

# COMPUTERS AND MINDS

By Benjamin Spiegel  
Proposed Fall 2019

<b>Banner ID:</b>	<b>B01423421</b>	<b>Graduation Year:</b>	<b>2021</b>	<b>Semester Level:</b>	<b>5</b>
<b>Email:</b>	<b>Benjamin_Spiegel@Brown.edu</b>				

<b>Degree:</b> <b><u>A.B. or Sc.B.</u></b>	<b>Sc.B</b>	<b>Is this a revision of a proposal submitted earlier?</b>	<b>Yes</b>
Will this concentration <b>replace</b> one previously declared		<b>Yes</b>	<b>If, yes, which?</b> <b>A.B. CS</b>
Will it be in <b>addition</b> to one previously declared			<b>If, yes, which?</b>

<b>Faculty Sponsor:</b>	George D Konidaris, Ph.D	<b>Sponsor's Department:</b>	Computer Science
<b>Sponsor's Title:</b>	Assistant Professor of Computer Science		

<b>Cocktail Party Description of the IC:</b>	<i>Computers and Minds</i> is a concentration focused on studying intelligence, specifically how it has manifested and is conceptualized in fields like linguistics and psychology for the ultimate purpose of modeling intelligence with computers.
--	--

The Curricular Resource Center maintains an archive of approved IC proposals to allow prospective ICers to get a sense of past work. Proposals are made available upon your graduation. <b>"X" here if you <u>DO NOT</u> want your IC proposal to be made available in the IC archive upon graduation.</b>	
--	--

<b>Signature:</b>	<i>Benjamin Spiegel</i>	<b>Date:</b>	12/14/19
-------------------	-------------------------	--------------	----------

# Computers and Minds

An Independent Concentration by Benjamin Spiegel, Class of 2021

Contents:

I. Overview.....

II. What.....

III. How.....

IV. Why.....

V. Courses.....

VI. Courses (Annotated).....

VII. Thesis/Capstone.....

VIII. Bibliography.....

IX. Faculty Sponsor Letter.....

**Before submitting this proposal**, please consult the [IC Website](#), attend the IC Coordinators' open hours at the Curricular Resource Center ([CRC website](#)), and look through the archive of approved IC proposals at the [IC Database](#).

**Applicants must submit their first IC proposal by the end of their 5<sup>th</sup> semester. All proposals must receive FINAL, full approval by the end of their 6<sup>th</sup> semesters. Proposals are only reviewed three times per semester on select dates.**

**Applicants must meet with the coordinators prior to submitting their proposal. This should be done well in advance of the first submission to allow time for thoughtful revisions.**

The word limits are generally suggestions, but please consult regularly with the IC Coordinators to ensure you are on track in terms of your writing. This rubric for evaluating IC proposals is [here](#). Please note the spelling, grammar, and coherence of the writing in the proposal is important and should be carefully reviewed before submitting.

Email your completed application to [independent\\_concentrations@brown.edu](mailto:independent_concentrations@brown.edu).

---

**1. Cover Letter/Overview** - Provide a summary of the proposed field's key questions/debates, its reason for existing as an IC at Brown, and a very brief summary of why you are proposing it. **~300 words**

**NOTE:** If you are **resubmitting** with revisions, please additionally provide a summary of the specific changes you have made

---

The goal of this concentration is to get a better picture of the nature of human intelligence and how we might apply our understanding of it in the pursuit of artificial general intelligence. What are the fundamental facets of the general intelligence that humans possess? What are the biological and computational underpinnings of human intelligence? What are the societal implications of the creation of AGI and how may we integrate it into society in an ethical fashion? *Computers and Minds* poses these questions and seeks to answer them by looking at relevant courses from multiple academic departments, combining them into a single and comprehensive study of modeling general intelligence and its implications.

- ◇ Cleaned up the **Learning Goals**.
- ◇ Fleshed out the **What** with an additional example of a major offered at another institution that mirrors this concentration. Added a paragraph explaining key terms and framing AI. Added an informational diagram.
- ◇ The **course selection** has been modified to better suit the goals of this concentration. Alternative courses have been listed for flexibility, see the **cover letter** for more details.
- ◇ The **course annotation** has been updated according to the course selection changes.
- ◇ The **Capstone/Thesis Concept** was further developed.
- ◇ A **cover letter** was written in defense of this concentration.

