Under the Paperwork Reduction Act of 1995 no persons are required to respond to a collection of information unless it displays a valid OMB control number.
PTO Form 1478 (Rev 09/2006)
OMB No. 0651-0009 (Exp 02/28/2021)

Trademark/Service Mark Application, Principal Register

Serial Number: 88895303
Filing Date: 04/30/020

The table below presents the data as entered.

Input Field	Entered
SERIAL NUMBER	88895303
MARK INFORMATION	
***MARK**	State of Mind AI (SOMAI)
STANDARD CHARACTERS	YES
USPTO-GENERATED IMAGE	YES
LITERAL ELEMENT	State of Mind AI (SOMAI)
MARK STATEMENT	The mark consists of standard characters, without claim to any particular font style, size, or color.
REGISTER	Principal
APPLICANT INFORMATION	
***OWNER OF MARK**	Pereptive Automata, Inc.
***MAILING ADDRESS**	Suite #33, 1250 Borregas Ave
***CITY**	Sunnyvale
***STATE** (Required for U.S. applicants)	California
***COUNTRY/REGION/JURISDICTION/U.S. TERRITORY**	United States
***ZIP/POSTAL CODE** (Required for U.S. and certain international addresses)	94089
***EMAIL ADDRESS**	XXXX
LEGAL ENTITY INFORMATION	
TYPE	corporation
STATE/COUNTRY/REGION/JURISDICTION/U.S. TERRITORY OF INCORPORATION	Delaware
GOODS AND/OR SERVICES AND BASIS INFORMATION	
INTERNATIONAL CLASS	009
***IDENTIFICATION**	Downloadable software, source code, machine learning algorithms, APIs, supporting tools and ser systems, including software using artificial intelligence, that process sensor data and output humar signals, such as intent-to-cross and awareness of human road users for automated vehicle applicati modulate its operating decisions and actions, all for automated systems and/or robotics developers,
FILING BASIS	SECTION 1(a)
FIRST USE ANYWHERE DATE	At least as early as 12/01/2018
FIRST USE IN COMMERCE DATE	At least as early as 12/01/2018

SPECIMEN FILE NAME(S)	
ORIGINAL PDF FILE	SPE0-2601646480e18060b390b47911678-20200430134303601403_._Perceptive_Automata_SOM
CONVERTED PDF FILE(S) (14 pages)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0003.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0004.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0005.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0006.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0007.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0008.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0009.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0010.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0011.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0012.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0013.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0014.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0015.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0016.JPG
ORIGINAL PDF FILE	SPE0-2601646480e18060b390b47911678-20200430134303601403_._Customer_Use_Case_Man
CONVERTED PDF FILE(S) (12 pages)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0017.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0018.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0019.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0020.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0021.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0022.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0023.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0024.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0025.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0026.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0027.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0028.JPG
ORIGINAL PDF FILE	SPE0-2601646480e18060b390b47911678-20200430134303601403_._SOMAI_Service_Release_
CONVERTED PDF FILE(S) (1 page)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0029.JPG
ORIGINAL PDF FILE	SPE0-2601646480e18060b390b47911678-20200430134303601403_._State_of_Mind__SOMAI_
CONVERTED PDF FILE(S) (3 pages)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0030.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0031.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0032.JPG

SPECIMEN DESCRIPTION	State of Mind (SOMAI) presentation to customers State of Mind (SOMAI) user manual for custom Download and Release Instructions SOMAI User Guide
INTERNATIONAL CLASS	042
*IDENTIFICATION	Research and development of automated vehicles; software design and development in the field of cloud-based and/or browser-accessible software, source code, machine learning algorithms, APIs, automated and/or robotic systems, including software using artificial intelligence, that process sens and/or behavior prediction signals, such as intent-to-cross and awareness of human road users for a automated system can use to modulate its operating decisions and actions, all for automated system manufacturers, suppliers, and end-users.
FILING BASIS	SECTION 1(a)
FIRST USE ANYWHERE DATE	At least as early as 12/01/2018
FIRST USE IN COMMERCE DATE	At least as early as 12/01/2018
SPECIMEN FILE NAME(S)	
ORIGINAL PDF FILE	SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._Perceptive_Automata_SC
CONVERTED PDF FILE(S) (14 pages)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0033.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0034.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0035.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0036.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0037.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0038.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0039.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0040.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0041.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0042.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0043.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0044.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0045.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0046.JPG
ORIGINAL PDF FILE	SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._Customer_Use_Case_Ma
CONVERTED PDF FILE(S) (12 pages)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0047.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0048.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0049.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0050.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0051.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0052.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0053.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0054.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0055.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0056.JPG

	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0057.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0058.JPG
ORIGINAL PDF FILE	SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._SOMAI_Service_Releas
CONVERTED PDF FILE(S) (1 page)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0059.JPG
ORIGINAL PDF FILE	SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._State_of_Mind__SOMA
CONVERTED PDF FILE(S) (3 pages)	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0060.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0061.JPG
	\\TICRS\EXPORT18\IMAGEOUT18\888\953\88895303\xml1\APP0062.JPG
SPECIMEN DESCRIPTION	Customer Presentation re: SOMAI State of Mind AI Customer Use Case Summary re: State of Mi User Guide SOMAI Customer Service User Guide and Download Instructions
ATTORNEY INFORMATION	
NAME	Peter Piotr Szymanski
ATTORNEY BAR MEMBERSHIP NUMBER	XXX
YEAR OF ADMISSION	XXXX
U.S. STATE/ COMMONWEALTH/ TERRITORY	XX
STREET	1900 Jefferson, 205
CITY	San Francisco
STATE	California
COUNTRY/REGION/JURISDICTION/U.S. TERRITORY	United States
ZIP/POSTAL CODE	94123
EMAIL ADDRESS	peter@siliconvalleycounsel.com
CORRESPONDENCE INFORMATION	
NAME	Peter Piotr Szymanski
PRIMARY EMAIL ADDRESS FOR CORRESPONDENCE	peter@siliconvalleycounsel.com
SECONDARY EMAIL ADDRESS(ES) (COURTESY COPIES)	pete@siliconvalleycounsel.com
FEE INFORMATION	
APPLICATION FILING OPTION	TEAS Standard
NUMBER OF CLASSES	2
APPLICATION FOR REGISTRATION PER CLASS	275
***TOTAL FEES DUE**	550
***TOTAL FEES PAID**	550
SIGNATURE INFORMATION	
SIGNATURE	/Peter Piotr Szymanski/
SIGNATORY'S NAME	Peter Piotr Szymanski
SIGNATORY'S POSITION	Attorney of Record, California bar member

SIGNATORY'S PHONE NUMBER	650-776-4826
DATE SIGNED	04/30/2020

Under the Paperwork Reduction Act of 1995 no persons are required to respond to a collection of information unless it displays a valid OMB control number.
PTO Form 1478 (Rev 09/2006)
OMB No. 0651-0009 (Exp 02/28/2021)

Trademark/Service Mark Application, Principal Register

Serial Number: 88895303
Filing Date: 04/30/020

To the Commissioner for Trademarks:

MARK: State of Mind AI (SOMAI) (Standard Characters, see mark)
The literal element of the mark consists of State of Mind AI (SOMAI). The mark consists of standard characters, without claim to any particular font style, size, or color.
The applicant, Pereptive Automata, Inc., a corporation of Delaware, having an address of
Suite #33, 1250 Borregas Ave
Sunnyvale, California 94089
United States
XXXX

requests registration of the trademark/service mark identified above in the United States Patent and Trademark Office on the Principal Register established by the Act of July 5, 1946 (15 U.S.C. Section 1051 et seq.), as amended, for the following:

International Class 009: Downloadable software, source code, machine learning algorithms, APIs, supporting tools and services for automated and/or robotic systems, including software using artificial intelligence, that process sensor data and output human state of mind and/or behavior prediction signals, such as intent-to-cross and awareness of human road users for automated vehicle applications, that an automated system can use to modulate its operating decisions and actions, all for automated systems and/or robotics developers, manufacturers, suppliers, and end-users

In International Class 009, the mark was first used by the applicant or the applicant's related company or licensee or predecessor in interest at least as early as 12/01/2018, and first used in commerce at least as early as 12/01/2018, and is now in use in such commerce. The applicant is submitting one(or more) specimen(s) showing the mark as used in commerce on or in connection with any item in the class of listed goods/services, consisting of a(n) State of Mind (SOMAI) presentation to customers State of Mind (SOMAI) user manual for customers SOMAI Customer Service Software Download and Release Instructions SOMAI User Guide.

Original PDF file:
SPE0-2601646480e18060b390b47911678-20200430134303601403_._Perceptive_Automata_SOMAI_Overview_for_Customers__1_.pdf
Converted PDF file(s) (14 pages)
Specimen File1
Specimen File2
Specimen File3
Specimen File4
Specimen File5
Specimen File6
Specimen File7
Specimen File8
Specimen File9
Specimen File10
Specimen File11
Specimen File12
Specimen File13
Specimen File14
Original PDF file:
SPE0-2601646480e18060b390b47911678-20200430134303601403_._Customer_Use_Case_Manual_for_SOMAI_Deployment.pdf
Converted PDF file(s) (12 pages)
Specimen File1
Specimen File2
Specimen File3
Specimen File4
Specimen File5

Specimen File6
Specimen File7
Specimen File8
Specimen File9
Specimen File10
Specimen File11
Specimen File12
Original PDF file:
SPE0-2601646480e18060b390b47911678-20200430134303601403_._SOMAI_Service_Release___Download_Page.pdf
Converted PDF file(s) (1 page)
Specimen File1
Original PDF file:
SPE0-2601646480e18060b390b47911678-20200430134303601403_._State_of_Mind__SOMAI__User_Guide_v1.04.pdf
Converted PDF file(s) (3 pages)
Specimen File1
Specimen File2
Specimen File3

International Class 042: Research and development of automated vehicles; software design and development in the field of automated vehicles; Non-downloadable cloud-based and/or browser-accessible software, source code, machine learning algorithms, APIs, supporting tools and services for automated and/or robotic systems, including software using artificial intelligence, that process sensor data and output human state of mind and/or behavior prediction signals, such as intent-to-cross and awareness of human road users for automated vehicle applications, that an automated system can use to modulate its operating decisions and actions, all for automated systems and/or robotics developers, manufacturers, suppliers, and end-users.

In International Class 042, the mark was first used by the applicant or the applicant's related company or licensee or predecessor in interest at least as early as 12/01/2018, and first used in commerce at least as early as 12/01/2018, and is now in use in such commerce. The applicant is submitting one(or more) specimen(s) showing the mark as used in commerce on or in connection with any item in the class of listed goods/services, consisting of a(n) Customer Presentation re: SOMAI State of Mind AI Customer Use Case Summary re: State of Mind (SOMAI) State of Mind AI (SOMAI) User Guide SOMAI Customer Service User Guide and Download Instructions.

Original PDF file:
SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._Perceptive_Automata_SOMAI_Overview_for_Customers__1_.pdf
Converted PDF file(s) (14 pages)
Specimen File1
Specimen File2
Specimen File3
Specimen File4
Specimen File5
Specimen File6
Specimen File7
Specimen File8
Specimen File9
Specimen File10
Specimen File11
Specimen File12
Specimen File13
Specimen File14
Original PDF file:
SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._Customer_Use_Case_Manual_for_SOMAI_Deployment.pdf
Converted PDF file(s) (12 pages)
Specimen File1
Specimen File2
Specimen File3
Specimen File4
Specimen File5
Specimen File6
Specimen File7
Specimen File8

Specimen File9
Specimen File10
Specimen File11
Specimen File12
Original PDF file:
SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._SOMAI_Service_Release___Download_Page.pdf
Converted PDF file(s) (1 page)
Specimen File1
Original PDF file:
SPE0-1-2601646480e18060b390b47911678-20200430134303601403_._State_of_Mind__SOMAI__User_Guide_v1.04.pdf
Converted PDF file(s) (3 pages)
Specimen File1
Specimen File2
Specimen File3

The owner's/holder's proposed attorney information: Peter Piotr Szymanski. Peter Piotr Szymanski, is a member of the XX bar, admitted to the bar in XXXX, bar membership no. XXX, is located at

1900 Jefferson, 205
San Francisco, California 94123
United States
peter@siliconvalleycounsel.com

Peter Piotr Szymanski submitted the following statement: The attorney of record is an active member in good standing of the bar of the highest court of a U.S. state, the District of Columbia, or any U.S. Commonwealth or territory.
The applicant's current Correspondence Information:

Peter Piotr Szymanski
PRIMARY EMAIL FOR CORRESPONDENCE: peter@siliconvalleycounsel.com
SECONDARY EMAIL ADDRESS(ES) (COURTESY COPIES): pete@siliconvalleycounsel.com

Requirement for Email and Electronic Filing: I understand that a valid email address must be maintained by the applicant owner/holder and the applicant owner's/holder's attorney, if appointed, and that all official trademark correspondence must be submitted via the Trademark Electronic Application System (TEAS).
A fee payment in the amount of $550 has been submitted with the application, representing payment for 2 class(es).

