
From: Joscha Bach [REDACTED]
Sent: Friday, November 8, 2013 2:12 AM
To: Kevin Slavin
Cc: Joi Ito; Martin Nowak; Epstein Jeffrey; Ari Gesher; takashi ikegami
Subject: Re: The benefits of deception

Kevin, thank you for your excellent input and inspired questions on this =opic!

> The broad idea is that we'll need to build slack into systems of =biquitous computing. That there's a paradox in that in order to have a =ohesive identity, certain aspects of our lives require omission (or =.g., deception).

There are two sides to this: first of all, many of our interaction =artners may have difficulties to accept (or even comprehend) the =interactions and social norms that we maintain in other contexts (think: =iving in a small town in the midwest vs visiting Burning Man; religious =community vs. less usual sexual interests; family life vs. work life; =rganizational role vs. friendships). As a result, we compartmentalize =ur social expression. In a culturally inhomogeneous society, this kind =f deception (in the sense of hiding parts of my personality depending =n the context) is a prerequisite of meaningful freedom.

But on a deeper level: while we may maintain the illusion of a coherent =elf, we do not possess a single social persona. Our social identity =merges not only over our personality, but also over the individual =elationships in which it manifests. We become who we are through the =eflection of others.

> The ethnographer Tricia Wang coined "The Elastic Self" after spending = lot of time with Chinese and American youth using various forms of =ocial software (...)

The notion of the "Elastic Self" might capture both aspects. On one =and, we may explore different possible ways of self-actualization by =eeping social contexts isolated from each other, on the other, we =sually maintain more than one persona.

> How would you build a system that provides a patient with the ability =o continue acting deceptively -- in order to maintain a cohesive =identity -- while providing the doctor with the ability to find -- and =ct on -- perfect information?

The simplest setup: the patient is a BDI (belief, desire, intention) =gent, and has accurate knowledge about himself. He also has beliefs =hat part about his beliefs, desires and intentions are acceptable to =he outside world (which here is a homogenous hospital context and does =ot need further differentiation), and which ones need to be adapted to =it the actual intentions. For perfect information, the doctor gets =irect access to both layers.

The BDI model of agency (i.e. the idea of describing an agent using a =et of beliefs, desires and intentions) is a somewhat crude =mplification. We may well hold different, contradictory sets of =eliefs and goals at different levels, and the one that becomes relevant =epends on the context, and the state the agent is in. Also, many =eliefs and intentions are not directly represented, but generated when =eeded, and their content and form may be highly dependent on the =nteraction. In reality, there is often no perfect information for the =octor to act upon, e.g., because some of the information will be =enerated during and depending on the actions of the doctor.

We might want to skip that kind of accuracy if we just want to get an =dea of the benefits of deception for the organization. If we treat the =patients (as well as doctors, nurses, ...) as straightforward BDI =gents, the "perfect information" would be the set of beliefs and =esires that the patient functionally acts upon. If we want to include =elf-deception, we might simply add another layer on top of that: the =et of beliefs that the agent has about his own beliefs,

desires and intentions at any given time. For functionally accurate deception of others, the patient needs to represent beliefs about the actual beliefs of the doctor, and actionable desired beliefs of the doctor about the beliefs, desires and intentions of the patient himself. (A perfectlyclairvoyant doctor would have access to all belief layers of the patient, and the relationships between them.) Only if we want to assess the benefits of deception to the agent himself do we need to implement a more detailed cognitive model.

Cheers,

Joscha

```
<?xml version=.0" encoding=UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN" "http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version=.0">
<dict>
    <key>conversation-id</key>
    <integer>272682</integer>
    <key>date-last-viewed</key>
    <integer>0</integer>
    <key>date-received</key>
    <integer>1383876705</integer>
    <key>flags</key>
    <integer>8590195717</integer>
    <key>gmail-label-ids</key>
    <array>
        <integer>6</integer>
        <integer>2</integer>
    </array>
    <key>remote-id</key>
    <string>359754</string>
</dict>
</plist>
```