In this final version of my proposal I will address the point about further discussion of the social implications of AI. While I have generally endorsed a positive outcome for artificial general intelligence on society, it is very possible that the rise of AGI may do more harm than good. If AGI were to fall into the wrong hands it could lead to a dystopian future where inequality would be greater than ever before. This is a risk that some are willing to take, and others are not, as many leaders in the technology space have publicly voiced concerns for AI while others have encouraged its rapid growth as a field. Regardless, we may only be able to see the social consequences of AI in decades, perhaps even centuries. My selection of courses on these topics serve to help me understand the fundamental groundwork of how societal structures operate so I may make my own inferences on how the rise of AGI will affect them. Their scope is limited to what Brown has to offer, and I have carefully selected courses that I believe will best complement the focus of this concentration.

With respect to how the above narrative coincides with my thesis, my concentration advisor, a leading researcher in the field of artificial intelligence, has recommended that I strictly pursue technical research into the field of Natural Language Understanding (NLU), an emerging subfield of AI. My purpose for creating this concentration was primarily to learn about the practical and theoretical uses and consequences of AI, thus a logical next step will be to pursue a Ph.D into a subfield of AI. In order to do this, I must demonstrate an ability to do in-depth technical research in NLU at the cutting edge. The addition of the ethical aspect of AGI into my thesis would detract from my ability to achieve a certain level of depth into the technical aspect, which would ultimately weaken my application for graduate school.

In regard to the course selection, I have included the possibility for alternate classes so that any future students who want to pursue this concentration have the flexibility to complete it even if certain classes are not being offered by professors in a given year. These alternate courses have similar content to their counterparts. In many departments, core courses are always offered, but the courses that are in this concentration are subject to the availability of the professors who teach them. Many professors in the computer science department go on sabbaticals or choose to do research full time in lieu of teaching their class.

---

**2. Learning Goals** - *Provide a clear description of the proposed concentration's academic goals and trajectory, reflecting on what skills a concentrator in this field would learn. ~250 words*

• *Think of this like the "Description" paragraph you see describing traditional concentrations on [Focal Point](#). You can also find examples of learning objectives for every concentration listed.*

---

This concentration is designed to unify the many interdisciplinary topics taught at Brown that are relevant to a course of study in artificial general intelligence (this will be defined in the **What** section).

In this concentration I will:

- Learn about the fundamental nature of human intelligence by studying its manifestation in human language and psyche.
- Learn about metaphysical aspects of the mind, including consciousness and happiness.
- Learn about the application of computer science in artificial intelligence.
- Understand the ethical implications of artificial general intelligence on society.


---

**3. What** - *Describe the context of the concentration outside of Brown and define key terms ~500-750 words*

- *If the concentration exists at other undergraduate liberal arts institutions comparable to Brown, describe the similarities and differences between the IC and those programs. Consider reaching out to members of those programs for details.*
  - *If the concentration does not exist elsewhere but is a clearly established field of study, describe the state of the field, bring in resources from active scholars, and discuss how your concentration will connect with the larger field. Consider consulting leading journal or contacting some of the relevant scholars for guidance.*
  - *If the concentration does not (yet) exist as a field of study, provide a very compelling description of its necessity, relevance, and academic merit.*
- 

Central to this concentration is the notion of intelligence, a widely debated topic in the scientific field. For the purposes of this document it's important to give rigid definitions to certain concepts, intelligence being one of them. These definitions are given **within the context of the field of artificial intelligence research**.

AI researchers often view **intelligence** with respect to an ability to solve tasks. These tasks can vary in diversity (think how different playing chess is to processing speech into text), and complexity (think about playing chess versus Connect 4). Historically, researchers would craft algorithms to be good at a particular task, but not necessarily good at other kinds of tasks that vary in diversity (imagine an algorithm that is really good at chess, but knows nothing about how to tell whether an image has a cat in it). This sort of algorithm is called a **narrow artificial intelligence** because it is only good at solving a narrow set of tasks within the task space. Recently, cutting-edge research has been on designing algorithms that are good at solving many different types of tasks, as opposed to just one type of task. In 2016, Google released one of these more **general** algorithms that learned to play seven different Atari 2600 games. A **general artificial intelligence (AGI)**, or an algorithm that can learn to do any task given sufficient time, is the Holy Grail of AI, and is expected by many futurists to significantly change the way the world works.



1. This diagram represents the task space (the set of all possible tasks) in two dimensions: complexity of the task and diversity, a spectrum of the types of tasks. The space is divided into areas that designate which tasks are do-able by which algorithm.