Declaration

☑ **Basis:**
If the applicant is filing the application based on use in commerce under 15 U.S.C. § 1051(a):

- The signatory believes that the applicant is the owner of the trademark/service mark sought to be registered;
- The mark is in use in commerce and was in use in commerce as of the filing date of the application on or in connection with the goods/services in the application;
- The specimen(s) shows the mark as used on or in connection with the goods/services in the application and was used on or in connection with the goods/services in the application as of the application filing date; and
- To the best of the signatory's knowledge and belief, the facts recited in the application are accurate.

And/Or
If the applicant is filing the application based on an intent to use the mark in commerce under 15 U.S.C. § 1051(b), § 1126(d), and/or § 1126(e):

- The signatory believes that the applicant is entitled to use the mark in commerce;
- The applicant has a bona fide intention to use the mark in commerce and had a bona fide intention to use the mark in commerce as of the application filing date on or in connection with the goods/services in the application; and
- To the best of the signatory's knowledge and belief, the facts recited in the application are accurate.

☑ To the best of the signatory's knowledge and belief, no other persons, except, if applicable, concurrent users, have the right to use the mark in commerce, either in the identical form or in such near resemblance as to be likely, when used on or in connection with the goods/services of such other persons, to cause confusion or mistake, or to deceive.

- ☑ To the best of the signatory's knowledge, information, and belief, formed after an inquiry reasonable under the circumstances, the allegations and other factual contentions made above have evidentiary support.
- ☑ The signatory being warned that willful false statements and the like are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001, and that such willful false statements and the like may jeopardize the validity of the application or submission or any registration resulting therefrom, declares that all statements made of his/her own knowledge are true and all statements made on information and belief are believed to be true.

Declaration Signature

Signature: /Peter Piotr Szymanski/   Date: 04/30/2020
Signatory's Name: Peter Piotr Szymanski
Signatory's Position: Attorney of Record, California bar member
Payment Sale Number: 88895303
Payment Accounting Date: 04/30/2020

Serial Number: 88895303
Internet Transmission Date: Thu Apr 30 14:09:43 ET 2020
TEAS Stamp: USPTO/BAS-XXXX:XXX:XXX:XXXX:XXXX:XXXX:XX
X:XXXX-20200430140943073671-88895303-710
5d947ae8a3c1f7679d99c01cccac50b635f4c2fa
13d932b15c5f747a274-CC-09419029-20200430
134303601403

State of Mind AI (SOMAI)


PERCEPTIVE
AUTOMATA
Human Intuition for Machines
Overview for
Customers

ANTICIPATING HUMAN BEHAVIOR IS ONE OF THE THE HARDEST PROBLEMS FOR AUTOMATED DRIVING

"It's the prediction piece that's still the great unknown. Humans are very good at predicting human behavior on the road. Machines will need to be able to predict and anticipate human behavior much better."

Gill Pratt
CEO of Toyota Research Institute
March 2018

"The choices made by driverless cars are critically dependent on understanding and matching the expectations of humans. This is the heart of the problem going forward."

Chris Urmson
Head of Waymo, 2009-2016
Co-Founder of Aurora, April 2017

"With radar and high-resolution cameras and all the computing power we have, we can detect and identify the objects on a street. The hard part is anticipating what they're going to do next. We have developed about 80% of the technology needed to put self-driving cars into routine use. But the remaining 20%, including developing software that can reliably anticipate what other drivers, pedestrians and cyclists are going to do, will be much more difficult."

Bryan Salesky
Co-Founder & CEO of Argo
July 2019

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

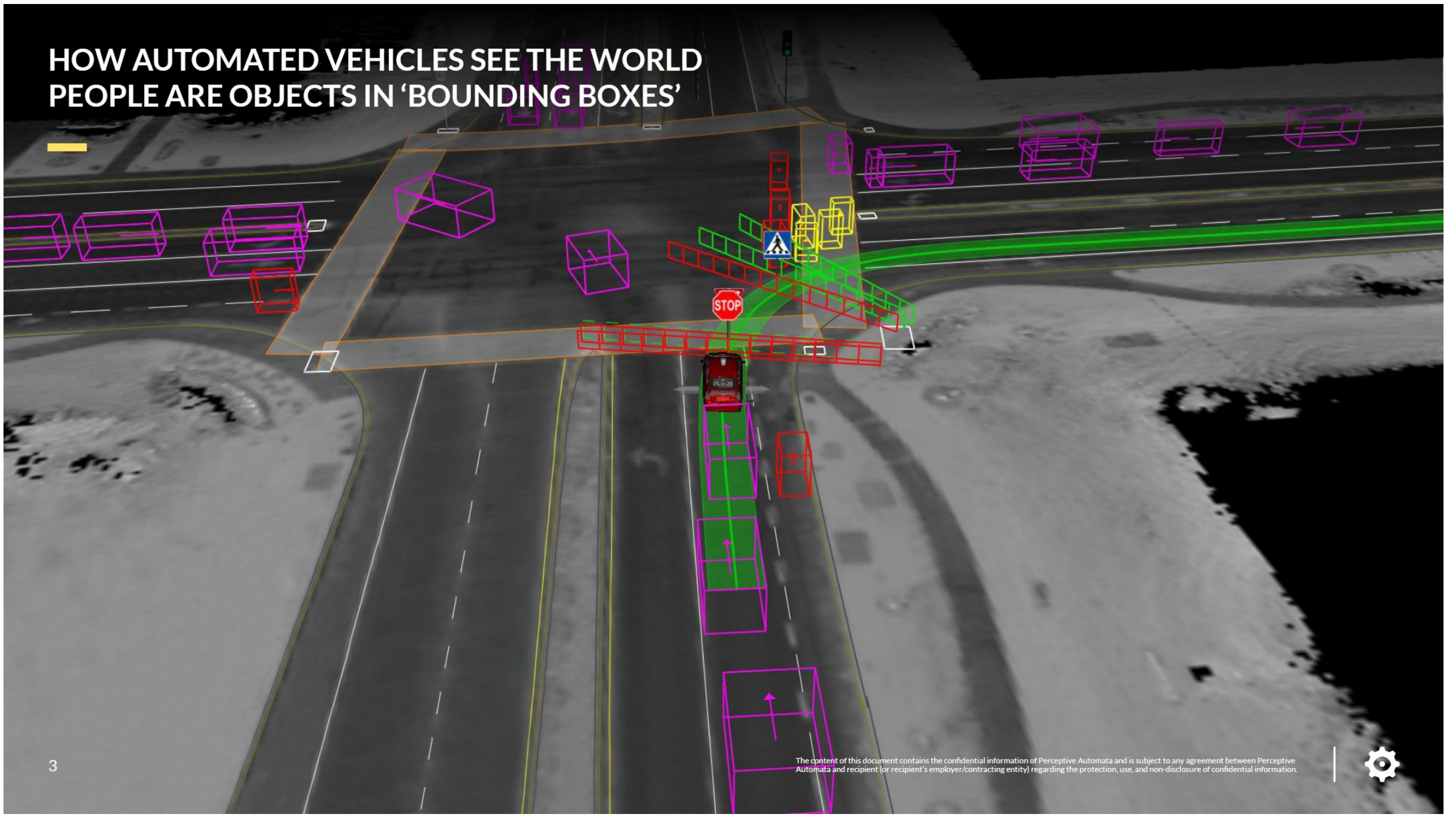
HOW AUTOMATED VEHICLES SEE THE WORLD
PEOPLE ARE OBJECTS IN 'BOUNDING BOXES'
STOP
3
The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.


BUT VIEWING PEOPLE SIMPLY AS BOXES LEADS TO “PARANOID DRIVING” BECAUSE CRITICAL INFORMATION IS MISSING - THE AV HAS NO CHOICE BUT TO STOP
The Physics-Only AV Perspective
4
The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

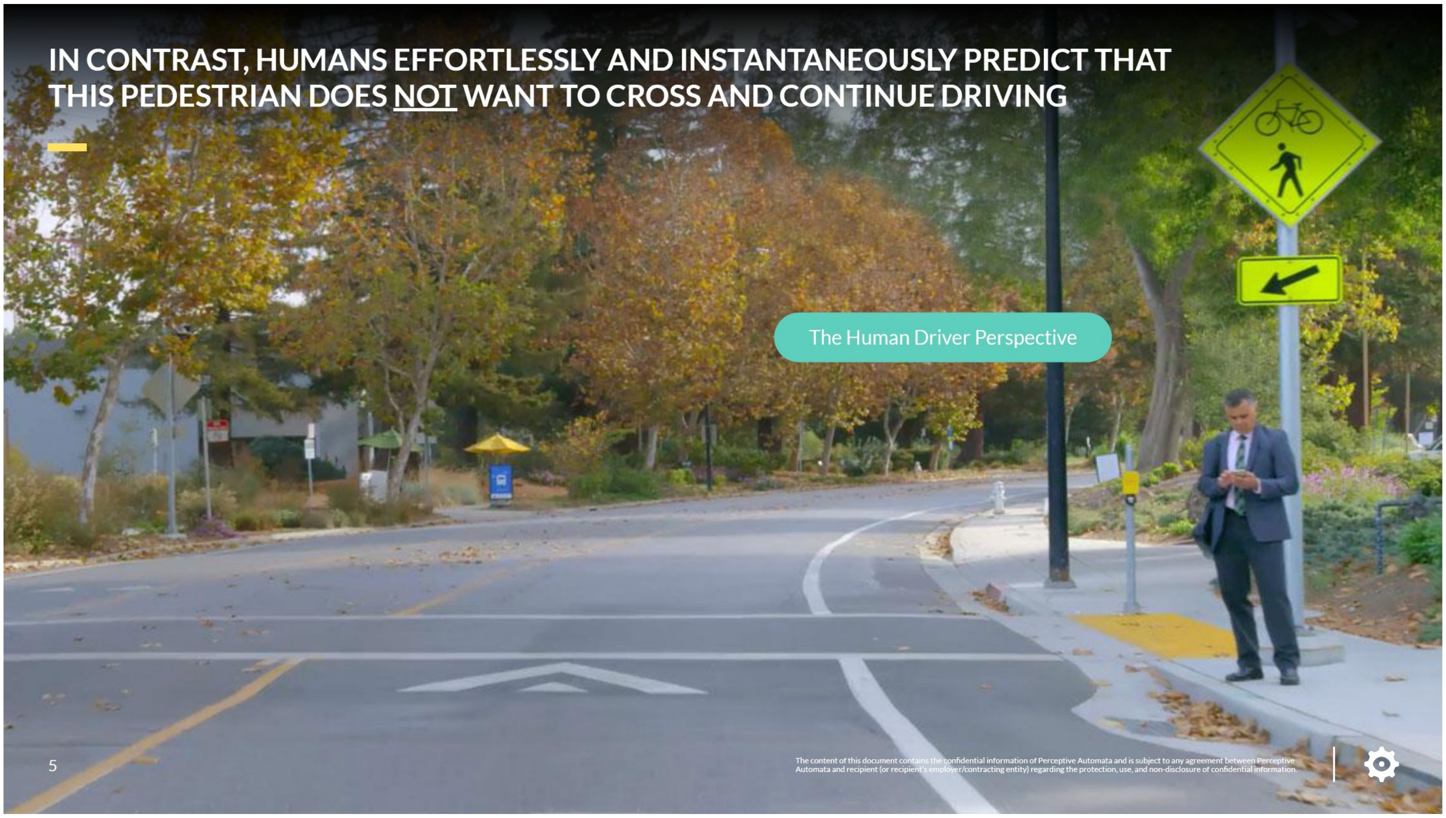
IN CONTRAST, HUMANS EFFORTLESSLY AND INSTANTANEOUSLY PREDICT THAT THIS PEDESTRIAN DOES NOT WANT TO CROSS AND CONTINUE DRIVING
The Human Driver Perspective
5
The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

