Technology Institute

# Mobile Innovations Forecast

## The elements of contextual intelligence

*Raman Chitkara, Global Technology Industry Leader*

Today's mobile users want to navigate their daily lives using a combination of highly personalised mobile devices, information and applications. Mobile operators covet the revenue these users will generate. Brands know they need to be prominent in this mobile bundle if they want to engage this audience. And no serious enterprise discounts the role mobile will play in its future.

However, the mobile user experience is inspired more by desktop metaphors than by the always present, highly personalised reality of mobile devices. Mobile interactions too often involve 'mobilised' desktop applications or websites shrunk to fit the mobile screen. Even native mobile apps have limited awareness of what a user wants without explicit input.

But the ground is shifting. As mobile markets mature, value creation is migrating to how well a device or service adapts its content and functionality to a user's needs and preferences. Having awareness of the user's contextual situation—where she is; what she likes; who she knows; how she has previously used a device or service and so forth—and having the intelligence to *act* on that knowledge is becoming a core driver of mobile innovation. As this process iterates, contextual awareness and intelligence will become a major source of growth for both the mobile and information technology industries.

### Synopsis

This is the fourth article in the Mobile Innovations Forecast Phase II: New technological capabilities. Here is a roadmap of the series:

- The Introduction argues that the dominant drivers of mobile innovation to 2018 will revolve around capturing and modeling the contextual situation of mobile users, and will transform the mobile device into an intelligent digital assistant.
- The second article examines how device and environmental sensors interact to capture information to model the user's physical context.
- The third article examines how communications networks enable interaction of the user's physical context data with information and applications in the cloud.
- **This fourth article describes the capabilities that enable the mobile device to generate contextually relevant information and services.**
- A concluding article will highlight the most significant new capabilities driving smart devices towards true digital companionship, setting the stage for new use cases and business models to follow. ■

This article examines the emerging capabilities that enable mobile devices and services to provide contextually intelligent information or take actions on behalf of the user in a natural, almost humanlike, fashion. PwC believes that contextual intelligence arises from coordination across three domains of technologies and capabilities found in the mobile device, the telecom network and the cloud:

a) **Conversational intelligence:** The natural language interfaces that converse with end-users on their mobile devices and deliver most contextual experiences to them;

b) **Sensor intelligence:** The software development environments that add contextual capabilities to mobile applications and services;

c) **Decision intelligence:** The machine learning capabilities that derive personal behaviour patterns to predict when and how to proactively engage the user with information, recommendations, and actions.

These domains interact at multiple levels and across multiple players. No single entity or industry is likely to own and control the end-to-end contextual value chain. Rather, the success of contextual technologies and services is likely to be ecosystem-led more than technology feature-led. With that in mind, it's time to explore each domain.

### Contextual awareness and intelligence defined

PwC defines contextual awareness and intelligence as a form of computing in which situational information about people, places and things augments the more slowly changing personal profiles of users and is leveraged to anticipate an end-user's immediate needs. Based on that knowledge, contextually intelligent services proactively offer enriched, situation-aware and actionable content, functions and user experiences. More importantly, every interaction adds to the depth and breadth of insight about a mobile user.

Since 2011, the most well known example of a contextually intelligent mobile service is Apple's Siri, a voice-activated virtual assistant that responds with information and services that grow in relevance the more the system is used. Other general-purpose virtual assistants such as Google Now or Microsoft's Cortana have launched to establish positions in the nascent market for mobile contextual services.

Whether embodied in a virtual assistant or integrated into a specific application, contextual awareness and intelligence is especially valuable for mobile user experiences. Unlike when she uses a desktop computer, the mobile user typically is not navigating documents or websites to meet her needs. The mobile user is more often navigating locations, choices and relationships that change dynamically as she progresses through her day.

The recurring question for mobile users as they encounter new information, people or choices seems to boil down to 'what is the best next step for me to take right now or in the near future?' The ability of a mobile device and service to be ready with contextually relevant information and services helps people decide what those next important steps should be.

## *Conversational intelligence*

Contextually aware and intelligent experiences require users to educate their mobile devices or services about their needs, preferences and desires. Some of this education will be explicit, such as a user giving a voice command to his device. Much more of this process will happen in the background.

