Recorded:
May 25, 2016

Watson Innovations Podcast #1:

From eyesight to insight: How computer vision is creating opportunities for companies

Watson

IBM

Episode Guide

Podcast

Watson Innovations Podcast

Episode 1:

From eyesight to insight: How computer vision is creating opportunities for companies

The digital universe is growing at a massive rate, and much of it is unstructured data in the form of images. In fact, in 2015 alone, consumers took 1 trillion pictures and uploaded them to social media. All this visual data provides organizations more opportunities than ever to transform their operations and deliver superior customer experiences at scale. However, without a way to quickly and expertly analyze the content of images, organizations are in the dark. In this podcast, Alyssa Simpson, Watson Senior Offering Manager, discusses why computer vision is important in the digital era and how the Watson Visual Recognition cognitive API is helping companies like OmniEarth, Inc. turn image insights into value. Learn how industries as diverse as manufacturing to marketing can benefit from computer vision as well as how you can get started today.

Listen here:

- Soundcloud: http://ibm.biz/wi-vision

Subscribe to the Watson Innovations Podcast using iTunes or your favorite podcatcher to have new episodes of the Watson Innovations Podcast delivered right to your mobile device or desktop:

- RSS: http://feeds.soundcloud.com/users/soundcloud:users:239204627/sounds.rss

Show notes

OmniEarth: Combating drought with IBM Watson cognitive capabilities

Drought is a complex problem that defies easy solutions, but that hasn't stopped one startup from tackling it. Environmental technology company OmniEarth helps cities and towns with aggressive water conservation targets by providing granular data on water consumption. Using advanced geoanalytics powered by the Watson API, Visual Recognition, OmniEarth processes, clarifies and fuses satellite and aerial imagery with other data sets to help local agencies quickly identify land parcels that need to reduce water usage. In this case study, find out how OmniEarth uses IBM Watson APIs and developer tools on the Watson Developer Cloud to build solutions that quickly make sense of vast amounts of unstructured data and enable the green data revolution.

https://ibm.biz/cs-omni

Watson Visual Recognition API

Visual Recognition uses deep learning algorithms to analyze images for scenes, objects, faces, text, and other subjects that can give you insights into your visual content. You can organize image libraries, understand an individual image, and create custom classifiers for specific results that are tailored to your needs.

http://ibm.biz/visualrecognition

Watson Developer Cloud

Watson Visual Recognition is part of the Watson Developer Cloud, a suite of flexible cognitive application programming interfaces (APIs) backed by the latest advancements in machine learning. Watson Deveoper Cloud makes it easy to accelerate application development with an array of software development kits, code examples, and tooling created for developers, by developers, and in use by tens of thousands across 20 industries, with over a billion API calls processed per month (as of publication date).

http://ibm.com/WatsonDeveloperCloud

IBM Watson Blog

Stories of how cognitive computing is transforming our world.

https://www.ibm.com/blogs/watson/

Transcript

Eve-Marie Lanza (EL): Welcome to the Watson Innovations Podcast brought to you by IBM. My name is Eve-Marie Lanza and today I am speaking with IBM Watson senior offering manager Alyssa Simpson. Alyssa oversees cognitive vision services for IBM Watson and is a frequent speaker at events like Exponential Two Thousand Sixteen, the Silicon Valley Forum's Women In Tech Festival and Launch Festival Two Thousand Sixteen where she recently presented Disrupt: How Innovative Companies Are Using Cognitive APIs to win.

Today, she'll be talking with us about computer vision: what it is, why it's becoming an increasingly important capability for organizations to have, and how Watson is helping organizations see more opportunities than ever before.

Hi Alyssa, thanks for joining us.

Alyssa Simpson (AS): Thanks for having me

(EL): Why is computer vision getting so much attention now?

(AS): The digital universe is growing at a massive rate—really unprecedented growth. Much of that data—as much as eighty percent is what we call unstructured, considered dark. It's something that traditional computing systems are blind to the content of this data and are unable to understand the information within that data.

Think about this in the context of any type of information that you can't put into an excel spreadsheet. A picture is as a really good example—doesn't fit nicely into rows and columns. Consider that in 2016 consumers took over one trillion pictures and uploaded them to social media. Countless more images were generated by devices like security cameras, satellite imagery, drones and all sorts of other cameras. We all carry a camera in our pocket.

There's a lot of potential insight locked into those images that organizations can benefit from. That's where computer vision comes in.

(EL): Could you tell us at a high level what computer vision is?

(AS): In short, we're really teaching computers to see. If you think about that, it's a pretty difficult issue.

Imagine that picture of you and your family at the beach: smiling, hanging out, having a beach day. Then, try to think of the top five words to describe that image. If you were describing it, you may say things like: family, people, my family at the beach. However, if you ask your mother she might say something different such as: beach day or sunset or happy family smiling. If you ask a department store what they would like to know about this image, if it was posted on social media, they may want to know: swimwear or striped women's bathing suites. If you ask a doctor, they may say: healthy. Psychologists may even say: excited. As you can see, it gets pretty complicated to understand what is true about an image based on who's asking and answering the question. It's all context-specific.