OUR UNIQUE PATENTED METHODS HARNESS THE POWER OF THE HUMAN PERCEPTUAL SUPERCOMPUTER

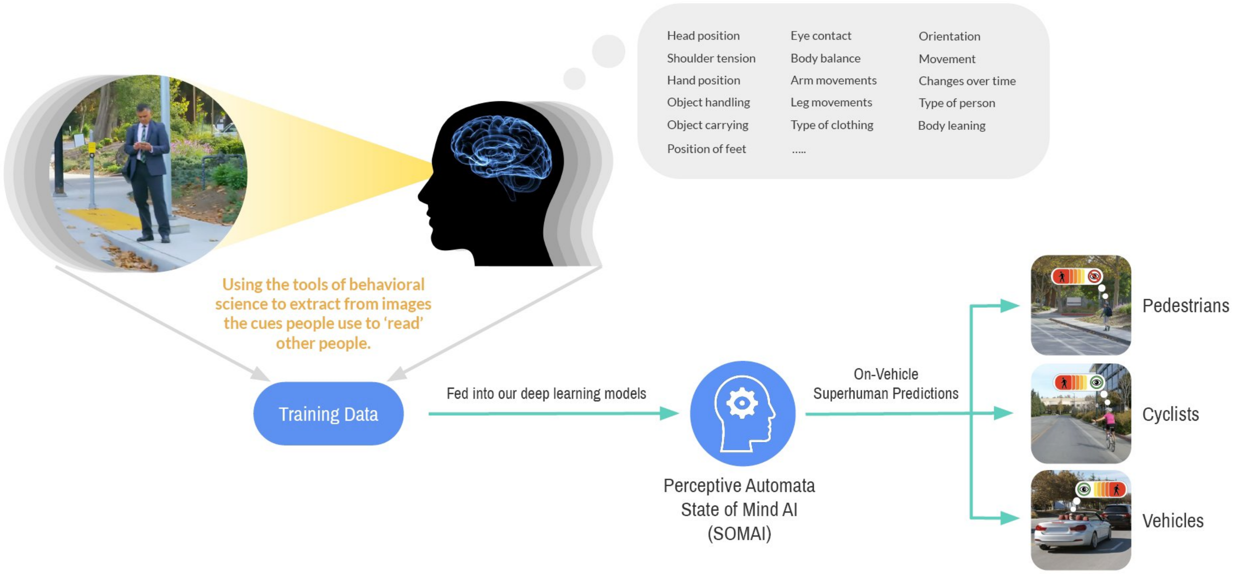

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

OUR LICENSABLE SOFTWARE PRODUCT IS THE SOMAI ("STATE OF MIND AI") SERVICE PACKAGE

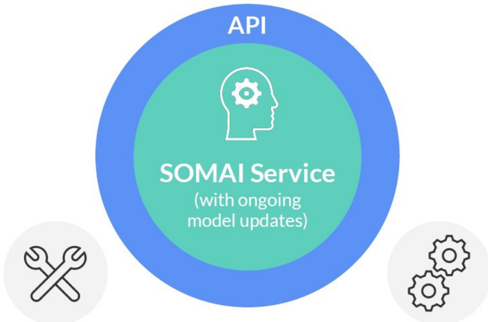

Supporting Tools including:

- State of Mind Analyzer
- Simulation Plug-In
- Visualization Library

Supporting Services including:

- Test Fleet Support
- Performance Analysis Support
- SOMAI Signal Usage Consulting

Pedestrians

Cyclists

Vehicles

Applications

- Robotaxis
- Last Mile Delivery
- Shuttle Services
- Automated Trucking
- Construction Equipment
- Warehouse Equipment
- Smart City
- ...

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

SIMPLE INTEGRATION WITH ON-VEHICLE AUTOMATION STACK WITHOUT REPLACING OR SLOWING DOWN CUSTOMER SUBSYSTEMS

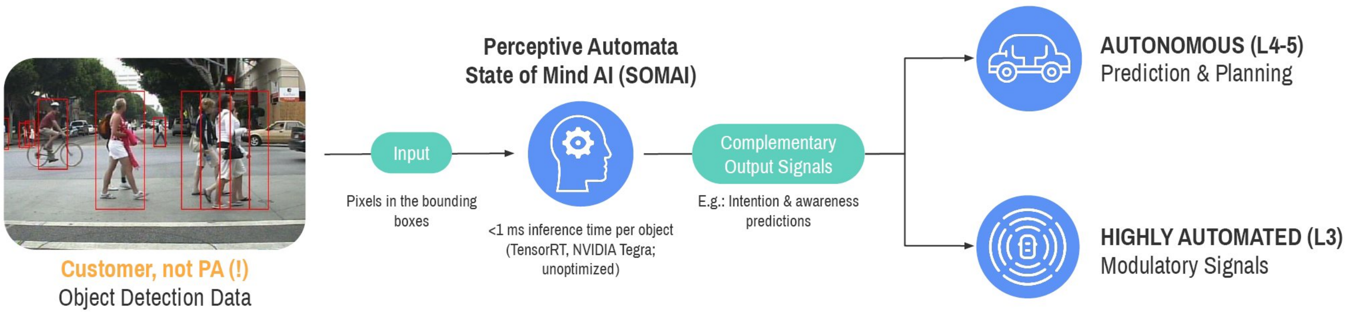

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

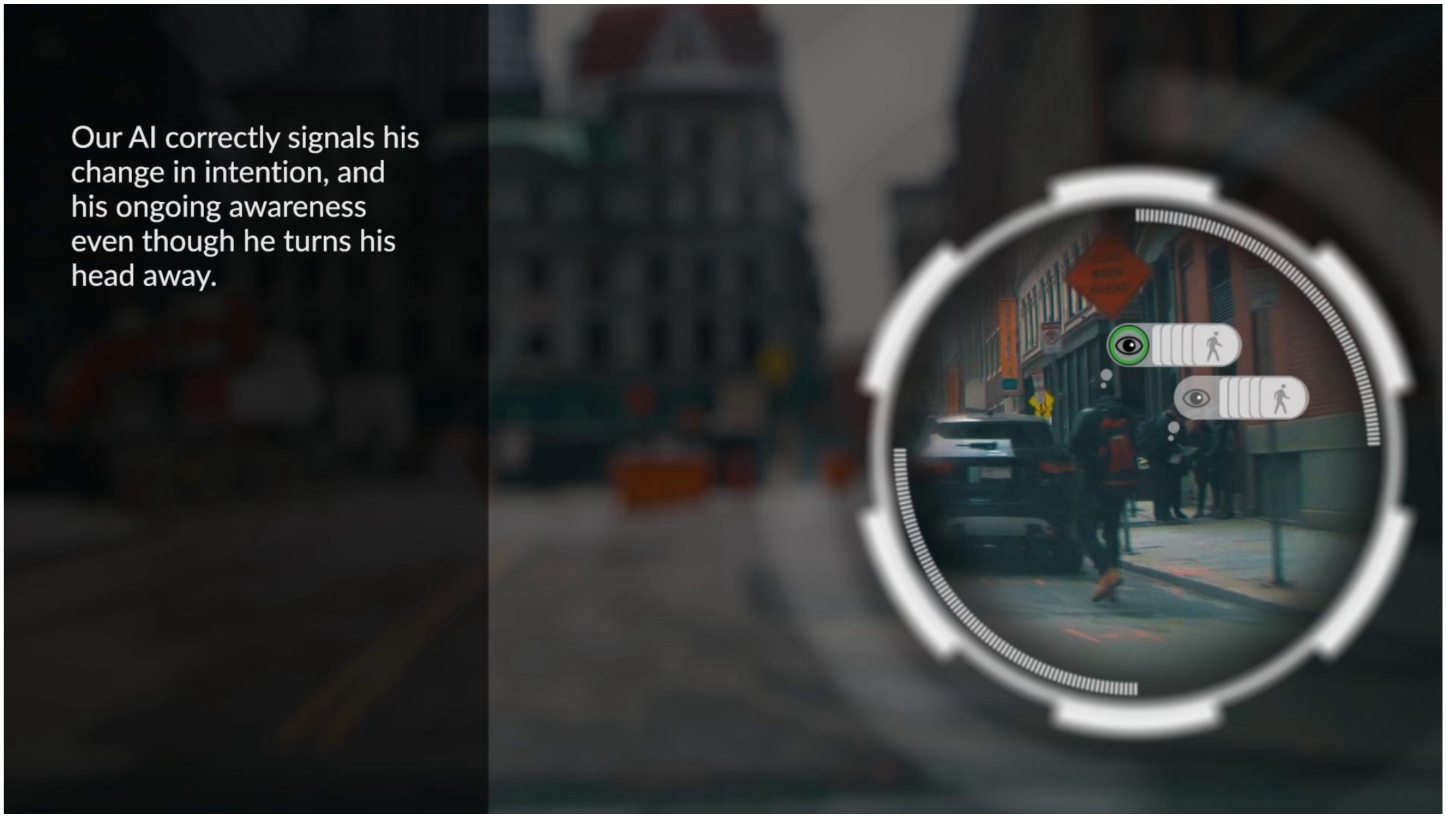
Our AI correctly signals his change in intention, and his ongoing awareness even though he turns his head away.

SOLVING FOR THIS CHALLENGE IS CRITICAL FOR REAL WORLD DEPLOYMENT

AUTONOMOUS - L4/5

- Enhanced Safety
  - detect and react sooner, before motion
  - predictable driving => reduce rear-endings
- Smoother Driving
  - passenger comfort
  - passenger trust
- Better Fleet/Network Economics
  - efficiency of vehicles completing trips
  - less vehicle wear and tear + fuel consumption

HIGHLY AUTOMATED - L3

- Enhanced Safety
  - increased situational awareness for driver
  - deploy in more complex road environments
- Driver Acceptance
  - reduce frustrating false positives
  - finer modulatory signals improve outcomes
- Enabling Transition from L3 to L4/5
  - focus on common platform
  - early enablement of L4/5 functionality

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

WAYMO: THE BLEEDING EDGE OF AI
IS THE ABILITY TO UNDERSTAND PEOPLE

“Anyone can buy a bunch of cameras and LIDAR sensors, slap them on a car, and call it autonomous. But training a self-driving car to behave like a human driver [...] is on the bleeding edge of artificial intelligence research. Waymo’s engineers are modeling not only how cars recognize objects in the road, for example, but how human behavior affects how cars should behave.”

The Verge, May 2018
Inside Waymo’s Strategy to Grow the Best Brains for Self-Driving Cars

“How can we make it adapt to the drivers that it’s sharing the road with? How do you tailor it to be more comfortable or drive more naturally? You really need a system that fricking works. They need reasoning to understand intentions of other drivers and pedestrians.”

Anca Dragan
Waymo, 2017

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient’s employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

GLOBAL LEADERSHIP POSITION FOR HUMAN BEHAVIOR PREDICTION

JAZZ VENTURE PARTNERS · SLOW · First Round · TOYOTA AI VENTURES · HYUNDAI CRADLE · PAVE · WORLD ECONOMIC FORUM · Gartner Cool Vendor 2018

$20M in capital (oversubscribed Series A)

Technical Founders: Neuroscience + Artificial Intelligence Experts

Sam Anthony, PhD
Co-Founder & CTO
20 years of cognitive science *and* computer science experience
HARVARD UNIVERSITY

David Cox, PhD
Co-Founder & Advisor
Director, MIT-IBM Watson AI Lab
Massachusetts Institute of Technology
IBM Watson

Walter Scheirer, PhD
Co-Founder & Advisor
Field-leading researcher in computer vision and deep learning
UNIVERSITY OF NOTRE DAME

SLATE · WIRED · QUARTZ · Bloomberg · NVIDIA · The Boston Globe · DAIMLER · FAST COMPANY

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

WE ACCELERATE OUR CUSTOMERS' TIME TO MARKET AND PROVIDE VALUABLE OPTIONALITY WITH A PROVEN PATENTED SOLUTION AT DRAMATICALLY LOWER COST

Faster Time
To Market

Certainty With
Proven Solution

Access to Unique
Talent & Protected IP

Dramatically
Lower Cost

Optionality &
Redundancy

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

Thank you!

James Gowers
VP Strategy & Business Development
james@perceptiveautomata.com

PERCEPTIVE AUTOMATA

Human Intuition for Machines

Customer Use Case Manual

For SOMAI Deployment

Table of Contents

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

COMPANY OVERVIEW

Perceptive Automata was founded by a team of Harvard and MIT neuroscientists, computer vision researchers, machine learning experts, and software engineers to solve what is often described as the hardest problem for highly automated driving: reading the state of mind of humans for the safe large-scale rollout of highly automated vehicles (L2-L5).

ANTICIPATING HUMAN BEHAVIOR
PROBABLY THE HARDEST PROBLEM FOR AUTOMATED DRIVING

> “It's the prediction piece that's still the great unknown. Humans are very good at predicting human behavior on the road. Machines will need to be able to predict and anticipate human behavior much better.”
>
> **Gill Pratt**
> CEO of Toyota Research Institute
> March 2018

> “The choices made by driverless cars are critically dependent on understanding and matching the expectations of humans. This is the heart of the problem going forward.”
>
> **Chris Urmson**
> Head of Waymo, 2009-2016
> Co-Founder of Aurora, April 2017

We enable automated vehicles to better understand what people might do next so they can navigate safely around humans, including pedestrians, bicyclists, and motorists. Our AI processes, in real-time, our customers' object data and outputs human state of mind attributes, such as 'intention' and 'awareness' as defined in the below graphic, to our customers' AV path planner or ADAS system.

