to its viability and performance. Therefore, we will not be interested here in the algorithms provided by AI, so we will not examine topics like techniques for searching the space of solutions, methods for logical inference, algorithms for planning, etc. Our interest is to show why the AI approach does its job, why it works, and when it doesn't. And, we want to know the relation of AI with the concept of cognition. In class T5 the main conceptual aspects of classical AI were already discussed.

Therefore, now we will start with the topic of understanding how AI relates to the conceptual design of systems.

Then, in the second topic, we will show that classical AI subsumes a methodology of symbols manipulations in a symbol system.

Finally, we will discuss the AI's view of the idea of agent.

Slides 5 and 6

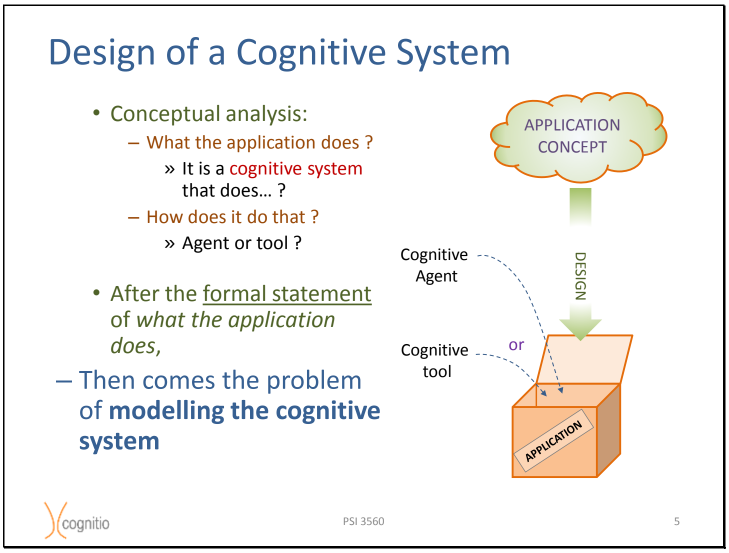

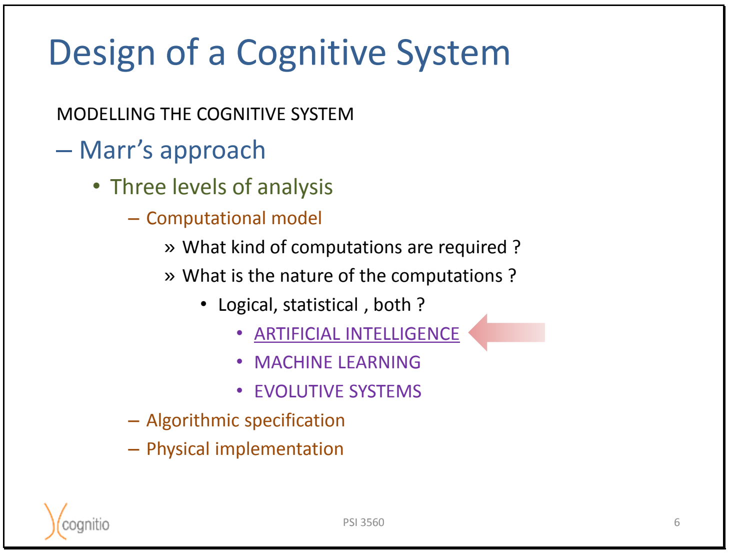

On class T4, we discussed some general aspects of the first step of a cognitive application's design, namely it's conceptual analysis. At this step you have to answer questions of the kind:

- What the application does?
- How does the application work, in terms of actions or output provided for a user?
- Is it an agent or a tool?

The answers should be then formally provided, in order to produce a specification of the application suitable for starting the next step of design process, which is the modelling step.

Thus, on class T4, we took Marr's theory and applied it to model the approach of investigation of a cognitive system, in order to derive a methodological strategy useful for the second step of the design procedure. What we learned was that it can be divided in three levels of analysis: (i) the computational modelling, (ii) the algorithm specification and, (iii) the physical implementation.