1

Term/Concept	Explanation and Definition
<b>Intelligence</b>	The ability to acquire and apply knowledge and skills. This definition is broadly accepted by cognitive scientists and computer scientists.
<b>General Intelligence</b>	The ability to integrate <i>multiple cognitive abilities</i> for the purpose of solving a novel problem. Humans have this sort of intelligence; humans are good at solving multiple different types of tasks.
<b>Artificial Intelligence (AI)</b>	The simulation of intelligent processes on machines.
<b>Narrow AI</b>	Systems that are good at learning how to solve a single type of task, but are not good at solving tasks outside of that category.
<b>General AI (AGI)</b>	Systems that are capable of learning multiple tasks of varying types. This sort of AI does not yet exist, and researchers are currently in a race to achieve it.

Majors and academic programs that have the same learning goals as *Computers and Minds* have been offered at other schools in the United States at least as far back as 1986. Below are three instances of pathways at other renowned institutions in the United States that very closely mirror *Computers and Minds*:

<sup>1</sup> Toward Tractable Universal Induction Through Recursive Program Learning - Scientific Figure on ResearchGate. Available from: [https://www.researchgate.net/figure/Approach-to-artificial-general-intelligence-Instead-of-trying-to-solve-complex-but\\_fig1\\_300646042](https://www.researchgate.net/figure/Approach-to-artificial-general-intelligence-Instead-of-trying-to-solve-complex-but_fig1_300646042) [accessed 17 Nov, 2019]

### *Carnegie Mellon University – Artificial Intelligence*

In 2018, CMU began offering the nation's first bachelor's degree in artificial intelligence. The program and its curriculum focus on how complex inputs — such as vision, language and huge databases — can be used to make decisions or enhance human capabilities. The curriculum includes coursework in computer science, math, statistics, computational modeling, machine learning and symbolic computation, as well as courses in ethics and social responsibility.

### *Stanford University – Symbolic Systems*

An interdisciplinary academic program created in 1986 that focuses on computers and minds, specifically the relationship between natural and artificial systems that represent, process, and act on information. Core requirements of the program include courses in symbolic logic, artificial intelligence, mathematics of computation, probability and statistics, programming, cognitive psychology, philosophy of mind, and interdisciplinary approaches to cognitive science. *Computers and Minds* shares the learning goals of this pathway.

### *University of Pennsylvania – Computer and Cognitive Science*

This pathway through the School of Engineering and Applied Science combines the application of theoretical insights from Computer Science, Linguistics, Neuroscience, Philosophy, and Psychology to the formal study of intelligence, perception, reasoning, and other properties of mind, and their application in the service of Information Technology. The degree combines a firm grounding in relevant aspects of Computer Science, from programming to algorithms to artificial intelligence. *Computers and Minds* shares the learning goals of this pathway.

For the sake of comparison, let's look at an example course-load that would satisfy Stanford's AI Concentration (this does not include Stanford's core requirement classes):

<b>Stanford's AI Concentration</b>	<b>Computers and Minds</b>
<b>CS 224N:</b> Natural Language Processing with Deep Learning	<b>CSCI 1460:</b> Computational Linguistics
<b>CS 229:</b> Machine Learning	<b>CSCI 1420:</b> Machine Learning
<b>CS 223A:</b> Introduction to Robotics	<b>CSCI 2951K:</b> Topics in Collaborative Robotics
<b>PHIL 20N:</b> Philosophy of Artificial Intelligence	<b>PHIL 1770:</b> Philosophy of Mind
<b>CS 221:</b> Artificial Intelligence: Principles and Techniques	<b>CSCI 1410:</b> Artificial Intelligence
<b>CS 107E:</b> Computer Systems From the Ground Up	<b>CSCI 0330:</b> Intro to Computer Systems
<b>MS&amp;E 234:</b> Data Privacy and Ethics	<b>DATA 0080:</b> Data, Ethics and Society

---

**4. How** - Discuss the scholarly methods to be used, the connections with any related disciplines, and how you plan to combine those disciplines using resources at Brown. **~500-750 words**

- Outline the analytic research methods or forms of inquiry of the concentration, highlighting the role of the methods course(s) in your course plan. These should come from and work with your methods course(s) and involve the practical tool for interacting with the information in your field (quantitative, qualitative, creative, etc.)