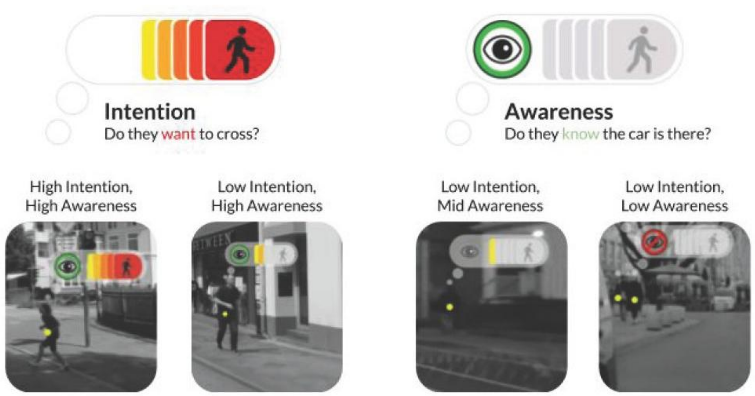

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

PEDESTRIAN, BICYCLIST, AND VEHICLE MODELS AND LONG-TAIL TOO

Pedestrians

Cyclists

Drivers

Our signals increase safety and roadmanship by enabling AVs to anticipate human behavior before actual body motion is detected. Our signals modulate otherwise paranoid automated driving behavior to be more predictable and natural, which reduces rear-endings and enables a smooth ride experience. For ADAS applications our signals give human drivers better situational awareness and increase driver acceptance of more advanced driver assistance functions.

These benefits are critical for increasingly sophisticated driver assistance functions to be embraced by human drivers - too many false positives lead to human drivers turning their vehicles' driver assistance systems off. They also critically enable autonomous vehicles to seamlessly integrate into human-dominated road environments and to deliver a pleasant ride experience for passengers.

OUR VALUE PROPOSITION PRODUCT

AUTONOMOUS - L4/5

- Safety
  - detect and react sooner, before motion
  - predictable driving => reduce rear-endings
- Smooth Driving
  - passenger comfort
  - passenger trust
- Fleet/Network Economics
  - efficiency of vehicles completing trips
  - less vehicle wear and tear + fuel consumption

HIGHLY AUTOMATED - L3

- Safety
  - increased situational awareness for driver
  - deploy in more complex road environments
- Driver Acceptance
  - reduce frustrating false positives
  - finer modulatory signals improve outcomes
- Enabling Transition from L3 to L4/5
  - focus on common platform
  - early enablement of L4/5 functionality

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

We raised $20M from top-tier investors and are working out of offices in Boston and Silicon Valley with OEMs, suppliers, and technology companies that are developing ADAS and/or autonomous driving systems.

GLOBAL LEADERSHIP POSITION FOR HUMAN BEHAVIOR PREDICTION

The value we deliver to our customers through our software licensing business model includes fewer years of technology development for faster market deployment, tens of millions of dollars in development cost savings, and access to scarce talent and intellectual property.

OUR VALUE PROPOSITION BUSINESS STRATEGY

THE CHALLENGE

People make effortless, intuitive judgments about what someone else wants or knows constantly. In any given thirty second interval at a busy urban intersection, there will be dozens of instances where one person looks at another and thinks, for example, “she isn’t going to cross,” or “he knows I’m here and is willing to yield,” or “she doesn’t see me, I should stop.”

Humans are incredibly good at silently communicating with each other. That communication is the key to safe, smooth and considerate driving. Until now machines have lacked this critical ability. They can't decipher our unspoken social communications nor intuit what's going on inside our heads. So how can we give machines like self-driving cars the ability to read human intentions, to know what humans know?

The best currently deployed approaches to interacting with the world using computer vision and motion planning are insufficient for the task of interacting with humans. Existing non-PA approaches use physics to locate people and to calculate their trajectory, but can't usefully anticipate what a pedestrian might do next. People aren't billiard balls — we aren't predictable based purely on our past motion.

Let's take the below example to illustrate this point. Many of today's approaches convert detected humans into 'bounding boxes', shown below as a yellow box in the image on the left, and then attempt to predict their likely actions based on calculated trajectory and context.

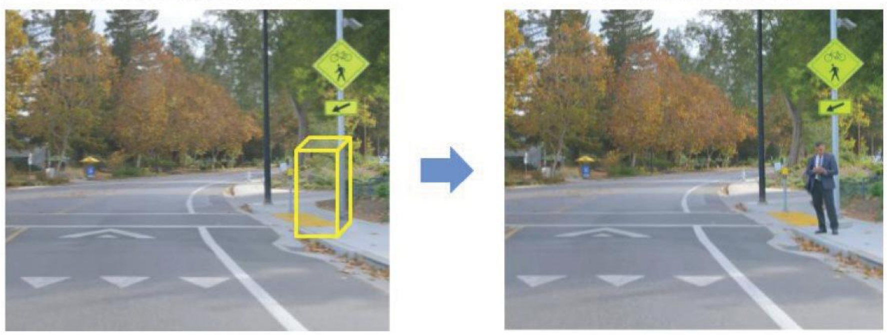

The Physics-Only AV Perspective   The Human Driver Perspective

In this example on the left, a pedestrian is detected at the edge of the road at/near a crosswalk with zero motion. The only safe decision for the AV is to stop in case the pedestrian wants to cross the road. This means that instead of smoothly responding to a situation the way a human driver would, self-driving cars and trucks act alternately 'paranoid' — timid, skittish, and easily startled — and oblivious.

Next, let's take a closer look at the image on the right showing what a human driver sees, and what a self-driving car *with* Perceptive Automata software can see too. This image on the right clearly includes much more decision-making information - not just location and trajectory shown in the left image, but also the full range of visual information needed to determine that this pedestrian, in all likelihood, has no intention to cross despite standing at/near a crosswalk.

The physical and social cues humans effortlessly pick up on include that he is leaning back slightly, carrying himself without much tension, attending to his smartphone, and showing no indication in his facial expression that he's trying to make a crossing decision. So if we polled, say, 100 human

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

drivers approaching this scene, the likely prediction, in contrast to the physics-only self-driving car in the left image, would be that this person has no intention to cross in front of the car.

A typical human driver would therefore *not* stop in front of the pedestrian, and, just as important, in the context of rear-endings, other human drivers behind the lead car also would not expect the lead car to stop. Self-driving cars with Perceptive Automata's software would come to precisely the same conclusion as those 100 human drivers and also not stop.

This is just one example, but existing self-driving vehicles frequently slow down or come to a complete stop in the presence of pedestrians and/or cyclists when there is no need to do so, only to start moving again at exactly the wrong time. This style of driving is nauseating for passengers and infuriating for other road users. It's the reason self-driving cars are often rear-ended and it can cause human drivers to behave erratically.

It is now widely understood that there will be no meaningful real-world deployment of self-driving vehicles without solving the problem of human understanding.

THE PERCEPTIVE AUTOMATA APPROACH

In an effort to solve this problem we've designed a State of Mind AI (SOMAI) that can use the whole spectrum of subtle, unconscious insights that we, as humans, use to make incredibly sophisticated judgments about what's going on in someone else's head. You could say that, in a sense, our models develop their own human-like 'intuition.'

We take monocular camera recordings from vehicles that show interactions with people (pedestrians and bicyclists today, other motorists in the near future). This pixel data is incredibly rich raw material; every time a person appears and interacts with a vehicle, they're giving off hundreds of signals that we as humans can innately use to understand their state of mind.

This method was developed by two of our co-founders, Sam Anthony and Walter Scheirer, over the course of a decade of research at Harvard. The fundamental method relies on the realization that supervised learning — the method by which machine learning models are trained using labeled data — relies in almost all cases on data that was labeled by humans. Human labels are generated via a process that involves showing a training sample to a human observer and asking them to record their response. This type of paradigm is per se a human behavioral experiment. In most cases the psychological and neuroscientific aspects of this labeling are afterthoughts. However, for Perceptive Automata they comprise the core of our models' inferential power.

USING SOMAI OUTPUTS

SOMAI modulatory signals can be injected at a number of points in the perception/action loop in order to improve use case performance. Below we describe some of the potential integration approaches based on what we've found to be helpful for customers in the past. Please note that

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

this should not be considered a comprehensive list. In practice, there are many more integration paths, and the paths listed below will generally require a degree of customization.

ADAS FALSE POSITIVE REDUCTION

In ADAS systems where the driver is in full control, SOMAI's modulatory signal can be used to reduce false positives. For instance, in the case of an automatic emergency braking (AEB) system, the activation threshold for the system can be modulated by a weighted linear combination of the SOMAI awareness and intention model outputs. The weightings and the amount of threshold modulation can be set by the customer.

In practice, what this means is that the activation threshold for the emergency braking system -- the certainty of evidence required before brakes are applied -- would be higher in cases where a pedestrian in the roadway did not intend to cross, and was aware of the car. When tuned correctly, this can effectively remove false positive emergency braking scenarios, for example when a pedestrian enters the roadway for reasons other than intending to cross the road.

An additional tuning parameter that's available in this case comes from the distributional nature of SOMAI's model outputs. What our models generate is not a scalar value for intention or awareness, but a prediction of the distributional characteristics of human responses to the presented human road user. This means that there is information not just about the central tendency -- the equivalent of the scalar mentioned above -- but also the variance, kurtosis, skew, heteroskedasticity, and even multimodality of the predicted human distribution.

In the false positive reduction case and many of the other cases described here, a useful tuning parameter is to use the kurtosis of the output distribution to threshold the model input to the AEB system. In cases where the kurtosis is very low, and the distribution approaches uniform, it is reasonable to assume that a measurement of central tendency would not be useful for prediction.

On the other hand, in cases with high kurtosis the central tendency is likely to represent substantial human agreement that the output intention or awareness value accurately reflects the internal mental state of the detected human road user.

KALMAN FILTER MODULATION

The outputs of the SOMAI intention model can be used directly to modulate the estimates of a Kalman filter. The prior state used in the Kalman filter estimation can be weighted by the current intention level. This will have the effect of 'tuning' the filter.
In the high intention case, in practical terms this means the filter will attend preferentially to motion (even if that motion is below the nominal noise floor of the motion sensor) that could potentially cross the ego vehicle path in cases where the pedestrian looks like they want to cross the street.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

In the low intention case, the filter will preferentially discount as noise motion readings even though those readings might be slightly above the nominal noise floor for the sensor.

WORLD MODEL MODULATION

In a lattice planner-based autonomous driving stack, it is essential to have a world model. A world model in autonomous driving can consist of multiple layers of perceptual inference about the state of the road around the vehicle. SOMAI model outputs comprise an additional layer of information that modulates the existing world model to make it more accurately represent the probabilistic likelihood of safe driving across the modeled grid.

The basic informational layer of a world model is an occupancy grid. An occupancy grid comprises a binary random variable for each point in a discretized topological rendering of the autonomous robot's local environment. The posterior estimate of occupancy at a given time step can be constructed with a bayesian estimator that takes into account not just the outputs of depth sensitive sensors like LIDAR, but also semantically-aware estimates of minimum safe distance around moving agents (e.g., the maximum distance a human agent on foot can be expected to be able to move).

When a vehicle is simulating proposed motion paths, the semantically-aware minimum safe distance estimates for moving agents can be biased by the SOMAI outputs. So, for instance, for a pedestrian with high intention, the minimum safe distance region would be stretched towards the current vehicle path and squeezed in other directions, preserving total area.

POMDP PRIOR BIASING

In most probabilistic autonomous robot planning systems, path planning is achieved via a partially observable Markov decision process. In the world model example above, a POMDP would be efficiently approximated to pick the optimal route. The estimation at each time step relies on an observation vector which is assumed to comprise a functionally sufficient set of information to estimate the underlying state of the world.

The SOMAI models output an observation that correlates with a hidden portion of the world-state, the internal goal and belief structure of human agents. Because of this, the SOMAI model output can be directly added to the observation vector of a POMDP.

When the POMDP is trained including the SOMAI state of mind vector in its observation vector, the motion paths selected by the estimator will be sensitive to the ‘best guess’ information about the state of mind of human road users, without further system redesign being necessary.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

FLEXIBILITY, EASE OF INTEGRATION, SPEED, AND SAFETY

SOMAI predictions are integrated with customer systems as a *subsystem* that mostly runs in a separate process. The subsystem is comprised of:

- A subsystem 'service' process that accepts inputs from the customer/system process and returns state of mind predictions and other optional outputs steps (e.g. visualizations);
- A lightweight 'interface' library that is compiled into the partner system and provides a simplified API that can be used to interact with the service.