Whether users actively or passively exchange information with their devices and services, their contextual interactions, whilst mobile, are likely to be conversational rather than based on formal commands or menus. "Siri, tell me which exit I should take to go to Northgate Shopping Mall" is a far more intuitive—not to mention safer—user experience than trying to input key words into a search box, and then read the results whilst driving a car. If the user's conversational experience is his initial exposure to contextual services, then the quality of the conversation significantly influences his decision to adopt contextual services in the first place.

Consequently, a core capability for contextually aware and intelligent mobile services is natural language processing (NLP). NLP refers to the ability of a computer to understand and converse in human language as human beings communicate with it. NLP is a hybrid discipline that combines expertise in computer science, artificial intelligence (AI) and linguistics. It is the technology and information foundation for the natural language interfaces used by virtual assistants like Apple's Siri, Google Now and Microsoft's Cortana. As such, NLP has a profound impact on the contextual experience for the vast majority of users.

Writing in the *Chronicle of Higher Education*[1] , Dr. Geoffrey Pullum, a professor of general linguistics at the University of Edinburgh, said that effective NLP results from the mastery of three core tasks by computer systems. Assuming that speech recognition technologies have extracted human words and phrases from the ambient environment, an NLP system must first understand **syntax** of a target human language to uniquely identify one sentence or phrase from another sentence or phrase. Second, an NLP system must possess **semantic** capabilities to extract the literal translation from a sentence or phrase and relate it to translation of other sentences and phrases. Finally, and most difficult, an NLP system must possess the **pragmatic** rules of thumb to infer the intent behind a human utterance, and thus discern what should be assumed or performed given the meaning of a sentence or group of sentences.

Although NLP has been a recognised subcategory of AI research since the 1950s, it was not until 2011 that NLP became well known to the general public. NLP made its most recent public debut in February 2011 when IBM's Watson supercomputer defeated two human champions at Jeopardy! While a flurry of positive and negative stories tried to untangle the ultimate impact of a computer's victory over humans in a contest conducted in natural language almost all observers were impressed by Watson's ability to interact smoothly with its human competitors and the game show host.

Then, in October 2011, NLP took another mass-market step with the launch of Siri on the iPhone 4S. Siri was the first mass market, general purpose[2] virtual assistant that employed NLP to create an interface in which users could speak in free-form natural language to query for factual output ("what was the score of last night's game between X and Y teams?") or perform voice-activated functions ("Siri, please send a message to the person I'm meeting at 3 p.m. that I will be 20 minutes late"). More recent virtual assistants, such as Google Now and Microsoft's Cortana, combine spoken and text-based interfaces based on NLP to expand the range of queries and tasks that can be executed by the user.

PwC believes that broad adoption of contextual intelligence pivots on the ability of users to engage with their devices and services in a conversational manner rather than in a command and menu approach. Just as the graphical user interface transformed the desktop computing experience for ordinary people, and hyperlinking defined how people experienced the World Wide Web, NLP enables a fundamentally new interface amongst people, information and technology.

These activities happen continuously at multiple levels of the mobile device, the network and stored data and/or functionality in the cloud. Depending on the nature of the contextual experience, these activities can execute across a number of service providers and ecosystems.

1 http://chronicle.com/blogs/linguafranca/2013/05/09/natural-language-processing/

2 Domain specific NLP systems for applications such as stock trading have existed for over a decade; Siri was the first to successfully remove the assumption of a specific semantic domain.

**Table 1: The elements of natural language processing**

| Element | Capability |
|---|---|
| Syntax | Uniquely identify one phase/sentence from another |
| Semantic | Extract literal translation |
| | Relate to other sentences/phrases |
| Pragmatic | Apply rules of thumb to phrase/sentence to discern what should be assumed or performed given the meaning of the phrase/sentence |

*Source: PwC*

## Common attributes of contextual systems

Contextually intelligent systems interact and learn about humans in real-time by incorporating the following information and processing:

1. **Acquiring information about a user and his or her environment.** The system typically draws its raw contextual data stream from physical sensors in the device or immediate surroundings, plus stored information on the device or in the cloud. Observation of contextual data might be episodic and initiated by the user, but most often is passive and continuous.

2. **Modeling the current circumstance or intention of the user.** A contextual system usually operates according to pre-programmed heuristic rules to build a predictive model of user intent. These rules are refined and augmented as the user repeatedly accesses contextual services.

3. **Reasoning about the next best step and taking action.** Based on its predictive model of user intent, a contextual system can access applications and data to offer and deliver contextually relevant experiences. The back-end resources that enable these experiences may or may not be known to the user.