Whilst we still be dealing with the conceptual analysis of the application, the first level of Marr's approach provides the elements we need to model the application at this stage. Some particular issues on modelling, however, are concerned with technical details of data and process representation that affects certain aspects of the design processes, which we will not treat in this course, because they respect just to the algorithmic and implementation points of view. Again, we will consider only the conceptual analysis of the application, neither its algorithms, nor its implementation.

Precisely, the issues on computational modelling level of analysis that affect our interests are:

- What kinds of computations are required?
- Which is the nature of these computations: logical, statistical or both?

The first one of the answers, namely the logical approach provided by classic AI, is the theme of this T6 class.

Classes T7 and T8 will deal with the second answer, the statistical, by examining machine learning. Then, classes T9 and T10 will approach the third answer that is a mixture of both aspects, logical and statistical, comprehending adaptive dynamic systems and evolutionary computation. Check the syllabus below (from class F0):

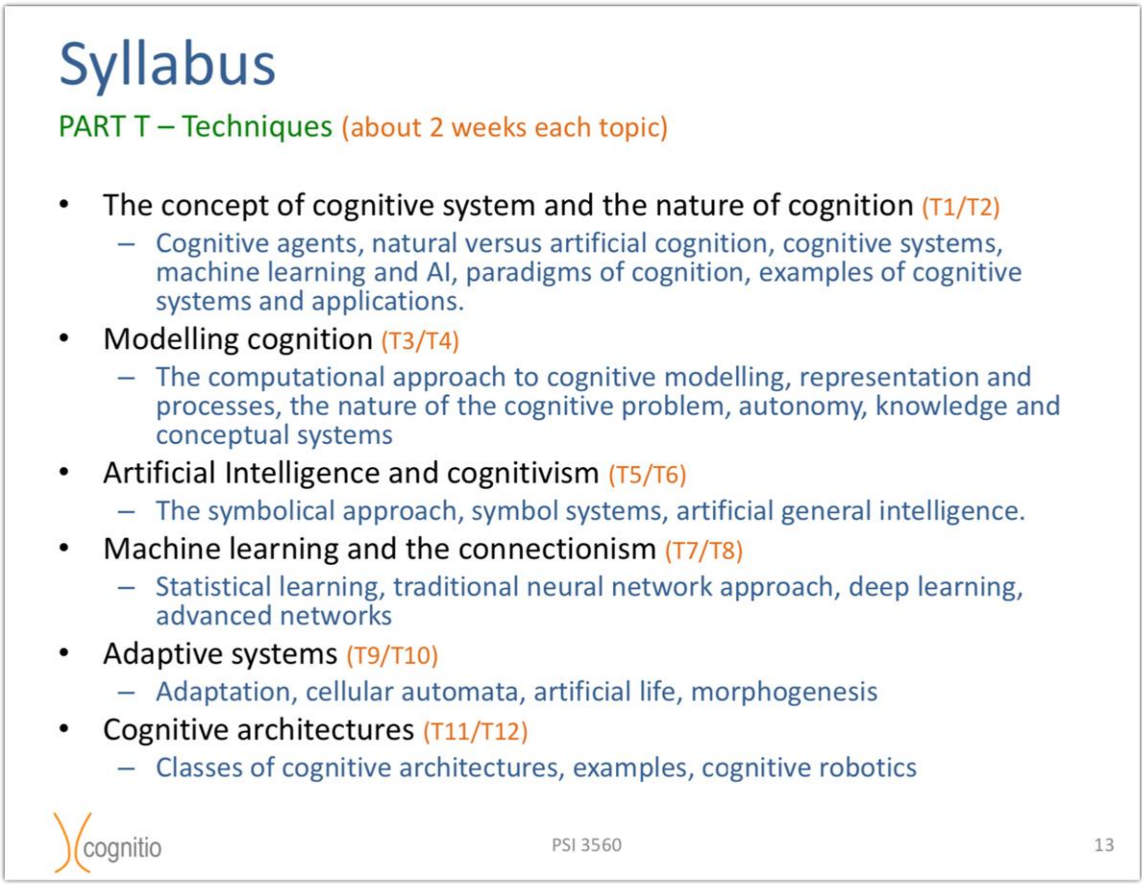

Slide 7

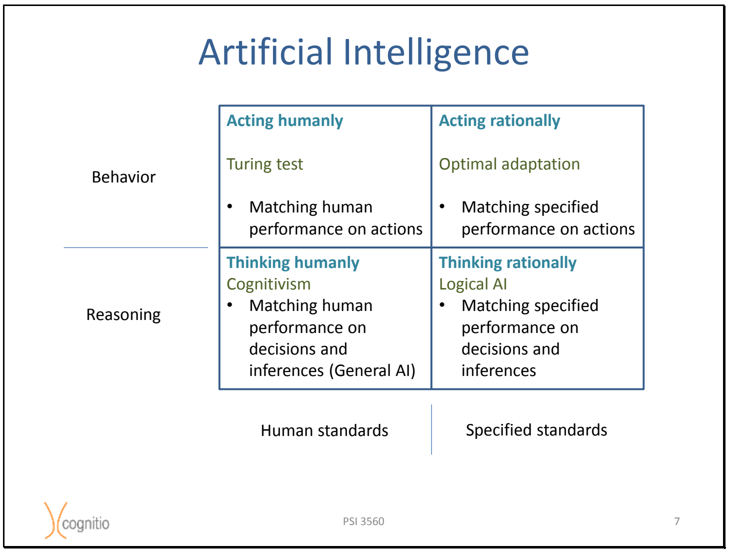

The issues of what is intelligence, how it differs from cognition, what is AI and how did it come up, were all treated in class T5, so we will consider now a more practical question that is how does AI accomplish its solutions. A summary of the given answers was elaborated by Russell and Norvig (1995-2010), which we present here with slight modifications.

Since AI takes the human intelligence as reference, one can achieve an artificial realization of it by specifying some standards that serve to guide the conception of applications that use AI. And, by providing such standards, one may observe two main aspects by which human intelligence manifests: human intelligent behavior and human intelligent reasoning.

Therefore, the table depicted above exams the two domains of realization, the natural (human) and the artificial, under the aspects of behavior and reasoning. Behavior relates to acting, while reasoning relates to thinking. In each case one has a corresponding practical target:

- Acting humanly – targets to matching human performance on actions. If the machine does as well as a human person, it can be taken as intelligent (Turing test, see class T5).
- Acting rationally – matching human performance, although interesting, in practice may not be achievable and maybe it is enough to obey just the rational principles that seems to guide human behavior, by giving the standards for practical rationality on acting.
- Thinking humanly – targets to matching human performance in making decisions and inferences. This was the point of view that guided cognitivism (see class T5).
- Thinking rationally – as in the case of human behavior, again it is more practical to comply with some specified standards for rationality of thinking, derived of characteristics of the human case.