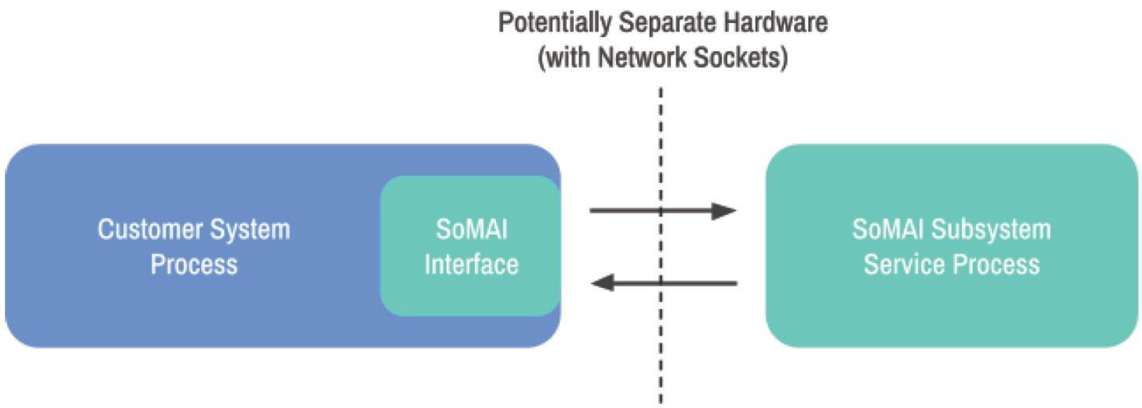

The interface library is responsible for launching the service and managing inter-process communication with the service. It is packaged as a C++ static library with a simplified API that is tailored to the specific needs of each customer.

This results in a limited API that is (a) easier to use and understand than a generic API, and (b) minimizes the footprint of Perceptive Automata code in the customer's core system processes, which can be significant from a functional safety standpoint.

The service process performs all the heavy lifting of the SOMAI subsystem - generating state of mind predictions and optionally performing a variety of other functions. As described in the following subsections, the subsystem service is designed to be ***flexible***, ***fast***, ***safe***, and ***upgradeable***.

FLEXIBILITY

With a number of optional calculations and a wide variety of alternatives for how data is transferred to and from the process, the SOMAI subsystem is designed with an emphasis on flexibility. The service can accept inputs from any number of sources, including sockets (network or UNIX), files, ROS, or other custom IO modules. Similarly, outputs can be published to one or more of a variety of destinations.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

SIMPLICITY - EASE OF INTEGRATION

This flexibility is typically hidden behind the simplified customer-tailored interface (so as not to introduce complexity for customers) but can be leveraged to quickly accommodate specific customer needs. Additionally, the specific calculations performed by the service (and therefore the inputs and outputs) can be easily controlled to suit different customer needs.

For example during evaluations and initial testing:

- Object detection can either be provided by the customer system (as inputs) or computed by the SOMAI subsystem (and optionally returned as outputs);
- Visualizations of the road user state of mind, which are often useful for evaluation, can optionally be generated by the subsystem.

SPEED

The SOMAI subsystem is fast. Forward inference with an *unoptimized* model takes <0.2 ms on a middle-of-the-road GPU (GTX 1060). Overall prediction by the subsystem takes a bit longer (due to data transfer, etc.) but is still about ten times faster than real-time for as many as 40 simultaneous pedestrians or cyclists.

The calculation is so light-weight that we anticipate real-time performance even without a GPU. By applying quantization and weight pruning optimizations, real-time performance on an ARM-based ECU should be well within reach.

Note that object detection (if requested) slows down the prediction time, but still easily achieves 30 frames per second with a GTX 1060.

SAFETY

Conscious of the functional safety concerns of our partners, we've developed the SOMAI subsystem with functional safety in mind from the start. For example:

- Packaging the SOMAI subsystem as a separate process allows the bulk of Perceptive Automata's code to be segregated from the core system process, minimizing the risk of subsystem problems affecting other parts of the system.

  PA code can even be kept entirely out of the customer process if need be - for example, by communicating inputs and outputs via ROS or another message passing library.

- We employ an increasing number of 'traditional' functional safety safeguards (including many of those described in ISO 26262) during the development of the subsystem.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

- Beyond ISO 26262, we've begun applying SOTIF techniques (including those described by ISO/PAS 21448) to our deep learning models and to the overall subsystem.

UPGRADEABILITY

Finally, the SOMAI subsystem architecture enables partner systems to benefit from ongoing improvements via over-the-air updates. The separation of different pieces of the system into separate files (and processes) supports updates of the subsystem executable and/or deep learning models on the fly, with no other updates to the system required.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

PERCEPTIVE AUTOMATA | Human Intuition for Machines

SOMAI Customer Service Release & Download Page

Dear Customer,

The following link is for the SOMAI Service package that includes support for Pascal, Turing, and Volta GPU architectures on Ubuntu 18.04.

SOMAI API Package v02.03

The configuration requirements are as follows:

The SOMAI Service needs to know

1. the frame size in order to properly initialize the detector at runtime;
2. the frame rate to optimize tracking and smoothing;
3. which of the supported GPU architectures it will be run on.

If you launch the Service manually using the launch script, you can set environment variables, 'SOMAI_FRAME_WIDTH', 'SOMAI_FRAME_HEIGHT', 'SOMAI_FRAME_RATE', and `GPU_ARCH` which will be picked up by the 'config.sh' script and used to configure the service. For example:

```
SOMAI_FRAME_WIDTH=3840 SOMAI_FRAME_HEIGHT=1920 SOMAI_FRAME_RATE=10.0
GPU_ARCH=pascal ./run_service.sh <args>
```

The values shown above represent the default configuration, which means that the environment variables will default to these values if left unset. To bypass setting these environment variables at the command line, you can also edit the 'config.sh' script directly to set the environment variables to the desired values.

Alternatively, if you launch the Service with the Service Daemon, you can set the 'frameWidth', 'frameHeight', 'frameRate', and 'gpuArch' parameters, which are members of the 'sStateOfMindAIConfig' struct. Then, when the Service is launched by the Daemon, it will be configured to use the appropriate frame parameters without having to use the environment variables.

Please note that due to the requirements imposed by the Driveworks 2.0 detection and tracking modules, this package has been built for Ubuntu 18.04 and includes many additional libraries in its 'lib' folder which may be removed in future releases.

Please submit any support requests to support@perceptiveautomata.com.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

PERCEPTIVE AUTOMATA

SOMA | State of Mind Analyzer

User Guide

Confidential

Table of Contents

Introduction

The ***State of Mind Analyzer*** (***SOMA***) is a tool that enables in-depth analysis of the outputs of Perceptive Automata's ***State of Mind AI*** models.

The tool's primary aim is threefold:

1. Provide a streamlined, responsive interface for analyzing ***State of Mind AI*** model outputs.
2. Render intuitive visualizations of those outputs to make them easy to understand at a glance.
3. Facilitate deeper, fine-grained analysis by enabling frame-by-frame inspection into each output.

In addition, SOMA provides controls for quickly viewing all the sequences within a video and selecting one to focus on, graphing the awareness and intention predictions of a cyclist or pedestrian over the course of an entire sequence, and recording an interesting sequence with accompanying visualizations to a separate video file.

Getting Started

System Requirements

- Ubuntu Linux 16.04 LTS
- NVIDIA CUDA 9.0
- FFmpeg (optional for recording video)

Quick Start

1. Untar the package file into target directory:
   cp soma-<version>.tar.gz /path/to/target/directory
   cd /path/to/target/directory
   tar xzvf soma-<version>.tar.gz
2. Install FFmpeg (optional):
   sudo apt-get update
   sudo apt-get dist-upgrade
   sudo apt-get install ffmpeg
3. Launch SOMA:
   cd soma-<version>
   ./soma
4. Load a video using the ***Open video*** toolbar button. This action is explained in greater detail in the next section.


PERCEPTIVE
AUTOMATA
Human Intuition for Machines
Overview for
Customers

ANTICIPATING HUMAN BEHAVIOR IS ONE OF THE THE HARDEST PROBLEMS FOR AUTOMATED DRIVING

"It's the prediction piece that's still the great unknown. Humans are very good at predicting human behavior on the road. Machines will need to be able to predict and anticipate human behavior much better."

Gill Pratt
CEO of Toyota Research Institute
March 2018

"The choices made by driverless cars are critically dependent on understanding and matching the expectations of humans. This is the heart of the problem going forward."

Chris Urmson
Head of Waymo, 2009-2016
Co-Founder of Aurora, April 2017

"With radar and high-resolution cameras and all the computing power we have, we can detect and identify the objects on a street. The hard part is anticipating what they're going to do next. We have developed about 80% of the technology needed to put self-driving cars into routine use. But the remaining 20%, including developing software that can reliably anticipate what other drivers, pedestrians and cyclists are going to do, will be much more difficult."

Bryan Salesky
Co-Founder & CEO of Argo
July 2019

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

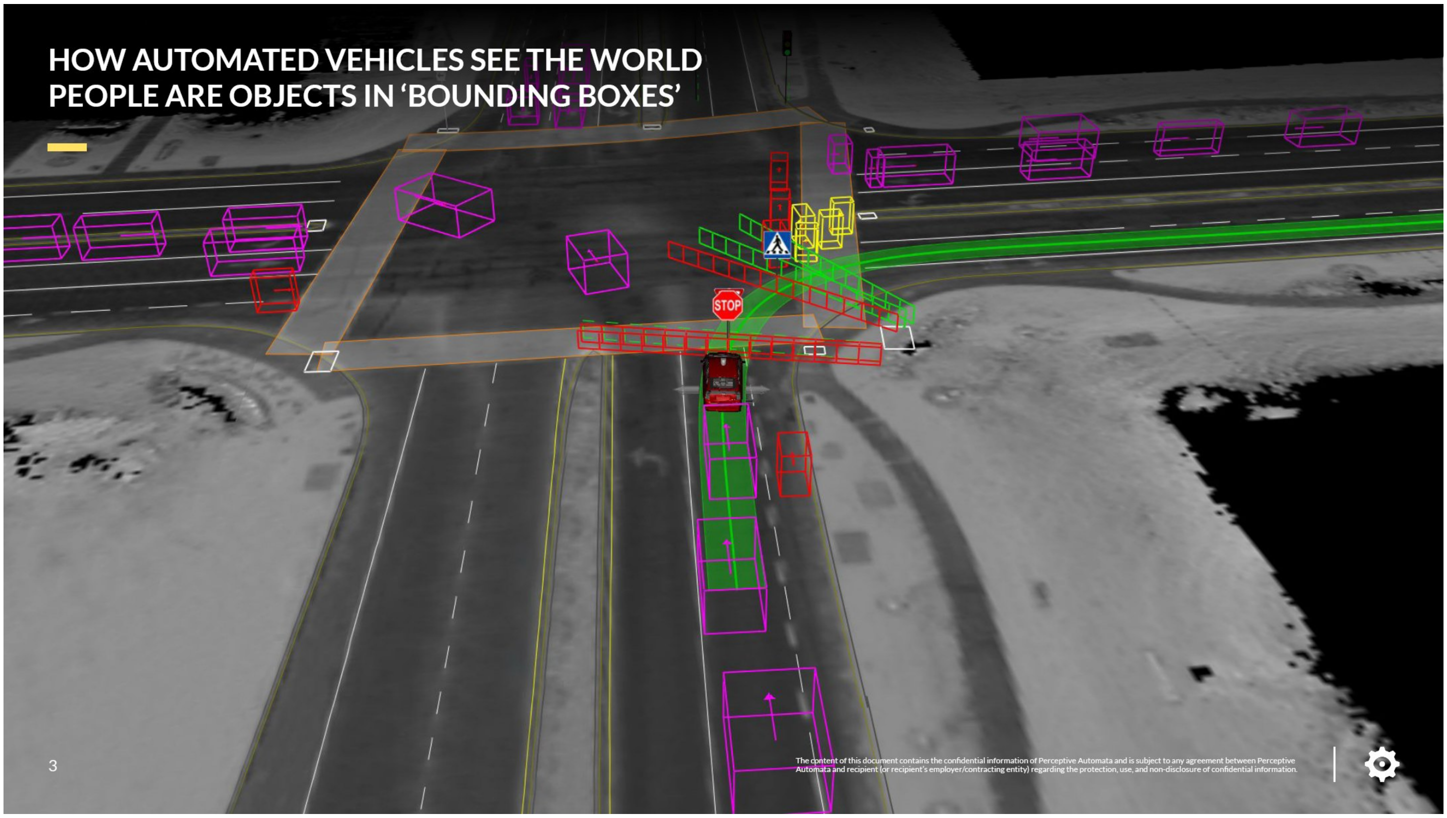
HOW AUTOMATED VEHICLES SEE THE WORLD
PEOPLE ARE OBJECTS IN 'BOUNDING BOXES'
STOP
3
The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.