4. **Notifying a user or other apps/ services of an action.** Contextual systems use conversational intelligence, based on NLP capabilities, to communicate with the user. The system also communicates with data sources and applications through APIs or similar means.

5. **Learning through feedback.** A contextual system continually analyses the conditions that trigger an event; the information and applications used to generate a contextualised response; plus the user's positive or negative reaction to the output offered by the system. These results feed back into the contextual system's configuration to provide increasingly personalised experiences to the user the more the system is used.

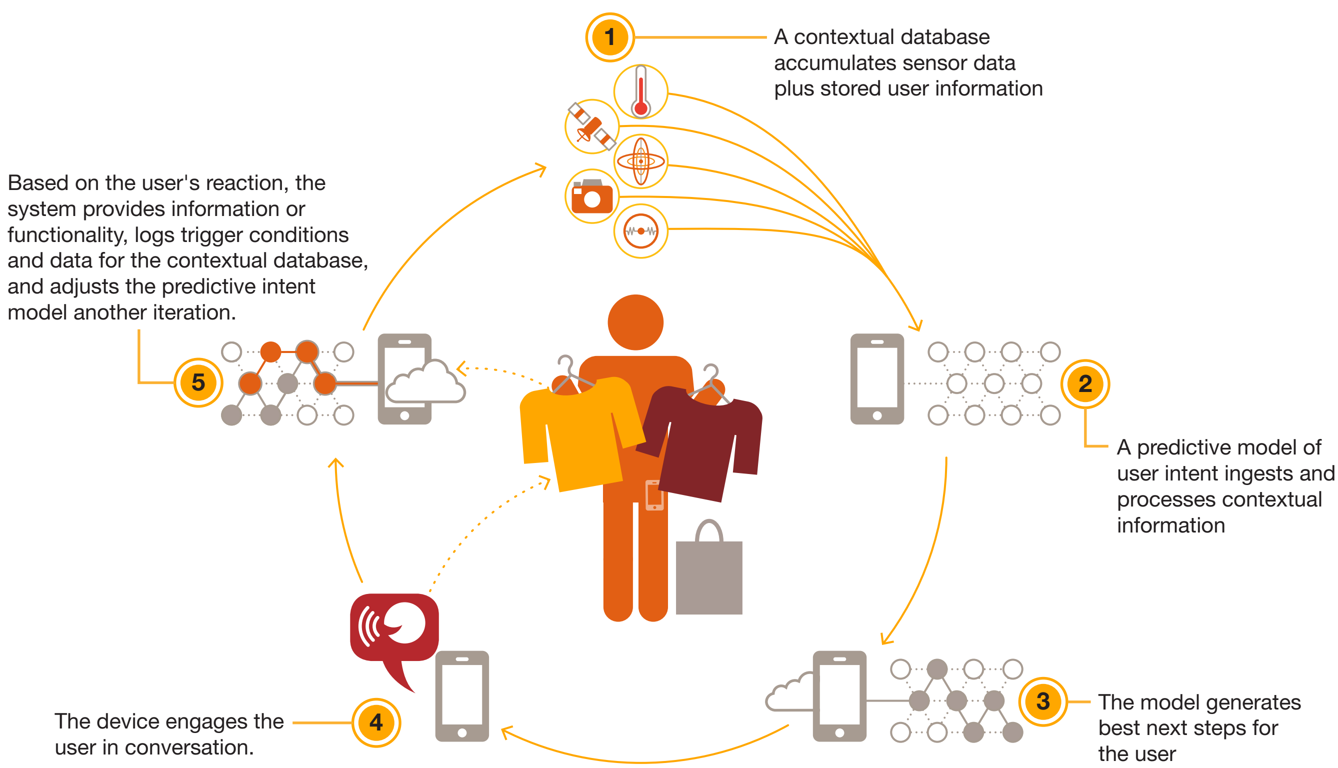


Source: PwC

*Contextually intelligent systems 'learn'* ***in real time*** *about humans and their environments through the exchange of information.*

## *Sensor intelligence*

If NLP provides a conversational interface to contextual experiences, then contextually intelligent applications use voice and sensor data to turn those conversations into action. This section explores the four primary contextual data inputs that help make a mobile application contextually aware and intelligent. It then provides a recent example of a context-focused software development kit (SDK) for mobile retail that packages these inputs so non-specialist developers can create contextually informed applications.