Slides 8 and 9

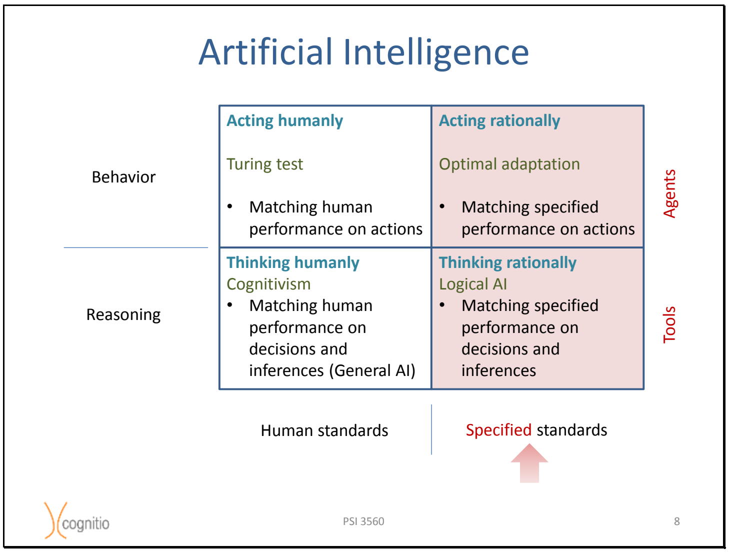

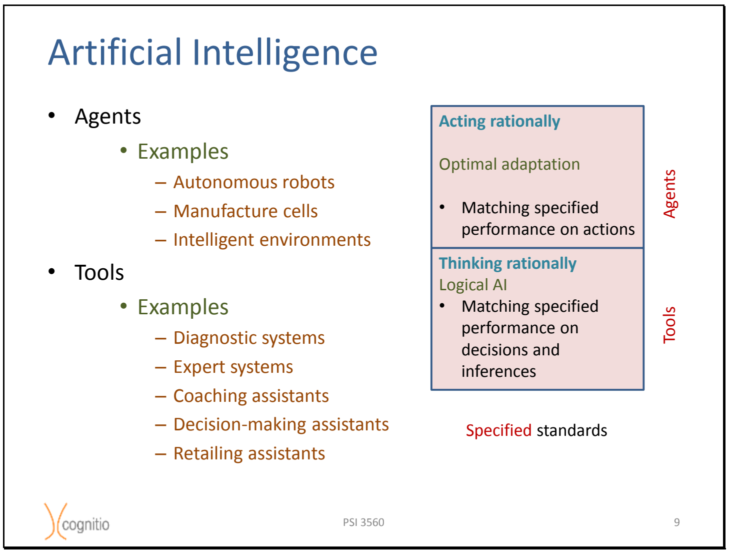

The left column of the table comprehends the themes considered in the F part of our course, where one is focusing on the foundational aspects of cognitive science. The right column are the aspects considered here at the T part. The approaches corresponding to each aspect, namely behavior and reasoning, are respectively:

- Optimal adaptation – matching specified performance on actions
- Logical reasoning – satisfaction of correctness and consistency

We have entitled these two lines of the table as “agents” and “tools”, as well as for the lists of examples of applications; however, this distinction was made in the terms that we usually have employed since the beginning of this course. In A.I. terminology, these terms are viewed differently, as we will show following.

Slide 10

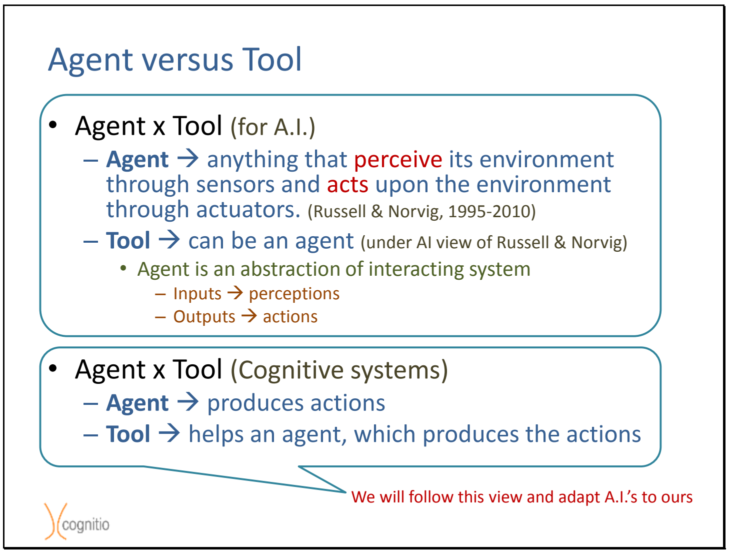

As put by Russell & Norvig (1995-2010), for the A.I. community agent is taken as an abstraction of a generic interacting system. Anything that receives inputs (perceptions, in the agent's terminology) and produces outputs (actions, in agent's terms), is considered an agent. For instance, a chatter robot software that receives inputs as queries from the keyboard and displays answers on the screen is taken as an agent.