BUT VIEWING PEOPLE SIMPLY AS BOXES LEADS TO “PARANOID DRIVING” BECAUSE CRITICAL INFORMATION IS MISSING - THE AV HAS NO CHOICE BUT TO STOP
The Physics-Only AV Perspective
4
The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient’s employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

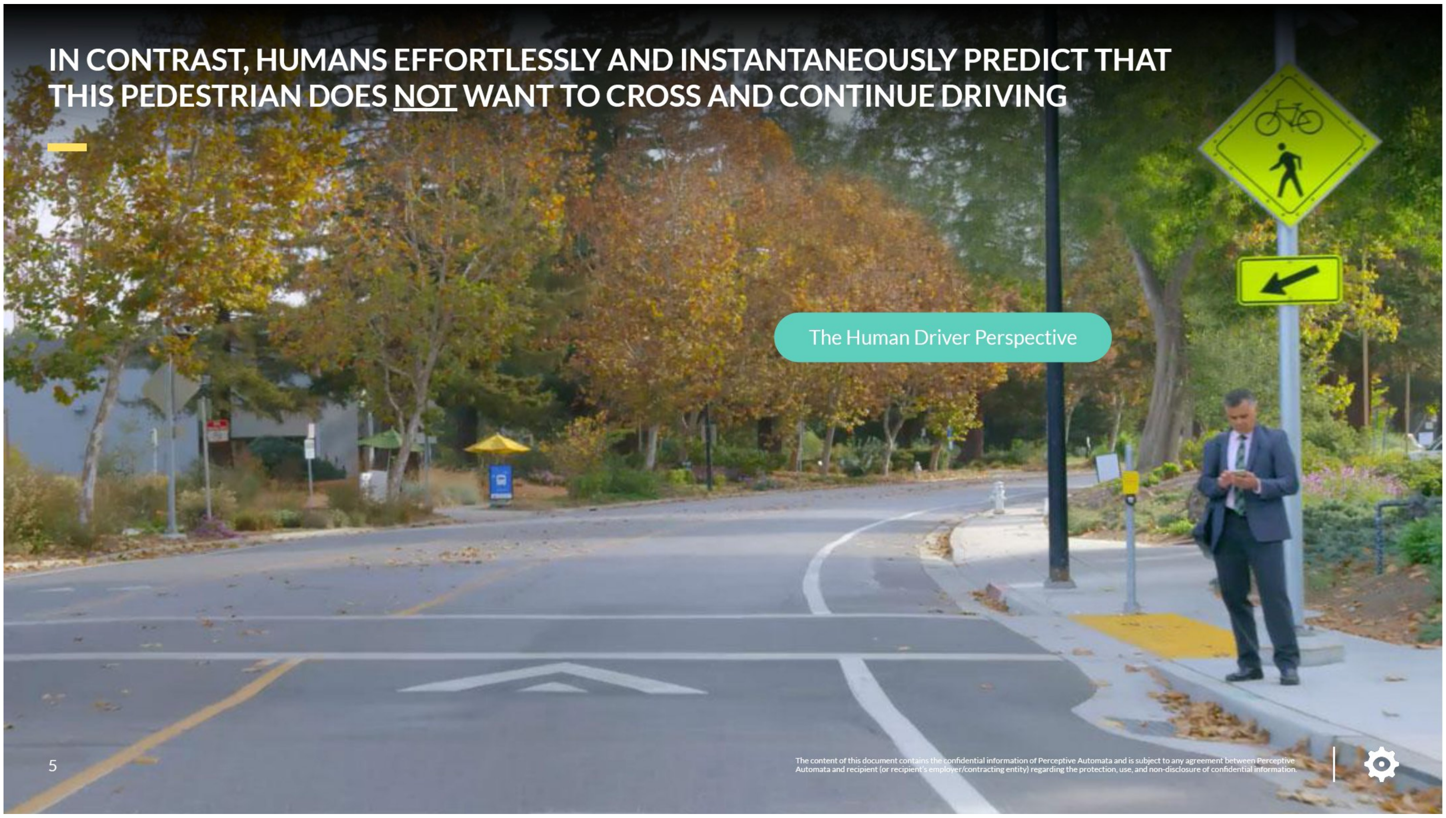
IN CONTRAST, HUMANS EFFORTLESSLY AND INSTANTANEOUSLY PREDICT THAT THIS PEDESTRIAN DOES NOT WANT TO CROSS AND CONTINUE DRIVING
The Human Driver Perspective
5
The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

OUR UNIQUE PATENTED METHODS HARNESS THE POWER OF THE HUMAN PERCEPTUAL SUPERCOMPUTER

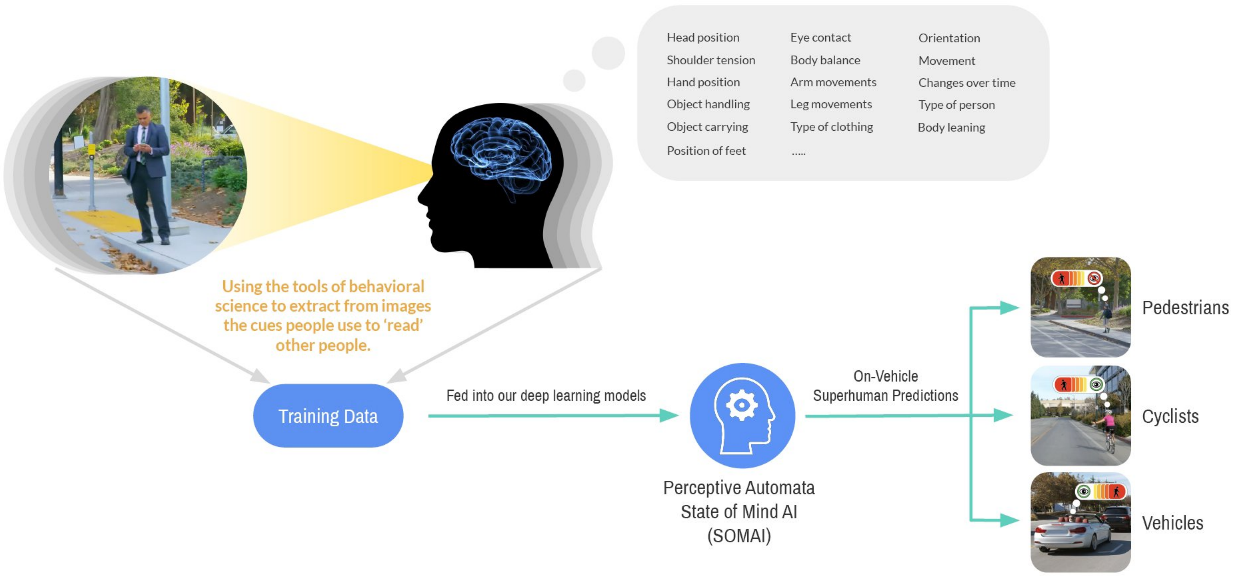

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

OUR LICENSABLE SOFTWARE PRODUCT IS THE SOMAI ("STATE OF MIND AI") SERVICE PACKAGE

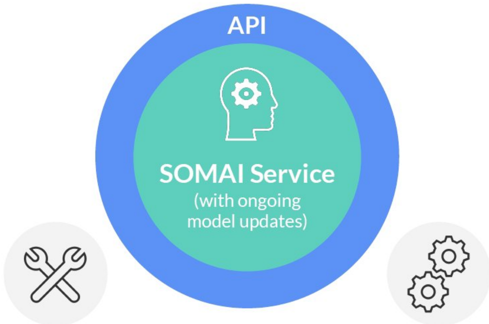

Supporting Tools including:

- State of Mind Analyzer
- Simulation Plug-In
- Visualization Library

Supporting Services including:

- Test Fleet Support
- Performance Analysis Support
- SOMAI Signal Usage Consulting

Pedestrians

Cyclists

Vehicles

Applications

- Robotaxis
- Last Mile Delivery
- Shuttle Services
- Automated Trucking
- Construction Equipment
- Warehouse Equipment
- Smart City
- ...

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

SIMPLE INTEGRATION WITH ON-VEHICLE AUTOMATION STACK WITHOUT REPLACING OR SLOWING DOWN CUSTOMER SUBSYSTEMS

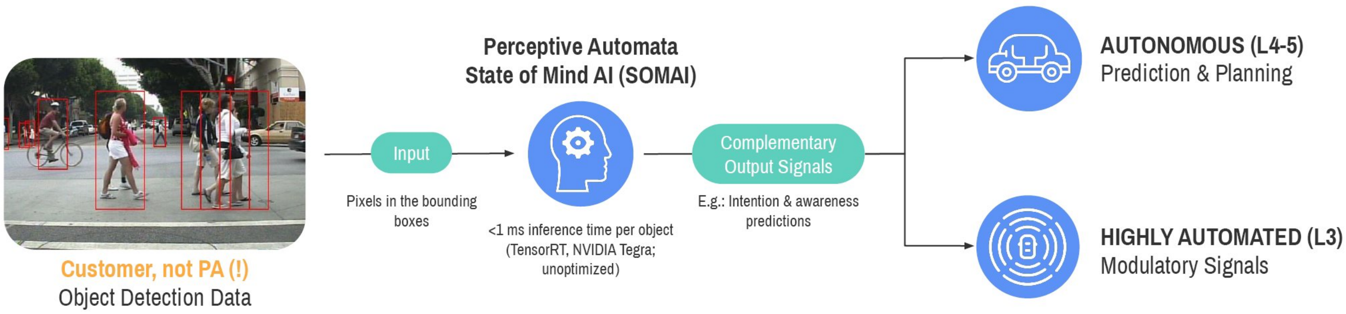

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

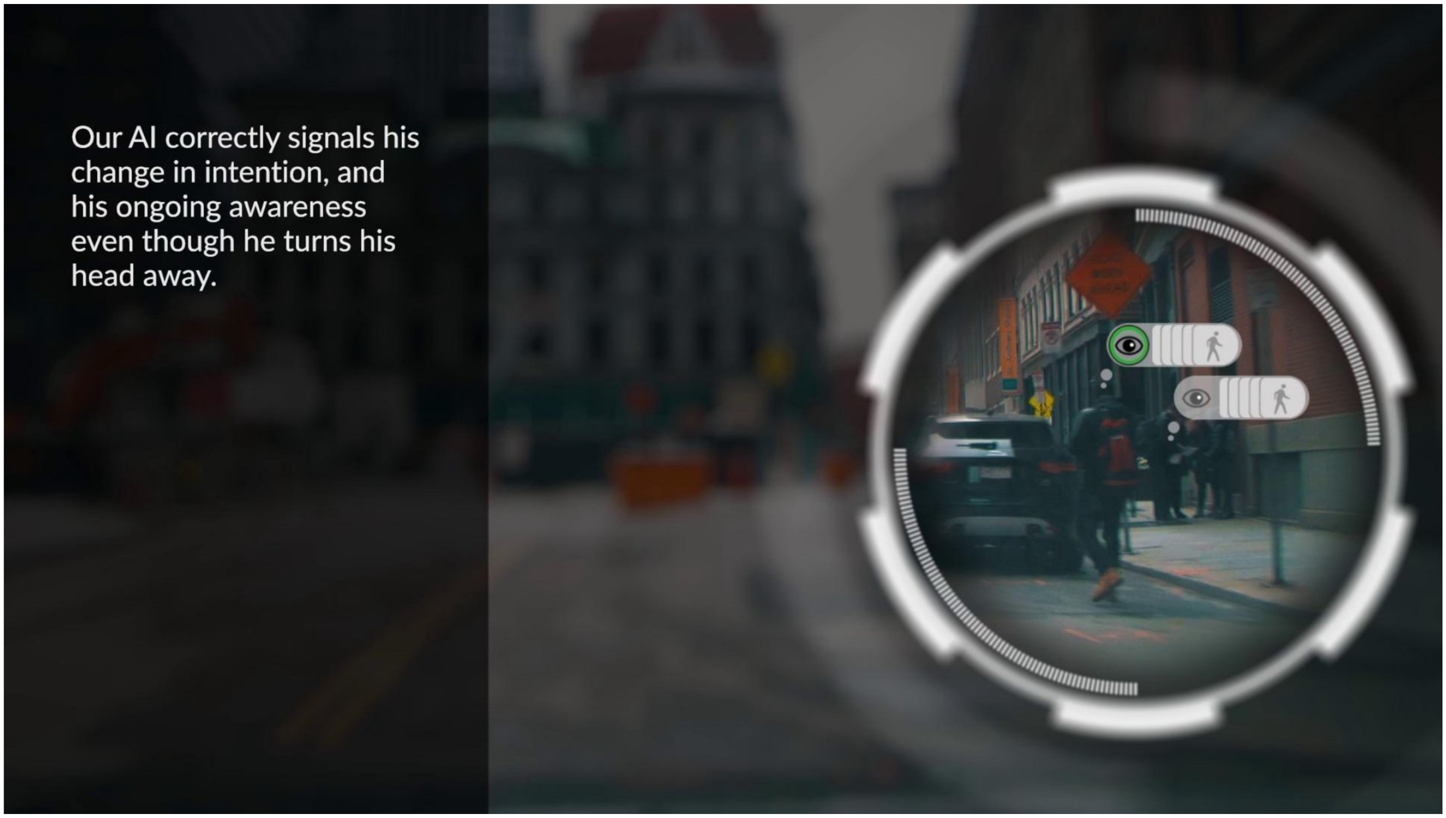
Our AI correctly signals his change in intention, and his ongoing awareness even though he turns his head away.