- Describe the connections with other academic fields, including how courses, texts, and techniques from other disciplines at Brown will be used to complete this concentration.

**NOTE:** If you are interested in participating in the [Engaged Scholars Program](#) as an independent concentrator, please include an additional description of your interest in the program and how it will relate to your proposed concentration.

---

This concentration studies intelligence and its implications within the scope of three pillars:

### **I. Cognitive Underpinnings and Manifestations of Human Intelligence**

What systems have emerged from human intelligence? Language, for one, is the culmination of thousands of years of combined effort to create an expansive system of communication. We are only now studying the shape and function of language retrospectively in the field of linguistics. The human mind is also a result of the development of human intelligence—studying cognitive science and psychology follow from this. Perhaps by studying the computational structure of the brain and the humanistic aspects of the psyche, we can glean some helpful clues about the nature of intelligence. This pillar examines the cognitive structures that give rise to general intelligence in humans alongside the historic manifestations of general intelligence in humanistic fields of study such as psychology and linguistics.

### **II. Societal Implications and Philosophical Aspects of General Intelligence**

How are intelligence and consciousness related? What are the societal implications of the advent of AGI? What will general automation do to our economy and how should we restructure society accordingly? What shape will AGI take within society? What distinguishes artificial from human general intelligence and what does this mean for the rights of individuals in a society? This pillar combines all of the theoretical aspects of artificial intelligence and their entailments, spanning consciousness and societal ethics. Classes in this category primarily fall into the philosophy and sociology departments. The goal of this pillar is to A) answer philosophical questions about whether artificial systems can be intelligent in the way humans are and B) theorize the possible impact of artificial intelligence on society in the future.

### **III. Technical Foundations and Applications of Artificial General Intelligence**

How can we leverage knowledge from the first two pillars in the application of AI models? What have researchers in the field of AGI done thus far? The purpose of this pillar is to grant us the tools to apply our knowledge to computer models of intelligence; it contains the breadth of the method courses for this concentration which mostly lie in the computer science department. Courses in this pillar will be conducive to a thesis in *Computers and Minds*.

---

### **5. Why – Provide a personal statement of justification for doing this concentration and an explanation of why the curriculum could not be pursued in a standard concentration at Brown. **~1000 words****

- Include a brief statement of personal interest in this topic and how you came to declare this IC.
- If the IC is similar to any other concentration(s), describe their differences - which should be significant enough to merit the creation of a new concentration. Why is this path necessary?
- Please refer explicitly to the existing concentrations as they appear on the [University Bulletin](#), which changes frequently! Be sure to reference standard concentrations accurately and mention advising conversations you



have had with other advisors from those programs.

- Read and consider various educational philosophies, including the [Liberal Learning Goals](#) and the philosophy outlined in [“Draft of a Working Paper for Education at Brown University”](#). Describe which educational goals the IC fulfills and how it will do so.
- 

Ultimately, I believe that Artificial General Intelligence will fundamentally change the way humans live. The integration of AGI into society has many potential risks, including the decimation of the entire human race, however it also has the potential to drastically improve quality of life across the board by eliminating the need for human labor, improving human education, and accelerating scientific research and technological innovation. The utilization of AGI in a society could resolve our presently unsustainable environmental practices and support the longevity of our civilization by giving us a chance at colonization of our species through space. Since the dawn of man, humans were fundamentally constrained by the need to acquire the natural resources necessary to sustain life. Societies formed to improve the overall efficiency of the acquisition of these resources, which in turn allowed for some of the population to focus on “human” endeavors like the arts and sciences. Advances in AGI could one day allow for the optimally efficient and sustainable acquisition of all of a society’s resources, thus freeing every human from nature’s fundamental demands.

I believe that a concentration studying AGI requires a holistic set of interdisciplinary classes that cannot be achieved in any other concentration offered at Brown. A few of the main contenders are listed and ruled out below.

#### *Computer Science – AI Track*

This concentration is designed to combine breadth in practical and theoretical computer science with depth in artificial intelligence. While this track could potentially cover two of the three pillars proposed by this concentration (Cognitive Underpinnings and Manifestation of Human Intelligence and Technical Application) it neglects to look in depth at the Cognitive Underpinnings of intelligence and the Philosophical Aspects of AGI. *Computers and Minds* is not a study of the application of computer science to a specific problem, it is the holistic study of general intelligence including its application, ethical and societal implications, cognitive underpinnings in humans and its manifestation in human fields.