Under our view of cognitive systems, however, an agent is an entity capable of executing and action, which effectively produces changes in the environment surrounding the agent. Therefore, under this view, the chatter robot may alter the inner state of its user, which could be considered a change in the robot's environment that includes the user. However, these actions are just conveying information that can be employed by the user to make a decision and issue some choice: the action is ultimately executed by the user. Consequently, under this view the chatter bot is no more than a tool for the user, who is the actual agent. Nevertheless, this is just a convention that facilitates in two points:

- Makes explicit the effective use of information (and knowledge, in particular) for making actions that can benefit the autonomy of the agent.
- Provides compatibility with the discourses of the several disciplines that cooperate to make cognitive science interdisciplinary.

For these reasons, we will preserve here our distinction and adopt it even in A.I. considerations.

Slide 11

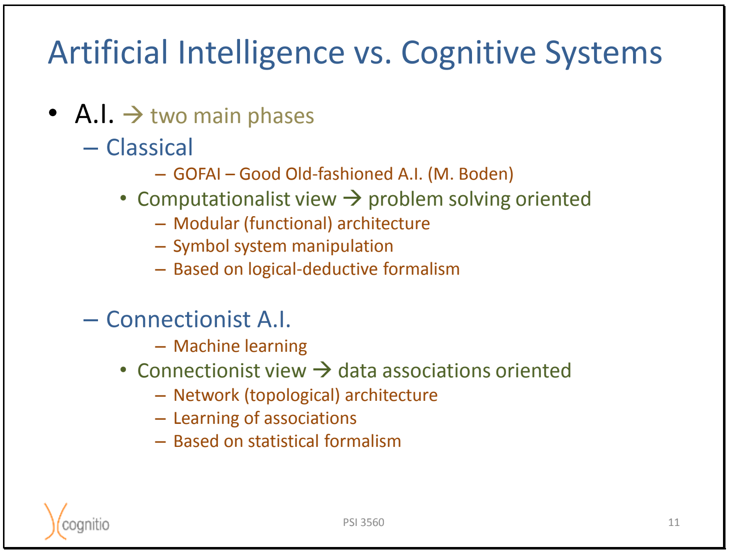

It was pointed on class T5 that A.I. as a scientific area lived two historic phases, classic A.I. and connectionist A.I.

Classic A.I., also called GOFAI – acronym for Good Old-Fashioned A.I., a term coined by Margareth Boden, a famous researcher on cognitive science – corresponds to the period since the appearance of the term A.I., to the beginning of the emergence of the successful use of machine learning methods with network-based architectures. This newer tendency was termed connectionist A.I., because of the use of networks as its architectural choice. However, the two eras differ in many other aspects.

The first aspect is the way by which the approaches consider the problems tackled with A.I. On classic A.I., the approach is the traditional problem-solving methodology of computer science, the branch of science dedicated to the structural study of the problems and their solutions. Under this view, problems can be grouped into distinct categories according to their structures and functional features. For this, solutions to them can be searched with correspondingly distinct strategies. Therefore, the mission of A.I. is to devise agents (and tools, in our terms), capable of solving the problem optimized to its structural characteristics, being a major challenge the structure identification.

Connectionist A.I., on the other hand, usually <u>ignore the structural aspects</u> of the problem, and consider problems as represented by solutions-data associations: which datum produces which answer. Therefore, the problem usually reduces to classification or regression.

The second aspect by which both differ is architectural. Classic A.I. systems reflects the organization of the problems they solve. On the other hand, connectionist A.I. systems are organized as networks of computing elements which topological ordering usually does not reflect the conceptual organizations of the problems they solve.