(EL): That does sound very complicated. How is Watson helping organizations understand their visual data?

(AS): Watson is really attempting to make it easy for organizations to see the way a human expert would, as in the previous example, with Watson visual recognition service, part of a cognitive suite of APIs available on the Watson Developer Cloud [WDC]. Officially speaking the WDC Vision service is the technology that enables our customers to find insight, drive value and take action on visual information of any kind but it's all in within the context of who's asking that question.

Visual Recognition understands the context of images, can help you organize them, create custom classifiers and develop applications that are tailored to your needs. We do this in a few different ways. At the most basic, we create tags from images, which are natural language which describe this image.

We can also extract which image off of a website is most relevant to a body of text. Consider, for example, a New York Times article. There's actually probably five or six images in that article and you may be only looking for one that's most visually relevant to the text of the article. We can also tell you if there's a person in the photo that can be used to identify them or crop it around that person.

The most interesting thing that we really drive is being able to teach Watson to see, trainable for learning around custom content in your image and that is unique to you.

(EL): Very interesting. So in general how is Watson approaching vision? What makes Watson Visual Recognition unique?

(AS): At Watson we approach computer vision like we do all of our other services, with an emphasis on making it easy for application developers, not just a fancy data scientist or an artificial intelligence expert, to put the latest advance and intimate machine learning to work for them to drive business value. In short, we want to make it really easy to use Watson and our cognitive services have a few basic components. They need to understand, they need to reason, and they need to learn.

When it comes to computer vision we focus on a few things. First, customizability. Each business is unique, every business challenge is unique, that means different businesses need to derive the different insights from visual data. So we've made it trainable—that means it can learn. That's really sort of the core pillar of Watson. End users don't receive a static response such as "person" from a photo of you and your family at the beach. You could train it to answer your particular business question. You could train it to say that's a happy picture of my family and the unique response that you would like.

Second, in talking with our customers, we've really learned that figuring out how to make sense of this data can be really difficult. So we focus a lot on designing our cognitive vision services to provide the data that you want and to make it really easy to get started. It's actually free to get started, and folks who don't have significant experience in computer vision or AI can use these APIs.

And finally, these APIs are open and flexible. They can be easily integrated into a lot of other services, because the real value is hooking up these things together like building blocks and combining them to solve a real business challenge.

(EL): That's great. What kind of insights can organizations get from images using the watching Visual Recognition Service and, more importantly, what can companies do with those insights?

(AS): So many things. I see companies using these services to answer every question imaginable and every time I talk to a new client, I learn something new and they're trying to answer a different question. Basically, “what is this a picture or video or an image of?” It answers the question, “what is this picture?” It can also answer more complex questions like, “how many pictures or videos do I have that are similar to this one are all within this type of a category?”

Getting a little bit more business-specific, understanding how one image in particular might interact with other data to really be relevant to the bottom line, whether or not the retail application or you're trying to listen for business intelligence, a lot of different drivers that people are trying to do. Even things like common themes in a batch of millions of photos. If you want to understand my personality a little bit better you may look at every image that I’ve posted to Instagram and Facebook and you can start learning things about my interests, such as: I like cooking and I am pretty outdoorsy and I like to go on hikes. Those are themes of, based on the imagery that I’ve posted to social media but may not be something that I’ve described specifically in text, so there's a lot of information that's really hidden in these images [and there’s value in] being able to extract it.

It's important to note that computer vision is not really just cool technology. It is, and I'm really excited about it, but the reason it's important is these insights are becoming a real necessary piece of delivering an awesome customer experience at scale. It's important for personalization, it can help companies make better decisions, enhance outcomes like really improving an online shopping experience or automating a mechanical process inside a manufacturing plant. Even conserving water in my home state of drought-stricken California.

(EL): You mention that visual recognition is being used today to combat drought. That’s a big issue for us in California. Can you tell me more about how Watson is helping?

(AS): You know, that’s something that I'm really concerned about here in California and a client of ours, OmniEarth, got some great press for this recently. Using our cognitive computing, they can basically understand from aerial images, combined with a lot of other data, patterns around water consumption. In California, the state has imposed water restrictions to regulate the usage and consumption of water, but the data that they rely on to make those decisions is really bad data—it’s really broad, it’s annual or multi-year averages, [using information] of that sort to understand the usage and set targets. In short, it’s really, really bad data and it's not particularly actionable.