The four primary contextual data inputs (location, identity, time and activity) not only answer questions of who, what, when and where; they provide the contextual data resource for local applications and decision engines in the cloud. For example, a contextual application might access location and time information from a context SDK operating on a mobile device. If the location data of a user in her home corresponds to the kitchen, and if the time stamp is 0700 local time, the statistical correlation score for the contextual situation of 'breakfast' is suitably enhanced. Based on that insight, the contextual application will return some type of information or take an action according to the design of its developer.

**Figure 1: The four primary contextual data inputs**

*Integration of primary and secondary data inputs creates contextual intelligence.*

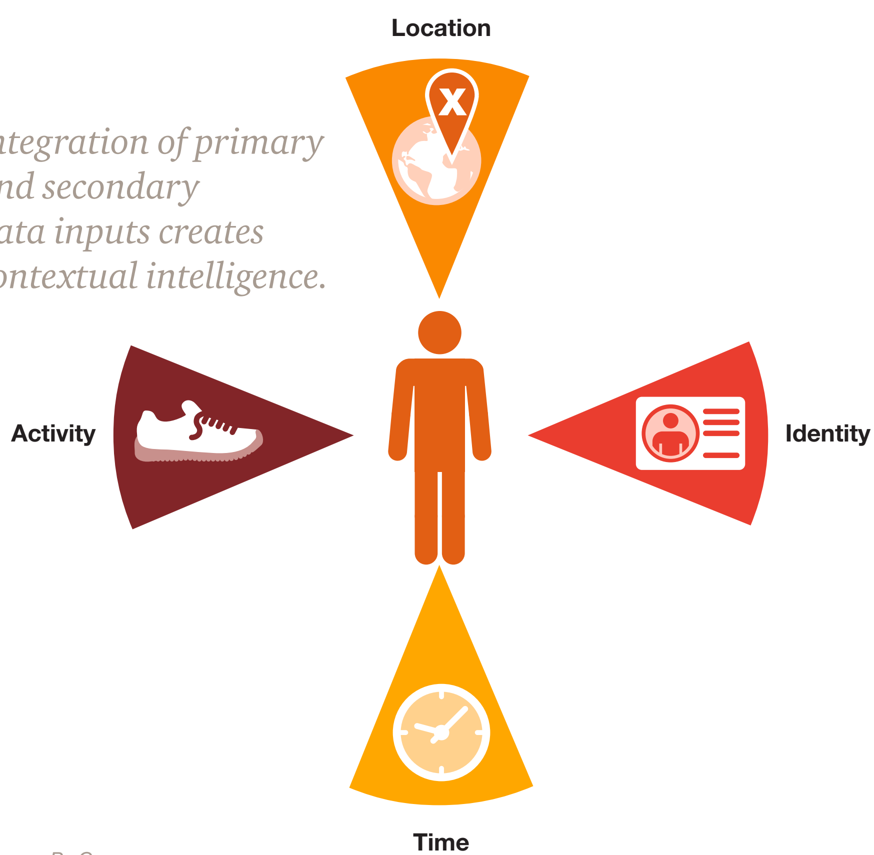


Source: PwC

Thus, it is useful to think of location, identity, time and activity as the primary colors of a contextual palette. Individually, these data types provide contextual value to an app depending on their presentation and use. More important, these inputs can be combined or mixed (location + activity) with secondary inputs, such as personal history, to create a range of options for adding further contextual value.

One of the first SDKs designed for building mass-market contextual applications is Gimbal from Qualcomm Retail Solutions. Gimbal uses location plus personalisation features to add contextual value to mobile apps focused around the experience of physical retail.

The platform is built around three core concepts:

1. **Location**—Where are the customers?
2. **Proximity**—What are the customers near?
3. **Interest**—Who are the customers and what matters to them?

In a typical scenario, a user downloads a Gimbal-enabled mobile app to share her location and other information on an opt-in basis in exchange for something in return. When the user's mobile device crosses a geofenced[3] boundary, that event triggers an action on the device or in the cloud that returns some type of

3 A geofence is a virtual perimeter around a real geographic location. Geofencing applications and services allow an administrator to set up triggers so when a mobile device enters (or exits) the boundaries defined by the administrator, a message is sent or a processing event occurs. Source: http://whatis.techtarget.com/definition/geofencing

*Machine-learning capabilities create the 'decision intelligence' that make the experience possible.*

information or digital token to the user. The output may be a standard offer for each opt-in participant or the output may be customised based on stored individual preferences or history. Simultaneously, the Gimbal platform records activity data for follow-on analysis by an application developer, a retailer or venue owner.