The third and fourth aspects by which the approaches differ are inter-related: the latter is the formal basis in which the former is founded. Classic A.I. rationale translates into a symbol-manipulation procedure, which finds its basis in the logical-deductive formalism. Connectionist A.I., on the other hand, has a rationale based on learning data-solutions associations ultimately founded on statistics and probability theory.

Slide 12

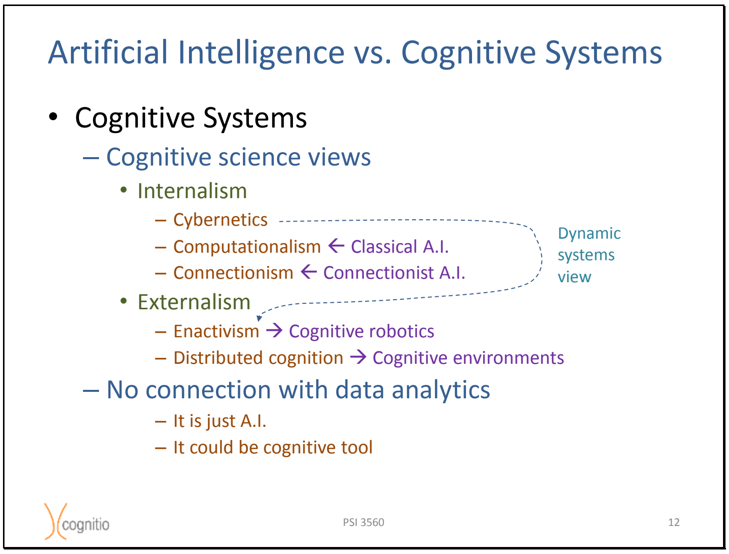

A latter aspect to consider briefly by now is the relation of A.I. with cognitive science. As it was examined on class T5, although intelligence and cognition are conceptually distinct terms, historically they were considered as equivalent for several practitioners. As was already discussed on previous classes, cognitive science and philosophy of mind are strongly related and they comprehend several schools of thinking. A first big divide is between the internalist and externalist points of view.

- Internalism – for internalists, the body – particularly the brain – is considered the site of cognition, and although interacting with the environment, they consider that the cognitive processing is a mechanism independent of the outside world, except for receiving inputs and producing actions. A.I. traditionally follows this view, both classic and connectionist.

- Externalism – for externalists, cognitive processes take place in every aspect of the world, so cognition extends beyond the brain and even the body, to tools and the environment. Although traditionally A.I. does not endorse this view, many areas of application of A.I. methods such as cognitive robotics and intelligent environments are being strongly influenced by the externalist views and have applied its principles to conceive their solutions.

Historically, A.I. is linked to the emergence of some of the schools of thinking in cognitive science. Computationalism, the school that defends that the mind reduces to computational processes is considered to have originated from classic A.I. and it is difficult to track the actual origins and identify which came up first.

Connectionism, the school of thinking in cognitive science that defends that the mind processes reduced to collective activity of neurons linked into circuits and networks has its origins both in neuroscience and connectionist A.I. . In this case, it is easy to trace the origin of connectionist A.I. also to neuroscience, because of the use of artificial neural networks as the prototypical model for connectionist architectures.

Another area in strong relation with A.I. is that of data analytics, also known as data science. This branch of study and application is devoted to the identification of problems in data, an innovative view distinct from the usual, which is to take a problem and find data that helps on its solution. The innovative view of data analytics is to study available data in some application domain and identify interesting problems that can be treated with the data and find their solutions. Potentially it is an area of strong interaction with cognitive science and A.I. However the approaches currently taken are closer to A.I. , mixing machine learning with classical logic-deductive methods. As we already discussed, their practitioners frequently label their activities as “cognitive applications”, while usually they aren’t. However, it is an area of potential interest for cognitive systems applications.