#### *Behavioral Decision Sciences*

This concentration covers descriptive questions like how people, institutions, and nations make judgments and decisions; normative questions about rationality, such as what constitutes the best judgments and decisions; and prescriptive questions, such as how the process of decision making can be improved to make actual decisions closer to optimal ones. Students will choose three additional courses in consultation with a concentration advisor that will constitute an integrated specialization in some area of decision science.

BDS does cover a lot of the bases that *Computers and Minds* would cover, but with a very different perspective. In this concentration, AGI is at the center of attention. Decision making is the headliner of this concentration, but the study of AGI entails so much more than making decisions. This concentration neglects to address the focus of my field of study and simplistically reframes my studies with a perspective on the flat BDS, as opposed to a multidimensional AGI.

#### *Double Concentration in CS & Linguistics/Cognitive Science/Psychology/Philosophy*

The fundamental issues with this set of dual-concentrations is that they are limited to topics that are covered in each respective concentration. If I did CS + Linguistics, I'd have trouble fitting in classes in Cog Sci, Psych, and Phil. *Computers and Minds* spans each and every one of these fields and has a central theme that is not represented by the unification of any set of two individual concentrations.

## 6. Courses - List the courses required for the concentration.

- AB concentrations should include **about 11 Core Courses**, including **1 methods course**, and **1 departmental independent study course (DISP)** for your **capstone** (students who wish to complete a thesis or and Sc.B. capstone will have a minimum of 12 courses, as 2 DISPs are required). Most ScB concentrations include **about 17 Core Courses**.
- The course list should include a mix of introductory, intermediate, and advanced courses, logically building on earlier knowledge essential to completing capstone-like work. Adapt the "categories" column to best explain the flow and relationship between courses.
- Consider the knowledge and skills that are integral to the field. You may find it useful to list as alternative courses that accomplish similar intellectual goals (For example, if you plan to take 2 from a list of 5 courses).
- The **methods** course should provide the tools to comprehend, conduct, and report on research in your field (critical for your capstone/thesis). It may be qualitative, quantitative, theoretical, creative, practical, etc. and more than one course may be required depending on your field.
- Specify the [Writing Requirement course](#) you intend to take as a junior or senior.

**NOTE:** If you wish to join the Engaged Scholars Program, remember to list 2 [Engaged Courses](#), the .5 credit ESP reflection seminar or equivalent, and your practicum.

Stud. Sem.	Acad. Sem, Yr	Course Code	Course Title	Instructor	Category
1	Fall 2017	CLPS 0050A	Computing as Done in Brains and Computers	James Anderson	Intro
1	Fall 2017	CSCI 0190	Accelerated Introduction to Computer Science	Shriram Krishnamurthi	Intro
2	Spring 2018	CLPS 0100	Learning and Conditioning	Ruth Colwill	Intro
2	Spring 2018	PHIL 1770	Philosophy of Mind	Adam Pautz	Advanced
3	Fall 2018	CSCI 0330	Introduction to Computer Systems	Thomas Doeppner	Intro
3	Fall 2018	CSCI 1410	Artificial Intelligence	George D Konidakis	Advanced
		CSCI 1420	Machine Learning (Alternate for either 1410 or 1470)	Stephen Bach	Advanced
4	Spring 2019	CLPS 0710	Psychology and Philosophy of Happiness	Joachim Krueger	Intro
4	Spring 2019	CSCI 2951K	Collaborative Robotics	Stephanie Tellex	Advanced
5	Fall 2019	CLPS 0701	Personality	Brian Hayden	Intro

5	Fall 2019	CLPS 1342	Formal Semantics	Pauline Jacobson	Advanced
5	Fall 2019	CLPS 1370	Pragmatics	Scott Anderbois	Advanced
5	Fall 2019	CSCI 1470	Deep Learning	Daniel Ritchie	Methods
6	Spring 2020	CSCI 1460	Computational Linguistics	Eugene Charniak	Methods
		CSCI 2952I	Language Processing (Alternate for 1460)	Ellie Pavlick	Methods
6	Spring 2020	DATA 0080	Data, Ethics and Society	Deborah Hurley	Intro
6	Spring 2020	SOC 0300	Organizations and Society	Rachel Wetts	Advanced
6.5	Summer 2020	TBD	TBD WRIT Course	TBD	WRIT
7	Fall	DISP 1000	Independent Study	TBD	Thesis
8	Spring	DISP 2000	Independent Study	TBD	Thesis

There are a total of 20 courses listed in this chart, however only **18** of them are needed for the concentration, including the WRIT course and the two Independent Study Thesis courses.