In addition to enabling application developers to draw macro-level geofences around physical locations such as a store or a shopping mall, Gimbal offers micro-location capabilities for areas 50 metres or less in radius by using proximity beacons[4]. These in-store devices are based on Bluetooth Low Energy (LE) to transmit ID codes that can be picked up by mobile devices that have downloaded retail applications. When a device is within physical proximity to the beacon and detects it, an app can notify the customer of location-relevant and individual-relevant content, promotions and offers. On the back-end, the app developer can set various interactive rules for the beacons such as activating during specific hours or engaging only with a certain shopper profile when she is within proximity.

The ability for non-specialist application developers to use an SDK like Gimbal to merge contextual inputs such as time and identity with fine-grained location capabilities for a target user experience is a major advance for contextual services. As demonstrated in other markets, innovations do not grow to mass-market status by staying rooted in the industry sector that spawned them. For contextual experiences to scale, developers that specialise in particular areas, such as retail, must be able to integrate sophisticated contextual capabilities into their applications without being contextual technology specialists. In that sense, general purpose contextual SDKs are a significant ingredient for success.

4 A proximity beacon marks a radio-based enclosure that a developer sets up around a smaller geographic area; for example, a department within a store, a specific street address, a section of a parking lot or a landmark in a plaza. Source: www.gimbal.com In addition to Gimbal, Estimote, Kontakt, and GeLo, to name a few, are offering proximity beacons.

## Decision intelligence

If the purpose of natural language interfaces and contextual applications is to package and present a contextual experience for a human or piece of mobile technology, then machine-learning capabilities create the 'decision intelligence' that make the experience possible.

Machine learning represents a significant departure from traditional system development methodologies. For most of computing history, programmes were built by distilling knowledge from human experts into a series of logical structures that enabled a system to respond in predictable, repeatable ways. If you wanted to build an accounting system, you started by interviewing human accountants to understand and create the rules that software engineers encoded into formal logic that could be understood by machines. So long as a target process lent itself to high levels of formalisation, the methodology worked reasonably well.

However, highly formal systems don't handle ambiguity or exceptions very well. Take NLP. Humans have tried and failed numerous times to develop a complete but manageable set of formal language rules that can handle the standard tasks and the exceptions of translation. Not only is human language rife with exceptions due to regional dialects and a host of other idiosyncratic factors, it is constantly evolving.

But by building a framework that enables software to start with some pre-programmed examples of previous, successful translations and then to compare those examples with a new sequence of words, a computer system might get closer to making a successful new translation. Add in scoring mechanisms for the system to track whether its current translation is closer or further from a target accuracy, and the system gains the ability to adjust its processing for the next translation attempt. Over time, the system will 'learn' to recognise statistically significant translation patterns that should grow in accuracy the more the system is used. In plain English, a machine learning system distills the rules it requires from the data on which it is exposed and trained, rather than having all that knowledge directly coded by the programmer.

Machine learning is foundational to contextual systems because it offers the ability to sift through vast data sets and classify preliminary patterns in a user's contextual data stream without direct human intervention. Sensor data logs, user transactions, check-ins, captured media, repeated location visits—all of these and many more will be sifted for patterns that fine tune predictive algorithms that anticipate, engage and perform actions for humans.

Based on these patterns, a key output of machine learning engines will be to place a human user into a contextual knowledge graph. This graph combines literal intelligence about the user's documented habits and emotional inferences about his typical states-of-mind (example, no alerts or interactions during sleeping hours). A knowledge graph is further enhanced through mapping relationships, classifications and genres derived from the four primary contextual inputs described above. This structured information is then made available to contextual applications for eventual presentation to the user.

The ability of machine learning to provide structure to unstructured information with minimal human involvement lies at the heart of its value proposition. For example, today's virtual assistants are highly responsive and accurate in providing information about movies playing at local theatres. The movie is playing at the theatre or it is not. However, analysing previous user interactions around movie content *plus* the preferences of her social network *and* her current location and time to make a recommendation about a new movie she might like, requires a level of processing only machine learning can provide.