OmniEarth developed a solution that they knew could sort of help the state monitor water usage using imagery and combining that with local consumption data. [They could do] mapping data but the flexibility and speed [to] train up customer classifiers and process significant parcel-by-parcel imagery was pretty limited, because they had to have humans look at this and, and verify it, at local government levels. OmniEarth is really interested in helping them, so what they did was they combined Watson Visual Recognition and they trained it to recognize topographical features in unstructured satellite images that come from the USGS. So, as an example they trained it to understand what a lawn looks like or what a pool looks like and how a pool is visually different from a pond in an aerial photograph,

and what they did was, if you take a large satellite image and you break it up into individual parcels of land right, this is just the image of my house and then you run it through visual recognition and you can understand, Hey do I have a lawn in my backyard or do I have a pool? And they can understand averages around what it takes to water a lawn to keep it green and how much water it takes to fill a pool, and they can combine it with my water usage data to start to understand, am I over consuming? So it's actually really cool. One of the local governments they were working with, the individual at the government that they were working with, actually found a leak in his backyard based on this, that he didn't know he had and so he, of course, went home and fixed the leak and then used this data to market to other people in his county. So, we're really excited about OmniEarth and, and the work that they're doing.

They're able to really improve the processing time about forty [times], [achieving] faster imaging process time than was possible from the method that they were using before, and then [increasing] capacity for analyzing terrain on a pretty broad massive scale. Every single image for every parcel in every county in California is not a trivial amount of analysis to do. And it can be done in a number of minutes rather than previously what took years to run an analysis.

(EL): I'm there with you, right there. That's all very, very intriguing. Can you share some other ways industries might put a Watson cognitive vision service to work?

(AS): Absolutely. We see customers using this in a lot of different industries, a few come to mind:

Manufacturing around visual auditing. For example, think of if a product is placed correctly on a manufacturing line or if there is a defect in a product or even in an airplane or a windmill using drones to fly out there and inspect, you know, a power line or inspect a windmill in a far off place and come back to understand, "hey is there something wrong? Is there a dent? Did it get struck by lightning? Or is there something wrong that I need to prioritize?" And those types of use cases.

In the retail space, it's really around personalizing your customer segmentation, your promotion, making sure that when you're showing me a product image to buy that it's similar to things that I may be more likely to buy, which would be really different from my mother, right? We are both interested in black t-shirts but we have really different opinions on which one will work for us, so that can drive use cases like conversion optimization, product searching, and really analyzing user generated photos.

Marketing and advertising have similar use cases around really personalizing advertisements and campaign creation and automating the process of finding the right image to put in your ad campaign that's really going to resonate with your audience.

(EL): That's really impressive. For organizations that are listening and are interested and getting excited about the possibility of using computer vision in their business, what are a couple of things that they can start thinking about today to get started on the path?

(AS): The hardest part about using computer vision for your business is getting really specific around what problem do you want to solve, what action do you want to take from this image?

In OmniEarth's case, they wanted to understand: is there a pool here or not? With retailers, they want to understand: what do I sell that is most similar to an image the user has posted on Facebook? They're very narrow and very specific targeted use cases. That is the hardest part.

The second hardest part is coming up with the training data. Typically speaking, our businesses are solving unique problems that in some cases have not been solved before, and it requires a level of training to organize this data. For example, if you are a large car manufacturer and you're looking to understand how many times and where there are photos of my car being shared online, you need to train the computer, you need to train the system to understand that this is the latest and greatest car model that we released this year and here is what it looks like visually, and you need to pretend that you're talking to someone who's never seen that car before, who is not an expert and the same way is that I don't happen to be a car expert. I cannot tell you the difference between a Toyota Camry and Honda Civic because they both get me from "point A to point B" successfully—you need to train me. You need to train me to become an expert in the visual distinction that you want to make.

Does that make sense?

(EL): Absolutely. It sure does.

Great. Well, this is all been really interesting, Alyssa. Thank you so much for your time today.

Where can listeners go to find out more about Watson cognitive vision services?

(AS): You bet. We'd love our listeners to come check us out. There are demos available by visiting Watson Developer Cloud Website: IBM.biz/visualrecognition.

(EL): Excellent. Thanks again for your time today, it's been a pleasure.

(AS): Anytime.

(EL): And to our listeners thank you for joining us today. You can find out more about this and other Watson cognitive technologies as well as get links to the resources mentioned on today's show on the IBM Watson blog at IBM.biz/thewatsonblog.

Don't forget to tune in to the next Watson Innovations Podcast. Bye for now.

About the speakers

Alyssa Simpson leads Offering and Product Management for Watson Vision as part of the IBM Watson Developer Cloud. Previously Alyssa was Sr. Offering Manager at IBM Commerce: Customer Analytics and Tealeaf Mobile. A veteran of all things analytics and mobile, Alyssa has a product design background in mobile UX/app design and development. Alyssa is a data nerd, technologist, and photographer. She holds a BA in American Studies from Trinity College and lives in San Francisco.

Eve-Marie Lanza is a Senior Marketing Manager for IBM Watson.

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
July 2016