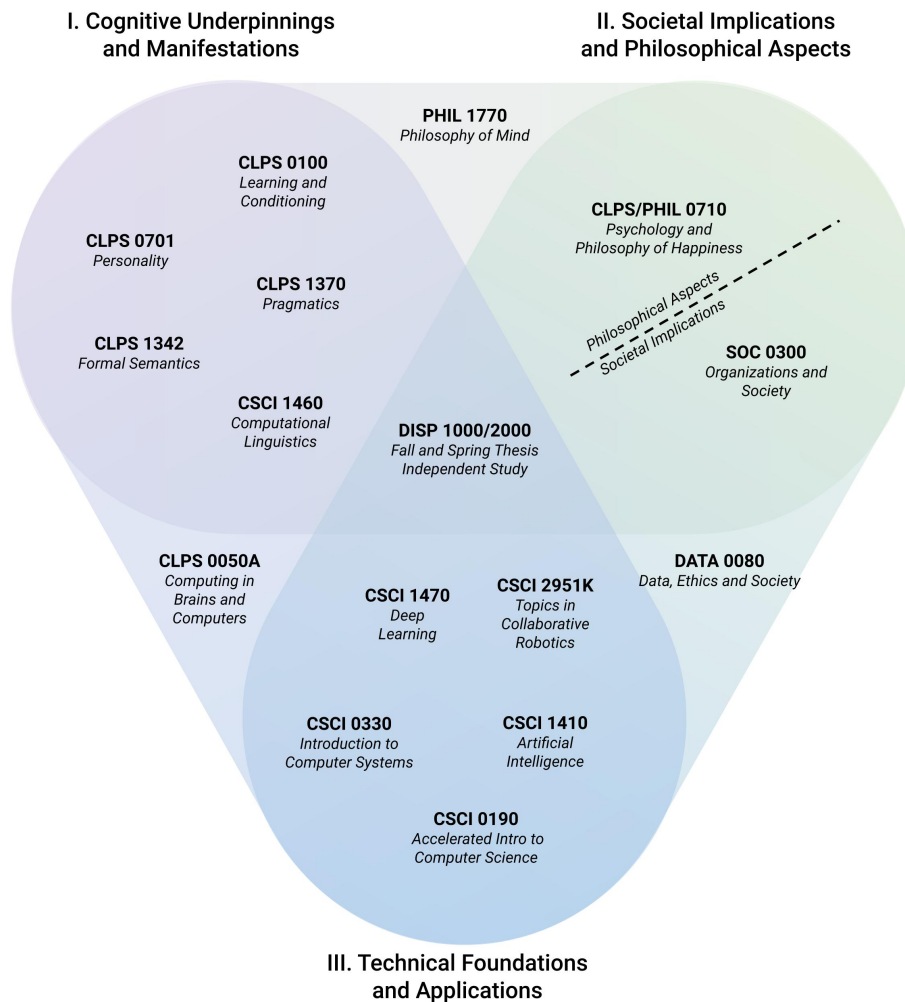
- CSCI 1420: Machine Learning may be substituted for either CSCI 1410: Artificial Intelligence or CSCI 1470: Deep Learning
- CSCI 2951I: Language Processing in Humans and Machines may be substituted for CSCI 1460: Computational Linguistics.

---

**7. Courses Annotations-** Describe how each of the courses supports the concentration ~**2-3 sentences each**

- Divide this section into categories that demonstrate the logical progression of the courses.
- Describe how each will contribute to the concentration's depth as well as breadth, relating the academic material to the concentration's field and **learning objectives**.

- Consider including a diagram (flow chart, Venn diagram, etc.) of courses to illustrate the coherence of the IC.



## I. Cognitive Underpinnings and Manifestations of Human Intelligence

The courses in this pillar examine the cognitive structures that give rise to general intelligence in humans alongside the historic manifestations of general intelligence in humanistic fields of study such as psychology and linguistics.

**CLPS 0050A - Computing as Done in Brains and Computers:** This course examines the similarities and differences between analog computing done in biological brains and digital computers.

**CLPS 0100 - Learning and Conditioning:** A psychology course focusing on how humans and other animals respond to stimuli with respect to a reward system.