SOLVING FOR THIS CHALLENGE IS CRITICAL FOR REAL WORLD DEPLOYMENT

AUTONOMOUS - L4/5

- Enhanced Safety
  - detect and react sooner, before motion
  - predictable driving => reduce rear-endings
- Smoother Driving
  - passenger comfort
  - passenger trust
- Better Fleet/Network Economics
  - efficiency of vehicles completing trips
  - less vehicle wear and tear + fuel consumption

HIGHLY AUTOMATED - L3

- Enhanced Safety
  - increased situational awareness for driver
  - deploy in more complex road environments
- Driver Acceptance
  - reduce frustrating false positives
  - finer modulatory signals improve outcomes
- Enabling Transition from L3 to L4/5
  - focus on common platform
  - early enablement of L4/5 functionality

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

WAYMO: THE BLEEDING EDGE OF AI IS THE ABILITY TO UNDERSTAND PEOPLE

“Anyone can buy a bunch of cameras and LIDAR sensors, slap them on a car, and call it autonomous. But training a self-driving car to behave like a human driver [...] is on the bleeding edge of artificial intelligence research. Waymo’s engineers are modeling not only how cars recognize objects in the road, for example, but how human behavior affects how cars should behave.”

The Verge, May 2018
Inside Waymo’s Strategy to Grow the Best Brains for Self-Driving Cars

“How can we make it adapt to the drivers that it’s sharing the road with? How do you tailor it to be more comfortable or drive more naturally? You really need a system that fricking works. They need reasoning to understand intentions of other drivers and pedestrians.”

Anca Dragan
Waymo, 2017

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient’s employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

GLOBAL LEADERSHIP POSITION FOR HUMAN BEHAVIOR PREDICTION

JAZZ VENTURE PARTNERS · SLOW · First Round · TOYOTA AI VENTURES · HYUNDAI CRADLE · PAVE · WORLD ECONOMIC FORUM · Gartner Cool Vendor 2018

$20M in capital (oversubscribed Series A)

Technical Founders: Neuroscience + Artificial Intelligence Experts

Sam Anthony, PhD
Co-Founder & CTO
20 years of cognitive science *and* computer science experience
HARVARD UNIVERSITY

David Cox, PhD
Co-Founder & Advisor
Director, MIT-IBM Watson AI Lab
MIT Massachusetts Institute of Technology
IBM Watson

Walter Scheirer, PhD
Co-Founder & Advisor
Field-leading researcher in computer vision and deep learning
UNIVERSITY OF NOTRE DAME

SLATE · WIRED · QUARTZ · Bloomberg · NVIDIA · The Boston Globe · DAIMLER · FAST COMPANY

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

WE ACCELERATE OUR CUSTOMERS' TIME TO MARKET AND PROVIDE VALUABLE OPTIONALITY WITH A PROVEN PATENTED SOLUTION AT DRAMATICALLY LOWER COST

- Faster Time To Market
- Certainty With Proven Solution
- Access to Unique Talent & Protected IP
- Dramatically Lower Cost
- Optionality & Redundancy

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

Thank you!

James Gowers
VP Strategy & Business Development
james@perceptiveautomata.com

PERCEPTIVE AUTOMATA

Human Intuition for Machines

Customer Use Case Manual

For SOMAI Deployment

Table of Contents

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

COMPANY OVERVIEW

Perceptive Automata was founded by a team of Harvard and MIT neuroscientists, computer vision researchers, machine learning experts, and software engineers to solve what is often described as the hardest problem for highly automated driving: reading the state of mind of humans for the safe large-scale rollout of highly automated vehicles (L2-L5).

ANTICIPATING HUMAN BEHAVIOR
PROBABLY THE HARDEST PROBLEM FOR AUTOMATED DRIVING

> “It's the prediction piece that's still the great unknown. Humans are very good at predicting human behavior on the road. Machines will need to be able to predict and anticipate human behavior much better.”
>
> **Gill Pratt**
> CEO of Toyota Research Institute
> March 2018

> “The choices made by driverless cars are critically dependent on understanding and matching the expectations of humans. This is the heart of the problem going forward.”
>
> **Chris Urmson**
> Head of Waymo, 2009-2016
> Co-Founder of Aurora, April 2017

We enable automated vehicles to better understand what people might do next so they can navigate safely around humans, including pedestrians, bicyclists, and motorists. Our AI processes, in real-time, our customers' object data and outputs human state of mind attributes, such as 'intention' and 'awareness' as defined in the below graphic, to our customers' AV path planner or ADAS system.

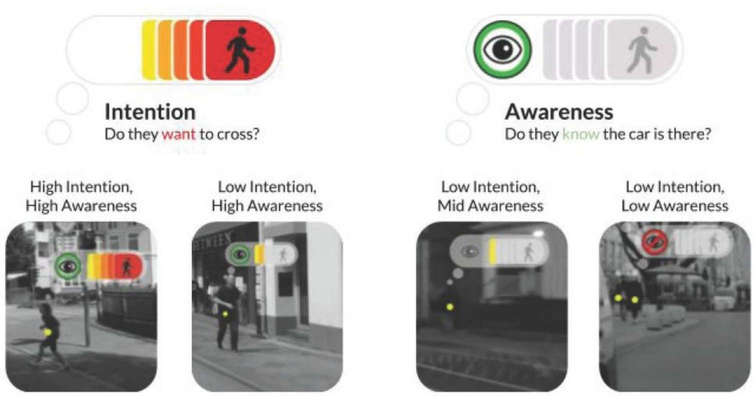

PEDESTRIAN, BICYCLIST, AND VEHICLE MODELS AND LONG-TAIL TOO

Our signals increase safety and roadmanship by enabling AVs to anticipate human behavior before actual body motion is detected. Our signals modulate otherwise paranoid automated driving behavior to be more predictable and natural, which reduces rear-endings and enables a smooth ride experience. For ADAS applications our signals give human drivers better situational awareness and increase driver acceptance of more advanced driver assistance functions.

These benefits are critical for increasingly sophisticated driver assistance functions to be embraced by human drivers - too many false positives lead to human drivers turning their vehicles' driver assistance systems off. They also critically enable autonomous vehicles to seamlessly integrate into human-dominated road environments and to deliver a pleasant ride experience for passengers.

OUR VALUE PROPOSITION PRODUCT

AUTONOMOUS - L4/5

- Safety
  - detect and react sooner, before motion
  - predictable driving => reduce rear-endings
- Smooth Driving
  - passenger comfort
  - passenger trust
- Fleet/Network Economics
  - efficiency of vehicles completing trips
  - less vehicle wear and tear + fuel consumption

HIGHLY AUTOMATED - L3

- Safety
  - increased situational awareness for driver
  - deploy in more complex road environments
- Driver Acceptance
  - reduce frustrating false positives
  - finer modulatory signals improve outcomes
- Enabling Transition from L3 to L4/5
  - focus on common platform
  - early enablement of L4/5 functionality

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

We raised $20M from top-tier investors and are working out of offices in Boston and Silicon Valley with OEMs, suppliers, and technology companies that are developing ADAS and/or autonomous driving systems.

GLOBAL LEADERSHIP POSITION FOR HUMAN BEHAVIOR PREDICTION

The value we deliver to our customers through our software licensing business model includes fewer years of technology development for faster market deployment, tens of millions of dollars in development cost savings, and access to scarce talent and intellectual property.

OUR VALUE PROPOSITION BUSINESS STRATEGY

THE CHALLENGE

People make effortless, intuitive judgments about what someone else wants or knows constantly. In any given thirty second interval at a busy urban intersection, there will be dozens of instances where one person looks at another and thinks, for example, “she isn’t going to cross,” or “he knows I’m here and is willing to yield,” or “she doesn’t see me, I should stop.”

Humans are incredibly good at silently communicating with each other. That communication is the key to safe, smooth and considerate driving. Until now machines have lacked this critical ability. They can't decipher our unspoken social communications nor intuit what's going on inside our heads. So how can we give machines like self-driving cars the ability to read human intentions, to know what humans know?

The best currently deployed approaches to interacting with the world using computer vision and motion planning are insufficient for the task of interacting with humans. Existing non-PA approaches use physics to locate people and to calculate their trajectory, but can't usefully anticipate what a pedestrian might do next. People aren't billiard balls — we aren't predictable based purely on our past motion.

Let's take the below example to illustrate this point. Many of today's approaches convert detected humans into 'bounding boxes', shown below as a yellow box in the image on the left, and then attempt to predict their likely actions based on calculated trajectory and context.

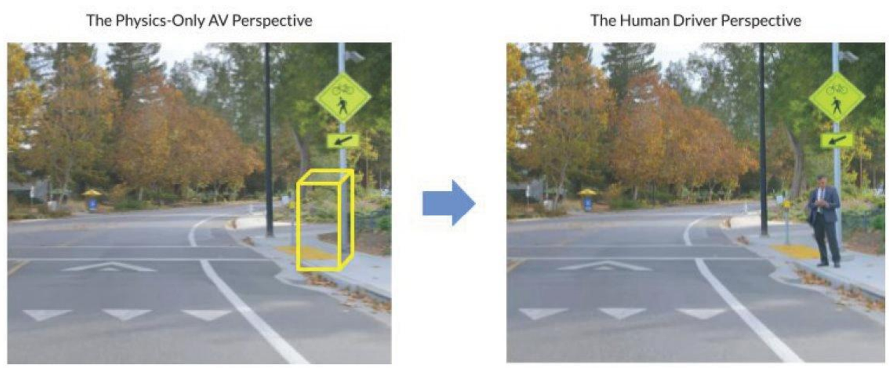

In this example on the left, a pedestrian is detected at the edge of the road at/near a crosswalk with zero motion. The only safe decision for the AV is to stop in case the pedestrian wants to cross the road. This means that instead of smoothly responding to a situation the way a human driver would, self-driving cars and trucks act alternately 'paranoid' — timid, skittish, and easily startled — and oblivious.

Next, let's take a closer look at the image on the right showing what a human driver sees, and what a self-driving car *with* Perceptive Automata software can see too. This image on the right clearly includes much more decision-making information - not just location and trajectory shown in the left image, but also the full range of visual information needed to determine that this pedestrian, in all likelihood, has no intention to cross despite standing at/near a crosswalk.

The physical and social cues humans effortlessly pick up on include that he is leaning back slightly, carrying himself without much tension, attending to his smartphone, and showing no indication in his facial expression that he's trying to make a crossing decision. So if we polled, say, 100 human

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

drivers approaching this scene, the likely prediction, in contrast to the physics-only self-driving car in the left image, would be that this person has no intention to cross in front of the car.

A typical human driver would therefore *not* stop in front of the pedestrian, and, just as important, in the context of rear-endings, other human drivers behind the lead car also would not expect the lead car to stop. Self-driving cars with Perceptive Automata's software would come to precisely the same conclusion as those 100 human drivers and also not stop.

This is just one example, but existing self-driving vehicles frequently slow down or come to a complete stop in the presence of pedestrians and/or cyclists when there is no need to do so, only to start moving again at exactly the wrong time. This style of driving is nauseating for passengers and infuriating for other road users. It's the reason self-driving cars are often rear-ended and it can cause human drivers to behave erratically.

It is now widely understood that there will be no meaningful real-world deployment of self-driving vehicles without solving the problem of human understanding.

THE PERCEPTIVE AUTOMATA APPROACH

In an effort to solve this problem we've designed a State of Mind AI (SOMAI) that can use the whole spectrum of subtle, unconscious insights that we, as humans, use to make incredibly sophisticated judgments about what's going on in someone else's head. You could say that, in a sense, our models develop their own human-like 'intuition.'

We take monocular camera recordings from vehicles that show interactions with people (pedestrians and bicyclists today, other motorists in the near future). This pixel data is incredibly rich raw material; every time a person appears and interacts with a vehicle, they're giving off hundreds of signals that we as humans can innately use to understand their state of mind.

This method was developed by two of our co-founders, Sam Anthony and Walter Scheirer, over the course of a decade of research at Harvard. The fundamental method relies on the realization that supervised learning — the method by which machine learning models are trained using labeled data — relies in almost all cases on data that was labeled by humans. Human labels are generated via a process that involves showing a training sample to a human observer and asking them to record their response. This type of paradigm is per se a human behavioral experiment. In most cases the psychological and neuroscientific aspects of this labeling are afterthoughts. However, for Perceptive Automata they comprise the core of our models' inferential power.