**Figure 2: Machine learning provides structure to unstructured information with minimal human involvement**

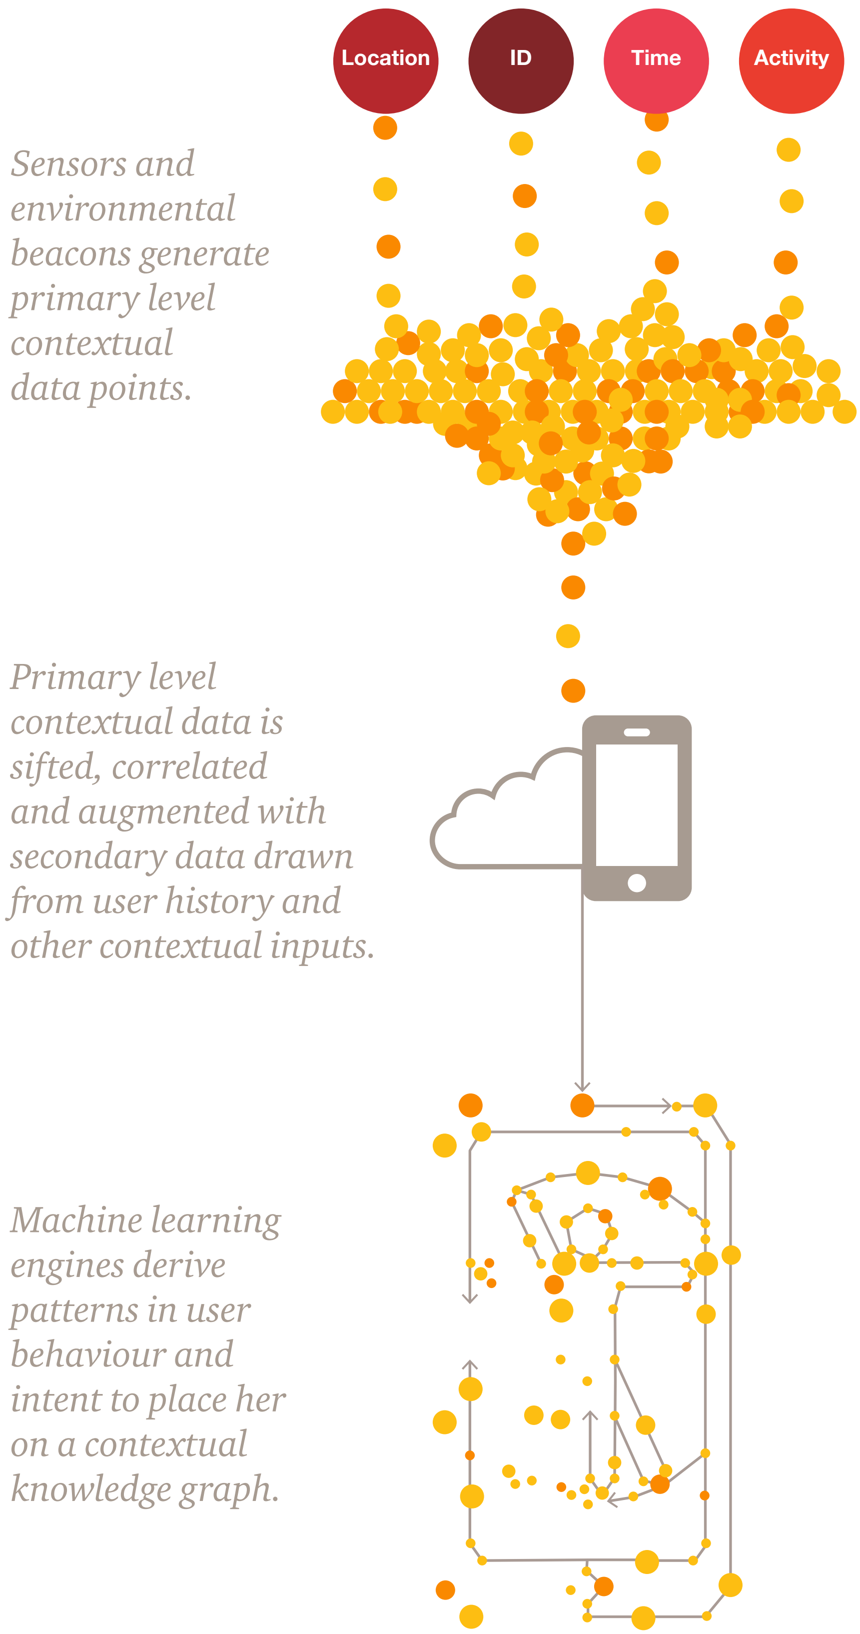

Machine learning in 2014 is focused mainly on improving NLP accuracy and performance. But the underlying principles of machine learning are applicable across the contextual value chain. Equally important, machine learning is the means through which contextual capabilities scale to the mass-market. Humans repeatedly prove that they will use any and all enhanced computing capabilities offered and then demand more, analogous to the way devices will use storage, application processing speed and other enabling technologies covered in **Phase** I of the Mobile Innovations Forecast. It follows that the more contextual intelligence powered by machine learning is added to an interaction, the more sophisticated user behaviour becomes, which creates more opportunities and demand to add value.