**CLPS 0701 - Personality:** A psychology course focusing on historic personality theorists and theories, covers theories of learning, being social, psychological needs. All of these topics have implications for how we might construct a model for general intelligence.

**CLPS 1342 - Formal Semantics:** A linguistics course focusing on the meaning of sentences and developing rules for direct compositionality of the syntax and semantics of sentences.

**CLPS 1370 - Pragmatics:** A linguistics course focusing on the meanings of utterances within a conversation context. This course plays nicely with conversational AI.

## II. Societal Implications and Philosophical Aspects of General Intelligence

The courses in this pillar combine multiple theoretical aspects of AGI spanning consciousness, ethics of data science, and societal structure. The goal of this pillar is to A) answer philosophical questions about whether artificial systems can be intelligent in the way humans are and B) theorize the possible impact of artificial intelligence on society in the future.

**CLPS/PHIL 0710 - Psychology and Philosophy of Happiness:** A joint psychology/philosophy course that focuses on happiness, a primary mental motivator for humans as well as all animals. Has implications for creating AI models that mirror the biological reward function.

**PHIL 1770 - Philosophy of Mind:** This course focuses on the nature of the mind with respect to consciousness and volition.

**DATA 0080 - Data, Society and Ethics:** This is a course on the social, political, and philosophical issues raised by the theory and practice of data science. Students in this course will examine the field of data science in light of perspectives provided by the philosophy of science and technology, the sociology of knowledge, and science studies, and explore the consequences of data science for life in the first half of the 21st century.

**SOC 0300 – Organizations and Society:** This course introduces the field of Organizational Studies, examining organizations as complex, multifaceted social settings. By investigating how organizations and society shape each other, students will build skills for informed, socially-responsible engagement in an increasingly organizational society. Understanding how organizations work is critical to understanding how AGI could disrupt existing organizational frameworks.

## III. Technical Foundations and Applications of Artificial General Intelligence

These courses serve as the majority of the methods courses for the concentration. They are in the computer science department and they lay the groundwork for the application of the first two pillars. **Note** that two of the listed courses serve as alternatives for other courses listed.

**CSCI 0190 - Accelerated Introduction to Computer Science:** An accelerated intro sequence to computer science.

**CSCI 0330 - Introduction to Computer Systems:** Covers high-level computer architecture and systems programming.

**CSCI 1410 - Artificial Intelligence:** A survey course that explores many of the groundbreaking and fundamental AI models that have laid the groundwork for modern AI research.

**CSCI 1410 – Machine Learning:** A course that dives deep into the mathematics of modern AI models.  
**Alternate for 1410 or 1470.**

**CSCI 1460 - Computational Linguistics:** Explores the application of computational methods to problems in natural-language processing, a prominent field in AI research.

**CSCI 2952I – Language Processing in Humans and Machines:** Like 1460, this class explores the application of computational methods to problems in natural-language processing, a prominent field in AI research.  
**Alternate for 1460.**

**CSCI 1470 - Deep Learning:** This course exposes students to the latest advancements made in the growing field of deep learning. Different types of neural networks and their underlying architectures are explained and developed.

**CSCI 2951K - Topics in Collaborative Robotics:** An independent research class where students write technical research papers in the field of AI.

**DISP 1000/2000 - Independent Study:** These independent study courses will allow for development of a thesis. The content of these courses will be developed by myself and my advisor.

#### **8. Capstone/Thesis Concept - *Outline the culminating project for the concentration ~250 words***

- All ICs must include a final project to tie together the knowledge and skills of the concentration. Describe the format, content, and methodologies of the intended Capstone or Honors Thesis (details on the [IC Info Sheet](#) and the [Thesis Guidelines](#))
- For sophomores, this should be a rough outline of what the final project might entail. Juniors should have a clearer idea of what specific questions the Capstone or Thesis will investigate and the resources necessary.

I have plans to do a year-long thesis incorporating the three pillars of *Computers and Minds*. I'm especially interested in the field of Natural Language Understanding, and plan to write a publication on this topic. I have been participating in an ongoing dialog with Professor of Computer Science Ellie Pavlick, Ph.D, a leading researcher in this field, who will be helping along my journey alongside Professor of Computer Science George D Konidaris, Ph.D, the advisor to my concentration. At the moment I expect it to look like this, because it is understudied and draws from a classical semantic parsing. Thinking of doing grad school, the way to do that is to write a technical paper. I need to be able to write with narrow depth.