USING SOMAI OUTPUTS

SOMAI modulatory signals can be injected at a number of points in the perception/action loop in order to improve use case performance. Below we describe some of the potential integration approaches based on what we've found to be helpful for customers in the past. Please note that

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

this should not be considered a comprehensive list. In practice, there are many more integration paths, and the paths listed below will generally require a degree of customization.

ADAS FALSE POSITIVE REDUCTION

In ADAS systems where the driver is in full control, SOMAI's modulatory signal can be used to reduce false positives. For instance, in the case of an automatic emergency braking (AEB) system, the activation threshold for the system can be modulated by a weighted linear combination of the SOMAI awareness and intention model outputs. The weightings and the amount of threshold modulation can be set by the customer.

In practice, what this means is that the activation threshold for the emergency braking system -- the certainty of evidence required before brakes are applied -- would be higher in cases where a pedestrian in the roadway did not intend to cross, and was aware of the car. When tuned correctly, this can effectively remove false positive emergency braking scenarios, for example when a pedestrian enters the roadway for reasons other than intending to cross the road.

An additional tuning parameter that's available in this case comes from the distributional nature of SOMAI's model outputs. What our models generate is not a scalar value for intention or awareness, but a prediction of the distributional characteristics of human responses to the presented human road user. This means that there is information not just about the central tendency -- the equivalent of the scalar mentioned above -- but also the variance, kurtosis, skew, heteroskedasticity, and even multimodality of the predicted human distribution.

In the false positive reduction case and many of the other cases described here, a useful tuning parameter is to use the kurtosis of the output distribution to threshold the model input to the AEB system. In cases where the kurtosis is very low, and the distribution approaches uniform, it is reasonable to assume that a measurement of central tendency would not be useful for prediction.

On the other hand, in cases with high kurtosis the central tendency is likely to represent substantial human agreement that the output intention or awareness value accurately reflects the internal mental state of the detected human road user.

KALMAN FILTER MODULATION

The outputs of the SOMAI intention model can be used directly to modulate the estimates of a Kalman filter. The prior state used in the Kalman filter estimation can be weighted by the current intention level. This will have the effect of 'tuning' the filter.
In the high intention case, in practical terms this means the filter will attend preferentially to motion (even if that motion is below the nominal noise floor of the motion sensor) that could potentially cross the ego vehicle path in cases where the pedestrian looks like they want to cross the street.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

In the low intention case, the filter will preferentially discount as noise motion readings even though those readings might be slightly above the nominal noise floor for the sensor.

WORLD MODEL MODULATION

In a lattice planner-based autonomous driving stack, it is essential to have a world model. A world model in autonomous driving can consist of multiple layers of perceptual inference about the state of the road around the vehicle. SOMAI model outputs comprise an additional layer of information that modulates the existing world model to make it more accurately represent the probabilistic likelihood of safe driving across the modeled grid.

The basic informational layer of a world model is an occupancy grid. An occupancy grid comprises a binary random variable for each point in a discretized topological rendering of the autonomous robot's local environment. The posterior estimate of occupancy at a given time step can be constructed with a bayesian estimator that takes into account not just the outputs of depth sensitive sensors like LIDAR, but also semantically-aware estimates of minimum safe distance around moving agents (e.g., the maximum distance a human agent on foot can be expected to be able to move).

When a vehicle is simulating proposed motion paths, the semantically-aware minimum safe distance estimates for moving agents can be biased by the SOMAI outputs. So, for instance, for a pedestrian with high intention, the minimum safe distance region would be stretched towards the current vehicle path and squeezed in other directions, preserving total area.

POMDP PRIOR BIASING

In most probabilistic autonomous robot planning systems, path planning is achieved via a partially observable Markov decision process. In the world model example above, a POMDP would be efficiently approximated to pick the optimal route. The estimation at each time step relies on an observation vector which is assumed to comprise a functionally sufficient set of information to estimate the underlying state of the world.

The SOMAI models output an observation that correlates with a hidden portion of the world-state, the internal goal and belief structure of human agents. Because of this, the SOMAI model output can be directly added to the observation vector of a POMDP.

When the POMDP is trained including the SOMAI state of mind vector in its observation vector, the motion paths selected by the estimator will be sensitive to the ‘best guess’ information about the state of mind of human road users, without further system redesign being necessary.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

FLEXIBILITY, EASE OF INTEGRATION, SPEED, AND SAFETY

SOMAI predictions are integrated with customer systems as a *subsystem* that mostly runs in a separate process. The subsystem is comprised of:

- A subsystem 'service' process that accepts inputs from the customer/system process and returns state of mind predictions and other optional outputs steps (e.g. visualizations);
- A lightweight 'interface' library that is compiled into the partner system and provides a simplified API that can be used to interact with the service.

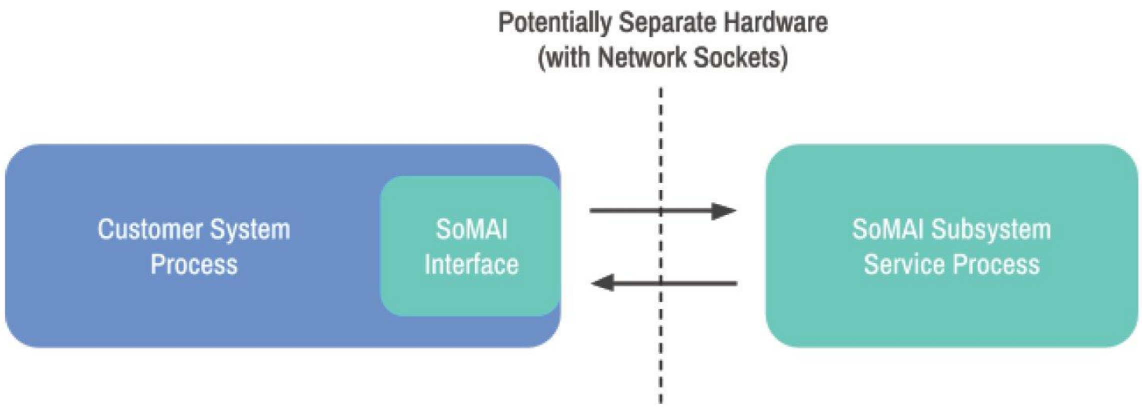

The interface library is responsible for launching the service and managing inter-process communication with the service. It is packaged as a C++ static library with a simplified API that is tailored to the specific needs of each customer.

This results in a limited API that is (a) easier to use and understand than a generic API, and (b) minimizes the footprint of Perceptive Automata code in the customer's core system processes, which can be significant from a functional safety standpoint.

The service process performs all the heavy lifting of the SOMAI subsystem - generating state of mind predictions and optionally performing a variety of other functions. As described in the following subsections, the subsystem service is designed to be ***flexible***, ***fast***, ***safe***, and ***upgradeable***.

FLEXIBILITY

With a number of optional calculations and a wide variety of alternatives for how data is transferred to and from the process, the SOMAI subsystem is designed with an emphasis on flexibility. The service can accept inputs from any number of sources, including sockets (network or UNIX), files, ROS, or other custom IO modules. Similarly, outputs can be published to one or more of a variety of destinations.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

SIMPLICITY - EASE OF INTEGRATION

This flexibility is typically hidden behind the simplified customer-tailored interface (so as not to introduce complexity for customers) but can be leveraged to quickly accommodate specific customer needs. Additionally, the specific calculations performed by the service (and therefore the inputs and outputs) can be easily controlled to suit different customer needs.

For example during evaluations and initial testing:

- Object detection can either be provided by the customer system (as inputs) or computed by the SOMAI subsystem (and optionally returned as outputs);

- Visualizations of the road user state of mind, which are often useful for evaluation, can optionally be generated by the subsystem.

SPEED

The SOMAI subsystem is fast. Forward inference with an *unoptimized* model takes <0.2 ms on a middle-of-the-road GPU (GTX 1060). Overall prediction by the subsystem takes a bit longer (due to data transfer, etc.) but is still about ten times faster than real-time for as many as 40 simultaneous pedestrians or cyclists.

The calculation is so light-weight that we anticipate real-time performance even without a GPU. By applying quantization and weight pruning optimizations, real-time performance on an ARM-based ECU should be well within reach.

Note that object detection (if requested) slows down the prediction time, but still easily achieves 30 frames per second with a GTX 1060.

SAFETY

Conscious of the functional safety concerns of our partners, we've developed the SOMAI subsystem with functional safety in mind from the start. For example:

- Packaging the SOMAI subsystem as a separate process allows the bulk of Perceptive Automata's code to be segregated from the core system process, minimizing the risk of subsystem problems affecting other parts of the system.

  PA code can even be kept entirely out of the customer process if need be - for example, by communicating inputs and outputs via ROS or another message passing library.

- We employ an increasing number of 'traditional' functional safety safeguards (including many of those described in ISO 26262) during the development of the subsystem.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

- Beyond ISO 26262, we've begun applying SOTIF techniques (including those described by ISO/PAS 21448) to our deep learning models and to the overall subsystem.

UPGRADEABILITY

Finally, the SOMAI subsystem architecture enables partner systems to benefit from ongoing improvements via over-the-air updates. The separation of different pieces of the system into separate files (and processes) supports updates of the subsystem executable and/or deep learning models on the fly, with no other updates to the system required.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

SOMAI Customer Service Release & Download Page

Dear Customer,

The following link is for the SOMAI Service package that includes support for Pascal, Turing, and Volta GPU architectures on Ubuntu 18.04.

SOMAI API Package v02.03

The configuration requirements are as follows:

The SOMAI Service needs to know

1. the frame size in order to properly initialize the detector at runtime;
2. the frame rate to optimize tracking and smoothing;
3. which of the supported GPU architectures it will be run on.

If you launch the Service manually using the launch script, you can set environment variables, 'SOMAI_FRAME_WIDTH', 'SOMAI_FRAME_HEIGHT', 'SOMAI_FRAME_RATE', and `GPU_ARCH` which will be picked up by the 'config.sh' script and used to configure the service. For example:

```
SOMAI_FRAME_WIDTH=3840 SOMAI_FRAME_HEIGHT=1920 SOMAI_FRAME_RATE=10.0
GPU_ARCH=pascal ./run_service.sh <args>
```

The values shown above represent the default configuration, which means that the environment variables will default to these values if left unset. To bypass setting these environment variables at the command line, you can also edit the 'config.sh' script directly to set the environment variables to the desired values.

Alternatively, if you launch the Service with the Service Daemon, you can set the 'frameWidth', 'frameHeight', 'frameRate', and 'gpuArch' parameters, which are members of the 'sStateOfMindAIConfig' struct. Then, when the Service is launched by the Daemon, it will be configured to use the appropriate frame parameters without having to use the environment variables.

Please note that due to the requirements imposed by the Driveworks 2.0 detection and tracking modules, this package has been built for Ubuntu 18.04 and includes many additional libraries in its 'lib' folder which may be removed in future releases.

Please submit any support requests to support@perceptiveautomata.com.

The content of this document contains the confidential information of Perceptive Automata and is subject to any agreement between Perceptive Automata and recipient (or recipient's employer/contracting entity) regarding the protection, use, and non-disclosure of confidential information.

SOMA | State of Mind Analyzer

User Guide

Confidential

Table of Contents

Introduction

The ***State of Mind Analyzer*** (***SOMA***) is a tool that enables in-depth analysis of the outputs of Perceptive Automata's ***State of Mind AI*** models.

The tool's primary aim is threefold:

1. Provide a streamlined, responsive interface for analyzing ***State of Mind AI*** model outputs.
2. Render intuitive visualizations of those outputs to make them easy to understand at a glance.
3. Facilitate deeper, fine-grained analysis by enabling frame-by-frame inspection into each output.

In addition, SOMA provides controls for quickly viewing all the sequences within a video and selecting one to focus on, graphing the awareness and intention predictions of a cyclist or pedestrian over the course of an entire sequence, and recording an interesting sequence with accompanying visualizations to a separate video file.

Getting Started

System Requirements

- Ubuntu Linux 16.04 LTS
- NVIDIA CUDA 9.0
- FFmpeg (optional for recording video)

Quick Start

1. Untar the package file into target directory:
   cp soma-<version>.tar.gz /path/to/target/directory
   cd /path/to/target/directory
   tar xzvf soma-<version>.tar.gz
2. Install FFmpeg (optional):
   sudo apt-get update
   sudo apt-get dist-upgrade
   sudo apt-get install ffmpeg
3. Launch SOMA:
   cd soma-<version>
   ./soma
4. Load a video using the ***Open video*** toolbar button. This action is explained in greater detail in the next section.