## *Conclusion and forecast*

Contextual intelligence enables a 'just-in-time' (JIT) mobile lifestyle that is becoming more prevalent. Cultural studies suggest that mobile users, especially younger demographics, are using their devices as much for organising their daily lives as for communication. Given that reality, it is clear that the days of the generic mobile user experience are numbered. Users will expect experiences specifically tailored for them that evolve in the face of new situations.

PwC believes the impact of contextual intelligence will be broad and deep—affecting every part of the mobile ecosystem. We anticipate three general trends that decision makers should consider as they develop strategies and capabilities for their organisations:

1) Moving beyond its roots in narrowly defined situations like stock trading, NLP transforms human computer interaction. Voice-controlled interfaces enabled by NLP will proliferate beyond smartphones to include a range of interactive situations (e.g., Amazon FireTV, connected cars, e-commerce, tele-medicine). Marketers and app developers will need to prepare for this fundamentally different model of customer relationship management.

2) Contextual SDKs like Gimbal will open up contextual app development to a much larger base of developers many of whom will create industry-specific contextual applications. Some early examples might be contextual apps for healthcare, education, travel, fitness and wellness. Technology companies with seemingly secure roles in these domains should be building out mobile contextual apps before new entrants challenge their market positions.

3) Machine-learning systems will encompass a range of contextual processing tasks in addition to NLP. Most likely, these systems will focus on domain-specific contextual knowledge to support industry-focused contextual SDKs. Over time, more individual domains will connect and integrate with one another to evolve general-purpose contextual assistants to more sophisticated and personalised levels. Companies with little or no machine-learning experiences should be exploring ways they can augment user experiences beyond the current standard of certainty (i.e., what is the purchase order number), and into areas of uncertainty (how reliable will this new supplier be).

The net result of these trends will be a veritable explosion of the contextual information market ecosystem, far beyond what is seen today. Activities, locations, transactions, preferences, emotional states—all of these will be logged, compared for patterns and archived to create the contextual experiences that drive JIT mobile markets.

Managing privacy will become one of the biggest and most important industry sectors for the contextual age. So too will be the industry players dedicated to managing and visualising the contextual information spaces of users.

**Daniel Eckert**
Managing Director,
Emerging Technologies
PwC

"There are significant opportunities for certain players with heavy infrastructure and analytic resources to catalyse entire ecosystems around contextual data," says Daniel Eckert, Managing Director of Emerging Technologies at PwC. "For example, data about weather is free from the government—but once the data is enriched by a contextual data provider for a particular audience—it can be utilised for many different services."

The opportunities are immense for contextual services as are the challenges, both technical and social. But the inherent value of enabling users to converse naturally with their devices and service to get what they want, when and how they want it is hard to deny.

## Let's talk

If you have any questions about the Mobile Innovations Forecast or would like to discuss any of these topics further, please reach out to us.

**Raman Chitkara**
Global Technology Industry Leader
PricewaterhouseCoopers LLP
raman.chitkara@us.pwc.com

**Daniel Eckert**
Managing Director, Emerging Technologies
PricewaterhouseCoopers LLP
daniel.eckert@us.pwc.com

## About PwC's Technology Institute

The Technology Institute is PwC's global research network that studies the business of technology and the technology of business with the purpose of creating thought leadership that offers both fact-based analysis and experience-based perspectives. Technology Institute insights and viewpoints originate from active collaboration between our professionals across the globe and their first-hand experiences working in and with the technology industry. For more information please contact Raman Chitkara, Global Technology Industry Leader.

## About PwC

PwC helps organisations and individuals create the value they're looking for. We're a network of firms in 157 countries with more than 184,000 people who are committed to delivering quality in assurance, tax and advisory services. Tell us what matters to you and find out more by visiting us at **http://www.pwc.com**