According to the Thesis Guidelines, students are required to "Think about potential topics; contact potential advisor and readers" by their 6<sup>th</sup> semester. I expect to have a clearer idea of my specific thesis idea by my 6<sup>th</sup> semester.

#### **9. Annotated Bibliography - *Provide an annotated list of scholarly works that are relevant to the concentration and support the contents of your proposal At least 10 works, 1-2 sentences each***

Last Updated 12/14/19 2:32 PM

- *Works should be either central to the concentration or necessary to fully understand its existence. Describe how each relates to the concentration. Do not simply provide a generic summary.*
- 

1. Feigebaum, A. *Computers and Thought*. McGraw-Hill, 1963.  
This book showcases the work of the scientists who not only defined the field of Artificial Intelligence, but who are responsible for having developed it into what it is today.
2. Jackendoff, Ray. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press, 2009.  
A deep dive into classical linguistics and the nature of various linguistic phenomena that AGI will be concerned with.
3. Jantsch, Erich. *The Self-Organizing Universe: Scientific and Human Implications of the Emerging Paradigm of Evolution*. Pergamon Press, 1992.  
A “big picture” systems-theory view of the mind and world, putting human intelligence in perspective.
4. Kellogg, Ronald T. *Fundamentals of Cognitive Psychology*. SAGE Publications, 2016.  
This text is a straightforward review of cognitive psychology, falls in line with the first pillar.
5. Kurzweil, Ray. *The Singularity Is near: When Humans Transcend Biology*. Duckworth, 2005.  
This book communicates a stunning and plausible future for AGI and its societal entailments. This falls in line with the second pillar.
6. RUSSELL, STUART NORVIG PETER. *ARTIFICIAL INTELLIGENCE: a Modern Approach*. PEARSON, 2018.  
This textbook closely examines traditional topics in narrow AI spanning language processing, reinforcement learning, machine learning, and planning.
7. Sutton, Richard S, and Andrew Barto. *Reinforcement Learning: an Introduction*. The MIT Press, 2018.  
This textbook focuses on reinforcement learning in artificial intelligence, most notably neural networks and control systems.
8. Tomasello, Michael. *Constructing a Language: a Usage-Based Theory of Language Acquisition*. Harvard University Press, 2005.  
This text reviews the concept of learning language as a social and embodied process. This has direct ties to the first pillar.
9. Wang, Pei, and Ben Goertzel. *Theoretical Foundations of Artificial General Intelligence*. Atlantis Press, 2012.  
This book lays down fundamental conceptual groundwork regarding AGI.
10. Wang, Pei. *Rigid Flexibility*. Springer, 2006.  
This book outlines a unique view on AGI and the underlying logical and control mechanisms.

---

#### **10. Faculty Sponsor Letter** - Attach a letter of support from the Faculty Sponsor

- The [Faculty Sponsor Info Sheet](#), including the instructions for this letter, can be found on the CRC’s website.
- 

See attached



**BROWN**  
Computer Science

**George Konidaris**

John E. Savage Assistant Professor  
Department of Computer Science  
<http://cs.brown.edu/people/gdk>  
[gdk@cs.brown.edu](mailto:gdk@cs.brown.edu)  
+1 413-335-1968

September 30<sup>th</sup> 2019

To Whom It May Concern:

I am writing to confirm my willingness to serve as Benjamin Spiegel's independent concentration faculty sponsor.

I have discussed Benjamin's independent concentration, in both intention and coursework detail, with him regularly since he took my class (Introduction to Artificial Intelligence) in the Fall of 2018. I believe that Benjamin is an unusually thoughtful and dedicated young student who has correctly identified the need for an inter-disciplinary approach to studying general intelligence, and who has carefully thought out a well-rounded curriculum to follow.

In terms of my own participation, I expect to meet with Benjamin three or four times a semester to discuss his progress throughout the concentration, and discuss his thoughts about the topic as a whole. We will also discuss career prospects and internship opportunities.

I will also serve as Benjamin's Honors thesis advisor during his senior year. Benjamin will enroll in two successive independent study courses with me, during which we will meet weekly. I will also supervise a short informal reading program the summer before Benjamin's senior year, where we will focus on reading cutting-edge research papers and developing a well-formulated research question. Our goal over the course of the Honors thesis will be to produce a publication in AI that both addresses Benjamin's interests and positions him well for PhD applications, if that is what he decides to do.

Please do not hesitate to contact me if you need further information.

Sincerely,

George Konidaris, Ph.D.